A Combined BEM/FEM Method for IC Substrate Modeling

Eelco Schrik

A Combined BEM/FEM Method for IC Substrate Modeling

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. J.T. Fokkema,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 11 september 2006 om 10:30 uur

door

Eelco SCHRIK

elektrotechnisch ingenieur
geboren te Voorburg.

Dit proefschrift is goedgekeurd door de promotor:
Prof. dr. ir. P.M. Dewilde

Toegevoegd promotor:
Dr. ir. N.P. van der Meijs

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof. dr. ir. P.M. Dewilde	Technische Universiteit Delft, promotor
Dr. ir. N.P. van der Meijs	Technische Universiteit Delft, toegevoegd promotor
Prof. dr. J.R. Long	Technische Universiteit Delft
Prof. dr. C.I.M. Beenakker	Technische Universiteit Delft
Prof. dr. W.H.A. Schilders	Technische Universiteit Eindhoven
Prof. dr. L.M. Silveira	Technical University of Lisbon, Portugal
Dr. W. Schoenmaker	Magwel NV, Leuven, België

Contents

1

Introduction

In present-day micro-electronic industry, product life-cycles are typically very short. With the advances in micro-electronic technology continuing fast and relentlessly, while the technology itself is becoming ever more expensive, the time to meet return on investment is becoming ever shorter. As a result, modeling and simulation as pre-fabrication verification tools of a circuit under design are becoming increasingly important. In this context, accurate and efficient modeling techniques have become an essential part in the design process for micro-electronic circuits.

The modeling techniques primarily aim at generating models that are representative of the physical situation, which typically includes parasitic effects ('parasitics') that may undermine the functionality and performance of the circuit. At the same time, the modeling techniques should be as fast and as accurate as possible in order to be valuable tools in the design process of micro-electronic circuits. However, as technology progresses, the impact and complexity of parasitic effects increases, placing higher demands on the speed and accuracy of the modeling techniques, but also on the compactness of the resulting models. To cope with these requirements, modeling techniques typically have to exploit the structure of the modeling problem to the limit.

This introductory chapter will explain that the structure of the modeling problem may be exploited by consistent problem partitioning, a concept that will be used in subsequent chapters.

1.1 Parasitics in Micro-Electronic Circuits

Figure 1.1 schematically represents two MOS transistors and a section of interconnect connected to their gates. For this illustration, we particularly consider digital technology, but the considerations presented in this section are similar for analog and mixed-signal technologies.

In the ideal case with respect to parasitics, the transistors would be ideal switches, the substrate would be an ideally insulating domain and the interconnect would be an ideal conductor in vacuum. In practice, however, physical parasitic effects influence the behaviour of the circuit. The transistors are physical devices with various non-idealities in their behaviour. The substrate, being a semiconductor, is mainly a resistive domain, possibly causing crosstalk between different segments of the chip. The interconnect is not an ideal conductor, which causes RC-delays along the interconnect. Furthermore, the interconnect is embedded in a dielectric with a higher permittivity than vacuum, which increases RC-delay. Additionally, field couplings through the dielectric may cause crosstalk among interconnect lines, and towards the substrate.

At low clock frequencies, these parasitic effects are typically negligible, because the switching delay in the transistors, as well as the RC-delay along the interconnect, are negligible compared to the speed of the signal. Furthermore, crosstalk between signals is negligible, because field couplings through the dielectric and the resistive substrate are weak at low clock frequencies.

However, through the years the clock frequencies in digital technologies have entered well into the GHz range, making delays and field couplings a dominant factor in IC-performance. As clock-frequencies increase, the signals carried by the interconnect become ever more aggressive, while the transistors switch at an ever faster rate. Under these circumstances, delays typically become critical in the performance of the circuit, and may cause synchronization problems. Furthermore, parasitic field couplings through the dielectric and the substrate increase, resulting in stronger crosstalk between signals, possibly causing signal-integrity problems. In general, the parasitics undermine the performance of the circuit, possibly to the extent that the entire circuit might not meet the requirements or, in the worst case, might not function at all. Therefore, models that are aimed at modeling parasitics should grow along with the requirements prescribed by the physics.

1.2 Parasitics Modeling

The objective of any modeling approach would be to generate a model that is as simple as possible, while capturing all relevant behaviour of the physical system. In our case, the physical system would be the whole of the micro-electronic circuit (i.e. its basic functionality and all parasitics). The modeling objective would then be an electrical circuit model which can be simulated with a regular circuit simulator like SPICE.

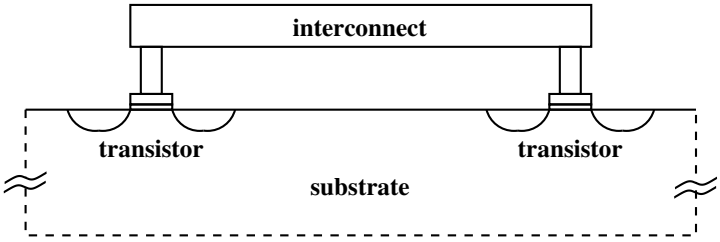

Figure 1.1: Schematic representation of interconnect, (MOS) transistors and substrate.

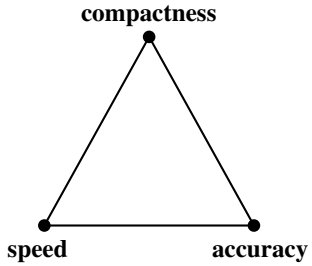

Figure 1.2: The triangle of trade-offs in modeling

As micro-electronic technology progresses, and the features are decreasing in size, while the operating frequencies are going up, the accuracy requirements on the models become ever more strict. At the same time, the overall complexity of the circuits continually increases, making it harder to obtain the models in a timely manner. Furthermore, as frequency increases, the phenomena relevant to the model are becoming ever more complex as well, making it harder to keep the models compact. This inherent triangle of trade-offs (see Figure 1.2) between speed, accuracy and compactness continually plays a role in modeling.

The speed-accuracy trade-off is primarily defined by refinements in the modeling approach. This is straightforwardly illustrated by mesh refinement in a mesh-based modeling technique. However, the choice for a particular modeling technique is also relevant. In this context, we find the contrast between 3D methods and 2D methods. The 3D methods model the problem in all three dimensions. They are generally applicable and are typically very accurate, but are computationally expensive and therefore usually slow. The 2D methods, on the other hand, model the problem in only two dimensions. They are not generally applicable, because they may only be applied to modeling problems where variations in the third dimension may be neglected. When applied to such modeling problems, these methods can be very accurate, while being computationally less expensive and therefore faster than the 3D methods. As a crossover between the 2D and 3D methods, there are also methods that are based on

a 2D method, but that partially take the third dimension into account. Such methods are commonly known as 2.5D methods; their behaviour with respect to accuracy and speed may be anywhere between the behaviour of the 2D and 3D methods.

In the compactness-accuracy trade-off, we typically find various model reduction techniques. These can roughly be classified into 3 different types: a-priori, on-the-fly and a-posteriori methods [Mei01]. A-priori methods aim at symplifying the model before it is generated, e.g. by neglecting effects that are known to have only little impact on the overall behaviour of the model. On-the-fly methods aim at symplifying the model while it is generated, e.g. by detecting and removing unnecessary detail. A-posteriori methods aim at symplifying the model after it has been generated, e.g. by detecting the most important factors that play a role in the overall behaviour of the model, and removing the less important factors.

In the speed-compactness trade-off, the model reduction techniques typically also play an important role. If the model requires compaction by an on-the-fly or a-posteriori model reduction procedure, the speed of the extraction typically decreases because of the additional time required for reducing the model. By efficient implementation of these procedures, the increase in extraction time can be kept as small as possible.

Considering the trade-offs from Figure 1.2, the 'ultimate' modeling approach seems utopia, but applying the appropriate modeling method in the right place, while fully exploiting all types of model reduction, should result in a useful overall trade-off between speed, accuracy and compactness of the model.

Fortunately, since modern IC's might consist of several (tens of) millions of transistors, actual design flows commonly apply divide-and-conquer strategies to (hierarchically) divide large problems into subproblems which allow to be solved more conveniently. Under these circumstances, parasitics modeling can efficiently be applied to smaller subcells of the original circuit. Even though the modeling and simulation of these smaller subcells cannot provide full insight into the behaviour of the overall chip, it may still identify the individual behaviour of relevant (possibly critical) subcircuits, which, in turn, may provide valuable insight into their contributions to the behaviour of the overall chip.

1.3 Problem Partitioning vs. Model Consistency

Similar to the hierarchical subdivision of the IC-design into subcells, it can also be very efficient to divide the modeling problem itself into subproblems. Particularly, if the overall modeling problem is very complex and a direct solution would have been prohibitively time-consuming, the subproblems may be solved more conveniently, resulting in partial solutions that together consistently approximate the solution of the overall modeling problem. Consistent partitioning of the modeling problem into sub-

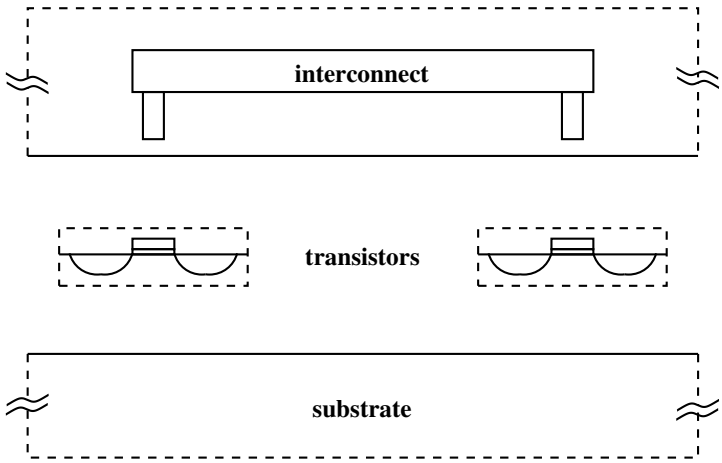

Figure 1.3: Partitioning of the modeling problem.

problems can be achieved by exploiting the structure of the underlying physics. In fact, as will be addressed in this thesis, *consistency* with the physics is even more important in a model than the highest accuracy.

The modeling problem for micro-electronic circuits can now roughly be subdivided into 3 categories, as schematically shown in Figure 1.3: interconnect, transistors, substrate.

Interconnect is usually considered an RCL modeling problem, where the R, C and L each present their own modeling challenges. Interconnect, however, is outside the scope of this thesis and will not be explicitly considered.

Building accurate transistor models – 'Compact Modeling', see e.g. [Gra90] – is a field of research in itself, and is therefore generally not explicitly considered in the field of parasitics modeling; off-the-shelf transistor models (like the SPICE BSIM3 models) are inserted, whenever a transistor is required in the model.

Under quasistatic assumptions, the substrate forms an RC modeling problem: the substrate is a semiconducting domain which has both resistive and capacitive characteristics. In other words, the domain parameters are determined by the resistivity and permittivity of the semiconductor. Depending on the frequency spectrum of the applied signals in the circuit, the permittivity of the semiconductor may, or may not, play a relevant role in the accuracy of the resulting model. In this thesis, for reasons of simplicity, the permittivity of the semiconductor will be assumed not to play a relevant role in the resulting model. In other words, the substrate will only be considered by its resistive characteristics. Note, however, that the concepts presented in this thesis may also be applicable to RC-modeling in the substrate.

With respect to the consistency between the model and the physics, as mentioned previously, omitting the substrate permittivity already introduces inconsistencies into the model. However, including this more advanced physical effect into the modeling technique proved to be a greater challenge than allowed by the scheduled time frame for this thesis. Therefore, the work presented in this thesis, which deals with consistent partitioning and efficient solution of the substrate modeling problem, should be considered as a *step towards* a fully consistent model.

The division into subproblems typically involves making assumptions. Certain assumptions may not fully be representative of the physical situation, but may still be allowed if the partitioned problem is a consistent approximation of the physical situation. As an example of this, consider the interconnect/substrate problem. From the physics, this is a fully coupled problem, where any field effects exist in both the capacitive and the resistive domains and cross the boundary between the domains. The partitioned approach considers the boundary between the domains to be a ground plane for the capacitance problem, and an isolating plane (except the contact areas) for the substrate problem. This allows the capacitance and substrate problems to be solved as efficiently as possible, while the combined model is an approximation of the physical situation. The validity of the approximation depends on the physical situation and should be carefully considered in cases which may reach, or cross, the validity-borders.

Furthermore, it is important to note that the subproblems should not be solved entirely independently; the coherence between the independent models is then not sufficient for overall consistency in the joint model. In other words, the subproblems should pass context information to each other, for sufficient coherence in the joint model. The importance of this observation is easily illustrated with interconnect RC modeling, where distributed effects along the length of the interconnect are captured in a lumped RC model, rather than a single overall resistor and a single overall capacitor.

The abovementioned concepts around modeling of micro-electronic circuits have been implemented in software: the SPACE layout-to-circuit extractor [Bee98, MSp]. The software can read and process designs, and eventually puts out an electrical circuit model, suitable for simulation.

1.4 Parasitics Modeling in the Substrate

In this thesis, we specifically focus on parasitics modeling in the substrate. The problem of substrate noise coupling and crosstalk has extensively been studied in literature, and has been compiled in several textbooks [Ver95b, Ara99, Don03].

Substrate noise is known to be a highly complex phenomenon in which many factors play a role. These factors include on-chip effects like noise injection through

aggressively switching digital transistors which are embedded in the substrate, or through capacitive coupling from interconnect carrying aggressive, high-frequency signals. Additionally, the substrate doping profile and the floorplan of the chip determine how the noise propagates through the substrate and where it is picked up. Finally, off-chip effects like inductance from the bond wires or the package may also play a role: the low-pass characteristics of the inductance may cause sharp switching spikes in the supply/ground lines to inject into the substrate through bias connections to the substrate or the wells. Even though this brief summary of substrate-related effects is not complete, it does indicate that any modeling approach aimed at predicting and/or analyzing substrate noise phenomena should be coherent and comprehensive.

Fortunately, some hierarchy is possible. At the device level, the noise behaviour/sensitivity of transistors is studied in the field of device physics, and is captured by accurate behavioral transistor models like the SPICE BSIM3 models. At the intermediate level, i.e. chip level, the noise behaviour of the chip, including injection from the interconnect, propagation through the substrate doping profile, and floorplan dependency, is studied in the field of physical verification, resulting in simulation, modeling and extraction tools like layout-to-circuit extractors. At the package level, the (inductive) characteristics of the bond wires and the chip package can be determined through package modeling and simulation tools.

When a common format (e.g. a SPICE netlist) is used for representing the models at the three different levels, then the separate models can be assembled into one comprehensive model that, in principle, allows for a representative simulation of the entire chip. However, as already mentioned previously, modeling and simulation of a full chip typically leads to complexity problems, and therefore the modeling and simulation is usually only applied to relevant subcircuits of the chip.

In this thesis, we will describe an improved technique for modeling substrate coupling effects at the chip level. However, before introducing the improvements, we will first properly place the substrate in the context of the partitioned problem.

1.4.1 Substrate Contact Definition

As indicated in Figure 1.3, the overall modeling problem can be partitioned into 3 subproblems, which, after their solution, can be joined to obtain the overall model. The 3 subproblems pass context information to each other, for sufficient coherence in the joint model.

With the substrate being treated as a separate partition of the overall problem, the coherence with the other partitions is achieved by strategically placing contact areas on the domain that represents the substrate [Gen01, Sch03]. In this case, 'strategically' means that the contact areas should be placed there where the circuit has close interaction with the substrate. It is well known [Ver95b, Ara99, Don03] that close interaction with the substrate takes place through the transistors themselves which are embed-

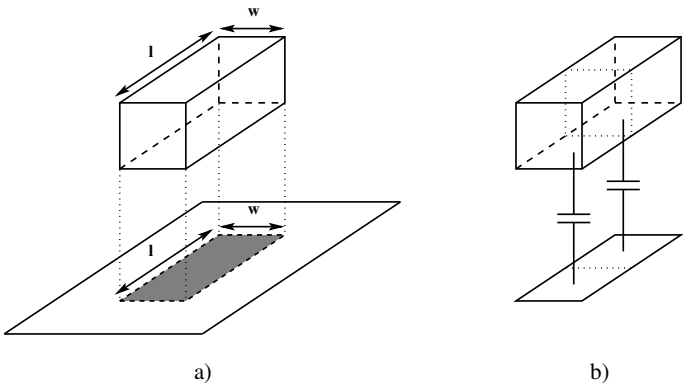

Figure 1.4: a) Shadow contact underneath interconnect wire, b) Lumped modeling approach for long wires.

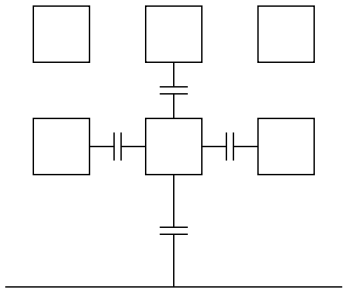

Figure 1.5: Cross-section of a wire with coupling to the substrate in a practical situation.

ded in the substrate, through noisy supply/ground lines (i.e. connected to the outside world through a parasitic package inductance) that are connected to the substrate and the wells, and through field couplings from the interconnect.

For transistors, the contact area to the substrate is well defined; it is situated immediately underneath the transistor itself. For a digital transistor, that is underneath the gate, source and drain areas; for an analog npn-transistor (as in e.g. Figure 1.7 on Page 10) that is typically underneath the collector. Similarly, biasing of the wells and the substrate takes place through via-connections between the supply lines and the substrate; via-connections also have a clearly defined contact area to the substrate.

Interconnect, however, influences the substrate more indirectly through capacitive field coupling, which requires the contact area to the substrate to be defined in a different way. Typically, the contact area from the interconnect is defined as the ”shadow”

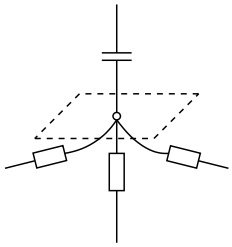

Figure 1.6: Single node model for a substrate contact area.

of the wire (see Figure 1.4a). For a single wire, the potential field would radiate in all directions and the interaction with the substrate would not be confined to just the shadow of the wire. However, in practical situations, the shadow approach proves to be adequate because the largest portion of the capacitance will be concentrated directly below the wire since (1) the distance to the substrate is shortest at that position and (2) in practical situations the surrounding of the noise generating wire will be shielded by other wires, see Figure 1.5. As such, it is assumed that the interaction between interconnect and substrate mainly takes place directly underneath the wire, which supports the shadow approach.

When a wire is relatively long, there will also be distributed effects along its length. Therefore, the substrate contact definition may require lumped modeling along the length of the wire (see Figure 1.4b). Thus, smaller portions of the interconnect capacitances are then connected to smaller substrate contact areas.

For each substrate contact area, a node is created in the final netlist, as shown in Figure 1.6. The substrate capacitance that is computed for the part of the interconnect that is above the area, is attached to this node. The substrate resistances that are computed for the substrate underneath the area, are also attached to this node. In this way, a representative and comprehensive electrical network is created from the partitioned problem. This electrical network (possibly after model reduction) can then be simulated using a simulator like SPICE.

1.4.2 Problem Partitioning – Traditional Approach

Also in the substrate modeling problem, problem partitioning is possible. Depending on the type of modeling, and the requirements on the resulting model, different partitioning approaches are possible. The traditional partitioning approach discussed with more detail later on in this thesis was aimed at maximum computational efficiency, at the loss of some accuracy. In particular, the approach would assume the substrate to consist of either one or two uniformly doped layers, such that a fast (BEM-based)

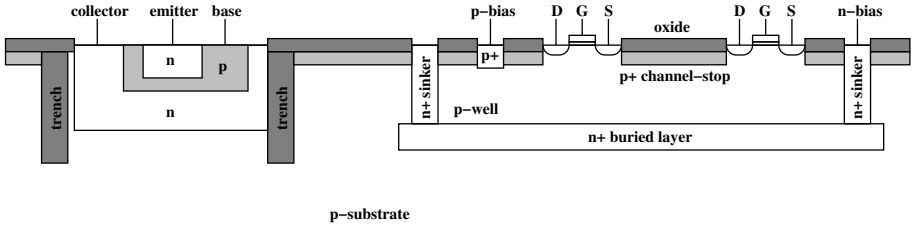

Figure 1.7: Advanced substrate structures

modeling technique could be applied. To a limited extent, this approach could take into account some influence from doping patterns like n-wells, but more advanced doping patterns, as are common in modern circuits, could not be captured consistently. Despite the accuracy limitations in this traditional approach, the method was straightforward, computationally efficient, and suitable for traditional modeling problems [Sme95a].

1.4.3 Improved Problem Partitioning – New Approach

Even though computational efficiency is still an important consideration, the advances in computing power and the reduced price of memory have allowed the modeling and partitioning approaches to aim more for accuracy rather than just computational efficiency. The traditional approach, though computationally very efficient, cannot consistently take into account specific doping patterns in the top layers of the substrate. However, modern (mixed-signal) technologies may typically include advanced substrate structures like channel-stoppers, buried layers, sinkers, trenches, and guard rings, as illustrated in Figure 1.7. Most of these structures are aimed at coping with substrate crosstalk at high frequencies, under ever larger scale integration. Therefore, it is important to include these substrate structures into the model, but this cannot be done through the traditional approach.

As will be detailed in this thesis, the advanced substrate structures can be included in the model by a more advanced approach to the partitioning of the problem. In particular, the new approach will take the doping patterns in the top few microns of the substrate fully into account through an accurate FEM-based modeling technique, and add a fast, BEM-based model for the deep substrate.

While this new approach is computationally more demanding than the traditional BEM approach, the resulting models will be more accurate and reliable. Furthermore,

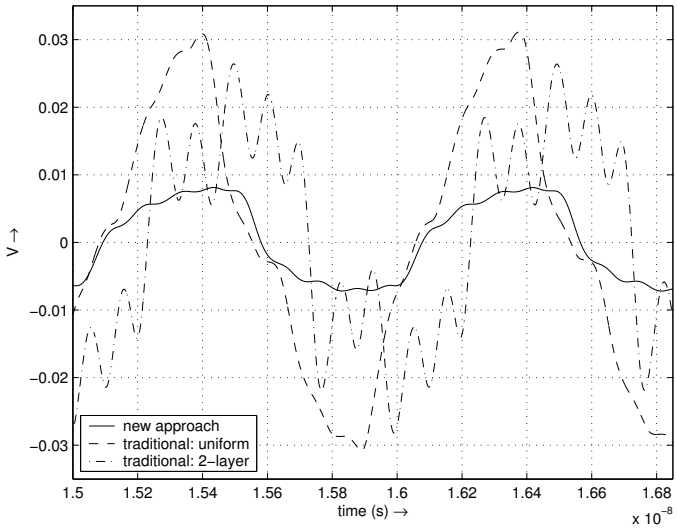

Figure 1.8: New approach vs. traditional approach to substrate modeling. Depending on the modeling approach, simulations with the resulting substrate model showed significantly different voltage waveforms on a substrate sensor node just outside a ring oscillator (more details follow in Section 6.3.2 and Figure 6.25).

we will show through a prototype implementation that the new approach is computationally more efficient than a FEM approach. Further optimizations to the method are reserved for future research.

1.4.4 Traditional vs. New Approach

Figure 1.8 compares the new approach to the traditional approach through simulations on a representative layout. The details of these simulations will be addressed in Section 6.3.2 and Figure 6.25. For the time being, it is sufficient to observe that the simulations are significantly different.

Unfortunately, we did not have access to actual measurements on a fabricated layout and therefore we did not have the opportunity to compare the results from the improved modeling approach to measurements from the physical situation. However, knowing that the physical situation can be modeled accurately with a 3D Finite Element Method, we will compare results from the improved modeling approach to those obtained with an independent 3DFEM. Together with the observation that the improved modeling approach stays close to the physics, this will be indirect experimental proof of the correct behaviour of the method.

1.5 Focal points & Structural Overview of this Thesis

From the topics discussed above, this thesis will focus on consistent modeling of substrates with specific patterns like in Figure 1.7, while aiming for low computational cost, low memory usage and the output of compact circuit models, suitable for simulation with SPICE, or an equivalent network simulator. The thesis will progress through the subject as indicated in Figure 1.9.

Chapter 2 will show that the physical situation in the substrate has progressed towards advanced substrate structures that are aimed at reducing or controlling substrate crosstalk. These advanced structures then impose new requirements on modeling techniques, particularly with respect to speed and accuracy. However, none of the available modeling techniques can comply well with the new requirements. In particular, the FEM is accurate but too slow, whereas the BEM is fast, but too limited with respect to the required accuracy. Nevertheless, based on the structure of the substrate modeling problem, a combined BEM/FEM method is identified as a possible new modeling method for the substrate.

Through an overview of the relevant properties of the BEM and the FEM in Chapter 3, Chapter 4 will introduce the combined BEM/FEM method. Based on the physics and the characteristics of the BEM and FEM, the substrate modeling problem is partitioned consistently into a BEM and a FEM part for the combined BEM/FEM approach. The combined BEM/FEM method proves to converge, and allows a sparse and reduced-order approach which, with respect to the new modeling requirements, has accuracy and speed benefits over the traditional BEM and FEM methods.

The implementation of the method allows practical evaluation of the method in Chapters 5 and 6, which will confirm the validity of the BEM/FEM method as a useful substrate modeling technique, after which Chapter 7 presents the conclusions.

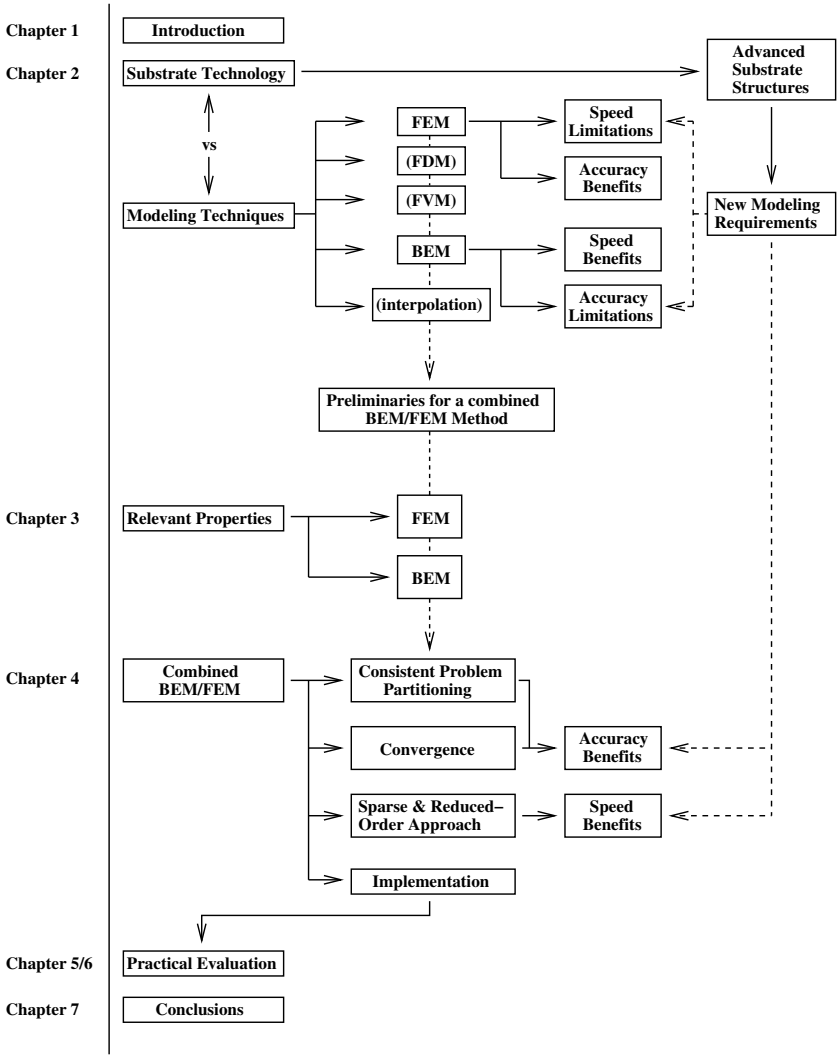

Figure 1.9: Structural overview of the thesis

2

Substrate Technology vs. Available Modeling Techniques

As described in the introduction, the overall parastics modeling problem may be partitioned into subproblems which allow more convenient solution than the overall modeling problem itself. In this chapter, we will focus on one of those subproblems: substrate modeling. The modeling problem and possible methods for its solution will be addressed through a summary of substrate technology and available modeling techniques (FEM, FDM, FVM, BEM and interpolation methods).

Each of the modeling techniques has its own advantages and limitations. However, as this chapter will show, the progress in micro-electronic technology imposes new requirements on the models and on the modeling techniques. Under these new requirements, the limitations in the modeling techniques are typically emphasized, which formulates the need for a new modeling strategy. Based on the structure of the substrate modeling problem, a combination of the BEM and FEM modeling techniques seems likely to result in the desired new modeling strategy. The next chapter will therefore describe the relevant properties of the BEM and FEM methods with respect to a combined BEM/FEM modeling strategy for the substrate.

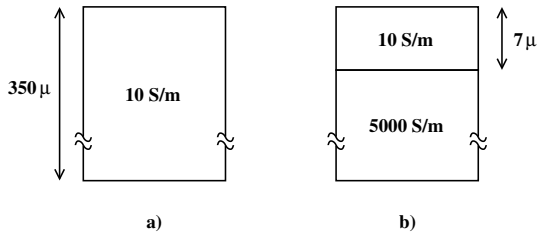

Figure 2.1: Schematic representation of common global doping profiles in the substrate: a) lightly doped substrate, b) heavily doped substrate. Numeric values are indications.

2.1 Global Doping Profile in the Substrate

The material characteristics of the substrate determine how noise/crosstalk propagates. In this context, the global characteristics of the substrate are determined by the global doping profile. Typical global doping profiles are p-type and may be classified into two main classes: lightly-doped and heavily-doped. The lightly doped substrate has a constant low doping concentration extending the full depth of the substrate (typically some $350\mu m$), whereas the heavily doped substrate only has a low doping concentration up to the first 5 – 10 μm, and then shifts to a much higher doping concentration.

The equations describing the physical behaviour of semiconductors are complicated, already in their most basic form. Furthermore, they are temperature-dependent. However, the full semiconductor equations are mostly relevant for *device* behaviour, whereas the global characteristics of the substrate can be captured with sufficient accuracy by taking into account only the dominant factors in semiconductor behaviour. In this context, the modeling approaches described in this thesis are aimed at modeling the global behaviour of the substrate and therefore do *not* take into account the full semiconductor equations. Instead, we aim for a modeling approach that generates a straightforward electrical network model in terms of general network elements.

Therefore, for a region with a given doping concentration, we will be using the equivalent conductivity in that area typically at a temperature of 300K [Nea92]. As mentioned in Section 1.3, it would also be possible to incorporate the capacitive characteristics of the semiconductor into the model, but this is typically only relevant for circuits operating at very high frequencies (i.e. high clock frequencies for digital, RF frequencies for analog). This type of circuits and frequencies is currently outside the scope of this thesis, and the methods proposed in this thesis are not evaluated in that context. However, including the capacitive substrate characteristics into the model may be a future extension to the proposed modeling methods.

Along these lines, Figures 2.1a and 2.1b show a schematic representation of the global doping profiles, together with numeric indications of the thicknesses and equiv-

alent conductivities. Thus, with respect to its global doping profile, the substrate may be considered to be a domain consisting of multiple layers with each their own conductivity.

2.2 Layout-Dependent Doping Patterns

Apart from the global doping profile, the substrate also contains layout-dependent (localized) doping patterns which determine crosstalk propagation and/or isolation through the substrate. In fact, these doping patterns are often used to reduce or control crosstalk.

Below, a selection of possible doping patterns will be discussed. Doping patterns are typically present in the top $4\mu m$ – $6\mu m$ of the substrate, and may be very different depending on the type of circuit. We will first highlight the n-well and the channel-stop layer, as these are very common and already representative of a wider range of doping patterns. After that, we will briefly summarize several other doping patterns. Even though the transistor itself may also be considered a 'doping pattern', it will not explicitly be discussed because it requires the full semiconductor equations to be taken into account, which is outside the scope of the modeling techniques described in this thesis.

2.2.1 N-Well

Figure 2.2 shows an example of an n-well in CMOS technology, and Figure 2.3 shows the first few microns of a typical n-well doping profile. Up to $5\mu m$ of depth, the domain is n-type. From $5\mu m$ to $6.5\mu m$ the doping shifts from n-type to p-type, which forms a junction. From 6.5 μm onwards to the full depth of the substrate ($350\mu m$ in this example) the doping is p-type and continues into the global doping profile.

Along the lines from Section 2.1, the p-type and n-type regions can be modeled as resistive domains. The junction areas could be modeled with a diode-model, but, assuming the biasing conditions in the substrate to be such that the diode is always reverse biased, the junction areas may also be modeled as a capacitance. This capacitance is non-linear, but assuming a digital application and a typical voltage swing, a linearized capacitance model may also be sufficiently accurate. Figure 2.4 shows the modeling approach for the n-well schematically.

2.2.2 Channel-Stop Layer

Figure 2.5 is the same as Figure 2.2, but supplemented with the channel-stop layer which is located underneath the oxide (LOCOS, or nowadays STI). Figure 2.6 shows the first few microns of the doping profile, where we observe that the channel-stop

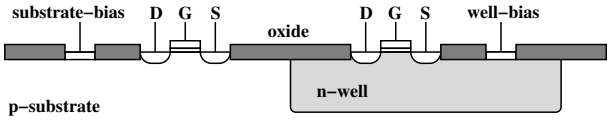

Figure 2.2: Schematic representation of the n-well as a local doping pattern.

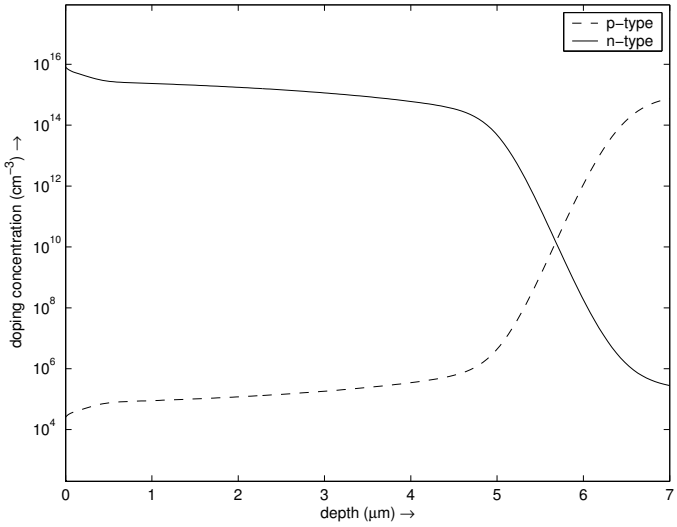

Figure 2.3: Doping profile for the n-well, embedded in a lightly doped substrate.

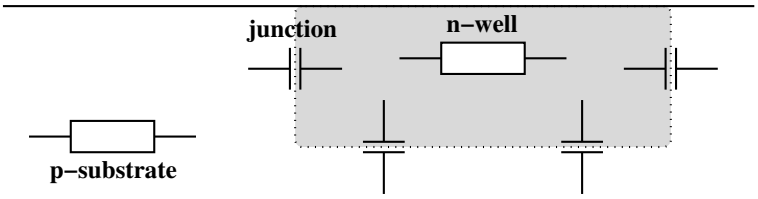

Figure 2.4: Modeling approach for the n-well.

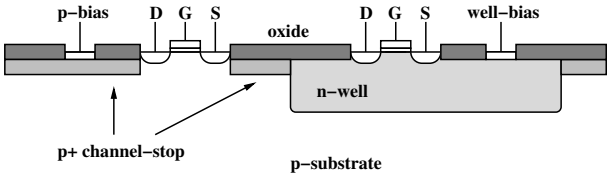

Figure 2.5: Schematic representation of common doping patterns in the substrate. The channel-stop layer is interrupted by transistors and wells.

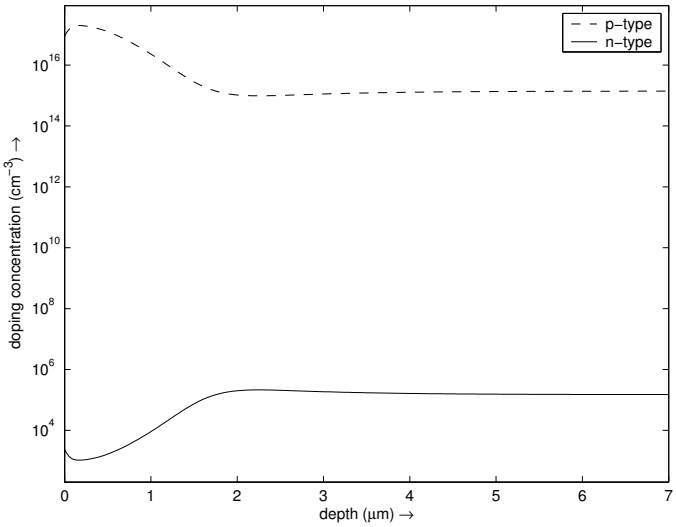

Figure 2.6: Doping profile for the channel-stop layer, in a lightly doped substrate.

layer is actually a thin layer with a high doping concentration. The high doping concentration in the channel-stop layer raises the treshold voltage of the parasitic FET formed by the lowest levels of interconnect and the underlying substrate; in effect this avoids that the parasitic FET is turned on.

However, since the threshold voltage of the actual transistors should not be modified, the channel-stop layer is not present underneath the transistors themselves. Furthermore, a p+ channel-stop layer does not extend into an n-well. This means that the channel-stop layer is a layout-dependent doping pattern. Since the channel-stop layer (typically 1000 S/m) is significantly more conductive than its surroundings (typically 10 S/m), it may either increase or decrease substrate crosstalk between different parts of the chip. This indicates that taking the channel-stop layer into account in a substrate model is important for accuracy.

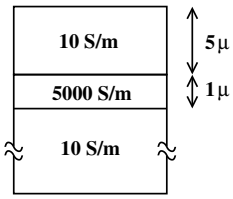

Figure 2.7: Schematic representation of the buried layer substrate. Numeric values are indications.

2.2.3 Other Doping Patterns

Below follows a selection of other possible doping patterns. Even though the selection is not complete, it does give an indication of the variety of doping patterns that may be present in the substrate.

Buried layers

The buried layer can be considered a cross-over between the two main *global* doping profiles. In particular, the lightly doped substrate (recall Figure 2.1a) was traditionally applied in analog circuits, because of its low-loss behaviour, whereas the heavily doped substrate (recall Figure 2.1b) was applied in digital circuits, because of its better behaviour with respect to avoiding latch-up. However, with the introduction of mixed-signal circuits, a substrate that combined low-loss behaviour and latch-up avoidance was required. As a result, the buried-layer substrate was introduced, which is mainly lightly doped for low losses, but has a heavily doped buried layer some $5\mu m$ below the surface for avoiding latch-up. Figure 2.7 shows this schematically.

If the buried layer would be a uniform layer through the whole substrate, it could be considered to be part of the global doping profile. However, dependent on the type of circuit, buried layer substrates may contain a buried layer that is not present everywhere, or contains n-type sections (see for example Figure 2.8). Under these circumstances, the buried layer becomes a layout-dependent doping pattern.

Since the buried layer provides a conducting layer in the substrate which may be present over large areas (somewhat similar to the channel-stop layer), it may have a significant impact on the substrate noise propagation in the substrate. Furthermore, if the buried layer is biased through a noisy supply, it may even facilitate the injection of noise into the substrate. In other words, it is important to take the buried layer into account in the substrate model.

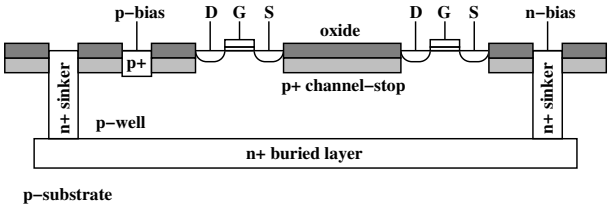

Figure 2.8: A triple-well structure

Triple-well structures

In some circuits, triple-well structures are applied; Figure 2.8 shows this type of structure. In this case, NMOS transistors are embedded in a p-well, which is entirely surrounded by an n-structure. As such, the NMOS transistors are isolated from the substrate by a double junction. This structure resembles a pnp transistor which can be 'turned off' under appropriate biasing conditions. The structure has several realization alternatives. In particular, the n+ sinkers may be n-wells themselves, with PMOS transistors embedded.

This structure is typically applied in mixed-signal technologies. The analog circuitry is embedded directly in the substrate (typically surrounded by a guard ring or a trench), while the digital circuitry is isolated by the triple-well structure.

Guard rings

Guard rings are typically placed around sensitive circuitry. Guard rings can have either the same doping type (i.e. p-type or n-type) as the substrate, or the opposite doping type and are typically biased appropriately. If the doping type is the same, the guard ring is intended as a 'noise-sink' (guard ring is typically shallow and biased through a low impedance path), if the doping type is opposite, the guard ring will force any substrate noise to approach the sensitive circuitry through the deep substrate (guard ring is typically deep), such that the impedance in the noise path increases. Note that this last approach is only effective at relatively low frequencies while the first approach requires low impedance (including low inductance if the frequency is high) in the bias connections.

Furthermore, it is important to note that if guard rings are biased through a noisy supply line, they may also *facilitate* the injection of noise into the substrate, having an opposite effect to what they were intended for. Guard rings should therefore be biased carefully.

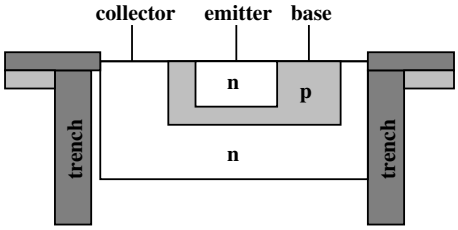

Figure 2.9: Trench around an analog transistor.

Trenches

In some cases it is more convenient to replace a deep guard ring with an isolation structure that does not need biasing. In such cases, a (deep) trench is usually applied. The trench (see Figure 2.9) consists of insulating material (typically oxide), and is placed around the sensitive circuitry, to force the substrate noise through the deep substrate, which increases the impedance in the noise path. Trenches may also be used specifically to break the buried layer. Similar to the guard ring, this structure is only effective at relatively low frequencies.

2.3 Substrate Modeling Techniques

Up to this point, we have indicated that the behaviour of the substrate with respect to substrate crosstalk is determined by both the global doping profile and local doping patterns. Therefore, any substrate modeling technique should consistently include the global resistive characteristics of the substrate, as well as the localized doping patterns.

Traditionally, there are 5 approaches available for substrate modeling:

- The Finite Element Method (FEM)
- The Finite Difference Method (FDM)
- The Finite Volume Method (FVM)
- The Boundary Element Method (BEM)
- Interpolation methods

Below, each of these methods will briefly be summarized, including references to literature and to commercial modeling tools.

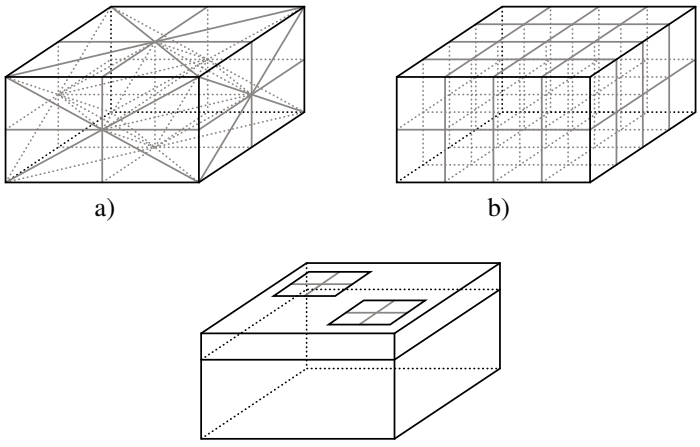

Figure 2.10: a) Example of a 3D tetrahedral volume discretization as would be used by the FEM; b) Example of a 3D cubic volume discretization as would be used by the FDM; c) Example of a discretization of contact areas on top of uniformly layered domain as would be used by the BEM.

2.3.1 FEM-based Substrate Model

The FEM applies a volume discretization to the entire domain. As an example, we consider a discretization using tetrahedral elements, as illustrated in Figure 2.10a. Along the edges in the discretization, mathematical relations between the nodes at the corners of the elements are defined. These mathematical relations are typically linear, but may also be of higher order. The mathematical relations between the nodes in the discretization may be accumulated into a large, but sparse, system of equations which may subsequently be solved by some (iterative) solution procedure. As will be explained with more detail in Chapter 3, the FEM solution procedure utilizes the property that the *exact* solution of the governing partial differential equation has minimum energy. In particular, the FEM solution procedure finds the best approximation of the exact solution by explicitly minimizing the energy in the approximate solution.

Each element in the discretization may have its own material properties assigned. Thus, the FEM is very flexible and would easily be able to take into account the doping patterns from Section 2.2. Furthermore, the FEM would be suitable to take into account the full semiconductor behaviour [Mil99], wich requires more advanced mathematical formulations than the Laplace equation. However, due to speed and memory limitations, the FEM can only handle relatively small structures. Furthermore, the FEM solution may not conserve current [Pol87]. For these reasons, the Finite Element Method is rarely used in (larger-scale) substrate and device modeling engines.

2.3.2 FDM-based Substrate Model

The FDM (see e.g. [Mit94]) discretizes the entire domain into cubic elements, as indicated in Figure 2.10b, where each element may have its own material properties assigned. With respect to the volume discretization, the FDM resembles the FEM, but with respect to the mathematical background the methods are different. In particular, the FEM uses the domain discretization to find an approximate solution of minimum energy under the governing partial differential equation, whereas the FDM domain discretization is actually a discretization of the partial differential equation itself and the solution procedure does not explicitly apply minimization of energy. The main restriction of the FDM over the FEM is that the domain has to be such that it can consistently be discretized with cubic elements. If this is the case, the meshing algorithms are more straightforward, the average degree of connectivity of the nodes is lower (improved sparsity), and the system of equations can be solved more efficiently. The FDM is widely applied; some field-solvers, e.g. Raphael [MRa], use the FDM instead of the FEM.

Similar to the FEM, the FDM is well capable of modeling the doping patterns from Section 2.2. It is computationally less demanding than the FEM and can therefore handle larger structures. In fact, the FDM has traditionally been a frequently used method for substrate modeling, and has been commercially available through the substrate modeling tool SubstrateStorm [MSu][1]. It was probably first applied in [Joh84] to obtain a resistance model for the substrate. Later, it was used in other books and papers [Ver93, Cle94, Ver95b, Pfo96] where the modeling was expanded with capabilities for taking into account the capacitive characteristics of the substrate by expanding the traditional resistive elements in the FDM with a capacitance in parallel. However, it was also noted that when different sections of the domain have different RC time constants, the solution of the system of equations results in higher-order elements as port-impedances. This then requires an Asymptotic Waveform Evaluation technique [Chi94] for a macromodel (used in [Ver93]), or a synthesis step based on Padé approximation and/or moment matching [Yen64, Chi94]. Fortunately, it is also noted that the dominant time-constants typically arise from the substrate resistance and *junction* capacitances, such that the majority of substrate problems can be handled by a purely resistive substrate model, supplemented with junction capacitances where necessary.

Traditionally, the FDM system of equations was solved using Gaussian elimination, but already [Pol87] and earlier work by the same authors proposed iterative methods. In the context of substrate modeling, [Sta94] addressed the application of the ICCG (Incomplete Choleski Conjugate Gradient) iterative method, and using appropriate preconditioning, while other papers report using multigrid (multilevel) methods [Sil99, Sil04] for efficient solution of the FDM system.

As shown in [Wem95], an FDM-like method is also possible with non-cubic elements. Similar to the FDM, the method uses stacked layers for discretizing the thick-

[1] SubstrateStorm has currently been withdrawn from the market.

ness of the substrate, but the horizontal discretization uses a Voronoi tesselation instead of rectangular elements. The result is a set of Voronoi prisms (Voronoi polygons supplemented with a thickness). Each prism is assigned a node in its center of gravity, and the Delaunay triangulation is used to define the horizontal connections between the nodes; the vertical connections are straigntforward. The main advantage of this method is that the mesh allows more general shapes, which may reduce the overall node-count. It can be argued that this method is actually an FVM method as described in the next section. However, this discussion will not be pursued here.

2.3.3 FVM-based Substrate Model

Similar to the FEM and FDM, the FVM also applies a volume-discretization to the domain. Each of the subvolumes in the discretization then has a node defined in its center of gravity, and for each of the subvolumes conservation of current is assumed and continuity conditions between adjacent subvolumes are assumed to hold. In this way, mathematical relations between adjacent nodes in the discretization are defined and a large, sparse system of equations is obtained. In obtaining the FVM solution, minimization of energy is not explicitly applied; in this respect, the FVM does not resemble the FEM. In fact, as explained below, the FVM resembles the FDM and has similar behavioural properties but is actually more general.

Contrary to the tetrahedral and cubic subvolumes in the FEM and the FDM discretizations described above, the shapes of the subvolumes in the FVM discretization are less constrained. In fact, the FVM system of equations may be formulated for non-uniform or even unstructured discretizations, as long as the entire volume is discretized consistently without overlap. As such, the FVM discretization allows more general shapes than both the FEM and FDM meshes, which may typically be exploited to reduce the overall node-count in the FVM system of equations. Also for regular meshes (as in Figures 2.10a and b) an FVM system of equations can be formulated.

Similar to the FEM and the FDM, the FVM is well capable of modeling the doping patterns from Section 2.2. Furthermore, if the FVM is applied such that it has a lower node-count than a FEM or an FDM, but (approximately) retains sparsity, it is computationally less demanding than both the FEM and the FDM and can therefore handle larger structures.

The FVM is used in the device modeling tool CURRY [Pol87], which also contains dedicated numerical (iterative) methods for solving the system of equations. The FVM is also used in device simulators like the 3D simulator Davinci [MDa], and its 2D counterpart Medici.

2.3.4 BEM-based Substrate Model

As will be explained with more detail in Chapter 3, the BEM only discretizes contact areas on the boundary of the domain, as indicated in Figure 2.10c, and uses a Green's function to capture the properties of the domain itself. A Green's function can be derived for uniformly layered domains, but including localized doping patterns into the Green's function is typically impractical, and may even be considered impossible for all but the simplest situations. Therefore, the BEM is suitable for modeling the global doping profile as a uniformly layered domain, but taking into account local doping patterns is not as straightforward as for the FEM. In fact, the effects from local doping patterns can only be incorporated in a limited way, merely as an addition to the BEM model of the global doping profile.

To retain the computational efficiency and model compactness of the BEM as much as possible, the n-wells and n-diffusions were traditionally the only doping patterns that were taken into account. The common approximative approach for incorporating the n-areas into the BEM model is schematically shown in Figure 2.11. Focusing on the n-well, it is lifted out of the substrate, and the well-bottom defines a BEM substrate contact, to which the well junction capacitances (including those from its sidewalls) are connected (in a similar way as in Figure 1.6). Depending on the resistive characteristics of the well, it is possible to model the well itself with a FEM, including a lumped model along the well-bottom. Despite the apparent limitations in BEM-based substrate modeling approach, the resulting models can still be useful [Sme95a].

2.3.5 Interpolation-based Substrate Model

The interpolation-based methods are a separate class of modeling techniques. Basically, these methods can be considered heuristic or low-order modeling techniques that are calibrated along 3D modeling data, and interpolate from the calibration points. The methods can be very fast and, when properly calibrated, can also be accurate depending on the type of modeling problem. In general, the methods are used to quickly obtain coarse, but representative, models.

A fast interpolation method for the substrate, based on a Delaunay triangulation between the contact areas is described in [Gen96b]; other interpolation approaches to substrate modeling were published in [Joa94, Ver95a]. However, these techniques will not explicitly be considered in this thesis and will therefore not be discussed in detail.

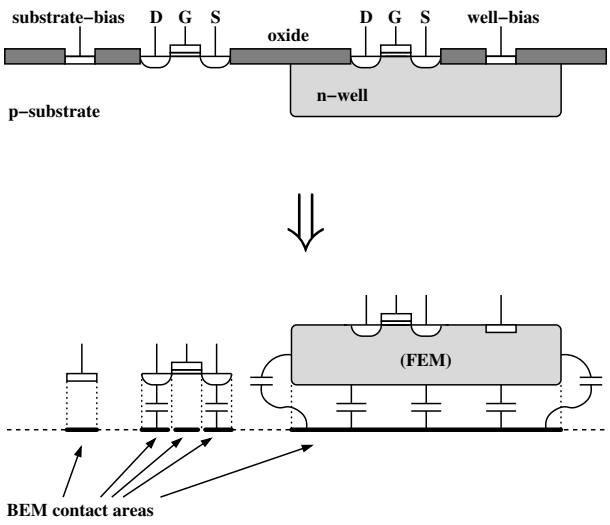

Figure 2.11: Schematic representation of a substrate structure (top) and the BEM modeling approach (bottom).

2.4 The Need for an Improved Modeling Strategy

In modern technologies, the role of the substrate has become increasingly important in overall chip-behaviour. In particular, the increasing frequency in the signals has caused the general behaviour of the chip to become more agressive. This is specifically shown by:

- Faster switching of transistors,
- Stronger field coupling between interconnect and substrate,
- Greater impact of package inductance.

Purely digital circuits can be considered relatively insensitive to substrate noise, even though e.g. a clock (typically PLL-based) might suffer from jitter due to substrate crosstalk, and the ever decreasing operating voltages do make the digital circuitry more sensitive. In purely analog (RF-) designs the substrate may play an important role in the behaviour of passive components in the circuit (e.g. substrate losses in an integrated inductor). However, technological trends have aimed at ever stronger integration by placing both analog and digital circuitry on a single chip (mixed-signal designs). In such designs, the analog circuitry may be very sensitive to the substrate noise induced by the digital circuitry, possibly causing a poor (or even unacceptable) Signal-to-Noise-Ratio in the analog signals.

In summary, technology has grown towards higher frequencies and the impact of substrate crosstalk has grown along with it. Therefore, substrate technology has mainly advanced towards avoiding substrate crosstalk. In general, the substrate has become an important factor in the physical behaviour of the chip, making a good substrate model important for reliable simulations. In this context, it has become increasingly important to develop modeling techniques that can generate consistent and accurate models for a wide range of different substrates.

As mentioned in Section 2.3, the FEM, FDM, FVM and BEM are the main available modeling techniques for substrate resistance extraction. However, with the ultimate objective of a modeling method efficiently calculating a sufficiently accurate and consistent substrate model for a full VLSI circuit, the FEM, FDM, FVM and BEM each have their limitations.

Where the FEM would easily reach the required accuracy and consistency, it is prohibitively slow and therefore restricted to relatively small domains. The FDM and FVM would also be able to reach the required accuracy and consistency, and they would be faster than the FEM, but they are still prohibitively slow for large modeling problems. Therefore, in order to be efficient, commercial substrate modeling tools based on, for example, the FDM (like SubstrateStorm [MSu]) tend to operate at some distance from the convergence point. In particular, these tools typically use coarse meshes and seem to use large error tolerances in the iterative solution procedure applied to the system of equations.

Where the BEM would reach the appropriate speed (especially using the windowing technique which will be described in Section 3.3.2), it is restricted to uniformly layered domains and, in order to stay as efficient as possible, can only include the effects of local doping patterns in a limited way. As such, the BEM only yields a consistent model in restricted situations and typically would not yield sufficient accuracy, if local doping patterns are part of the modeling problem.

2.5 Reconsidering the Modeling Problem

According to the discussion in Section 2.4, the advanced substrate structures and the new modeling requirements basically emphasize the accuracy limitations in the BEM and the speed limitations in the FEM, FDM and FVM. Therefore, the methods seem to be unsuitable for the problem under consideration. However, by considering again the modeling problem itself, we may find a new angle on the way the modeling techniques can be applied.

Consider Figure 2.12 which is similar to Figure 2.5, but with relevant cross sections indicated. Figure 2.13 contains the accompanying doping profile for each cross section, drawn in a single figure. Section 5.2 will explain these figures in more detail, but for the time being it is sufficient to observe that the doping patterns continue into

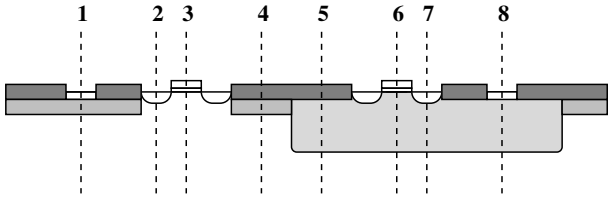

Figure 2.12: Relevant cross sections through the doping patterns from Figure 2.5.

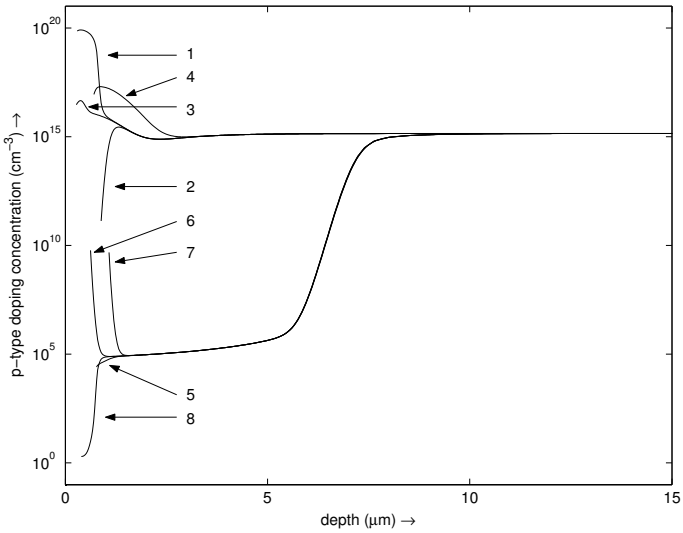

Figure 2.13: Example of a full doping profile, drawn in a single figure; loosely based on a standard CMOS technology in a lightly doped substrate provided by SubstrateStorm. Indicated numbers refer to cross-sections from Figure 2.12.

the global doping profile beyond the first few microns. According to the characteristics of the BEM and the FEM (where the FEM is chosen as the representative of the three previously described volume discretization methods), the BEM would be most suitable for modeling the global doping profile, while the FEM would be most suitable for modeling the doping structures in the top of the substrate. A separation of the problem into a BEM and a FEM part, and some consistent connection between them suggests a new modeling approach that will be detailed in the following chapters.

3

Relevant Properties of the BEM and FEM Modeling Techniques

As explained in the previous chapter, the structure of the substrate modeling problem is such that a combination of the BEM and FEM modeling techniques results in a new substrate modeling strategy. To determine the feasibility of such a combined BEM/FEM modeling technique, this chapter will outline the relevant properties of the BEM and the FEM that may allow a computationally efficient combination of the methods. In the next chapter, the combined BEM/FEM method will then be introduced.

3.1 Mathematical Formulation of the Problem

The most fundamental mathematical formulation of electromagnetic phenomena is given by Maxwell's equations. However, Maxwell's equations would be too general (and unnecesarrily complicated) for the situations described in this thesis. Therefore, we will limit the scope of the problem as follows:

- The fields are assumed quasistatic (i.e. the potential is defined).
- The domain is assumed to be purely resistive (conductive).
- The domain does not contain fixed charges or current sources.

Maxwell's equations then allow a more straightforward formulation as described below.

3.1.1 The Laplace Equation

Under the conditions summarized above, Maxwell's equations may be simplified to the following formulation [Web50, Hoo75]. Let p be a point with spatial co-ordinates (x_p, y_p, z_p) in a 3D domain Ω. In point p, let $\Phi(p)$ be the potential, $\sigma(p)$ be the conductivity, $J(p)$ be the current density vector, and $E(p)$ be the electric field vector. Then, a distributed formulation of Ohm's law can be written as follows:

$$J(p) = \sigma(p)E(p) \tag{3.1}$$

$$E(p) = -\nabla\Phi(p) \tag{3.2}$$

Furthermore, the law of conservation of current prescribes:

$$\nabla \cdot J(p) = 0 \tag{3.3}$$

Substituting Equation 3.2 into Equation 3.1 and substituting the result into Equation 3.3, we obtain the following differential equation:

$$\nabla \cdot (\sigma(p)\nabla\Phi(p)) = 0 \tag{3.4}$$

If the conductivity in the domain is homogeneous (that is, $\sigma(p)$ is constant in the entire domain), this equation reduces to a well-known equation from potential theory:

$$\sigma\nabla^2\Phi(p) = 0 \tag{3.5}$$

Equation 3.5 is known as the Laplace equation.

For inhomogeneous domains either Equation 3.4 should be used, or, if the domain may be divided into subdomains with each a constant conductivity, we may consider the Laplace equation to hold in each subdomain, and assume continuity relations along the interfaces between the subdomains.

3.1.2 Boundary Conditions

Since the substrate is considered to be predominantly resistive, we will approach the modeling problem as a problem that aims to find a resistance network between contacts on a (layered) resistive domain. This basically requires to solve the Laplace equation under appropriate boundary conditions.

The boundary conditions are chosen such that current can enter or leave the domain through the contact areas on the boundary, while the remainder of the boundary

has insulating properties. This requires to define Dirichlet boundary conditions on the contact areas, and homogeneous Neumann boundary conditions on the remainder of the outer boundary of the domain. In mathematical terms, this may be formulated as follows.

Let S be the outer boundary of the domain Ω. Then Let $S_D \subset S$ be the areas where Dirichlet boundary conditions hold, and $S_N \subset S$ be the areas where homogeneous Neumann boundary conditions hold, such that $S_D \cup S_N = S$ and $S_D \cap S_N = \emptyset$.

Then, the homogeneous Neumann boundary conditions are as follows

$$\frac{\partial \Phi}{\partial n} = 0 \quad \text{on } S_N \tag{3.6}$$

Using an externally applied potential ϕ_a, the Dirichlet boundary conditions are as follows:

$$\Phi = \phi_a \quad \text{on } S_D \tag{3.7}$$

Another property on S_D is that the normal component J_n of the current through the contact is proportional to $\frac{\partial \Phi}{\partial n}$ with σ as the proportionality factor. In particular:

$$\frac{1}{\sigma} J_n = \frac{\partial \Phi}{\partial n} \quad \text{on } S_D \tag{3.8}$$

Assuming the potential on contact areas to be constant, the tangential component of the current through the contact areas will be zero.

Under these boundary conditions, a resistance network between the contacts on the domain can be obtained from the solution of the Laplace equation.

3.1.3 Energy Functional

A convenient property of the field obtained from the Laplace equation and its boundary conditions is that it has minimum energy (Thomson's theorem), where the energy $\mathcal{E}$ is given by the following energy functional:

$$\mathcal{E} = \int_{\Omega} \sigma(p) ||\nabla \Phi(p)||^2 dp \tag{3.9}$$

This property is particularly useful in studying convergence of approximate field solutions. In particular, any approximate field solution that has energy close the minimum, is also close to the exact field solution in a point-wise sense.

3.2 The Finite Element Method

The Finite Element Method (FEM) is a well-known numerical method for solving the Laplace equation. It has been studied and applied extensively in many textbooks

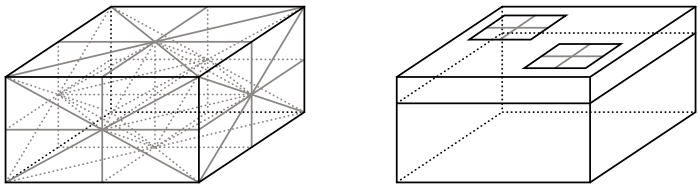

a) b)

Figure 3.1: a) Example of a 3D tetrahedral volume discretization as would be used by the FEM; b) Example of a discretization of contact areas on top of uniformly layered domain as would be used by the BEM. Images taken from Figure 2.10, but repeated here for convenience.

and scientific papers in many different fields of mathematics and engineering. As background knowledge for this thesis, we have used the more practical approach to the FEM presented in [Sil96], and some of the theory presented in [Str73, Nor73, Zie83].

3.2.1 General Approach

The FEM subdivides the entire domain into triangular (2D) or tetrahedral (3D) elements (see Figure 3.1a). Then, an appropriate basis function is applied to the elements. The basis function provides the mathematical relation between the potentials at the corners (the FEM nodes) of the elements. Even though higher-order basis functions are possible, we will use a straightforward piecewise linear (or piecewise planar for 3D) basis function. The FEM elements are combined into a global system of equations which is then used to minimize the energy functional from Equation 3.9. In this case, the energy functional is used as an alternative formulation for the Laplace equation.

The system of equations in the FEM is inherently sparse, and typically large. Despite its sparsity, the size of the system of equations typically makes the solution procedure computationally intensive, resulting in long computation times while requiring a lot of memory. Therefore, over the years, many efforts have been made to improve the speed, memory-usage and general efficiency of the FEM. The performance of the FEM can be improved by structured meshing, that is, coarser meshing in areas that do not require much detail and fine meshing in areas that do. This approach is commonly applied in many FEM-based software tools, e.g. FEMLAB [MFe], and mainly aims at a-priori reducing the size of the FEM system of equations, without losing accuracy. With respect to the solution procedure itself, dedicated solvers have been developed, e.g. based on Krylov iterative techniques [Saa03], multigrid techniques [Sha03], or a combination of both.

The 3D discretization will automatically incorporate any inhomogeneities of the domain into the model, because each element in the discretization can have its own

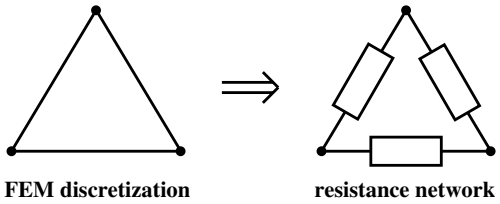

Figure 3.2: When using a piecewise linear basis function in the FEM, the mathematical relations along the edges in the FEM discretization already represent resistances.

material properties assigned. As such, the FEM is very accurate and flexible, making it a suitable method for a very wide range of mathematical and engineering problems.

3.2.2 The FEM discretization as a Resistance Network

Taking a slightly different approach to the FEM, the resistance network can be found without having to calculate the field solution. This approach is based on the observation that when choosing a piecewise linear basis function, the mathematical relations along the edges in the FEM discretization already represent resistances [Hal87, Gen88, Gen91]. In other words, the FEM discretization is then equivalent to a resistance network, as schematically shown in Figure 3.2. The proof for this approach is presented in [Gen91], and will not be repeated here. When the resulting network is used for subsequent circuit simulation, the minimization of the energy functional will automatically (though implicitly) be ensured by the circuit simulator itself.

The network obtained in this way has many nodes and a sparse structure (the number of resistances is approximately linear in the number of nodes). However, only a small number of the FEM nodes are actually terminal nodes, while the rest of the FEM nodes are internal nodes. Therefore, any internal nodes can be eliminated through node-by-node star-delta transformation (i.e. Gaussian elimination), through the Schur complement (i.e. block-wise Guassian elimination) or through a different solution procedure applying an iterative technique from e.g. [Saa03].

3.3 The Boundary Element Method

The Boundary Element Method (BEM) is a well-known numerical method for solving the Laplace equation. It has been extensively studied and applied in literature, in various fields of mathematics and engineering. Among the more recent books on this topic is, for example, a treatment on the BEM in the field of electrical engineering [Pol05]. However, [Bre78] is typically seen as a key reference.

3.3.1 General Approach

The Boundary Element Method (BEM) is based on an integral form of the Laplace equation. In the following, we assume the Laplace equation to hold in our domain Ω, and the boundary conditions to be as described in Section 3.1.2.

Green's Function

Let p and q be two points on the boundary of the domain, let r be the radius (distance) between them $r = \sqrt{(x_p - x_q)^2 + (y_p - y_q)^2 + (z_p - z_q)^2}$, and let $G(p, q)$ be a Green's function. In the present case, the Green's function is a fundamental solution to the Laplace equation, and, as such, it automatically ensures the minimization of the energy functional from Equation 3.9. The Green's function "encodes" the characteristics of the medium, and it can be interpreted physically as "the potential in point p due to a current injected at point q".

For general axisymmetric situations (of which the layered substrate is a special case) a Green's function can be derived through a Hankel transform. However, the Hankel transform is computationally an expensive procedure, even for the simplified situation of a layered substrate. Nevertheless, Hankel transform implementations with improved efficiency are available, for example [Kol04] which is available in the SPACE layout-to-circuit extractor [MSp]. The Hankel transform is applied to the substrate modeling problem in e.g. [Li 02, Bra04].

To avoid the computational complexity of the Hankel transform, the Green's function in uniformly layered substrates can alternatively be evaluated through a method which does not explicitly evaluate the Hankel transform: the method of images. Unfortunately, the method of images is still computationally expensive if there are many layers involved (see e.g. [Zha98]). However, for uniformly layered media that consist of only 2-3 layers, the method of images can efficiently be used to evaluate the Green's function. For the situation of a substrate consisting of 2 layers, a Green's function in terms of the method of images is discussed in [Sme95b].

Unfortunately, formulating a Green's function for domains that contain lateral conductivity variations is already cumbersome for variations of very simple shape and can be considered impossible for variations of more general shape. Therefore, localized conductivity patterns of arbitrary shape are typically impossible to capture in a Green's function. As such, localized doping patterns in the substrate are typically impossible to capture when applying the BEM to the substrate modeling problem.

The potential in any point of the domain can only be defined with respect to some reference point of fixed potential. This reference point is quite arbitrary [Ram94], and may conveniently be chosen based on the modeling problem. For the substrate modeling problem, the domain size is very large compared to the typical size of the contact areas. Therefore, it is convenient to define the potential at infinity as zero, and

use infinity as the reference point. As such, the presented BEM technique assumes the substrate to be a semi-infinite domain. The validity of this approach for modeling finite domains deserves a closer study, which is presented in Section 3.3.4. In Section 3.3.3, an expansion of the BEM technique is presented which approximates the finite dimensions of the domain, resulting in a pseudo-finite approximation of a finite domain. In general, however, the semi-infinite domain is typically considered valid, because, in typical substrate modeling problems, the lateral dimensions and the thickness of the substrate are large compared to the size of the layout-features under consideration.

For the 3D homogeneous case of the Laplace equation, the Green's function is

$$G(p,q) = \frac{1}{4\pi\sigma r} \tag{3.10}$$

We observe that the potential indeed goes to zero as r goes to infinity.

Integral Equation

Now, let $k(q)$ be the current density distribution function on S_D. Then, using an appropriate Green's function for the substrate modeling problem in a semi-infinite domain [Sme95b], the potential in point p can be written as

$$\Phi(p) = \int_{S_D} k(q)G(p,q)dq \tag{3.11}$$

Equation 3.11 is the integral formulation of the Laplace equation on which the BEM is based. It is basically found by applying Green's second identity (e.g. [Kre93]) to the Laplace equation, and applying a-priori knowledge with respect to the boundary conditions and the domain properties ('encoded' by the Green's function).

Since Equation 3.11 integrates over the whole contact area including the 'observation point' p, the Green's function from Equation 3.10 (or a more advanced version) also needs to be evaluated in a situation where the source point and observation point are one and the same. Under these circumstances, a singularity occurs, because r becomes 0. Fortunately, this singularity is integrable according to [Wil84].

Discretization and Solution

According to Equation 3.11 the BEM only has to discretize those parts on the boundary of the domain (see Figure 3.1b) where Dirichlet conditions hold (i.e. S_D). The discretization usually (but not necessarily) utilizes rectangular panels and applies a piecewise constant basis function, such that the current distribution is assumed to be constant on each panel and each panel forms an equipotential region. As such, the

discretization allows a piecewise constant approximation of the continuous current density distribution.

Alternatively, it is also possible to use a piecewise linear basis function in the BEM, combined with a triangular mesh. Such an approach is, however, significantly more complex with respect to implementation. Therefore, for the development of the prototype BEM/FEM method it was more convenient to use a straightforward BEM with piecewise constant basis function. As a result, the BEM with piecewise linear basis function will not explicitly be considered in this thesis.

Based on the discretization, the Method of Moments [Har68] allows to find the piecewise constant approximation of the current density distribution from a linear system of equations. In this case, we define **P** as the vector of panel potentials, **K** as the vector of (unknown) panel currents and *G* as the influence matrix. An entry G_{ij} in the influence matrix describes the potential at panel i due to a unit current injected through panel j (i.e. the Green's function is evaluated for each panel-pair). We then obtain:

$$\mathbf{P} = G\mathbf{K} \tag{3.12}$$

The BEM then continues by defining an incidence matrix *F* relating panels and contacts. An entry F_{ij} in the incidence matrix is '1' when panel i is on contact j, and is '0' otherwise. If each BEM contact is covered by only a single BEM-panel, F is an identity matrix. However, since each BEM contact will typically be covered by multiple BEM panels due to refinements in the discretization, matrix F is usually non-square. By denoting **V** as the contact potential vector and **I** as the contact current vector, we can write

$$\mathbf{P} = F\mathbf{V} \tag{3.13}$$

$$\mathbf{I} = F^T\mathbf{K} \tag{3.14}$$

Combining Equations 3.12, 3.13 and 3.14, we obtain:

$$\mathbf{I} = F^T G^{-1} F\mathbf{V} = Y\mathbf{V} \tag{3.15}$$

where *Y* is an admittance matrix for the resistive substrate with the substrate contacts as ports.

Network Structure

The admittance matrix resulting from the BEM defines a network which has a dense structure: each node in the network is a port-node that is connected to every other node, there are no internal nodes. An illustration of such a network is shown in Figure 3.3. For this type of network, the number of resistances is quadratic in the number of nodes, which may slow down the simulation of larger networks. Therefore, model order reduction should be used to simplify this network. The windowing technique described in Section 3.3.2 is a suitable technique for model order reduction in the BEM.

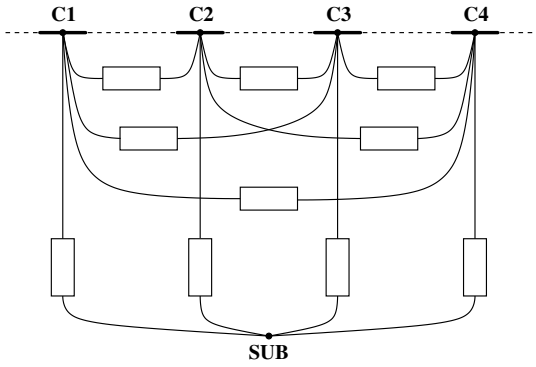

Figure 3.3: BEM resulting circuit model for a configuration of four contacts on top of the substrate and the reference node 'SUB' at infinity.

Reference Node

Apart from the port-nodes, the BEM network also contains a reference node which is typically referred to as 'SUB', as illustrated in Figure 3.3. This reference node respresents the fixed-potential reference point at infinity with respect to which the potentials in our modeling problem were defined.

3.3.2 Windowing Technique

In the context of model order reduction and the BEM, it is possible to extract a reduced-order model through the Schur algorithm for approximate matrix inversion [Dew87]. The algorithm is schematically represented in Figure 3.4. The main principle applied in the Schur algorithm is that it requires the influence matrix to be known only partly, in a (staircase) band around the main diagonal. The approximate inversion then (implicitly) estimates the matrix-entries outside the band (indicated by the sections Z in Figure 3.4) in a maximum-entropy sense, such that the resulting inverse matrix contains zeroes in those areas.

The Schur algorithm can *only* be applied to matrices with a staircase band structure. A BEM influence matrix with a staircase band structure suggests interactions between closely coupled BEM-panels; a windowing technique as in Figure 3.5 would then be a straightforward implementation of this concept. Unfortunately, for BEM modeling problems with 2D and 3D geometries, the influence matrix typically does not have a (staircase) band structure, but a multiple band structure [Nel89b]. In some cases, such a multiband matrix can be permuted into a (staircase) band matrix, but

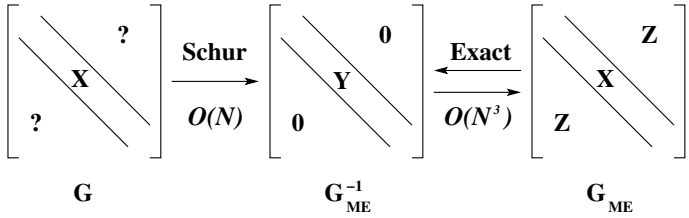

Figure 3.4: Schematic interpretation of the Schur algorithm for approximate matrix inversion.

in general this is not possible. Therefore, [Nel88] presents a hierarchical approach which uses the Schur algorithm to invert band-structured submatrices of the original multiband matrix. This hierarchical approach can consistently be combined with the windowing technique from Figure 3.5 [Mei92].

The hierarchical Schur algorithm for sparse approximate inverses of partially specified matrices [Nel89a] has been implemented in the SPACE layout-to-circuit extractor together with the windowing technique for the BEM. The main benefit of the Schur algorithm in conjunction with the BEM is that it improves linearity in the number of BEM-panels:

1. The number of evaluations of the Green's function becomes linear.

2. The inversion process itself becomes linear.

3. The resulting number of resistances becomes linear.

In other words, the computation of the BEM becomes more efficient, while the resulting network becomes sparser. In particular, the SUB reference node representing the potential at infinity already provides a sparse representation of global couplings through the deep substrate. Together with the windowing technique, this results in an overall sparser model.

3.3.3 Pseudo-Finite Domain BEM

The previous sections described the BEM for a semi-infinite domain. However, domains with finite lateral dimensions can be modeled by the BEM through an image-method for the domain sidewalls [Sme95b]. This image-method is schematically represented in Figure 3.6. Basically, images of the boundary element under consideration are defined with respect to the sidewalls of the domain, so as to approximate a Neumann condition.

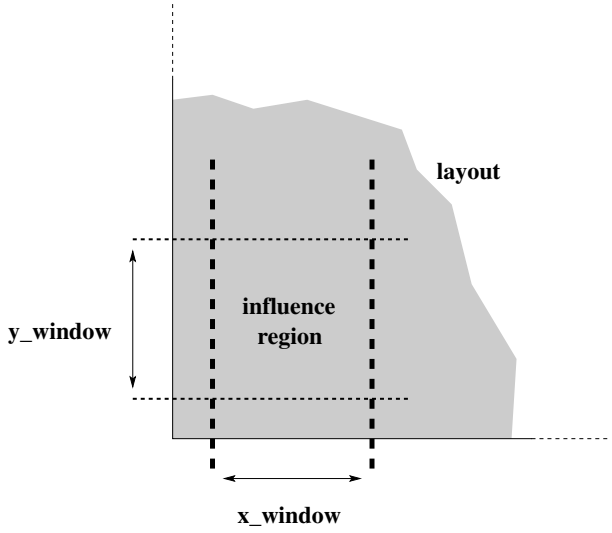

Figure 3.5: Windowing technique. The BEM influence matrix is only calculated inside the influence region bounded by the x- and y-windows. Through consistent (step-wise) sliding of the influence region, the whole of the layout is eventually covered.

Even though the image-series in the sidewall approximation actually diverges [Sme95b], an approximation of the sidewall effects can still be obtained by taking into account only the first image related to the sidewall(s) closest to the boundary element under consideration.

The sidewall-images technique allows to take the sidewalls into account only for boundary elements that are closer to the edge of the domain than specified by a user-defined distance parameter. If this parameter is chosen as (more than) half the lateral dimensions of the domain, the sidewalls are taken into account for all boundary elements.

A pseudo-finite BEM domain can now be achieved through a combination of sidewall-images and a double-layer BEM approach (as previously described in Section 3.3.1). In particular, finite lateral dimensions of the domain can be approximated by sidewall images, while finite thickness of the domain can be approximated by a double-layer BEM where the top-layer has the thickness and conductivity of the domain under study, while the bottom layer has a much lower conductivity, such that the bottom layer can be considered non-conductive. The combination of these techniques then results in a pseudo-finite BEM domain.

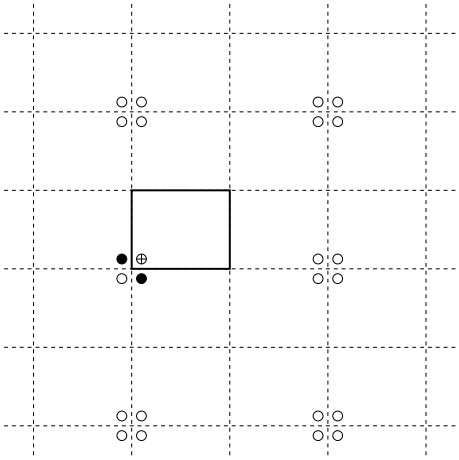

Figure 3.6: Image technique for taking domain sidewalls into account with the BEM. Solid rectangle indicates domain boundary. Symbol $\oplus$ represents the boundary element under consideration, empty circles represent images, black circles represent images actually used in the approximation.

3.3.4 Validity of semi-infinite approximations of finite domains

In the general BEM approach, the substrate is treated as a semi-infinite domain. This might suggest that the resulting models with the reference node at infinity are inappropriate for actual substrates which have finite dimensions. However, we will show below that the substrate as a semi-infinite domain is actually an appropriate approximation of finite substrates.

Consider the situation of two BEM contacts $C1$ and $C2$ on top of a lightly-doped semi-infinite substrate. A schematic representation of the field distribution in the substrate is shown in Figure 3.7. For increasing distance d between the contacts, we will observe that the port resistance saturates. This behaviour may be made plausible as follows. Consider the contacts to be equally sized (say, $5\mu m \times 5\mu m$) and to be covered by only a single BEM panel. Even though this may be considered a coarse discretization, the approximation actually becomes more accurate as the distance d between the contacts goes to infinity, since the contacts then resemble two point sources in space. The interaction between the contacts can be found through the Green's function, which, for a lightly-doped semi-infinite substrate is similar to the Green's function shown in Equation 3.10. Knowing that the the conductivity σ is constant, we may introduce the constant k as follows

$$k = \frac{1}{4\pi\sigma} \tag{3.16}$$

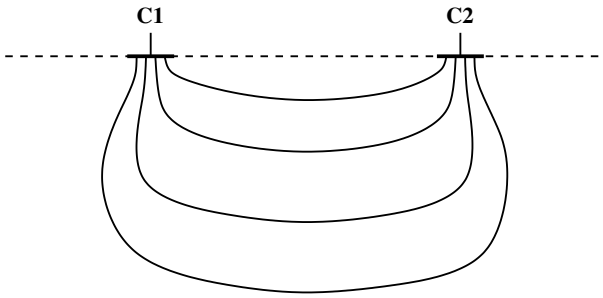

Figure 3.7: Schematic representation of field lines between contacts on a lightly doped substrate. The field distributes itself evenly.

such that the Green's function from Equation 3.10 rewrites to

$$G(p,q) = k\frac{1}{r} \tag{3.17}$$

The interaction between contacts $C1$ and $C2$ is then $\sim \frac{k}{d}$. The interaction of a contact with itself is less straightforward to compute due to a singularity in the Green's function when r becomes 0, but, as already mentioned previously, this singularity is integrable according to [Wil84]. Defining the value a as the result of the integration over $\frac{1}{r}$ using the technique from [Wil84], the interaction of a contact with itself is then $\sim ka$. Even without evaluating the integral, we can understand that a increases with the size of the contact, that it is independent of the distance d between the contacts, and that it has the same dimension as $\frac{1}{d}$, namely m^{-1}. Then, the influence matrix G looks as follows:

$$G = k \cdot \begin{bmatrix} a & \frac{1}{d} \\ \frac{1}{d} & a \end{bmatrix} \tag{3.18}$$

The resulting admittance matrix is found by straightforward inversion:

$$G^{-1} = \frac{1}{k(a^2 - \frac{1}{d^2})} \cdot \begin{bmatrix} a & -\frac{1}{d} \\ -\frac{1}{d} & a \end{bmatrix} \tag{3.19}$$

Figure 3.8 shows the corresponding resistance network obtained by direct inspection from the admittance matrix. The SUB node may now be eliminated to obtain the port resistance between contacts $C1$ and $C2$. The resulting port resistance, R_p is found to be:

$$R_p(C1, C2) = \frac{2k(a^2 - \frac{1}{d^2})}{a + \frac{1}{d}} \tag{3.20}$$

Increasing the distance d to infinity, we observe saturation of the port resistance to a value of $2ka$:

$$\lim_{d\to\infty} R_p(C1, C2) = \lim_{d\to\infty} \frac{2ka^2 - \frac{2k}{d^2}}{a + \frac{1}{d}} = \frac{2ka^2}{a} = 2ka \tag{3.21}$$

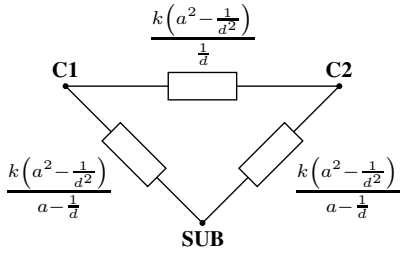

Figure 3.8: Circuit model for a finite distance between contacts C1 and C2 on a uniform, semi-infinite substrate. Resistance values obtained by direct inspection from the admittance matrix in Equation 3.19.

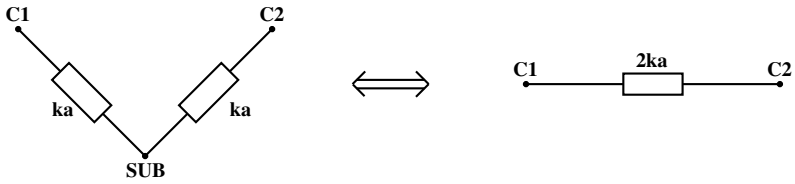

Figure 3.9: Circuit model for an infinite distance between contacts C1 and C2 on a uniform, semi-infinite substrate. The port resistance between the contacts saturates to a constant value.

Also from the resistance network in Figure 3.8, this behaviour can be observed. In fact, the resistor between $C1$ and $C2$ goes to infinity as d goes to infinity. The resistors to the SUB node converge to ka as d goes to infinity. This is shown in Figure 3.9, which also shows that a resistance of $2ka$ remains when the SUB node is eliminated. This confirms the observation from Equation 3.21.

For a heavily doped substrate (recall Figure 2.1b) the saturating behaviour of the port resistance is also observed. A schematic representation of the field distribution in the substrate is shown in Figure 3.10. In particular, the highly conductive deep substrate presents an (almost) ideally conducting domain, such that it may be considered to be an equipotential region. The reference point at infinity may be considered to be part of this equipotential region, such that the SUB reference node in the resulting network becomes representative for the highly conductive deep substrate. In this case, however, the saturating behaviour of the port resistance will already be observed for smaller distances between the contacts because, for increasing distance between the contacts, the resistive path will predominantly be determined by the vertical distance between the contacts and the highly conductive deep substrate already at medium distances between the contacts.

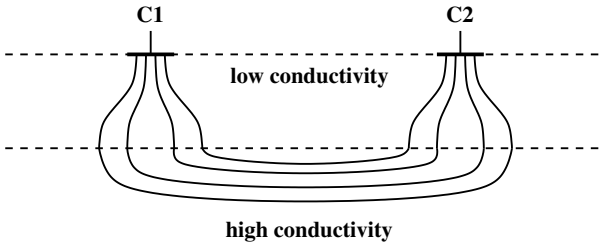

Figure 3.10: Schematic representation of field lines between contacts on a heavily doped substrate. The field is mainly vertically oriented between the contacts and the high-conductivity layer, and horizontally oriented inside the high-conductivity layer.

These observations are also confirmed in practice. Assume the two contacts $C1$ and $C2$ (each sized $5\mu m \times 5\mu m$) on top of a lightly doped semi-infinite substrate and on top of a heavily doped semi-infinite substrate as previously shown in Figure 2.1. Plotting the port resistance between the contacts as a function of the distance between the contacts, we obtain Figure 3.11. The saturating behaviour is clearly visible.

Now, consider Figure 3.12, where we consider the same 2 contacts (each sized $5\mu m \times 5\mu m$) on top of a $4000\mu m \times 4000\mu m$ substrate, with a thickness of $350\mu m$. The dimensions of this finite domain are representative of actual substrates. We will now study whether the resistance between the contacts on this finite substrate also exhibits saturating behaviour.

Plotting the resistance value between the contacts as a function of the distance d for the substrate of finite and infinite dimensions, Figure 3.13 is obtained. The modeling approaches apply an infinite domain BEM and the pseudo-finite domain BEM from Section 3.3.3.

This last method can be used as a pseudo-finite domain BEM by defining the substrate as a double-layer domain with a highly resistive layer at the bottom, of which the top is located at the prescribed depth of $350\mu m$. We observe saturating behaviour in both curves. Furthermore, we observe that the resistance values for both configurations are very close (within 4%), and that only as the contacts reach the edges of the finite domain, the resistance in the finite domain suddenly increases strongly. A detail view of this situation is shown in Figure 3.14 which reveals that the sudden increase is actually more gradual than might initially have been observed from Figure 3.13. An additional observation we make from Figure 3.13 is that the resistances calculated for the infinite domain are all somewhat smaller than the resistances calculated for the finite domain. This is correct because the field lines in the modeling problem are not confined to the finite domain, but may expand into the infinite domain, which introduces additional 'parallel paths' outside the finite domain, thus reducing the resistance.

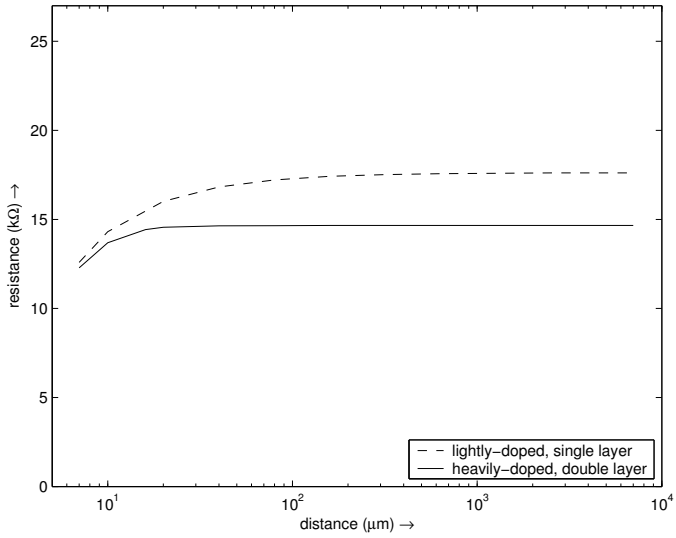

Figure 3.11: Resistance values for increasing distance between the contacts, using a single-layer and a double-layer BEM approach for the lightly-doped and heavily-doped substrates as shown in Figure 2.1. The resistance value saturates in both cases.

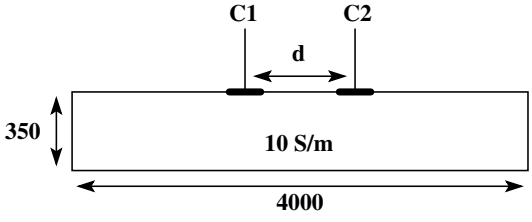

Figure 3.12: Test configuration of the substrate as a finite domain. Contacts C1 and C2 are each sized $5\mu m \times 5\mu m$. Distance values are in μm.

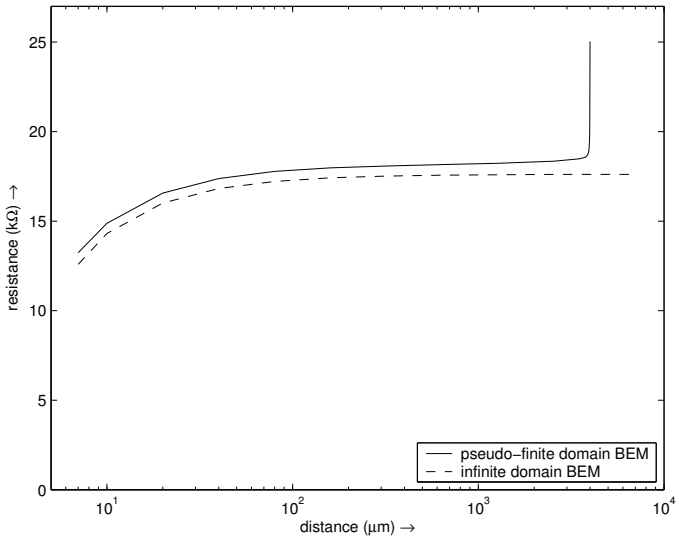

Figure 3.13: Resistance values for increasing distance between the contacts, using a pseudo-finite and an infinite domain BEM approach. The curves stay within approximately 4 % of each other, until the contacts get close to the edges of the (pseudo-)finite domain.

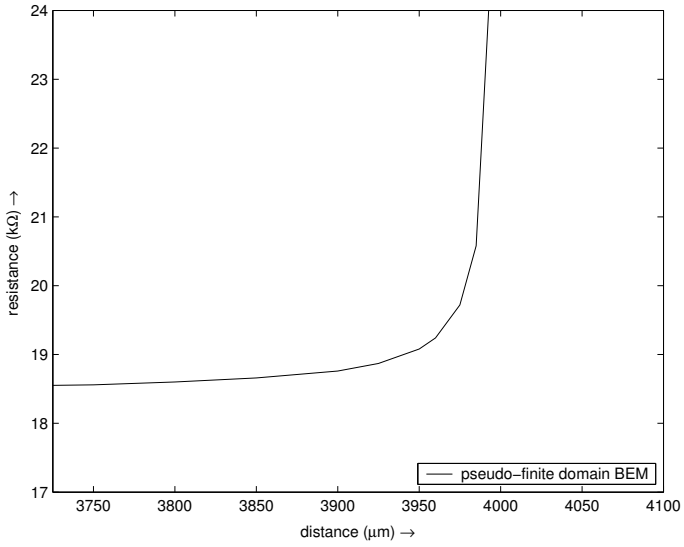

Figure 3.14: Detail view of the resistance curve for the pseudo-finite domain BEM, near the edge of the domain. The transition is gradual. Note: the horizontal axis is now linear whereas in Figure 3.13 it was logarithmic.

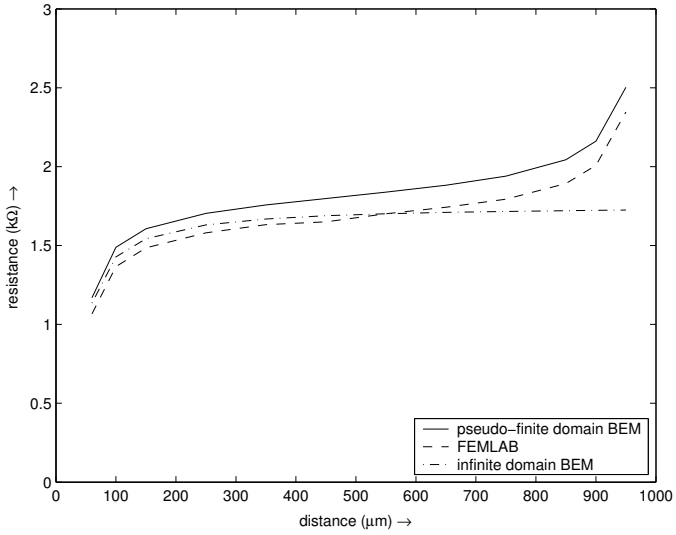

Figure 3.15: Resistance values for increasing distance between the contacts, using FEMLAB, and the infinite and pseudo-finite domain BEM. The curves for FEMLAB and the pseudo-finite domain BEM are close to each other over the whole distance range, while they diverge from the infinite-domain BEM as the contacts get close to the edges of the (pseudo-)finite domain. The finite domain had dimensions $1000\mu m \times 1000\mu m \times 350\mu m$, with $50\mu m \times 50\mu m$ contacts on top.

To verify whether the observed saturation and near-edge behaviour of the resistance in the finite domain is correct, we do a similar simulation with the 3DFEM modeling tool FEMLAB [MFe]. Unfortunately, FEMLAB could not model a situation equivalent to Figure 3.12 because the relative size of the contacts is extremely small ($5\mu m \times 5\mu m$) compared to the domain size ($4000\mu m \times 4000\mu m \times 350\mu m$); this introduced problems with meshing and convergence. Therefore, the problem was modified to a domain of $1000\mu m \times 1000\mu m \times 350\mu m$ with contact areas of $50\mu m \times 50\mu m$. This situation is perhaps no longer representative of actual substrate configurations, but the saturation and near-edge behaviour of the resistance value in the finite domain may still be observed. The results are shown in Figure 3.15, which, apart from some convergence-related difference between the curves, confirms the saturation and near-edge behaviour of the resistance in the finite-domain substrate. Furthermore the figure confirms (again) that the infinite-domain BEM provides a valid approximation of the finite-domain situations.

From the observations in the previous paragraphs, we conclude that the semi-infinite domain is an appropriate approximation of the typical situation in a lightly doped substrate, as long as the contact areas stay at an 'appropriate' distance from the edge. For the situation shown in Figures 3.13 and 3.14, the appropriate distance would

be about 2 % of the lateral dimensions of the domain. In Figure 3.15 that would be 30–40 %, but this situation is less representative of actual substrates.

3.4 Summary and concluding remarks

In this chapter we identified the Laplace equation as the mathematical formulation of the substrate modeling problem. Then, we summarized the Finite Element Method (FEM) and the Boundary Element Method (BEM) as the two main techniques for solving the Laplace equation. In particular, we focused on the relevant properties of the BEM and the FEM with respect to a combined BEM/FEM method as a new modeling strategy for the substrate.

In this context, the most relevant property of the FEM is as follows.

- When using a piecewise linear basis function, the FEM discretization can be interpreted as a resistance network (see Section 3.2.2).

With respect to the BEM, Section 3.3.4 provides a very relevant observation: the infinite-domain BEM is a valid approximation of a finite-domain substrate with relatively large dimensions, as long as the contact areas are not nearer to the edges of the finite domain than about 2 % of its lateral dimensions.

The most relevant properties of the BEM may now be summarized as follows:

- The 'SUB' reference node at infinity is a valid representation of the deep substrate in a finite domain of relatively large dimensions (see Section 3.3.4).
- The port resistance between contacts on a relatively large domain saturates for increasing distance between the ports (see Section 3.3.4).
- A windowing technique applied to the BEM significantly improves sparsity in the BEM model and computational efficiency of the BEM technique at the cost of only little accuracy (see Section 3.3.2).

Note (again) that the first two BEM properties refer to substrates of relatively large dimensions, which is a valid assumption for typical substrate modeling problems. Furthermore, these two properties hold for both lightly-doped and heavily-doped substrates.

The second BEM property implies that the 'horizontal' resistances between the contacts in the BEM model quickly become large for increasing distance between the

contacts, whereas the 'vertical' resistances towards the SUB reference become constant. In other words, the vertical resistances towards the reference node are relatively more important.

Based on these properties, the next chapter will show that the BEM and FEM methods may efficiently combine into a BEM/FEM method.

4

A Combined BEM/FEM Method

As mentioned in the previous two chapters, the Finite Element Method (but also the Finite Difference and Finite Volume Methods) and the Boundary Element Method both have their advantages and limitations. Generally speaking according to Section 2.4, the FEM (FDM, FVM) is accurate and flexible, but typically slow, whereas the BEM is typically faster, but less flexible and only accurate in more restricted situations. When applied to modern substrate technologies, however, the speed limitations in the FEM and the accuracy limitations in the BEM are typically emphasized.

In this context, much research from literature has aimed at increasing the speed of the FEM through efficient meshing and solution techniques which may even take the FEM to some distance from its convergence point. Our approach, however, aims at making the BEM more flexible by combining it with a FEM. In particular, the FEM can be applied in specific subdomains that require improved flexibility or accuracy, while the BEM can be applied in the remaining subdomain that does not have these requirements. The combination of the models then takes place on the boundary between the BEM and FEM subdomains.

Using the relevant properties of the FEM and the BEM from the previous chapter, this chapter will explore the possibilities for applying a combined BEM/FEM approach to the substrate modeling problem. The next two chapters will then present a practical evaluation of the BEM/FEM method.

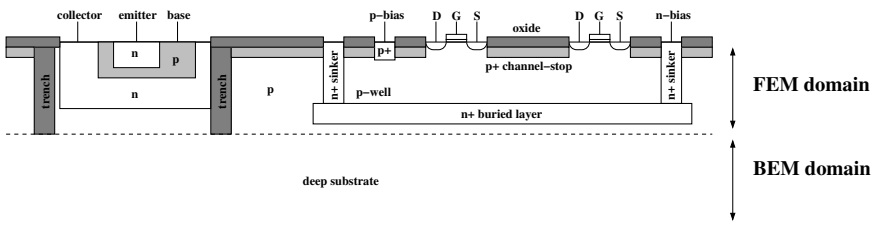

Figure 4.1: Separation of the substrate modeling problem into a FEM and a BEM domain.

4.1 General concept

Along the lines from Section 2.5, the substrate modeling problem can be separated into a BEM and a FEM part based on the doping profile. The separation will take place as follows. The layout-dependent doping patterns from Section 2.2 are typically located in the top few microns of the substrate. Under these circumstances, the top few microns of the substrate can be modeled with the FEM, while the underlying substrate (typically some $350\mu m$) can be handled by the BEM. This suggests an approach as indicated in Figure 4.1.

With the FEM and the BEM each applied to their own domains, the methods share the potential distribution along the common interface. Therefore, the whole of the interface will have to be discretized, such that a piecewise continuous approximation of the interface potential distribution is obtained. However, the typically piecewise linear basis function of the FEM and the typically piecewise constant basis function of the BEM approximate the interface potential distribution with different order and impose different requirements on the interface discretization. In particular, the FEM side of the interface is discretized into triangular elements, while the BEM side of the interface is discretized into BEM contact areas which may have arbitrary shape as long as their union completely covers the interface without overlap. The FEM nodes are then located at the corners of the FEM triangles, while the BEM nodes are located in the centers of gravity of the individual BEM contact areas. This is an inherent mismatch between the node locations of the FEM and the BEM.

Previous work [Now96, Now97] handled the mismatch between the BEM and FEM nodes through an energy conservation approach by introducing ideal tranformers between the BEM and FEM nodes. Here, however, we propose to choose the BEM contact division on the interface as the Voronoi dual of the FEM mesh [Sch02a]. In that case, the BEM and FEM nodes will coincide, allowing a straightforward combination of the BEM and FEM models through an incidence strategy.

In short, after [Pre85], the Voronoi dual is obtained from an arbitrary triangular shape by choosing a point halfway each of the branches in the triangle and placing a new line there, perpendicular to the branch. The new lines are extended, until they

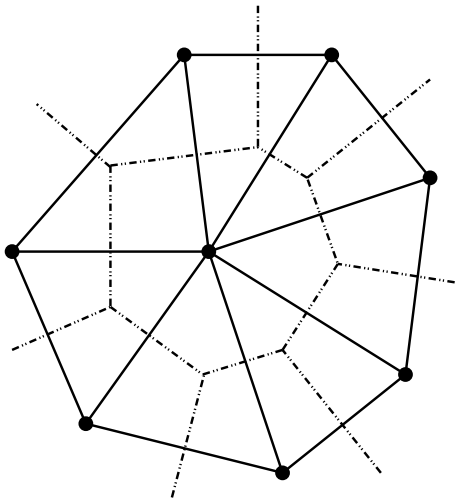

Figure 4.2: Voronoi dual (dash-dotted) of an arbitrary triangulation (solid)

reach each other in a point. Applying this approach to an arbitrary triangulation, the Voronoi dual is as shown in Figure 4.2. In general, with each node in the (FEM) triangulation, a Voronoi polygon (BEM contact area) is associated. If the triangulation is structured like in Figure 4.3, the Voronoi dual consists of hexagons.

In effect, the abovementioned approach results in a consistent covering of the interface with FEM nodes and corresponding BEM contact areas. The basic BEM mesh would now consist of at least one BEM panel per contact area, but a refined mesh with multiple BEM panels per contact area is also possible.

The fact that the BEM contact-division of the interface can be a Voronoi dual of some triangulation is related to the fact that the BEM is currently assumed to use a piecewise constant basis function. Alternatively, the BEM may also use a piecewise linear basis function, in which case the BEM contact areas are required to be triangular, and the BEM nodes would be at the corners of the triangles instead of in the center of gravity. In this case, it would be straightforward and valid to choose the BEM mesh equal to the FEM mesh. However, for reasons of simplicity, we have chosen in this thesis to work with the piecewise constant BEM. Conceptually, however, a piecewise linear BEM could also be used.

Also a mesh with rectangular elements, as would typically be the case in an FDM mesh, has a dual. This dual would typically consist of rectangular elements as shown in Figure 4.4. Compared to the general triangular/hexagonal case, the rectangular meshing approach allows a more straightforward implementation.

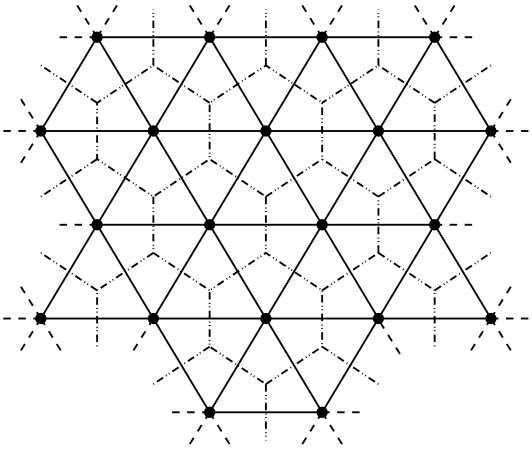

Figure 4.3: The triangular FEM mesh (solid) and its Voronoi dual defining hexagonal BEM contacts (dash-dotted).

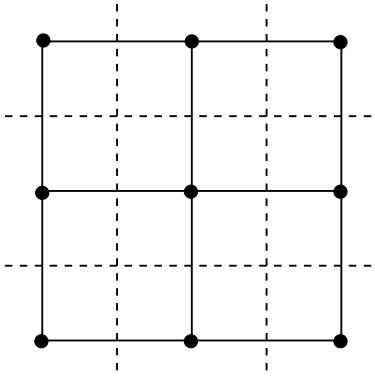

Figure 4.4: Solid line: FDM mesh. Dashed line: BEM contact division. The meshes are dual, such that the BEM and FDM nodes coincide.

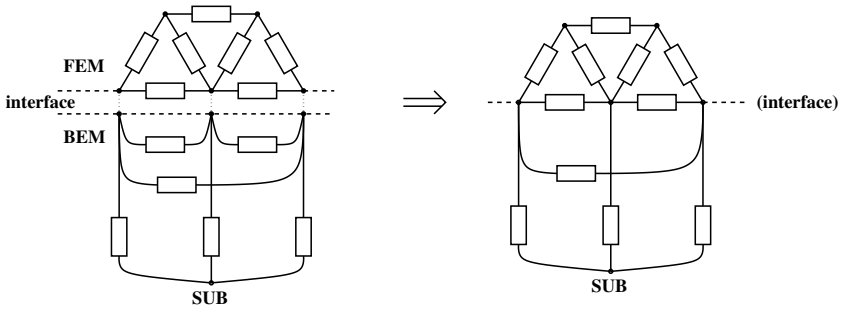

Figure 4.5: Combining the FEM and BEM models. Dual meshing ensures coinciding nodes along the interface, allowing a straightforward connection of the networks.

Once the BEM and FEM interface discretizations have been chosen properly with respect to each other, the nodes on either side of the interface coincide. Knowing that the interface has no physical meaning in the original modeling problem (Figure 4.1) and should therefore allow any current to pass through without a voltage drop, the potential distributions on either side of the interface should be equal. In terms of the BEM/FEM modeling approach, this implies that the potentials in the coinciding BEM and FEM nodes on the interface should be equal. The combination of the BEM and the FEM models is then straightforward: they can be connected through the coinciding nodes. Figure 4.5 shows this schematically.

4.2 Convergence

The convergence of the BEM/FEM method was already addressed in the work by Nowacka et al. [Now96, Now97], but the way in which the BEM/FEM method was used is different from the way we use it here. In particular, Nowacka applied the BEM/FEM method to capacitance modeling, whereas this thesis applies the BEM/FEM method to resistance modeling. Furthermore, for the combination between the BEM and FEM models along the BEM/FEM interface, Nowacka observed that the BEM and FEM nodes would not coincide and applied generalized transformers to mutually associate the non-coinciding BEM and FEM nodes, whereas this thesis applies dual meshing to ensure coinciding BEM and FEM interface nodes for straightforward combination of the models.

Despite these differences, the fundamental problem behind both capacitance and resistance modeling is a generalized Laplace equation. Therefore, the convergence proof for the BEM/FEM method developed by Nowacka will also be applicable to resistance extraction, but will be reformulated to our current approach using dual meshing instead of generalized transformers.

First, we will briefly address the convergence behaviour of the BEM and FEM methods. Then, we will formulate the convergence proof along the lines set out by Nowacka, but tailored to our situation.

4.2.1 General Remarks on Convergence

There are various possible formulations for convergence, mainly based on different formulations of the error. In the situations below, we assume the discretization error (i.e. the overall discretization of the problem) to be the main contributor to the overall error in the solution. This implies that we assume additional errors, for example resulting from the level of convergence in the numerical approximation of the solution and from finite-precision arithmetic in the computations, to be relatively unimportant.

The convergence of an approximate field solution towards the exact field solution can be quantified through the minimum-energy criterion according to Equation 3.9. In particular, the exact field solution for given boundary conditions has minimum energy. Any approximate field solution that (approximately) satisfies the given boundary conditions and that has minimum energy, is close to the exact solution, also in a pointwise sense. As such, minimization of energy will be used as an alternative formulation for point-wise convergence. In principle, the minimum energy criterion can be applied explicitly (i.e. minimize the energy in the discretized problem), or, equivalently, relative to the exact solution (i.e. minimize the energy in the difference-function between the exact solution and the discretized solution).

4.2.2 Convergence of the individual BEM and FEM methods

The FEM solution scheme basically ensures convergence by explicit minimization of the energy in the discretized problem [Sil96]. The convergence proof and more formal discussions of the approximative properties of the FEM can be found in literature, e.g. [Nor73, Str73, Zie83], and will not be discussed here. The rate of convergence in the FEM is dependent on the basis function. The simplest form of the FEM is the situation where it uses a piecewise linear (planar) basis function. That is, if h represents the mesh granularity, the basis function is $O(h)$. Under these circumstances, the convergence rate of the FEM under refinements in the mesh (i.e. $h \rightarrow 0$) can be found to be $O(h^2)$.

The BEM solution is based on a fundamental solution of the Laplace equation, so minimization of the energy and convergence are inherently ensured; contrary to the FEM, no explicit energy minimization is required here. The simplest form of the BEM is the situation where it uses a piecewise constant (i.e. $O(1)$) basis function. The rate of convergence in the BEM can then be found to be $O(h)$. The BEM also allows a piecewise linear basis function, in which case the BEM has a similar convergence rate as the FEM: $O(h^2)$. However, as already mentioned in Section 3.3.1, the BEM with piecewise linear basis function will not explicitly be considered in this thesis.

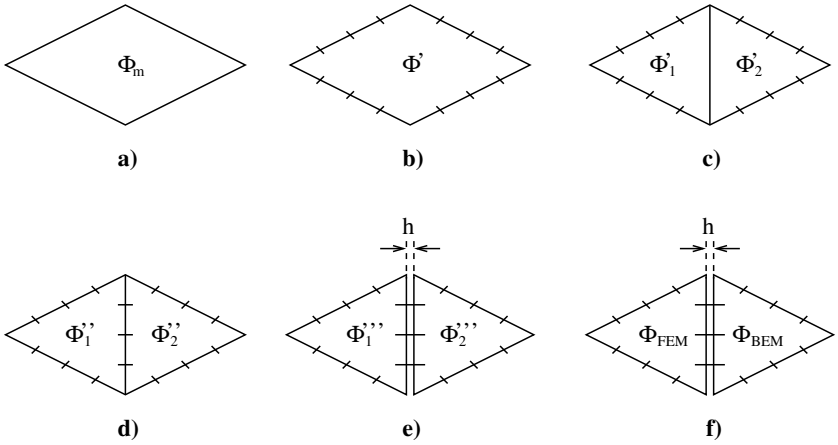

Figure 4.6: Consecutive steps in the convergence proof

4.2.3 Convergence of the combined BEM/FEM method

Here, we briefly reproduce the convergence proof from [Sch02b, Sch04a], which is similar to the convergence proof from [Now96, Now97], but tailored to the way we currently apply the BEM/FEM method. Knowing the work in [Now96, Now97] to be a theoretically consistent description of the proof for the BEM/FEM method, the theoretical details will not be repeated here. Practical support for the convergence of the BEM/FEM method will be provided through the convergence experiments in Chapter 5.

The steps in Figure 4.6 are used as a guide through the convergence proof. Note that the proof is not formal, according to the rules from traditional mathematics; it should merely be considered as an indication of how the formal proof can be constructed.

Figure 4.6a – The key behind the proof is the fact that a field (Φ_m) that satisfies the Laplace equation and that satisfies the continuous boundary conditions imposed on the domain will have minimum energy. This is the true, physical field.

Figure 4.6b – When the boundary conditions are discretized, the resulting field (Φ') will be such that its energy is minimal, consistent with the discretized boundary. Using the observation from Section 4.2.1 that fields that approximately satisfy the boundary conditions and that have energy close to the minimum are also close to the exact field in a point-wise sense, the approximate field Φ' can get arbitrarily close to the exact field Φ_m for refining discretization. In terms of the energy $\mathcal{E}$, this is formulated as $\mathcal{E}(\Phi') = \mathcal{E}(\Phi_m) \pm \varepsilon$ for some small ε.

Figure 4.6c – The domain is now divided with an interface. If the interface pre-

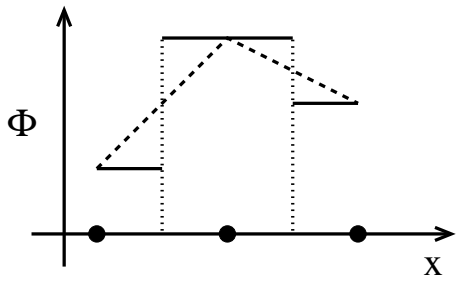

Figure 4.7: Potential approximations on the interface. The dots along the x-axis represent the coinciding BEM and FEM nodes. Legend: solid line = BEM approximation, dashed line = FEM approximation.

serves energy, there will be a continuous potential match along the interface. The energy will be such that $\mathcal{E}(\Phi_1') + \mathcal{E}(\Phi_2') = \mathcal{E}(\Phi')$.

Figure 4.6d – When the interface is discretized, there will be an exact, but *discrete* potential match along the interface, while the total field is still consistent with the boundary conditions on the outer boundary. However, because the discretized potential distribution on the interface *approximates* the *exact* distribution, we know that the total field necessarily has larger energy than the total field from Figure 4.6c. In terms of energy, this implies: $\mathcal{E}(\Phi_1'') + \mathcal{E}(\Phi_2'') = \mathcal{E}(\Phi') + \varepsilon$, for some other small ε.

Figure 4.6e – If the fields separated by the interface are solved with numerical methods that apply different basis functions, the total field will contain a discontinuity along the interface. To study this discontinuity with more detail, the interface is substituted with an h-thin layer, across which an interpolation between the basis functions is defined. Then, the energy contained in the h-thin layer is studied while $h \to 0$.

Figure 4.6f – As illustrated in Figure 4.6f, we assume a FEM field to the left and a BEM field to the right of the interface. Figure 4.7 schematically shows the situation on the interface, when choosing the simplest available basis functions for both methods (piecewise linear for the FEM and piecewise constant for the BEM), and assuming the BEM and FEM nodes to coincide (achieved by choosing the BEM contact-division on the interface as the Voronoi dual of the FEM interface mesh as explained in Section 4.1).

A difference field will now exist across the h-thin layer. This field is defined by a linear interpolation between Φ_{BEM} and Φ_{FEM} across the distance h. In particular, let x be a point on the interface and let k be a variable in the continuous interval $[0, h]$, the field Φ_i inside the h-thin layer is then defined as:

$$\Phi_i(k, x) = \Phi_{FEM}(x) + \frac{\Phi_{BEM}(x) - \Phi_{FEM}(x)}{h} \cdot k \tag{4.1}$$

In terms of the energy functional from Equation 3.9, we need the gradient of Φ_i. Knowing that Φ_i in the x-direction is $O(x)$ on the FEM-side and $O(1)$ on the BEM-side of the h-thin layer, the partial derivative of Φ_i with respect to x is $O(1)$. The partial derivative with respect to k is found from Equation 4.1. We remark that the granularities of the BEM and FEM interface discretizations and the thickness of the h-thin layer should be proportional. In this way, the k-derivative of the linear interpolation will be kept in proportion while $h \to 0$. Furthermore, knowing that the linear interpolation between the piecewise linear and piecewise constant basis functions of the FEM and the BEM is $O(h)$ (indeed linear), the k-derivative of the linear interpolation is then found to be:

$$\frac{\partial \Phi_i}{\partial k} = \frac{\Delta \Phi_i}{\Delta k} = \frac{\Phi_i|_{k=h} - \Phi_i|_{k=0}}{h - 0} = \frac{\Phi_{BEM} - \Phi_{FEM}}{h} = \frac{O(h)}{h} = O(1) \tag{4.2}$$

In short, we observe that:

$$\nabla \Phi_i = O(1) \tag{4.3}$$

Furthermore, the energy functional from Equation 3.9 requires the conductivity function to be known. Since the h-thin layer was introduced as part of constructing this proof, the conductivity function in the h-thin layer has no physical meaning and it will disappear as the thickness h of the h-thin layer approaches zero. Therefore, the conductivity function may be chosen appropriately within the context of this proof. A reasonable choice would be that in the direction of k the conductivity function is constant and in the direction of x (i.e. along the BEM/FEM interface) the conductivity function is equal to the average between the conductivity distributions on the FEM-side and the BEM-side of the BEM/FEM interface. Alternatively, the conductivity function in the direction of x may be equal to either of the conductivity distributions on the FEM-side or BEM-side of the BEM/FEM interface. In any of these presented cases, the conductivity function σ_i in the h-thin layer is independent of the thickness h of the layer itself.

Because the h-thin layer has the same thickness h everywhere, we can now rewrite the energy functional as follows:

$$\int_{V_i} \sigma_i ||\nabla \Phi_i||^2 dv \quad \Longrightarrow \quad \int_{S_i} \sigma_i ||\nabla \Phi_i||^2 \, h \, dS \tag{4.4}$$

where V_i is the volume of the h-thin interface layer and S_i is the surface of the interface layer. With Equation 4.3, and letting h approach zero while knowing that σ_i is independent of h and is therefore of $O(1)$ with respect to h, this last expression rewrites to

$$\lim_{h \to 0} \int_{S_i} O(1) \, h \, dS = 0 \tag{4.5}$$

which converges to 0 with a linear dependence on h, showing that the energy contained in the linear interpolation disappears as $O(h)$.

Concluding Observations

The BEM and FEM fields both minimize energy (by definition) and the energy contained in the linear interpolation disappears as $O(h)$, contributing a negligible part to the total energy. As such, for the given set of discretized boundary conditions on the outer boundary, the total energy in the BEM/FEM field from Figure 4.6f under strong enough refinement in the BEM/FEM interface discretization, can always become smaller than the energy in Figure 4.6d. Additionally, the energy in Figure 4.6f will never become smaller than the energy in Figure 4.6b, because Figure 4.6b gives the exact, minimum energy solution for this set of boundary conditions. The energy in Figure 4.6f is now enclosed between the energies in Figure 4.6d and Figure 4.6b, which shows the convergence. Based on the $O(h)$ convergence rate in the BEM with piecewise constant basis function, the $O(h)$ convergence rate in the interface, and the $O(h^2)$ convergence rate in the FEM with piecewise linear basis function, the overall convergence rate in the BEM/FEM method is $O(h)$.

Note that if the BEM would use a piecewise linear basis function, the BEM and FEM nodes along the interface would already coincide, and the duality between the BEM and FEM discretizations would not be necessary. Furthermore, the h-thin layer presented above would not be necessary, because the piecewise linear basis function and the coinciding nodes would produce a matching field along the interface; no discontinuity due to differing basis functions is introduced. The rate of convergence in the piecewise linear BEM is of the same order as the FEM: $O(h^2)$. The rate of convergence in the overall BEM/FEM method would then be $O(h^2)$ as well.

4.3 A sparse and reduced-order BEM/FEM approach

In the context of substrate resistance modeling, the BEM/FEM approach allows a sparse and reduced-order approach based on the physics.

4.3.1 A Sparse BEM

In general, the substrate modeling problem comprises a dense configuration of contact areas on top of a substrate of relatively large dimensions. Under these circumstances, local couplings between the contact areas concentrate near the surface of the substrate while the global couplings fan out through the deep substrate. The deep substrate therefore carries a relatively low current density.

Considering this situation in the context of BEM/FEM substrate modeling, the local couplings will be modeled by the FEM, while the global couplings will be modeled by the BEM. As a result, the BEM part in the overall BEM/FEM model may neglect local couplings because they play a negligible role compared to those captured by the

FEM. The global couplings, on the other hand, are captured by the 'SUB' node, which represents the deep substrate in the BEM (see Section 3.4). This implies that the BEM extraction may use a small window (recall Sections 3.3.2 and 3.4) without introducing large errors in the overall BEM/FEM model.

This last observation may now be more strongly supported by the saturation curves shown in Section 3.3.4 of the previous chapter. Lateral resistances go to infinity for increasing distance, while resistances towards the deep substrate become constant (recall also Section 3.4).

Furthermore, due to the relatively low current density in the deep substrate, relatively little current will pass through the BEM contacts on the BEM/FEM interface, such that the current density distribution on the BEM contacts can be taken into account with a coarse BEM mesh, without introducing large errors in the overall BEM/FEM model.

Based on these observations, the BEM part of the BEM/FEM extraction may use a small window and a coarse mesh without introducing large errors in the overall model. As a result, the BEM extraction is fast, and results in a sparse resistance network that is attached to the already sparse FEM network. The combined BEM/FEM model is then sparse, but still contains internal nodes which may be removed through some elimination procedure.

4.3.2 A Reduced-order 3DFEM

Consider Figure 4.8 which shows the equipotential lines and current density for two contacts attached to a highly conductive (channel-stop) layer on top of a significantly more resistive bulk. The current flow mainly takes place through the highly conductive layer. We observe that any vertical field components in this layer are most significant near the contact areas, whereas the field will be mainly horizontally oriented at some lateral distance from the contacts. In the areas where the field is horizontally oriented, the vertical field component will be negligible, causing an (almost) equipotential field distribution in the vertical direction.

This observation can be exploited when generating the 3DFEM discretization in such a way that the nodes in different layers are aligned vertically, as Figure 4.9 illustrates[1]. Basically, this vertical alignment can be achieved through a stacked horizontal 2DFEM mesh, where vertical resistances are placed between the 2DFEM layers in the stack. The result is a 3DFEM with prismatic elements, which is a reduced-order approach compared to a traditional 3DFEM with tetrahedral elements.

Similar to the traditional 3DFEM, the 3DFEM with prismatic elements defines a volume discretization, where the granularity of the mesh depends on the structures

[1] It can be argued that the discretization in Figure 4.9 resembles an FDM, but this makes no difference conceptually.

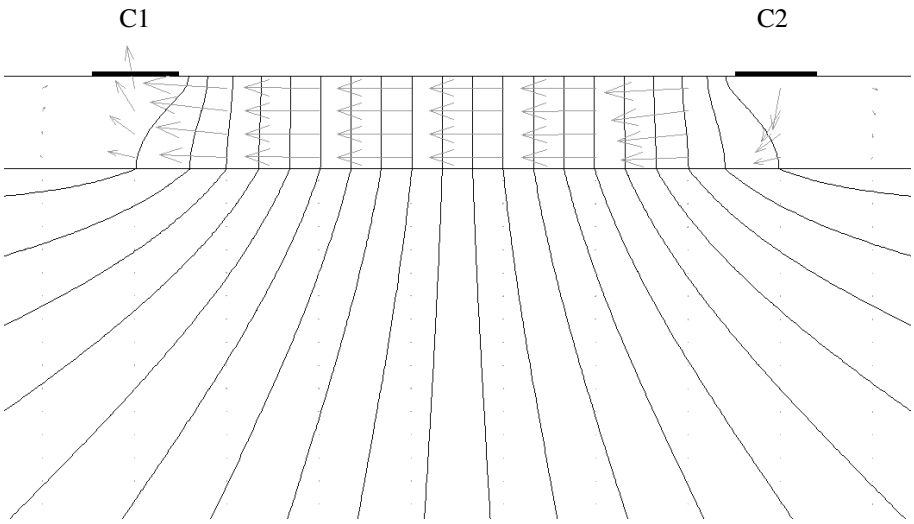

Figure 4.8: FEMLAB-generated plot of equipotential lines and current-density-arrows between contacts C1 and C2 attached to a highly conductive layer (1000 S/m) on top of a poorly conductive (10 S/m) bulk. The current mainly flows through the top layer; the current density in the bulk is comparatively so small that the arrows are hardly visible.

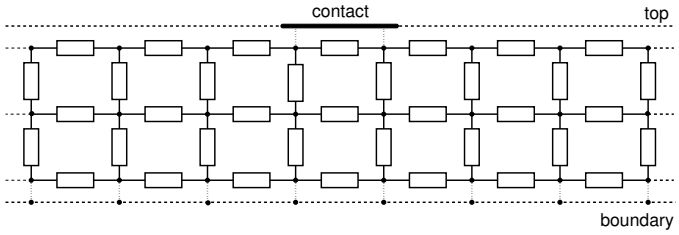

Figure 4.9: Side view of a 3DFEM with vertically aligned nodes.

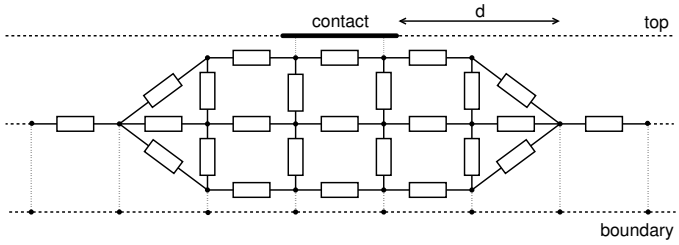

Figure 4.10: Side view of a 3DFEM around a contact area with node contractions beyond distance 'd'.

that are present in the domain. For the situation illustrated in Figure 4.1, for example, the prismatic discretization would have to be such that all structures in the indicated FEM domain are taken into account consistently.

However, in areas where the field distribution is equipotential in the vertical direction, as was illustrated for the highly conductive layer in Figure 4.8, the vertically stacked nodes will be equipotential, and they can be contracted into a single node without introducing any error. This is conceptually equivalent to [Gen96a]. Assuming the FEM mesh from Figure 4.9 to be applied for a highly conductive layer, the corresponding contraction approach is illustrated in Figure 4.10: at some (user-defined) distance 'd' from a contact area on the substrate, the vertically aligned nodes in the discretization may be contracted into a single node. This approach reduces the order of the 3DFEM with prismatic elements even further.

The presented 3DFEM approach with prismatic elements is generally applicable (it actually resembles a 3DFDM and has similar characteristics), but the node contractions are particularly useful for the channel-stop layer, which is a highly conductive layer at the top of the substrate on which contact areas are located.

4.4 Implementation

In this section, we briefly introduce the implementation of the combined BEM/FEM method in the SPACE layout-to-circuit extractor. Examples of how the new components in the implementation can be called are given in Appendix A. The implementation should still be considered a prototype that may undergo future optimizations and improvements, but the current prototype will prove to be sufficient for practical evaluation in Chapters 5 and 6.

4.4.1 Meshing

In order to simplify implementation, the discretization approach in SPACE is primarily based on rectangular elements. These elements are commonly known as 'tiles' [Mei92] of which the dimensions were originally only defined by the layout-features. In the layered FEM and BEM/FEM modeling approaches, however, the tile dimensions are additionally constrained by (user-defined) x and y size parameters, as illustrated in Figure 4.11. The tile division forms the basis both for the horizontal mesh in the FEM approach, and for the horizontal discretization on the BEM/FEM interface in the BEM/FEM approach. In particular, the edges of the tiles form the basis for the FEM (FDM) mesh, while the tiles themselves define BEM contact areas on the BEM/FEM interface. An example of the tile division (BEM contact division) for a simple layout is shown in Figure 5.13 in Chapter 5.

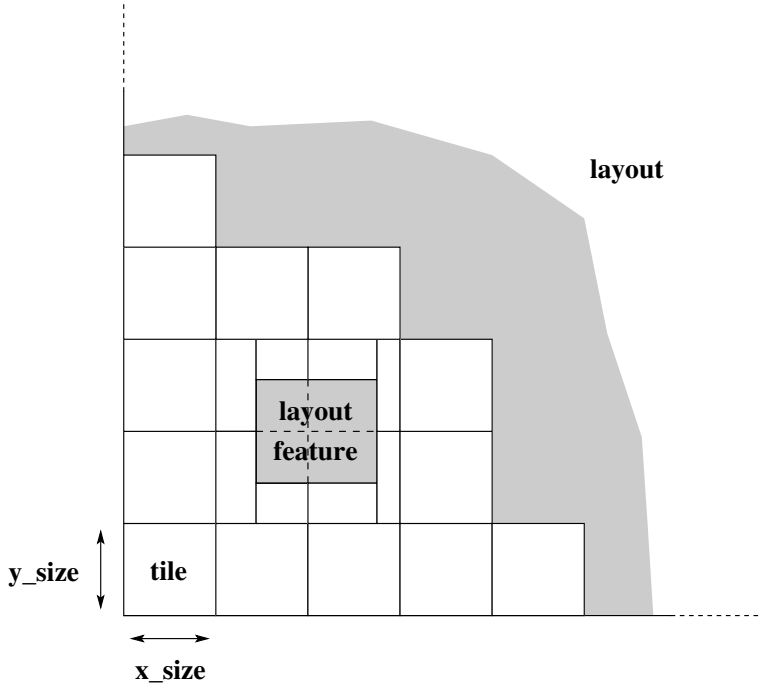

Figure 4.11: Basic discretization approach. Division of the layout into tiles with at most the dimensions of (user-defined) x and y size parameters. Layout features also have influence on the tile division, and may also contain tiles internally.

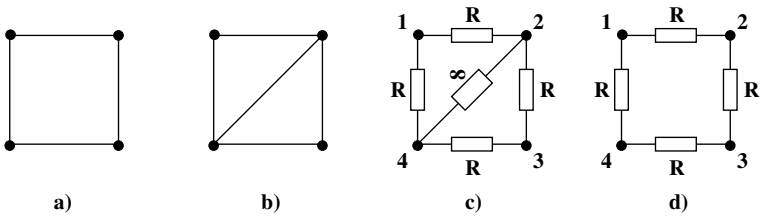

Figure 4.12: FEM meshing approach and resistor assignments for a square tile. a) tile, b) FEM mesh, c) node and resistor assignments, the diagonal resistor is infinite [Gen91], d) resulting network. The diagonal element in the FEM mesh may already be omitted during the meshing procedure because the associated diagonal resistor is infinite.

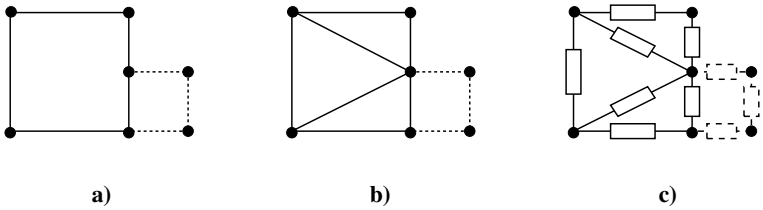

Figure 4.13: FEM meshing approach and resistor assignments for adjacent tiles of different size. a) tiles, b) FEM mesh, c) resulting network. For the side tile (dashed), the observations from Figure 4.12 are applied. For the main tile the diagonal resistors do not cross the full panel and are therefore *not* infinite [Gen91]; they are retained in the resulting network.

As suggested in Section 4.3.2, the FEM with vertically aligned nodes is constructed from a stacked 2DFEM approach with vertical resistances between the layers in the stack (more details about this follow in Section 4.4.2). A 2DFEM meshing approach applied to a single tile is illustrated in Figure 4.12. The diagonal element across this tile results in an infinite resistance [Gen91], and is therefore usually omitted during the generation of the 2DFEM mesh. However, in a situation where the diagonal elements would not cross the full tile, as illustrated in Figure 4.13, the resulting resistances are *not* infinite, and therefore the diagonal elements may not be omitted from the mesh. An example of the 2DFEM mesh for a simple layout is shown in Figure 5.14 in Chapter 5.

Whether any diagonal resistances in the FEM are infinite or not, we observe that each tile primarily has 4 FEM nodes at its corners. The main approach for handling the BEM/FEM interface is based on this situation. For a tile on the BEM/FEM interface, the approach is illustrated in Figure 4.14. Each interface tile has a FEM mesh along its edges and simultaneously defines a BEM contact. As a result, there are 4 FEM nodes at its corners and a BEM node in its center of gravity. The BEM and FEM nodes do not

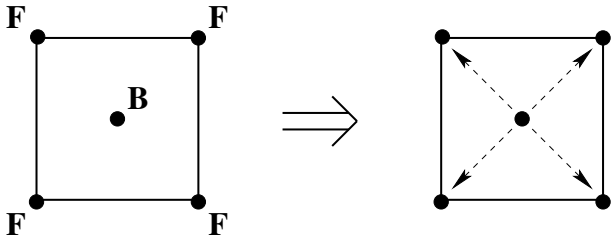

Figure 4.14: BEM/FEM interface tile. The tile has 4 FEM (FDM) nodes at its corners and a BEM node in its center of gravity. The BEM conductances attached to the BEM node are proportionally divided over the FEM nodes.

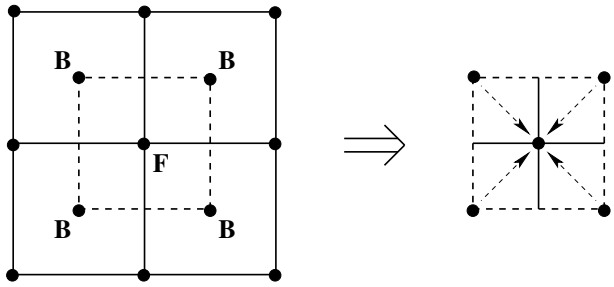

Figure 4.15: Expanded situation with multiple interface tiles, defining multiple BEM contacts around a FEM node. The division principle from Figure 4.14 associates a 'virtual' BEM contact (shown by dashed line) with the FEM node. This approximates the situation from Figure 4.4.

coincide, but the transition is made by dividing the BEM conductances proportionally over the 4 FEM nodes at the corners of the tile. In this way, the connection of the BEM network to the FEM network is straightforward.

Expanding this situation to multiple interface tiles, the approach looks as in Figure 4.15. In this case, proportional sections of the 4 BEM contacts are associated with the FEM node. This can be interpreted as a 'virtual' BEM contact that is associated with the FEM node. This is almost equivalent to the approach based on dual meshing, where the BEM/FEM nodes coincide as in Figure 4.4.

Based on these observations, the BEM/FEM results generated by SPACE are expected to converge in almost the same way as described by Section 4.2.3. Strictly speaking, however, the convergence proof from 4.2.3 is not valid for this situation, because the way in which the BEM and FEM models are combined is different here. Nevertheless, in Chapter 5 we will demonstrate the convergence in practice.

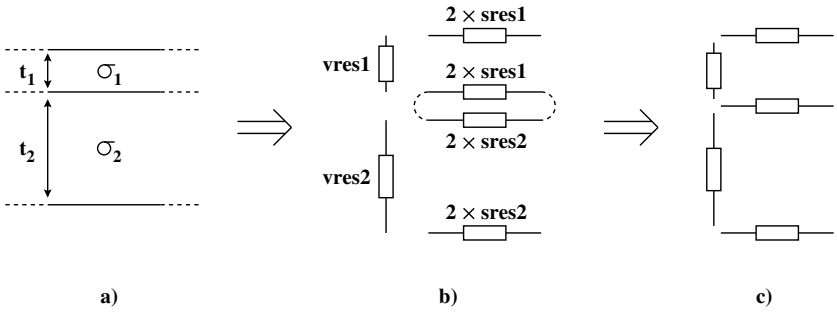

Figure 4.16: Construction of the layered FEM. Sheet resistances 'sres' in Ω/sq, vertical resistances 'vres' in $\Omega \cdot \mu m^2$

4.4.2 Layered FEM

As already mentioned previously in Section 4.3.2, the 3DFEM with vertically aligned nodes may be constructed from stacked 2DFEM layers connected by vertical resistances. The result is a layered FEM. The principle for generating this layered FEM is similar to the technique from [Gen88] (see also Section 3.2.2), but now expanded into the z-direction. This principle may be illustrated as follows.

Consider Figure 4.16a, where two substrate layers with conductivities σ_1 and σ_2 and thicknesses t_1 and t_2 are stacked on top of each other. For each layer, a sheet resistivity and a vertical resistivity are determined. Then, these resistivities are stacked as in Figure 4.16b; which reduces to 4.16c by placing the sheet resistivities along the interface between σ_1 and σ_2 in parallel.

A similar principle can be applied if a single resistive substrate layer would be divided into multiple FEM layers.

4.4.3 Node Contraction

The node contraction from Figure 4.10 can be implemented through a straightforward 'halo' principle around relevant areas in the layout, as indicated in Figure 4.17. The halo can be derived from the mask data that potentially induces relevant vertical currents in the FEM region. This mask data may be, for example, a via which defines an explicit contact from the interconnect to the substrate, or polysilicon interconnect which may have significant interaction with the substrate. The halo defines the area where the full FEM is applied, and can be grown over some user-defined distance, to allow control over the distance 'd' from Figure 4.10.

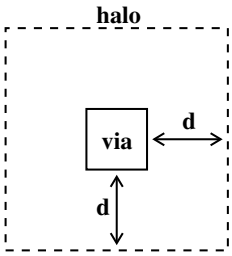

Figure 4.17: Halo around a via mask in the layout. The halo initially is an exact copy of the relevant mask (which does not necessarily have to be a via), and is subsequently grown over user-defined distance 'd'.

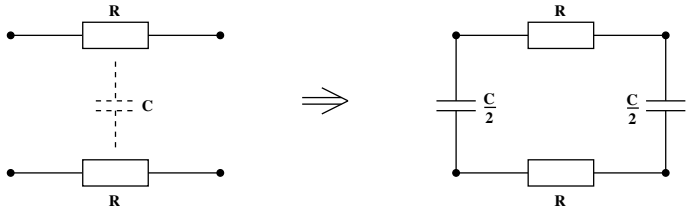

Figure 4.18: Distribution of junction capacitances in the FEM grid. Situation simplified to 2D for ease of illustration.

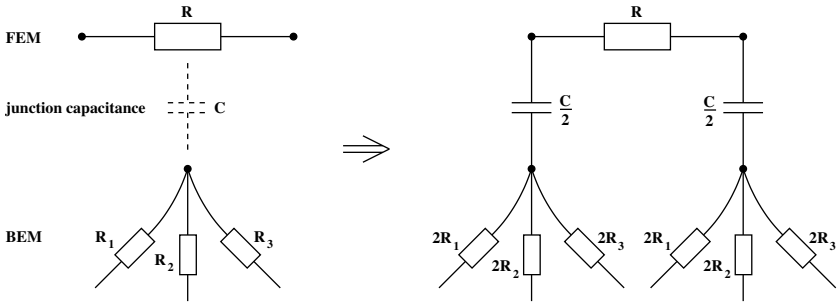

Figure 4.19: Distribution of a junction capacitance along the BEM/FEM interface. Notice also the distribution of the BEM resistances. Situation simplified to 2D for ease of illustration.

4.4.4 Junction Capacitances

As mentioned in Section 2.2.1 and illustrated in Figure 2.4, junctions in the substrate may be modeled by linear capacitances. For any thick n-area (like the n-well) the junction along its bottom and sidewalls has to be taken into account (for *thin* n-areas, the sidewalls may be neglected). The junction capacitances are inserted into the mesh based on the previously described tile-division of the layout. For the bottom of an n-area, the capacitances are related to the area of the tiles in the tile-division and proportionally divided over the corners of the tiles and attached to the FEM nodes there. If the junction is located in the FEM domain, the capacitance is inserted in the FEM grid as illustrated in Figure 4.18. The technique is similar for the sidewalls of the n-area, where capacitances may be divided proportionally over the FEM-layers in the layered FEM, based on the thickness of the layers. If the junction is located on the BEM/FEM interface, the capacitance is inserted between the BEM and FEM nodes, as illustrated in Figure 4.19. As shown in the figure, the proportional distribution of the junction capacitances is similar to the proportional distribution of the BEM resistances.

4.4.5 Elimination Method in Layered FEM and BEM/FEM

Both the layered FEM and the combined BEM/FEM networks contain terminal nodes and internal nodes, of which the internal nodes can be removed through some elimination procedure. This elimination procedure is currently based on straightforward node-elimination (Gaussian elimination) in the network.

Note, however, that the present implementation of the elimination procedure in SPACE is grossly inefficient for 3DFEM situations (both in the layered FEM and in the BEM/FEM case). This is because it was optimized for 2DFEM approaches in interconnect structures by using minimum-degree ordering of the nodes [Mei95] and exploiting the geometrical properties of the application [Gen96a]. These acceleration techniques are not (or only partly) applicable for large 3DFEM regions.

Thus, in the practical evaluation of the methods presented in Chapter 5 we will be unable to draw conclusions about absolute performance. Instead we will be satisfied with conclusions about the relative performance of different methods using the same suboptimal elimination method (solver). This is deemed sufficient for the practical evaluations in chapter 5. Future optimizations may be achieved by applying an iterative solver [Saa03], but a study into iterative methods has not been performed in the context of this thesis.

4.5 Summary and Concluding Remarks

In this chapter, we introduced a BEM/FEM method as a hybrid combination of the BEM and FEM methods. We have addressed the separation of the substrate modeling problem domain into a FEM and a BEM part. The FEM would be applied in the top few microns of the substrate, capturing local couplings, while the BEM would be applied in the deep substrate, capturing global couplings.

Subsequently, we discussed appropriate discretization on the interface between the BEM and FEM subdomains, such that the resulting models could be attached in a straightforward way. Furthermore, we showed convergence of the BEM/FEM method, based on the interface discretization.

With respect to the substrate modeling problem including localized doping patterns, the FEM is typically 'slow and accurate', while the BEM is typically 'fast and not accurate enough'. A hybrid combination of these two methods that inherits the unfortunate characteristics of both methods (i.e. a combined method that is 'slow and not accurate enough'), has no value in a practical context. However, we have shown that the combination of BEM and FEM can actually yield a method that is fast and accurate. In particular, we introduced two straightforward, physics-inspired techniques for sparsification of the BEM network and order-reduction in the FEM network. The result is a combined BEM/FEM approach that captures global couplings with a coarse (i.e. fast and sparse) BEM, and that captures local couplings with a FEM that contracts equipotential nodes wherever it is allowed.

Conceptually, the BEM/FEM method should now be faster than a full FEM method, because the full FEM requires to define layers in the deep substrate, whereas the BEM does not. The top layers can essentially be captured with the same discretization as the full FEM, but the deep substrate can be captured with a fast and sparse BEM; the resulting model is inherently smaller and retains sparsity. Furthermore, the performance of the BEM/FEM method can be improved by the contraction of equipotential nodes in the FEM network.

Finally, we addressed the implementation of the BEM/FEM method. In the next two chapters, we will study the behaviour of the BEM/FEM method in practice.

5

Basic Behaviour of the BEM/FEM Method

As proposed in the previous chapter, the BEM/FEM method is a hybrid combination of the BEM and FEM methods. As mentioned in Chapters 2 and 3, the FEM and the BEM each have their own typical characteristics: the FEM is accurate and flexible, but slow, whereas the BEM is typically much faster, but has less flexibility and may therefore not always be accurate. Through a study of the practical characteristics of the BEM/FEM method when applied to the substrate modeling problem, the present and the following chapter will show that the BEM/FEM method inherits good characteristics from its consituting parts. That is, the BEM/FEM method inherits accuracy and flexibility from the FEM, and speed from the BEM, resulting in a practically applicable method with an improved speed-accuracy trade-off compared to the BEM and the FEM individually.

Furthermore, the present and the following chapter will show that consistency between the model and the physical situation is more important than the highest accuracy in the model. This observation may be used for the benefit of efficiency in the BEM/FEM technique.

In this context, the present chapter will evaluate the basic behaviour of the BEM/FEM method by applying it to a simple situation. In the following chapter, the BEM/FEM method will be applied to a more realistic situation.

5.1 Objectives and Overview

In this chapter, the practical aspects of the usage of the BEM/FEM method will be evaluated. The main objective of this practical evaluation will be to test the implementation of the BEM/FEM method and to confirm the practical value of the BEM/FEM method with respect to performance, accuracy and consistency with the physical situation. This objective will be approached through a number of separate sub-objectives:

- Test the implementation and confirm convergence of the layered FEM, which forms a prerequisite for correct behaviour of the overall implementation of the BEM/FEM method.
- Test the implementation and confirm convergence of the BEM/FEM method.
- Confirm consistency of the resulting models with the physical situation.
- Identify behavioral characteristics of the BEM/FEM method with respect to accuracy, speed and memory usage.
- Identify relative characteristics of the BEM/FEM method with respect to the individual BEM and FEM methods.

This chapter will progress as follows. Section 5.2 will introduce the doping profile to be used with the case studies in this (and the following) chapter, while Section 5.3 explains how the BEM, FEM and BEM/FEM modeling techniques may be applied to the doping profile. Then, Section 5.4 will introduce the basic layout structure considered in this chapter: a 9-section ring-oscillator. After these preliminaries, the subsequent sections will further progress as follows.

Concentrating first on the simplified situation where only the channel-stop pattern in the ring-oscillator is present, Sections 5.5 and 5.6 will address the 5 objectives stated above. In particular, the implementation of the layered FEM and BEM/FEM methods will be tested, and convergence will be confirmed. Furthermore, each method will be compared to an independent 3DFEM reference, with the goal of showing that the layered FEM and BEM/FEM methods can reach good accuracy and that the resulting models are consistent with the physical situation. Additionally, the efficiency and accuracy of the BEM/FEM method will be studied when using a coarse and sparse BEM and when using node contractions in the layered FEM. Finally, the consistency between the 3DFEM reference (representing the physical situation), the BEM/FEM method and a plain BEM method will be studied with the goal of showing that the BEM/FEM method yields a model that is consistent with the physical situation in the presented example, whereas a plain BEM method does not.

The observations and conclusions will eventually be summarized in Section 5.7.

5.2 Substrate Doping Profile

Any substrate modeling problem starts from the physical situation in the substrate which is defined by the substrate doping profile. In this section we will introduce the substrate doping profile which will be used in the extractions and simulations that follow in later sections. Based on this doping profile, the FEM, BEM, and BEM/FEM modeling techniques can be applied, as will be illustrated in Section 5.3.

The case-studies presented in this chapter use an example technology that is loosely based on a standard CMOS technology in a lightly doped substrate provided by the dedicated substrate modeling tool SubstrateStorm [MSu]. In this context, consider Figures 5.1, 5.2 and 5.3 which were previously shown in Chapter 3, particularly Sections 2.2.2 and 2.5, but repeated here for convenience. Basically, Figure 5.2 defines relevant cross sections through Figure 5.1 which are each associated with a doping concentration curve in Figure 5.3. Depth zero was chosen on top of the oxide indicated in Figure 5.1, which explains the depth offset at the beginning of the curves.

We may distinguish 4 basic 'classes' in the cross sections:

- contacts: cross sections 1 and 8;
- transistors: cross sections 2, 3, 6 and 7;
- channel-stop: cross section 4;
- n-well: cross section 5.

These classes will be discussed in more detail below, but before doing so, it is important to note that due to the lack of measurements on an actual circuit in an actual substrate we took some liberty for simplifications in the modeling approach. Even though these simplifications may not be entirely valid with respect to the physical situation, they do allow the discussion below to be more straightforward, while the results are expected to expand to the full, unsimplified situation in a relatively straightforward way. Therefore, the behaviour of the BEM/FEM method is presently evaluated in an easily comprehensible situation, while the comparison of the BEM/FEM results to actual measurements is reserved for future research.

Considering the doping concentration curves for the contacts (i.e. curves 1 and 8) in Figure 5.3, the doping concentration near the surface is very high ($\sim 10^{20}$, comparable to 100,000 S/m), and therefore the contacts are assumed ideal. It may be observed from Figure 5.3 that doping concentration curve '1' for the p-bias contact does not join doping concentration curve '4' for the channel-stop layer, but actually joins curves '2' and '3' representing the lower doping concentration immediately underneath the transistor. As such, the channel-stop layer is interrupted by the presence of a contact, but for historical reasons we have assumed any contact areas to be immediately connected to the channel-stop layer, as illustrated in Figure 5.2. This may

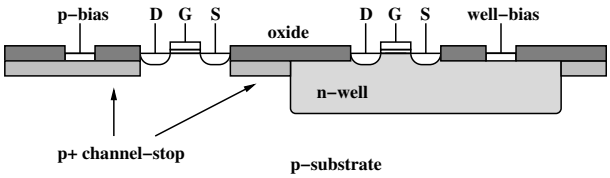

Figure 5.1: Repetition of Figure 2.5. Schematic representation of common doping patterns in the substrate. The channel-stop layer is interrupted by transistors and wells.

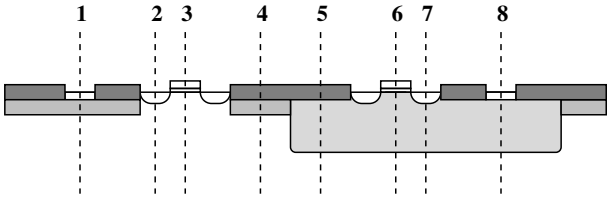

Figure 5.2: Relevant cross sections through the configuration in Figure 5.1, which are associated with the doping profile in Figure 5.3.

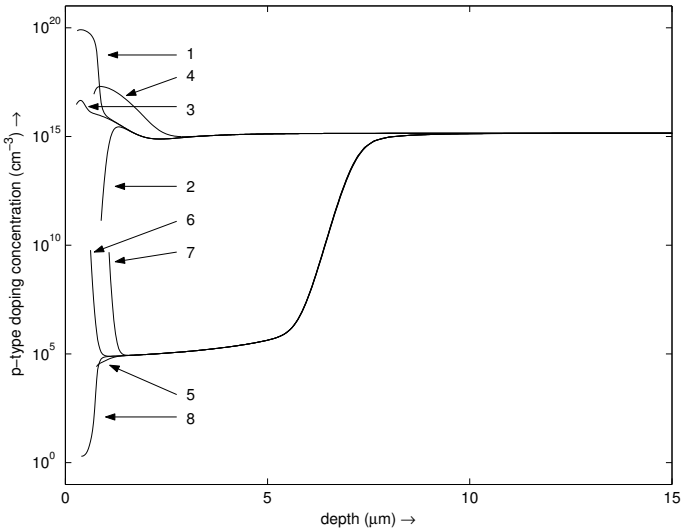

Figure 5.3: Example of a full doping profile, drawn in a single figure; loosely based on a standard CMOS technology in a lightly doped substrate provided by SubstrateStorm. Indicated numbers refer to cross-sections from Figure 5.2.

introduce some small modeling errors with respect to the physical situation, but since the contact areas are typically very small, these errors are considered negligible. In general, this modeling detail is of no consequence to the study of the behaviour of the BEM/FEM method itself.

Because the transistors require the full semiconductor equations to be taken into account, which is not the topic of this thesis, we will use the traditional way of capturing the active devices by compact behavioural models (like BSIM3 [MB3]) with suitable substrate connections underneath them. However, because the transistor models are only valid in a very small domain around the transistors, the doping profiles *underneath* the transistors should be taken into account by the modeling technique applied to the remainder of the substrate. Thus, for the transistor *inside* the n-well, the doping profile underneath the transistor is assumed to be the same as in the remainder of the n-well. For the transistor *outside* the n-well, the doping profile underneath the transistor is assumed to have a constant doping concentration, equal to the doping concentration in the deep substrate; this implements the interruption of the channel-stop layer.

Under the circumstances presented in the previous 2 paragraphs, the substrate modeling problem is reduced to taking into account the channel-stop layer and the n-well. Both curves have approximately the same depth offset in Figure 5.3 (0.70 μm for the channel-stopper, 0.77 μm for the n-well). The offsets are removed, such that both curves now start at zero depth; the difference of 0.07 μm between the offsets is ignored. Consequently, the channel-stop and the n-well are the only layout dependencies of the doping profile to be modelled, as described below.

Figure 5.4 focuses on the p-type doping concentration near the surface of the substrate for the channel-stop situation. We observe the high doping concentration near the surface; this is the channel-stop layer. The deeper regions have a significantly lower doping concentration. The peak doping concentration in the channel-stop layer is about $2 \cdot 10^{17} cm^{-3}$, which corresponds to a conductivity of approximately 1000 S/m, whereas the deep substrate has a doping concentration of about $1.4 \cdot 10^{15} cm^{-3}$, which represents a conductivity of approximately 10 S/m [Nea92].

Figure 5.5 shows both the n-type and p-type doping concentrations for the n-well. In the extractions and simulations in the following sections, the n-well will be considered an ideally conducting domain (contrary to Figure 2.4). Since the n-well in this example doping profile has an n-type doping concentration of about $4 \cdot 10^{15} cm^{-3}$, which represents a conductivity of about 100 S/m, while the n-well is separated from the remainder of the substrate by its junction capacitance, this choice may be justified. Indeed, as Section 6.3.1 in the next chapter will show, neglecting the well-resistance is a valid simplification that does not introduce inconsistencies in the presented examples. The junction along the bottom and sidewalls of the n-well will typically be taken into account by distributed capacitances (recall Section 4.4.4).

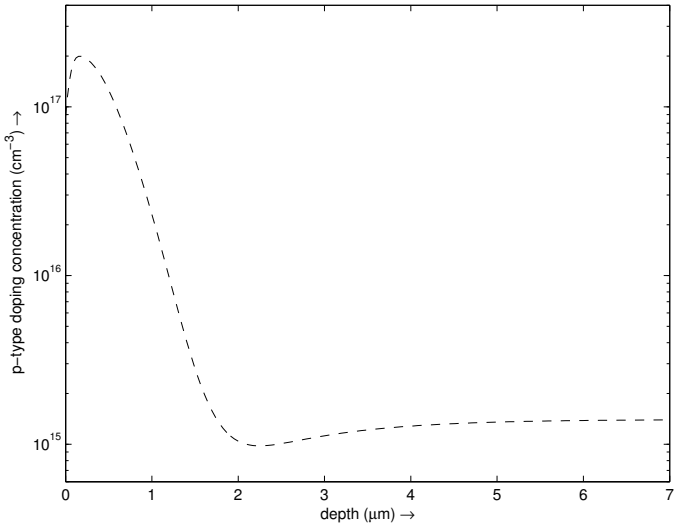

Figure 5.4: Channel-stop doping profile in a typical p-substrate. Peak conductivity: approx. 1000 S/m, deep substrate conductivity: approx. 10 S/m. The doping profile actually extends to the entire thickness of the wafer (in this case 375 μm).

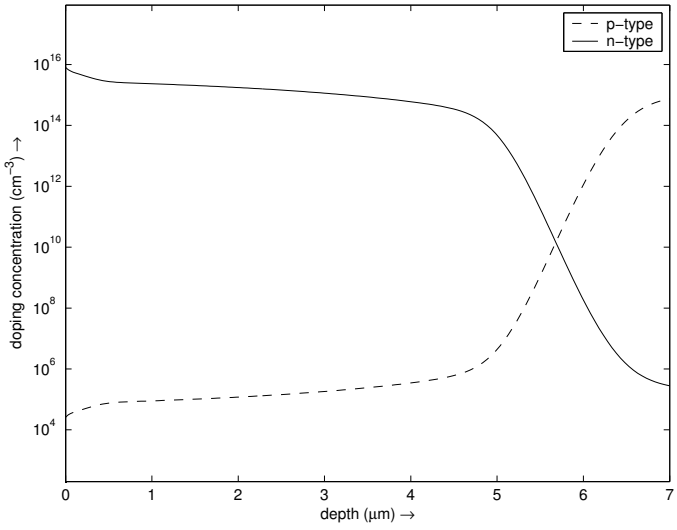

Figure 5.5: Doping profile for the n-well, embedded in a lightly doped substrate. Repetition of Figure 2.3.

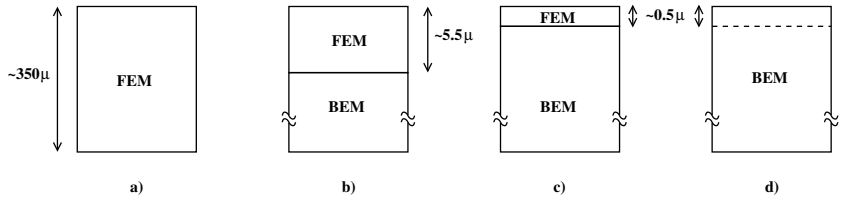

Figure 5.6: a) full FEM, b) regular BEM/FEM, c) reduced BEM/FEM, d) BEM (dashed line indicates boundary in double-layer approach)

5.3 Modeling Approaches

Based on the doping profile, the FEM, BEM/FEM and BEM modeling techniques may now be applied. Based on the previous chapters from this thesis, we can choose 4 basic modeling approaches as schematically shown in Figure 5.6. To determine the value of BEM/FEM as a new modeling method, we will consider and mutually compare all 4 modeling approaches.

5.3.1 FEM Approach

The FEM is applied to the whole substrate as in Figure 5.6a. The doping profiles from Figures 5.4 and 5.5 are approximated through the layered FEM approach with vertically aligned nodes. Basically, this requires to divide the substrate into a number of layers with constant doping concentrations, or associated conductivities, such that they form a piecewise constant approximation of the doping pattern.

Typically, we are looking for an approximation that is sufficiently accurate, while using as few layers as possible. However, to determine an 'optimal' layer division requires a separate study that was not explicitly performed. Therefore, while recognizing that it might not be optimal, the conductive layer division is based upon the division provided by the FDM-based dedicated substrate modeling tool SubstrateStorm [MSu]. The approximation of the doping profile is given in Figure 5.7, the accompanying layer division is shown in Table 5.1. The division takes into account the whole doping pattern including the n-well, which is why layers 5 and 6 are thin, to properly take into account the well-bottom. The deep substrate is handled with 4 layers of quickly increasing thickness towards the deep substrate. From the data in Table 5.1, the FEM division with vertically aligned nodes can be determined according to the approach from Figure 4.16. In this way, from the 10 conductive layers in Table 5.1, an 11-layer FEM is obtained (see also Appendix A.2.1).

Underneath the transistors, where no channel-stop layer is present, the thickness of the conductive layers is equivalent to those proposed in Table 5.1, but the conductivity of the top layers is 10 S/m instead of the $\sim$1000 S/m for the top two layers repre-

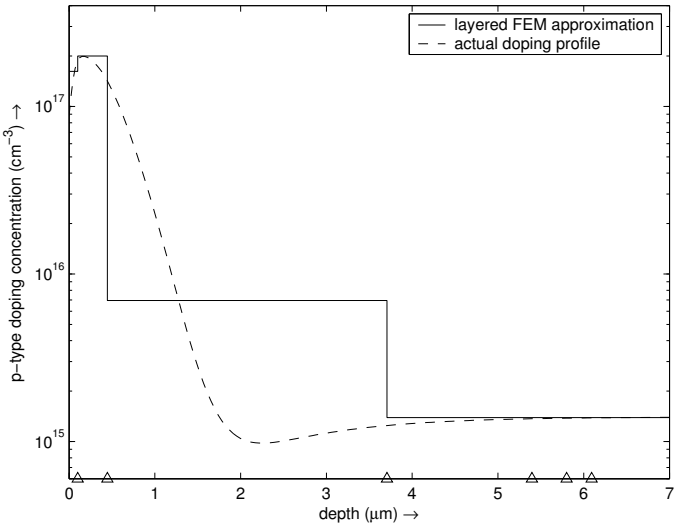

Figure 5.7: Piecewise constant approximation of the doping profile in Figure 5.4. The triangular tick marks along the horizontal axis are transition points between the layers defined in Table 5.1.

Table 5.1: Conductive layer division for Figure 5.7. Note that the mapping from doping concentrations to conductivities is slightly non-linear [Nea92]. Original wafer thickness is 375 μm, but total cumulative thickness of layer stack is 374.533 μm due to round-off errors.

layer	thickness (μm)	approximate p-type doping concentration (cm^{-3})	approximate conductivity (S/m)
1	0.100	1.62e17	790
2	0.346	2.00e17	975
3	3.26	6.95e15	50
4	1.69	1.39e15	10
5	0.404	1.39e15	10
6	0.293	1.39e15	10
7	5.54	1.39e15	10
8	20.2	1.39e15	10
9	73.7	1.39e15	10
10	269	1.39e15	10

senting the channel-stop layer and the 50 S/m for the area immediately underneath the channel-stop layer. In effect, this implements the interruption of the channel-stop layer due to the presence of a transistor.

Based on the division from Table 5.1, the bottom of the n-well can be chosen underneath layer 4 (at 5.4 μm depth), layer 5 (at 5.8 μm depth), or layer 6 (at 6.1 μm depth). These depth-samples are over a relatively small interval and, considering Figure 5.5, any of these depths would be feasible, but, considering the overall approximations applied in the modeling approach, the *exact* depth of the well will probably not be critical in obtaining an appropriate model. Therefore, the well-bottom is assumed to be immediately underneath layer 4, which will also prove to be convenient for the BEM/FEM method in the next section, because as few as possible FEM layers are then applied in the BEM/FEM context. Under the present circumstances, the junction capacitance at the bottom of the n-well couples across layer 5.

As mentioned in Section 5.2, the n-well itself is considered an ideally conducting domain. Therefore, no FEM layers will be defined inside the n-well. The doping profile *underneath* the n-well is captured by layers 6–10 from Table 5.1.

5.3.2 Regular BEM/FEM Approach

The BEM/FEM approach is equivalent to the FEM approach, except that the deep substrate is handled with the BEM instead of the FEM. Therefore, the BEM/FEM approach may be derived from the FEM approach in a straightforward way. In the BEM/FEM approach, the FEM is typically applied up to the depth of the deepest doping pattern present. In the present situation, the deepest doping pattern is the n-well. As presented in the previous section, the bottom of the n-well is located underneath layer 4 from Table 5.1. Thus the interface between the BEM and FEM domains is defined immediately underneath layer 4 (see also Appendix A.2.2), as schematically indicated in Figure 5.6b. The junction capacitance at the bottom of the n-well then immediately couples to the BEM. In this way, layers 5–10 from Table 5.1 are substituted with the BEM. As a result, the original 11-layer FEM approach is now captured by a BEM / 5-layer FEM approach.

The present approach defines the interface towards the BEM immediately underneath layer 4, such that as few FEM layers as possible are applied and the well junction-capacitance couples immediately to the BEM. However, considering the layered FEM approach presented in the previous section, where the junction capacitance was defined across layer 5, it can be argued that the interface towards the BEM should actually be defined underneath layer 5. Nevertheless, since layer 5 is thin, its presence, or absence, is not expected to play a significant role in the resulting BEM/FEM models, while its presence would increase the computation time. Under these considerations, layer 5 was omitted from the BEM/FEM approach.

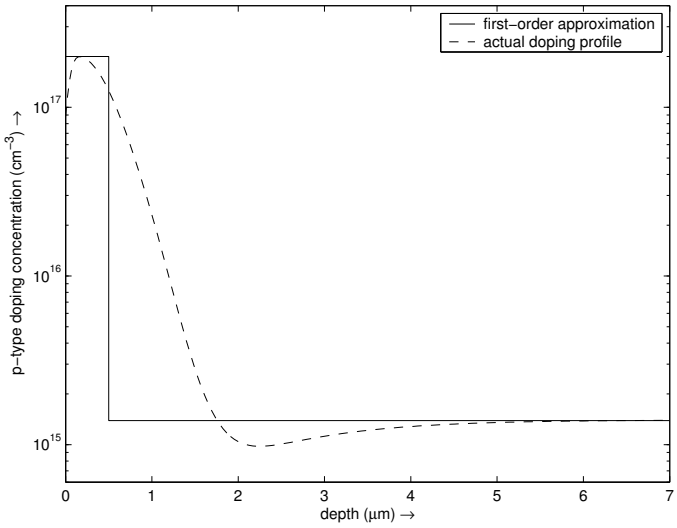

Figure 5.8: Low-order approximation of the doping profile in Figure 5.4. Peak conductivity: approx. 1000 S/m, deep substrate conductivity: approx. 10 S/m.

5.3.3 Reduced BEM/FEM Approach

For simplified structures where there is no n-well present (or when the effect of the n-well is approximated with lower order to reduce extraction time), the deepest doping pattern is the channel-stop layer, which is 0.5 μm thick. In this case, the FEM is only applied to the top 0.5 μm of the substrate, as indicated in Figure 5.6c. This representation will mainly be applied in the simplified simulations from Section 5.6.

In this case, the approximation of the doping profile is as indicated in Figure 5.8. Basically, this introduces a coarse piecewise continuous approximation of the doping profile, consisting of only 2 layers: the channel-stop layer and the deep substrate. The approximation is based on the peak doping concentration in the channel-stop layer, and the typical doping concentration in the bulk. From the doping profile, these concentrations are known to be $2 \cdot 10^{17} cm^{-3}$ and $1.4 \cdot 10^{15} cm^{-3}$, respectively. To complete the definition of the 2-layer piecewise continuous approximation, the thickness of the channel-stop layer needs to be determined. This can be done by integrating the doping concentration in the actual doping profile over the thickness of the substrate and requiring that the piecewise constant approximation of the doping profile integrates to the same value. The thickness of the channel-stop layer in the approximation can then be determined by fitting. In this way, the thickness of the channel-stop layer is found to be approximately 0.5 μm. This value actually confirms a common rule-of-thumb for the thickness of the channel-stop layer.

5.3.4 BEM Approach

The situation in Figure 5.6d represents the plain BEM approach. Even though the method allows, to a limited extent, to take into account n-wells and n-diffusions (see Figure 2.11 in Section 2.3.4), it is not capable of consistently taking into account general layout-dependent doping patterns. For example, with respect to the channel-stop layer, the uniformly layered approach by the BEM assumes the channel-stop layer to be present either *nowhere*, or *everywhere*. In the former case, the BEM assumes a uniformly doped substrate with a conductivity of 10 S/m. In the latter case, the BEM uses a double-layer approach where the channel-stop layer is assumed to be a uniformly doped layer, 0.5 μm thick, with a conductivity of 1000 S/m on top of a uniformly doped substrate with a conductivity of 10 S/m (see also Appendix A.3).

5.3.5 Preliminary Comparison Between the Methods

To get a global idea about the practical implications of the ways in which the different modeling approaches presented above take into account a doping pattern like the channel-stop layer, let us first consider a simple situation and compare the results from the different methods to an independent reference. Here, only the situation and the results will be posed. Deeper analysis of more advanced situations will follow in subsequent sections.

Consider Figure 5.9, which shows two contacts, each surrounded by an n-well guard ring which interrupts the channel-stop layer. When applying the plain BEM, it will model the n-wells as being outside of the substrate (recall Section 2.3.4 and Figure 2.11) and then model the channel-stop layer as being either absent (single-layer BEM) or present everywhere (double-layer BEM). Clearly, the plain BEM approach cannot consistently capture the modeling problem from Figure 5.9. The BEM/FEM and layered FEM approaches, however, *can* consistently capture the modeling problem.

Applying the single-layer BEM, double-layer BEM, BEM/FEM and layered FEM approaches to the modeling problem, we eventually obtain 4 networks which may be compared to a reference calculated by the independent 3DFEM modeling tool FEMLAB [MFe]. Even though the networks contain a few capacitances (resulting from the n-well junction capacitance towards the substrate), the resistive part of the networks is currently the most relevant part to study, because it is most directly affected by the 3 resistance-oriented modeling methods.

For the single-layer BEM, double-layer BEM and BEM/FEM methods, the resulting resistance networks contain 3 nodes: C1, C2 and the reference node SUB. Only by eliminating the SUB node from these networks, can they be compared to the layered FEM and to FEMLAB which do not contain the SUB node (see also Section 5.4.3). The results can be found in Table 5.2.

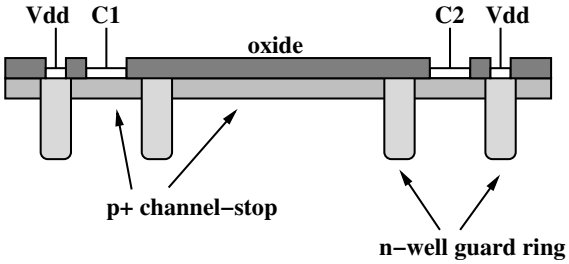

Figure 5.9: Two contacts, each surrounded by an n-well guard ring which interrupts the channel-stop layer. Distance between the terminals: $34\mu m$. Terminal size: $1\mu m \times 1\mu m$

Table 5.2: Resistance values in $k\Omega$ for the situation of Figure 5.9. R_p on the bottom row is the resistance between C1 and C2 after elimination of the SUB node. The BEM/FEM and layered FEM results are close to the independent FEMLAB result, whereas the single- and double-layer BEM methods are not.

	1-layer BEM	2-layer BEM	BEM/FEM	layered FEM	FEMLAB
$R(C1, C2)$	756.7	4.387	5923	82.89	97.56
$R(C1, SUB)$	19.19	1.509	43.40	n/a	n/a
$R(C2, SUB)$	19.19	1.509	43.40	n/a	n/a
$R_p(C1, C2)$	36.53	1.788	85.55	82.89	97.56

We observe that the BEM and BEM/FEM methods produce networks that are not comparable to each other, while only the BEM/FEM method compares well to the layered FEM method and to FEMLAB.

The remaining differences between the BEM/FEM, layered FEM and FEMLAB results are caused by a difference in convergence. The BEM/FEM and layered FEM methods used nominal extraction settings, while FEMLAB used a fine mesh due to the relatively small structures in the layout (thin channel-stop layer; narrow guard ring closely around contact area).

Knowing that the doping patterns in the layout may sometimes be specifically designed to avoid crosstalk, it is important that the doping patterns are taken into account consistently. From the results in Table 5.2, we conclude that only the BEM/FEM and FEM methods are actually capable of generating a model that is consistent with the physical situation represented by the 3DFEM reference.

5.4 Experiment Setup – Ring Oscillator

Below, we will address the setup of the experiments with respect to the layout under consideration, the applied modeling tools and the preliminaries for a valid comparison between the methods.

5.4.1 Layout

The case-studies presented in this chapter are all based on the same basic layout. The considered layout is the 9-section ring oscillator from Figure 5.10. The unit-length λ in this layout is 0.03125 μm (i.e. 1/32 μm). The inner ring consists of polysilicon interconnect which carries the actual oscillating signal. The outer ring is a ground line that is situated in the first metal layer. The outer ring is also connected to the substrate through a via halfway along the left edge. All transistors have 250 nm gatelength. The transistors connected to the outer ring are n-channel enhancement MOSFETs; they have a direct connection to the substrate immediately underneath their gate, but also underneath their source and drain junctions. The transistors connected to the inner supply line, V_{dd}, are p-channel enhancement MOSFETs and they are embedded in an n-well. These transistors have a connection to the substrate through the junction capacitance between the n-well and the p-substrate.

The layout uses the substrate technology from Section 5.2. A channel-stop pattern as shown in Figure 5.11 is then also defined by the layout. As already mentioned previously, the channel-stop pattern typically has a thickness of 0.5 μm, and a conductivity of 1000 S/m.

5.4.2 Applied Modeling Tools

The following modeling tools will be used in subsequent sections:

- *SPACE [Bee98, MSp]: Layout-to-circuit extractor with implementations of BEM, layered FEM and combined BEM/FEM.*
- *FEMLAB [MFe]: FEM-based multiphysics modeling tool.*

SPACE and FEMLAB run on different machines, an HP J5600 (dual PA_RISC processor at 552 MHz) and a dual Pentium IV 2.8GHz, respectively, each having 1Gb of memory. In the following sections, relevant performance comparisons are all done using SPACE on the HP, while FEMLAB is independently used as a reference.

The technology data for the FEM, BEM/FEM, and BEM methods shown in Figure 5.6 may be passed to SPACE through a high-level, or a low-level technology de-

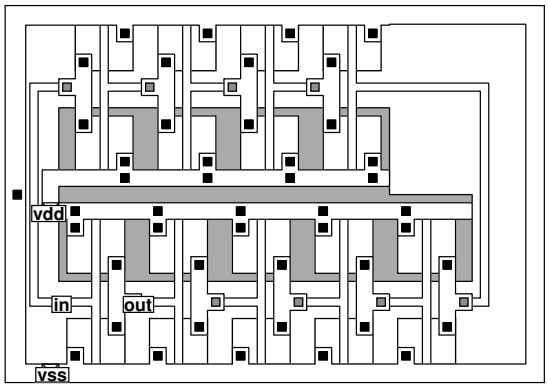

Figure 5.10: Layout of a 9-section ring oscillator

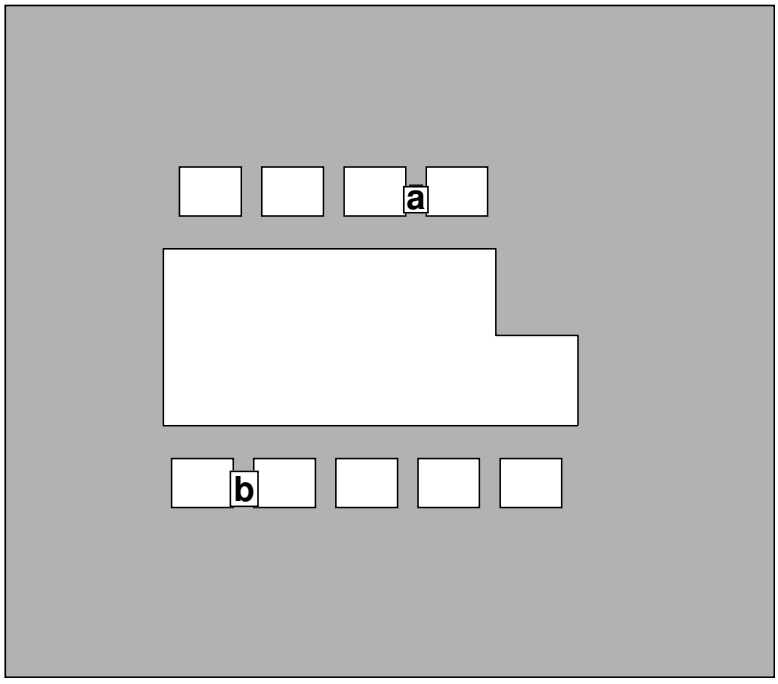

Figure 5.11: Channel-stop pattern for the ring oscillator. Size is 20.5 $\mu m \times$23.5 μm. Markers are the terminals for R_{ab} in Tables 5.4 – 5.9.

scription language as described in Appendix A. FEMLAB may be used through its graphical user interface, or its description language.

Unfortunately, the commercial substrate modeling tool SubstrateStorm proved to be unsuitable as a reference. When applied to straightforward terminal configurations on a standard substrate, the results from SubstrateStorm did not match the results calculated by 2 other tools, while these tools did match each other [Sch04b]. The same problem occurred in several different test-cases. Furthermore, SubstrateStorm allowed insufficient control over the mesh to study convergence under mesh refinement. This supports the observation from Section 2.4 that FDM-based substrate modeling tools tend to operate at some distance from the convergence point in order to be efficient.

5.4.3 Preliminaries for a valid comparison

In subsequent sections, the BEM, FEM and BEM/FEM methods will be compared to each other. The fundamental difference in domain-size between the BEM and the FEM requires to take specific care in a comparison between the methods. Starting from the BEM and the FEM, a valid comparison between the methods can be constructed in either of two ways:

- The FEM can approximate the BEM by defining the FEM domain as large as possible.

- Conversely, the BEM can approximate the FEM by using sidewall-images or a pseudo-finite BEM domain (recall Section 3.3.3).

In a comparison with the BEM/FEM method these approaches can also be applied to the BEM and/or FEM subdomains.

Recall from Section 4.4.5 that the elimination procedure is unintended for solving the layered FEM, such that the computations by SPACE for solving the layered FEM (but also for solving BEM/FEM) are very inefficient. Therefore, in the subsequent sections, conclusions about absolute performance of the methods cannot be drawn, but we will be satisfied with establishing the relative performance of different methods using the same suboptimal solver.

Note that the SUB reference node has to be eliminated from the BEM, or BEM/FEM, network for a valid comparison with a FEM network which does not have this reference node.

5.5 Convergence of the Layered FEM Method

To verify the implementation of the layered FEM, this section presents a practical study of the convergence of the layered FEM by applying it to a simplified situation which does not yet take into account the whole ring oscillator, but only concentrates on the channel-stop pattern. This convergence study is relevant because the layered FEM forms an essential part of the BEM/FEM implementation. Furthermore, verification of the implementation is relevant because the layered FEM flexibly applies the capabilities of the SPACE layout-to-circuit extractor in a way that was originally unintended.

The convergence of the layered FEM can now be studied by refining the horizontal and vertical discretization. However, due to the inefficiencies in the solver, a convergence experiment for the layered FEM using strong refinements in both horizontal and vertical discretization is limited to a very small structure. The structure under consideration is shown in Figure 5.12. It represents the channel-stop pattern from Figure 5.11 with a small segment of the underlying substrate. The structure basically consists of three conductive layers, loosely based on the doping profile approximation from Figure 5.7 and Table 5.1: the top layer is 0.5 μm thick and has a conductivity of 1000 S/m, the intermediate layer is 3.25 μm thick and has a conductivity of 50 S/m, while the bottom layer is 15 μm and has a conductivity of 10 S/m.

For the vertical FEM discretization, each of these conductive layers may now be divided into a number of sublayers from which a layered FEM can be derived. The division into sublayers takes place according to the following systematic approach:

- The highly-conductive top-layer (where a uniform current is assumed) is divided uniformly.
- The intermediate and bottom layers are divided progressively according to the row '1 – 1.5 – 2.25 – 3.375 – ...'

Table 5.3 shows this approach for different levels of refinement. We observe in the bottom row that (according to Figure 4.16), the presented layer divisions result in a 7-layer, 10-layer, 13-layer and 16-layer FEM respectively.

As explained in Section 4.4.1, the horizontal FEM discretization is based on tiles with at most the dimensions of (user-defined) x and y size parameters. Convergence with respect to the horizontal discretization may be studied by reducing the maximum tile dimensions.

Table 5.4 shows the resistance between terminals a and b from Figure 5.11 for increasing refinements in the horizontal and vertical discretizations. The empty locations in the table were not computed due to excessive computation times[1] resulting

[1] The 7-layer FEM at $3\lambda \times 3\lambda$ required 12 days of computation time. Extractions which could be expected to take more time were not performed.

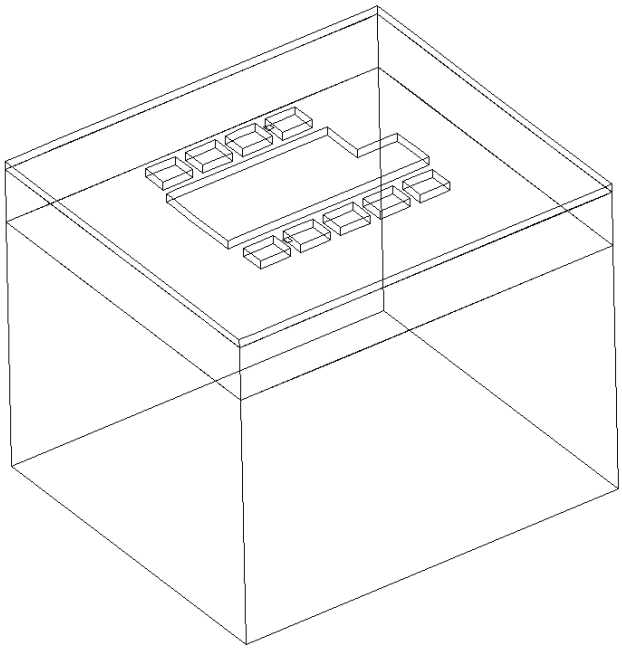

Figure 5.12: Channel-stop pattern for the ring oscillator, and the underlying substrate layers based on Table 5.3. Dimensions: 20.5 $\mu m \times$23.5 $\mu m \times$18.75 μm. Image taken from the FEMLAB user interface.

Table 5.3: Layer divisions at different levels of refinement in the vertical direction

total thickness, conductivity	layer division, thicknesses in μm			
0.5 μm, 1000 S/m	0.25	0.167	0.125	0.1
	0.25	0.167	0.125	0.1
		0.167	0.125	0.1
			0.125	0.1
				0.1
3.25 μm, 50 S/m	1.3	0.684	0.4	0.246
	1.95	1.03	0.6	0.370
		1.54	0.9	0.555
			1.35	0.832
				1.25
15 μm, 10 S/m	6	3.16	1.85	1.14
	9	4.74	2.77	1.71
		7.11	4.15	2.56
			6.23	3.84
				5.76
total number of conductive layers (N)	6	9	12	15
resulting number of FEM layers (N+1)	7	10	13	16

Table 5.4: R_{ab} resistance values ($k\Omega$) for increasing refinement of the FEM mesh in horizontal (i.e. reduced maximum tile dimensions) and vertical (i.e. increasing number of FEM layers according to Table 5.3) direction. Refinements in the vertical direction do not have much influence, but convergence due to refinements in the horizontal direction is apparent. Empty locations in the table refer to situations that were not computed due to excessive computation times resulting from the non-optimized solver; the 7-layer FEM at $3\lambda \times 3\lambda$ required 12 days of computation time.

maximum tile dimensions	FEM			
x_size (λ) $\times$ y_size (λ)	7-layer	10-layer	13-layer	16-layer
50×50	4.21	4.19	4.16	4.14
40×40	4.49	4.48	4.46	4.42
30×30	4.95	4.94	4.92	4.89
20×20	5.56	5.57	5.57	5.55
10×10	5.93	5.94	5.93	
5×5	6.09	6.11		
3×3	6.15			

from the non-optimized solver. We observe that, for this example, refinements in the vertical discretization are not critical in reaching convergence. In fact, for given maximum tile dimensions, the resistance values calculated at each different vertical discretization are almost identical. This is most probably caused by the fact that the majority of the current will flow through the highly conductive top layer, which is apparently already captured with sufficient accuracy by a division into two conductive layers, such that the introduction of more conductive layers has hardly any effect on the computed resistance. Nevertheless, for refinements in the horizontal discretization we clearly observe convergence. This is particularly visible for the 7-layer FEM, where the last three refinement steps in the horizontal discretization result in a growth of R_{ab} by 6.7%, 2.7% and 0.99% respectively. This clearly indicates convergence. Unfortunately, even further refinement in the horizontal discretization was not possible due to excessive computation times caused by the non-optimized solver.

Using the resistance value of 6.15 $k\Omega$, which is the best converged value under the limitations of the non-optimized solver, we will now identify how good this value is as an approximation by comparing it to an independent reference calculated by FEMLAB. For the shown structure in Figure 5.12, using strong mesh refinement in the domain and on the contact areas, FEMLAB calculated R_{ab} = 6.25 $k\Omega$. The mesh was chosen as fine as possible given the available memory of the machine on which FEMLAB runs. Between the layered FEM value of 6.15 $k\Omega$ and FEMLAB's reference value of 6.25 $k\Omega$ we observe a difference of 1.6%. This is considered reasonably satisfactory.

In summary, the layered FEM implementation shows converging behaviour, and compares well to the independent FEMLAB reference.

5.6 Convergence of the BEM/3DFEM method

This section presents a practical study of the behaviour and convergence of the BEM/FEM method by applying it to a simplified situation comparable to the one used in Section 5.5, which only concentrates on the channel-stop pattern. As mentioned in Section 5.5 we would ideally compare the BEM/FEM method to the layered FEM method for comparable situations, but the layered FEM is unfortunately not suitable as a general reference due to the yet unoptimized implementation which makes it infeasible to choose a large FEM domain as an approximation of the infinite BEM domain. Therefore, as in Section 5.5, FEMLAB, which provides optimized solvers and allows to choose a larger FEM domain, will be used as an independent reference.

The situation for these extractions is the reduced BEM/FEM approach from Section 5.3.3 (Figure 5.6c), but note that the FEM domain is considered to have a conductivity of zero, except in the areas where the channel-stop layer is present. The BEM and FEM horizontal discretizations can then be limited to the channel-stop pattern itself, as illustrated in Figures 5.13 and 5.14. In effect, the channel-stop pattern is then situated on top of a 10 S/m domain representing the deep substrate. The extractions below again calculate the port-impedance R_{ab} between terminals 'a' and 'b' from Figure 5.11.

5.6.1 Independent 3DFEM Reference

We have used FEMLAB to calculate a reference for R_{ab}, choosing the domain as indicated in Figure 5.15. The dimensions of the channel-stop pattern itself are fixed, but to obtain a value for R_{ab} that becomes (almost) independent of the domain size, as well as properly converged with respect to mesh refinements, the size of the block underneath was increased (under fixed aspect ratios) and the mesh was chosen with as much refinement as possible, given the available amount of memory in the machine on which FEMLAB runs. The eventual dimensions of the block are 120 $\mu m\times$ 123 $\mu m\times$ 35 μm, for which, under strong mesh refinement, FEMLAB calculates 7.2 $k\Omega$. Having observed the behaviour of FEMLAB in reaching this point, this value is considered within 1% accurate. In accordance with Section 5.4.3, this large domain approximates the BEM situation of a domain extending to infinity, which allows an as fair as possible comparison between the FEMLAB 3DFEM and the BEM/FEM method.

The contribution of the deep substrate to the calculated port-impedance can be determined by ignoring the deep substrate and calculating the port impedance for only the channel-stop pattern. For the isolated channel-stop pattern, FEMLAB's 3DFEM method calculates 8.7 kΩ. From these results, we conclude that the deep substrate is indeed relevant for accurate results; comparing the situation *with* the deep substrate (7.2 kΩ) to the situation *without* (8.7 kΩ), a difference of approximately 20% becomes apparent.

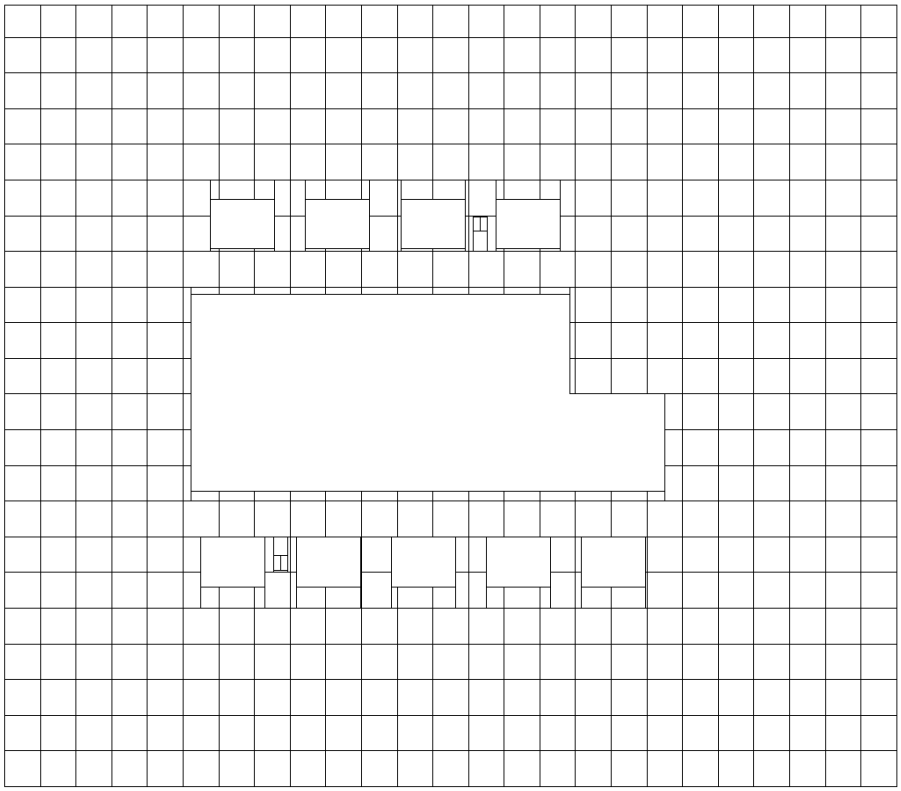

Figure 5.13: Example of the tile division (i.e. BEM contact-area division) along the BEM/FEM interface using the approach from Figure 4.11. Parameters x_size and y_size are both 30λ.

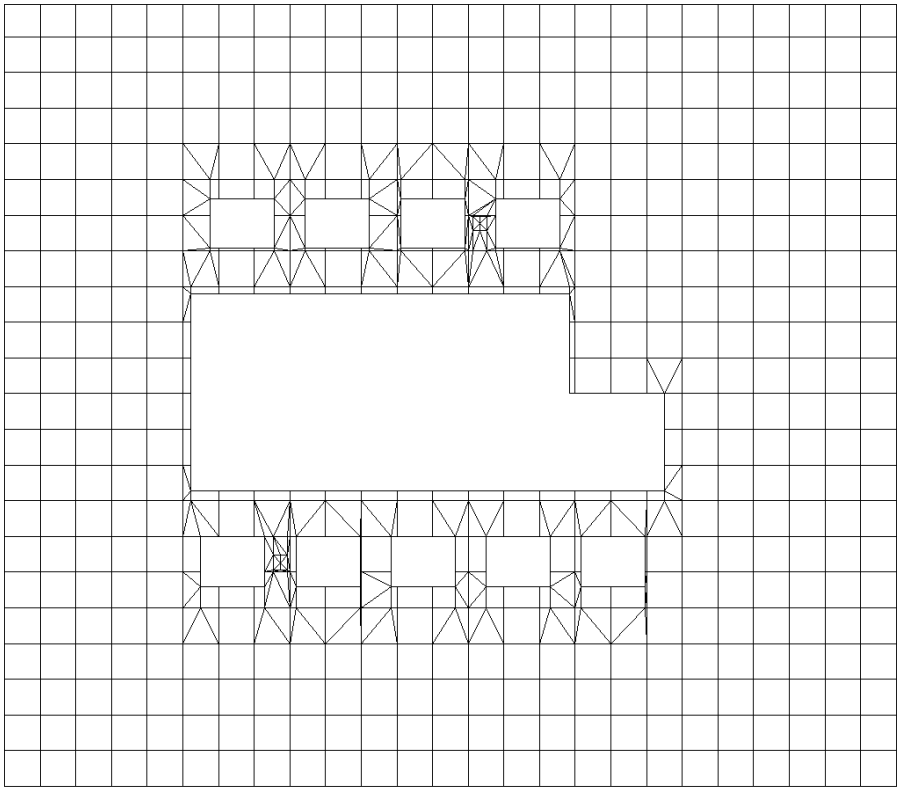

Figure 5.14: Example of the FEM mesh corresponding to the tile division (BEM contact-area division) from Figure 5.13.

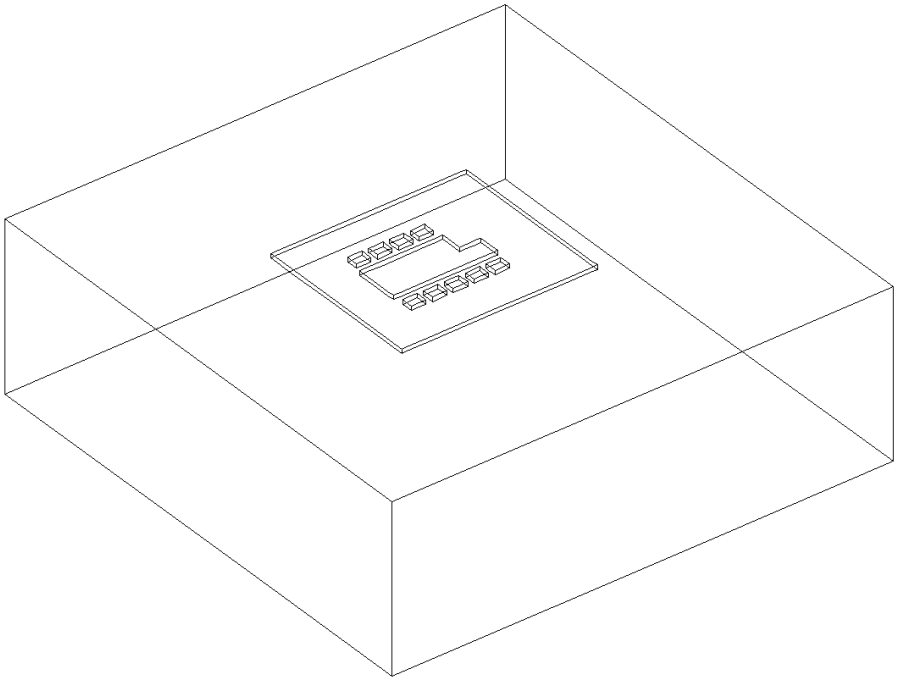

Figure 5.15: The channel-stop pattern on top of a 10 S/m domain. Image taken from the FEMLAB interface.

5.6.2 Convergence of the BEM/3DFEM method

This section will study the convergence of the BEM/3DFEM method through the parameters that determine its behaviour. We will choose a nominal setting for all parameters, and vary only individual parameters, while keeping the other parameters at their nominal settings.

Recall that the unit-length in the design is $\lambda = 0.03125\ \mu m$ (1/32 μm). The 4 main parameters that determine the behaviour of the BEM/FEM method are explained below, the nominal settings are therefore given without further explanation:

- Maximum tile dimensions in BEM/FEM horizontal discretization = $30\lambda \times 30\lambda$
- Number of FEM layers = 3
- Size of BEM window = $5\mu m$
- BEM mesh = 1 BEM panel per interface tile

A fifth parameter, the node contraction distance from Figure 4.10, will also be addressed, but, strictly speaking, it has no influence on the convergence of the BEM/FEM method itself. It merely provides a transition between a BEM/2DFEM and a BEM/3DFEM situation which may each individually converge, but this is independent of the contraction distance. Therefore, the convergence simulations for the 4 main parameters do *not* use node contractions.

Note that, according to Section 5.4.3, the SUB reference node, which is inherent in the applied BEM technique, is eliminated from the networks by Gaussian elimination. The obtained resistance values are then comparable to the resistance values obtained with the FEM (either FEMLAB or the layered FEM), which does not have this reference node.

Maximum Tile Dimensions in BEM/FEM Horizontal Discretization

As explained in Section 4.4.1, the BEM/FEM horizontal discretization is defined by tiles with dimensions of at most the (user-defined) x_size and y_size parameters. Reducing the maximum tile dimensions in the BEM/FEM horizontal discretization should show convergence in the resulting model.

Table 5.5 and Figure 5.16 suggest converging behaviour of R_{ab} for reducing tile dimensions, but the resistance value at $10\lambda \times 10\lambda$ is still some 6.8% larger than the resistance value at $20\lambda \times 20\lambda$, indicating that the converged value has not yet been reached. Unfortunately, further refinement is not possible due to memory limits of the system on which the modeling software runs.

However, an estimation of the converged value can be obtained by fitting a curve to the resistance values that we were able to obtain. This will be discussed with more detail in Section 5.6.3

Number of FEM layers

The layered FEM approach is based on a stack of 2DFEM layers connected by vertical resistances. The case where there is only one FEM layer actually represents the 2DFEM. Increasing the number of FEM layers will introduce the third dimension into the model and improve the accuracy of the FEM. Consequently, the overall accuracy of the BEM/FEM model improves. Table 5.6 shows R_{ab} as a function of this parameter. Convergence is apparent, as shown in Figure 5.17.

Size of BEM window

The BEM allows sparsification through the previously described windowing technique, which can make the method more efficient at the cost of some accuracy. Table 5.7 and Figure 5.18 show that, in a combined BEM/FEM situation, less sparsification (i.e. a larger BEM window) has little influence on R_{ab} while the extraction times and memory usage increase significantly.

BEM mesh

The initial BEM mesh consists of a single BEM panel per BEM contact. In the context of BEM/FEM modeling, the BEM contact division on the BEM/FEM interface is equivalent to the tile division on the interface, such that the initial BEM mesh consists of 1 BEM panel per interface tile. The BEM mesh then allows refinement by introducing multiple BEM panels per interface tile. Table 5.8 and Figure 5.19 show that, in the combined BEM/FEM situation, such refinements have very little influence on R_{ab} while they are relatively costly with respect to extraction time and memory usage.

Node Contractions

The contraction distance 'd' (Figure 4.10) plays no explicit role in the convergence behaviour of the Combined BEM/FEM method: the calculated resistance will be bounded by the resistance values calculated by the BEM / 2DFEM method (i.e. full contraction) and the BEM / 3DFEM method (i.e. no contraction), provided that equivalent horizontal discretizations have been used. Table 5.9 and Figure 5.20 show this behaviour for R_{ab}. The computation time and memory usage of the BEM / contracted 3DFEM method is only slightly larger than BEM / 2DFEM, while the accuracy of BEM / 3DFEM is still obtained.

Table 5.5: Extraction statistics and resistance values for increasing refinement of the BEM/FEM horizontal discretization. Converging behaviour in the resistance value is apparent, as shown in Figure 5.16.

maximum tile dimensions x_size (λ) $\times$ y_size (λ)	# interface nodes	time (s)	mem (Mb)	R_{ab} ($k\Omega$)
50 $\times$ 50	295	2.32	4.61	5.05
40 $\times$ 40	353	4.18	4.79	5.29
30 $\times$ 30	608	16.0	13.3	5.86
20 $\times$ 20	1172	83.2	39.9	6.51
10 $\times$ 10	4269	8923	321	6.95

Table 5.6: Extraction statistics and resistance values for an increasing number of FEM layers in the BEM/FEM model. Convergence in the resistance value is apparent, as also shown in Figure 5.17.

# FEM layers	time (s)	mem (Mb)	R_{ab} ($k\Omega$)
1 (2D)	5.97	6.62	5.31
2	11.5	12.1	5.64
3	16.0	13.3	5.86
4	23.7	14.1	5.95
5	30.0	14.4	6.01
6	53.3	20.1	6.04

Table 5.7: Extraction statistics and resistance values for increasing size of the BEM window. The non-monotonous behaviour of the resistance is caused by meshing phenomena induced by the window. The resistance value changes little, as also shown in Figure 5.18.

BEM window size (μm)	time (s)	mem (Mb)	R_{ab} ($k\Omega$)
1	9.74	6.29	5.49
3	9.85	9.99	5.85
5	16.0	13.3	5.86
10	31.6	24.0	5.73
25 (full)	43.3	36.8	5.70

Table 5.8: Extraction statistics and resistance values for increasing refinement in the BEM mesh. The resistance value changes little, as also shown in Figure 5.19.

#BEM panels / interface tile	time (s)	mem (Mb)	R_{ab} ($k\Omega$)
1	16.0	13.3	5.86
4	25.7	16.9	5.84
8	55.1	30.7	5.82
16	286	107	5.81

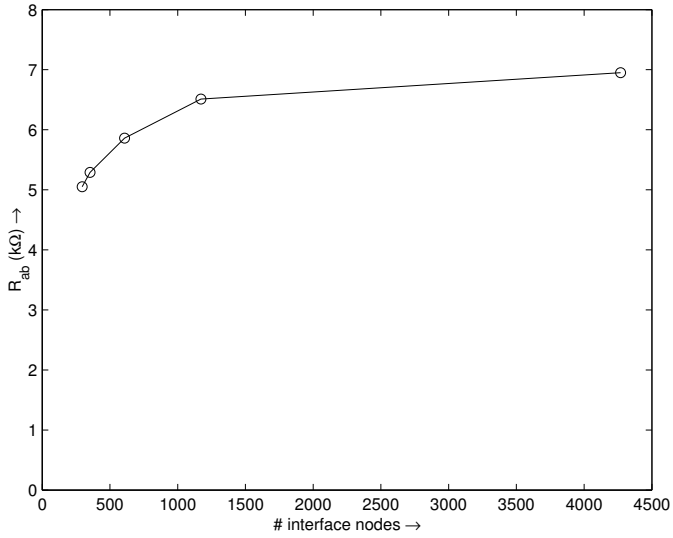

Figure 5.16: Graphical representation of Table 5.5. Resistance R_{ab} converges for an increasing refinement of the BEM/FEM horizontal discretization, represented by the increasing number of interface nodes along the horizontal axis. Convergence is apparent.

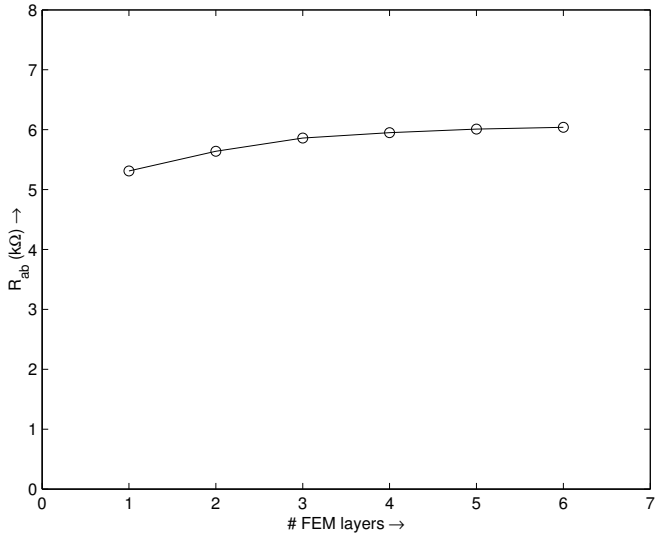

Figure 5.17: Graphical representation of Table 5.6. Resistance R_{ab} converges for an increasing number of FEM layers.

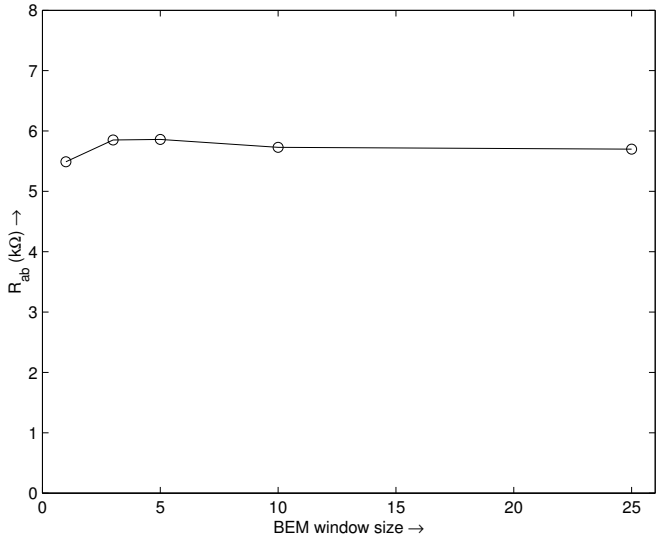

Figure 5.18: Graphical representation of Table 5.7. Resistance R_{ab} changes little for increasing BEM window size.

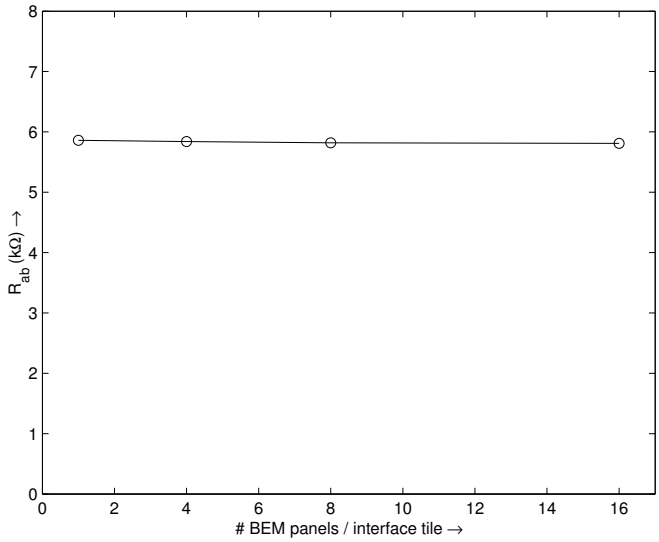

Figure 5.19: Graphical representation of Table 5.8. Resistance R_{ab} changes little for an increasing number of BEM panels per interface tile.

Table 5.9: Extraction statistics and resistance values for various settings of the contraction distance 'd'. The resistance value is enclosed between the BEM / 2DFEM and BEM / 3DFEM situations. The extraction uses the nominal settings, but some mesh refinement near the contacts is necessary to accommodate the halo that guides the contraction: there are now 640 interface nodes instead of the 608 from Table 5.5. See also Figure 5.20 for a graphical representation.

contraction distance	time (s)	mem (Mb)	R_{ab} ($k\Omega$)
– (BEM / 2DFEM)	9.19	8.46	5.83
2 λ	9.21	8.87	6.09
4 λ	9.21	8.87	6.21
8 λ	9.26	8.88	6.46
16 λ	9.28	8.88	6.57
32 λ	9.30	8.89	6.61
∞ (BEM / 3DFEM)	21.4	13.9	6.62

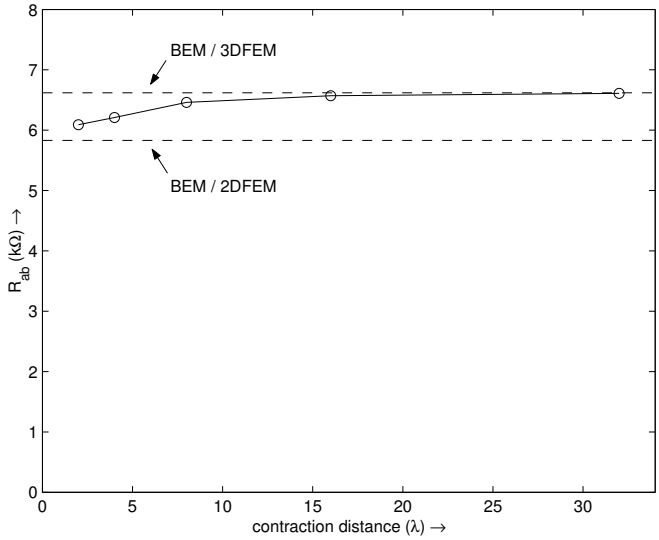

Figure 5.20: Graphical representation of Table 5.9. Resistance R_{ab} is enclosed between the BEM / 2DFEM and the BEM / 3DFEM situations and converges towards the BEM/3DFEM situation for increasing contraction distance.

Concluding Remarks

Tables 5.5 – 5.9 and the associated Figures 5.16 – 5.20 show that the BEM/FEM method behaves as expected with respect to computation time and memory usage which increase for refined parameter settings.

However, only for the BEM/FEM horizontal discretization and the number of FEM layers can we observe significant influence on the calculated resistance value. The resistance value is almost insensitive to the size of the BEM window and refinements in the BEM mesh. This observation confirms the remarks from Section 4.3.1: a BEM extraction with a relatively coarse mesh and a small window is sufficient for the combined BEM/FEM results to be accurate. Thus, the BEM extraction parameters can be chosen such that the BEM contributes only little to the computational complexity of the overall combined BEM/FEM extraction, while appropriate accuracy is still achieved. Note, however, that these results may be biased due to the structure of the example, where most of the current between the contacts is confined to the channel-stop layer and the deep substrate indeed plays only a limited role in the overall model. Therefore, a second evaluation of this observation is done for a more realistic situation in Section 6.2.4.

Furthermore, the node contractions may strongly reduce the computation time, while high accuracy is still achieved. Note, however, that also these results may be biased due to the structure of the example. In fact, there are only 2 terminal areas involved around which the 3DFEM is applied. Therefore, the contractions strongly reduce the computation time because they take place over large areas. In actual designs, however, the 3DFEM will typically be required over larger areas, and the computation time reduction owing to the node contractions is expected to be smaller. Nevertheless, Section 6.2.1 will show for a more realistic design that the node contractions can indeed provide a significant reduction in the extraction time, while the accuracy in the overall model is retained.

5.6.3 Full Convergence and Comparison to Reference

Having established the stable behaviour of the BEM/FEM model with respect to its individual parameters in the previous section, this section presents a more fully converged situation which is obtained by choosing (strongly) refined settings for all parameters. The result is compared to the independent FEMLAB reference from Section 5.6.1.

For this extraction, Tables 5.5 – 5.9 were used to choose (strongly) refined settings for all parameters in the BEM/FEM method. In particular:

- Maximum tile dimensions in BEM/FEM horizontal discretization = $10\lambda \times 10\lambda$
- Number of FEM layers (where applicable) = 6

Table 5.10: Full convergence of the BEM / 3DFEM, BEM / contracted 3DFEM and BEM / 2DFEM methods. The error is calculated w.r.t. the independent FEMLAB reference value of 7.2 kΩ. The BEM / 3DFEM and BEM / contracted 3DFEM methods both show a very small error w.r.t. the reference.

BEM/	3DFEM	contracted	2DFEM
time (hrs:mins)	8:21	4:08	3:56
memory (Mb)	960	960	960
R_{ab} (kΩ)	7.03	7.05	6.06
error (%)	-2.4	-2.1	-16

- Size of BEM window = 10 μm
- BEM mesh = 4 BEM panels per interface tile
- Contraction distance (where applicable) = 16 λ (0.5 μm)

The calculated values and extraction statistics for the BEM/FEM method are given in Table 5.10, together with the error compared to the FEMLAB reference of 7.2 kΩ. The BEM / 2DFEM method gives a significant error of 16%, whereas the BEM / 3DFEM and BEM / contracted 3DFEM match the FEMLAB reference within 2.5%. The large amount of memory in all extractions is caused by the heavy BEM settings.

Even though the BEM / 3DFEM resistance values are close to the FEMLAB reference, the limitations in the system resources previously encountered in the construction of Table 5.5 prohibit an explicit practical evaluation of the convergence. However, a theoretical estimate of the fully converged value may be obtained through a curve fit on the data in Table 5.5 and its graphical representation in Figure 5.16. Indeed, let x be the number of interface nodes from Table 5.5, and performing a nonlinear least-squares fit on $\{1, x^{-1}, x^{-2}\}$, the resistance R_{ab} is described by the following approximate function within 1% accuracy:

$$R_{ab}[k\Omega] = 7.18 - \frac{917}{x} + \frac{85.7e3}{x^2} \tag{5.1}$$

The original data and the curve fit are shown simultaneously in Figure 5.21. Equation 5.1 indicates an estimated converged value of 7.18 $k\Omega$. This means that the 6.95 $k\Omega$ at $10\lambda \times 10\lambda$ tile-size is $\sim$3% from convergence. The other parameters (i.e. those represented in Tables 5.6 – 5.9) are estimated to be within 0.5% from convergence, and actually have a tendency to cancel each other. Thus, we may estimate that the R_{ab} value of 7.03 $k\Omega$ for the BEM/3DFEM can still grow $\sim$3% to 7.24 $k\Omega$ when smaller tile dimensions would be chosen. This would mean a remaining difference of less than 1% when compared to the 7.2 $k\Omega$ resulting from FEMLAB. This is considered satisfactory.

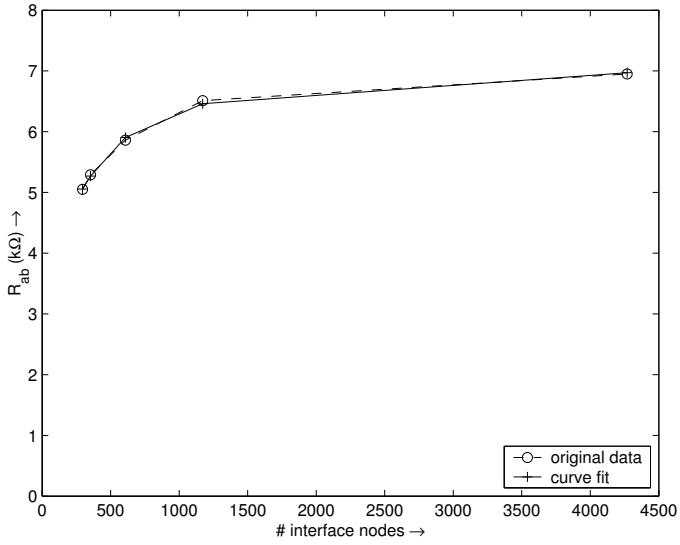

Figure 5.21: The original data from Figure 5.16 (Table 5.5) and the curve fit from Equation 5.1 drawn in a single figure. The match is within 1%.

In summary, we observe that the BEM/FEM method compares well to the reference. At the same time, we observe that the node contractions may strongly reduce the computation time, while the accuracy of the BEM/3DFEM method is still achieved. However, as already mentioned at the end of the previous section, the present example allows the contractions to take place over large areas which may not be a representative situation for an actual design. The study in Section 6.2.1 of the next chapter will reveal the practical value of the node contractions for a more representative situation.

5.6.4 Comparison to plain BEM

In the previous sections the BEM/FEM method showed converging behaviour and compared well to an independent reference. In this section, the improved consistency of the BEM/FEM method over the plain BEM method will be established by comparing the BEM/FEM method, a single-layer BEM and a double-layer BEM to a 3DFEM reference calculated with FEMLAB. By choosing two different basic configurations for the test terminals in the channel-stop pattern (see Figure 5.22) we will show that only the BEM/FEM method yields results that are consistent with the FEMLAB reference, whereas the plain BEM approaches may result in large errors depending on the situation.

FEMLAB used similar refined extraction settings as before (see Section 5.6.1), the BEM/FEM method used the nominal extraction settings, while the plain BEM used

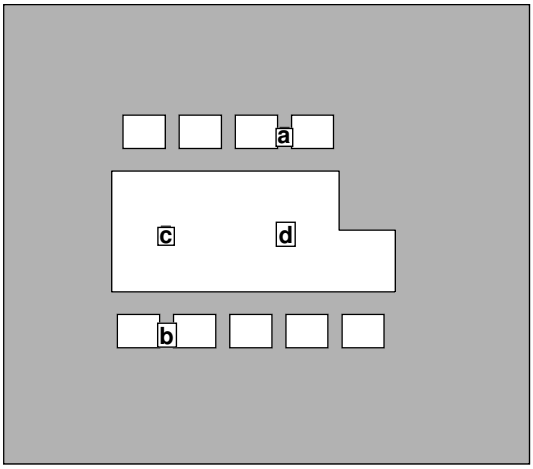

Figure 5.22: Channel-stop pattern for the ring oscillator. Dimensions are equal to those from Figure 5.11: $20.5\mu m \times 23.5\mu m$. Markers are test-terminals for the simulations in Table 5.11.

Table 5.11: Resistance values calculated by the various modeling approaches. R_{ab} and R_{cd} calculated between the two corresponding markers in Figure 5.22. The BEM/3DFEM results are consistent with the FEMLAB reference, whereas the plain BEM results are not.

method	channel-stop	R_{ab} ($k\Omega$)	R_{cd} ($k\Omega$)
1-layer BEM	absent	231	227
3DFEM ref. (FEMLAB)	pattern	7.2	214
BEM/3DFEM	pattern	5.86	245
2-layer BEM	uniform layer	3.61	3.28

its own nominal settings. Table 5.11 shows the numerical values of the comparison. Columns R_{ab} and R_{cd} have been generated by doing separate extractions for terminal configurations ab and cd from Figure 5.22. Terminal configuration ab was already used in the convergence experiments from the previous sections; the terminals are inside, but close to, the border of the channel-stop region. Terminal configuration cd has been chosen such that the terminals are in locations where no channel-stop is present.

In column R_{ab}, we observe that 3DFEM, combined BEM/FEM and double-layer BEM are near each other, whereas the single-layer BEM gives a large error. However, in column R_{cd} we observe that 3DFEM, combined BEM/FEM and single-layer BEM are near each other, whereas the double-layer BEM gives a large error. Even

though the errors between BEM/FEM and the reference are still significant (15 – 20%), they are mainly caused by the difference in extraction settings: 'nominal' for the BEM/FEM method, versus 'refined' for the reference. Taking note of this difference, Table 5.11 confirms that the results of the BEM/FEM method are consistent with the 3DFEM reference. The single-layer and double-layer BEM, on the other hand, may calculate reasonable resistance values for either of the terminal configurations, but present an arbitrarily large error in the other.

In terms of the usability of these models, only models that are consistent with the physical situation underlying the modeling problem are actually usable in practice, while the accuracy that is achieved in the models is relatively less important. Unfortunately, due to the absence of measurements on a physical channel-stop pattern, it is not possible to verify the 4 resulting models with a physical situation. However, assuming that a 3DFEM as used in FEMLAB is capable of accurately modeling the physical situation that is currently under study, FEMLAB may be used as a reference. With this reference, only the BEM/FEM method is consistent, and is therefore a better usable model than the 1-layer and 2-layer BEM. In other words, the BEM cannot be the universal modeling tool in case of lateral doping variations. Actually, this was already concluded from Section 5.3.5.

5.7 Summary and Concluding Remarks

This chapter started by identifying the objectives of the practical evaluation. Then, after introduction of the substrate doping profile, appropriate settings for the FEM, BEM/FEM and BEM methods were chosen, aiming for a sufficiently accurate approximation of the substrate doping profile. Subsequently, a ring-oscillator with its accompanying channel-stop pattern was introduced as the main structure under consideration.

The behaviour and convergence of the layered FEM and BEM/FEM methods was studied by applying them to a simplified situation where only the channel-stop pattern and the underlying substrate were present. Both methods showed converging behaviour and the results compared favourably to an independent 3DFEM reference. Furthermore, it was observed that a coarse and sparse BEM provides sufficient accuracy to the overall BEM/FEM model, while it also improves the computational efficiency of the overall BEM/FEM method. It was also observed that node contractions can improve the computational efficiency even further, without sacrificing accuracy. Finally, it was shown that the studied example could be modeled more consistently by the BEM/FEM method than by a plain BEM approach.

The next chapter will evaluate the behaviour of the BEM/FEM method when applied to a more realistic situation.

6

Behaviour of the BEM/FEM Method in a Realistic Situation

With the convergence and consistency of the BEM/FEM method being confirmed for the simplified situation from the previous chapter, the method will now be applied to a more realistic situation. The present chapter will again confirm that the BEM/FEM method inherits the good characteristics from its constituting parts, and that consistency of a model with the physics is more important than the highest accuracy in the model. Using these observations, the next chapter will present the conclusions.

6.1 Introduction

This chapter is a continuation of the previous chapter and therefore uses the objectives, substrate doping profile and modeling approaches from Sections 5.1, 5.2 and 5.3, respectively.

Sections 6.2 and 6.3 will again address the objectives from Section 5.1, but now by applying the layered FEM and BEM/FEM methods to a more realistic case involving the full ring-oscillator. In these sections, the first objective will be to confirm whether the BEM/FEM method can be faster than the layered FEM while reaching comparable accuracy. The second objective will be to confirm whether consistency between the model and the physical situation is more relevant than the highest accuracy in the

model. If this is the case, this observation can be exploited to the benefit of the efficieny of the BEM/FEM method. In fact, it may also be beneficial to other (substrate) modeling techniques.

Section 6.4 will consider the effect of capacitive coupling between the interconnect and the substrate on the noise levels in the substrate, with the objective of showing that not only consistency in the substrate model itself is relevant, but also consistency of the overall model with the overall physical situation including the interconnect is very relevant.

The observations and conclusions will ultimately be summarized in Section 6.5.

6.2 Ring Oscillator with Guarded Sensor Node

A typical scenario for substrate crosstalk would be the situation where the switching activity in a digital circuit injects noise into the substrate which is picked up by sensitive circuitry elsewhere. A possibility for protecting the sensitive circuitry would be to place a guard ring around it. In this context, as shown in Figure 6.1, a sensor node surrounded by an n-well guard ring is added to the layout at some distance from the ring oscillator. If the guard ring is properly biased, the noise reaching the sensor node will be reduced. The channel-stop pattern is now as shown in Figure 6.2.

Contrary to the BEM/FEM convergence simulations from the previous section, it is not sufficient to apply the FEM only to the channel-stop pattern, because we now also have to take into account the n-well, which, in our case, is assumed to be 5.4 μm thick. Therefore, we now use the regular BEM/FEM modeling approach from Section 5.3.2 (Figure 5.6b). A schematic side-view of the present situation is shown in Figure 6.3.

The nominal BEM/FEM extraction parameters in this example will be as follows:

- Maximum tile dimensions in BEM/FEM horizontal discretization = $50\lambda \times 50\lambda$
- Number of FEM layers = 5
- Size of BEM window = 50λ (1.5625 μm)
- BEM mesh = 1 BEM panel per interface tile

Note that the 5 FEM layers were chosen as described in Section 5.3.2. Furthermore, the BEM settings have been chosen coarse according to the observations from Sections 4.3.1 and 5.6. In particular, the BEM window has been chosen as the minimum (i.e. equal to the maximum tile dimensions), and the BEM discretization has been chosen as coarse as possible.

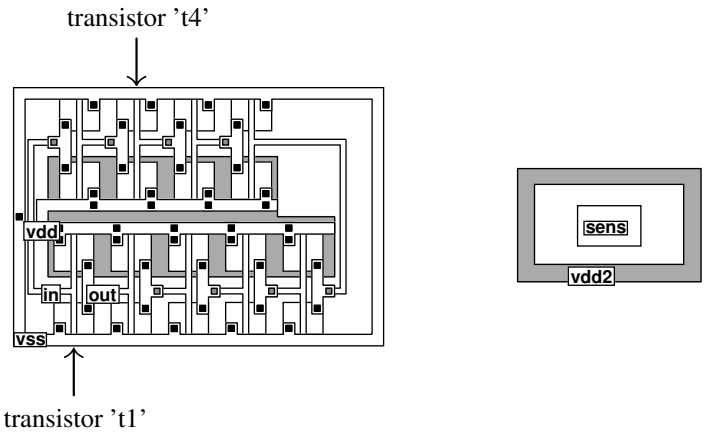

Figure 6.1: The ring oscillator with guarded sensor node. Marked transistors will be used later on in the chapter.

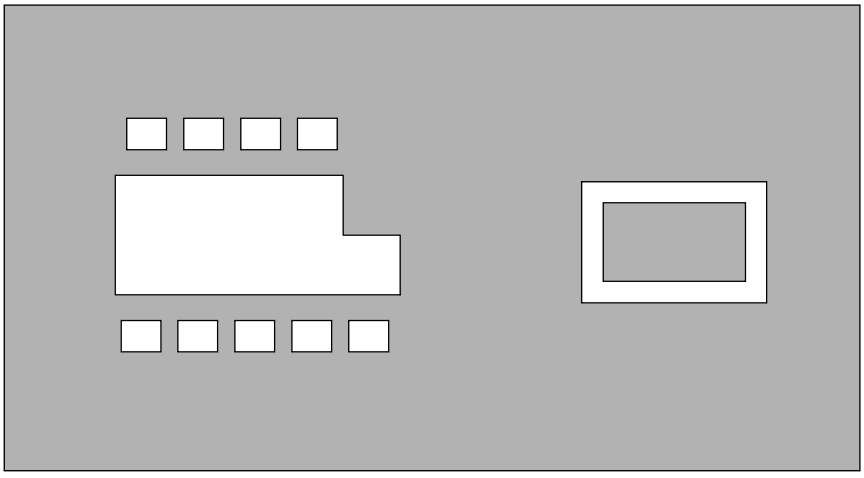

Figure 6.2: Channel-stop pattern for the ring oscillator with guarded sensor node.

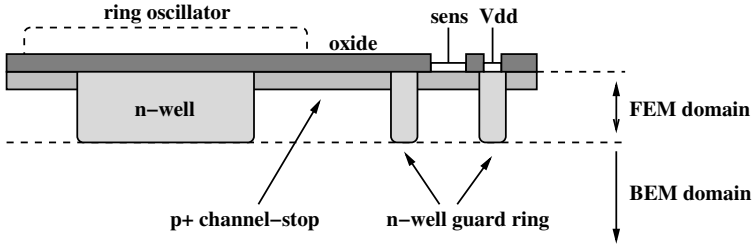

Figure 6.3: Schematic side view of the situation. Note that the BEM domain starts immediately underneath the n-well, making the BEM necessary to capture resistive coupling between ring oscillator and sensor node.

Note (again) that, according to Section 5.4.3, the SUB reference node, which is inherent in the applied BEM technique, is eliminated from the networks by Gaussian elimination. The obtained resistance values are then comparable to the resistance values obtained with the FEM (either FEMLAB or the layered FEM), which does not have this reference node.

Note also that the situations in this, and subsequent, sections present an as-complete-as-possible model with respect to the overall partitioned problem from Figure 1.3. That is, the overall model takes into account the interconnect resistance as well as the capacitive coupling between interconnect and substrate. However, to avoid too much computation time being 'lost' to the calculation of the interconnect to substrate capacitance, a fast 2.5D interpolation method will be used. In Section 6.4, we will also study in what way the consistency of the model with the overall partitioned problem influences the simulation results.

Unfortunately, it proved impossible to use FEMLAB as a reference tool in this example. The large number of small contact areas on top of the domain due to the transistors and the interconnect capacitances caused meshing and convergence problems in FEMLAB. The meshing algorithm frequently failed, or convergence in the iterative solver could only be reached for a large error tolerance or after a very large number of iterations. Even in a simplified situation where the interconnect capacitances were not taken into account, the meshing and convergence problems persisted.

Therefore, the layered FEM will be used as a reference, despite its yet unoptimized implementation. Initially, in Section 6.2.1, the 11-layer FEM will be used in a comparison with its corresponding BEM / 5-layer FEM. Since both methods use the same unoptimized solver, this allows a meaningful comparison between the performances of the BEM/FEM and layered FEM methods. Subsequently, in Section 6.2.2, a more deeply converged layered FEM reference is generated consisting of a 20-layer FEM with a refinement in the horizontal mesh that is as strong as possible without introducing an impractical extraction time. In Section 6.2.3 the BEM / 5-layer FEM is compared to this reference.

For a proper comparison between the layered FEM and BEM/FEM methods (see also Section 5.4.3), the domains in which they are applied should be comparable. Using the layered FEM as a reference, it is not feasible with respect to computation time to choose a large FEM domain for a representative comparison with the infinite-domain BEM part of the BEM/FEM method. However, the layered FEM does use a vertical discretization that represents the full 375 μm thickness of the substrate. As a result, the relatively small lateral dimensions of the layout cause domain of the substrate modeling problem to be 'thin' and 'high'. To match this domain with the BEM/FEM method, the BEM part in the BEM/FEM simulations below will use sidewall-images to take into account the small lateral dimensions of the domain. The comparison between the layered FEM and BEM/FEM methods is then valid.

6.2.1 BEM / (contracted) FEM vs. equivalent layered FEM

As already suggested in Section 4.5, the BEM/FEM method is expected to be faster than the full FEM because the BEM/FEM method does not require to discretize the deep substrate, but can suffice with a coarse BEM there. Furthermore, it is expected that the performance of the BEM/FEM method significantly improves when using node contractions. In this example, we will show that the BEM/FEM method indeed lives up to these expectations. In particular, this section will compare the 11-layer FEM to its corresponding BEM / 5-layer FEM.

In this situation we will be applying the BEM/FEM method with *no* contraction, *partial* contraction and *full* contraction. The partial contraction will furthermore be applied with two different settings for the halo in which the 3DFEM is applied: a minimal case and a nominal case. In both cases the contraction distance will be 8λ (0.25 μm), but in the minimal case the 3DFEM is only applied near the via connections to the substrate (V_{ss} and the sensor node), whereas in the nominal case, the 3DFEM is also applied underneath the polysilicon interconnect (which is relatively close to the substrate). In areas where the contraction takes place, the top three FEM layers representing channel-stop will be contracted into a single layer, such that the 5-layer FEM will be contracted into a 3-layer FEM. Note that the contractions only take place in the channel-stop layer, because only there the conditions hold under which the contraction is allowed (see Section 4.3.2).

The full FEM uses the same maximum tile dimensions for its horizontal discretization as the BEM/FEM interface discretization. Furthermore, the halos that guide the node contractions (see Section 4.4.3) also have influence on the horizontal mesh. Therefore, in this set of simulations, they are present in all cases, but only actually used in the contraction simulations. In this way, the full FEM and the FEM segment of the BEM / (contracted) FEM approach all use *exactly* the same horizontal mesh, allowing a fair comparison between the methods.

Table 6.1: Selected resistance values and extraction statistics. We compare the 11-layer FEM to the combined BEM / (contracted) FEM method. The BEM/FEM method may be approximately 4 times faster than the full FEM without loss of accuracy, and approximately 8 times faster with some loss of accuracy.

	3DFEM	BEM/FEM			
contraction	n/a	none	partial	partial	full
3DFEM halo	n/a	n/a	nominal	minimal	n/a
FEM situation	11-layer	5-layer	hybrid	hybrid	3-layer
time (min:sec)	117:26	55:56	25:16	13:55	10:48
memory (Mb)	59.9	107	50.2	47.1	37.2
R(vss, sens) ($k\Omega$)	183.4	181.3	179.6	320.3	220.8
R(vss, t1) ($k\Omega$)	11.59	11.73	11.37	8.719	6.012
R(sens, t1) ($k\Omega$)	382.2	386.3	391.0	410.0	494.6

The networks were ensured to be topologically identical. Note that the resulting networks contain 11 nodes: 'sens', 'Vss' and 9 transistor backgates. Each network then contains 55 resistances in total. Table 6.1 shows the extraction details and a selected set of resistances from the resulting networks. When comparing the BEM / 5-layer FEM to the 11-layer FEM, we observe that the resistance values show only marginal differences, while the BEM / 5-layer FEM is about twice as fast. Figure 6.4 shows the simulations of the waveforms on the sensor node indicated in the layout from Figure 6.1 with the networks generated by these two methods; the noise waveforms are comparable, but still somewhat different. Some simple experimentation showed that the most important reason for this difference is a difference in the handling of RC effects along the well-bottom, which is located on the BEM/FEM interface. In particular, the transition from FEM mesh to BEM mesh on the interface in the BEM/FEM method is somewhat different from the full FEM method, where the FEM mesh continues without the transition. The BEM resistances are then attached to the junction capacitances in a different way than the FEM resistances, as already illustrated previously in Figures 4.18 and 4.19. In general, however, the results match very well.

Furthermore, Table 6.1 shows that the BEM / nominal hybrid FEM is about twice as fast as the BEM / 5-layer FEM while the differences in the resistance values are still only marginal. In turn, the BEM / minimal hybrid FEM and the BEM / fully contracted FEM are both approximately twice as fast as the BEM / nominal hybrid FEM, but the resistance values are now significantly different.

Having ensured that the networks are topologically identical, the BEM/FEM networks can be compared to the full FEM network by plotting the resistances from BEM/FEM networks against the full FEM resistances, as shown in Figure 6.5. Despite some significant differences, we observe clear correlation between the networks. In other words, we observe that the networks are mutually *consistent*.

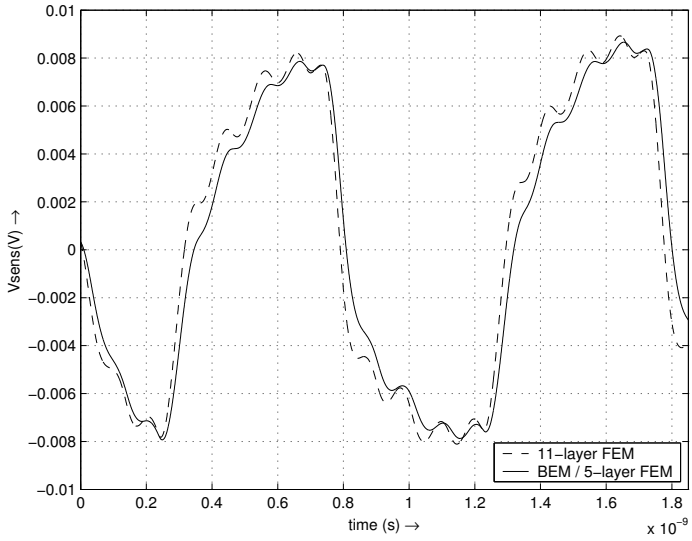

Figure 6.4: Simulations on the sensor node using networks generated with the 11-layer FEM and the BEM / 5-layer FEM methods. There is only a small difference between the waveforms.

The overall accuracy of the BEM/FEM networks with respect to the full FEM can be calculated in a straightforward way. Let $R_{BF_{ij}}$ be a resistance value in the BEM/FEM network and let $R_{F_{ij}}$ be the corresponding resistance value from the full FEM network. For each resistance value, we can then define the ratio A_r as follows:

$$A_r = \frac{R_{BF_{ij}}}{R_{F_{ij}}} \tag{6.1}$$

The ratios for all resistances are shown graphically in Figure 6.6. The numbering of the resistances along the horizontal axis is as follows:

- 1–10: Resistances connected to the sensor node. In particular, resistance 1 represents R(vss,sens) and resistance 2 represents R(sens,t1) from Table 6.1.
- 11–19: Resistances connected to the V_{ss} node (except R(vss,sens), which was already captured under 1–10). In particular, resistance 11 represents R(vss,t1) from Table 6.1.
- 20–55: Remaining resistances in the network, all connected to the backgates of the transistors.

We observe that the BEM / 5-layer FEM and the BEM / nominal hybrid FEM both have a ratio of approximately 1 in all resistances, indicating a very good match with

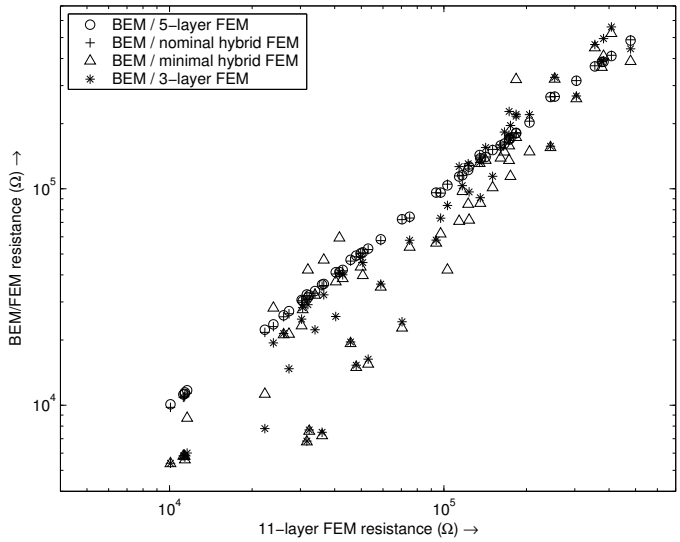

Figure 6.5: BEM/FEM resistances versus 11-layer FEM resistances. The differences may be significant, but correlation between the networks is apparent.

the full FEM network. For the BEM / minimal hybrid FEM and the BEM / 3-layer FEM the match with full FEM is poor.

Knowing that the resistance values for BEM/FEM using node contractions for the 2-terminal case from Table 5.9 were enclosed between the BEM/3DFEM and BEM/2DFEM, we might have expected that the data points for the BEM / hybrid FEM approaches for the multi-terminal case in Figures 6.5 and 6.6 would have been enclosed by the data points for BEM / 5-layer FEM and BEM / 3-layer FEM. However, this behaviour is not apparent from these figures. In fact, this behaviour is observed when the network is studied in terms of *port impedances* rather than the resistance values. The port impedance ratios (calculated in a similar way as in Equation 6.1) are shown in Figure 6.7, where the BEM / hybrid FEM situations are indeed enclosed. The numbering for the impedances in this figure is identical to the numbering in Figure 6.6. In particular, in comparing the BEM / minimal hybrid FEM to the BEM / 3-layer FEM we clearly observe that the port impedances are equivalent, except for port impedances 11–19. These port impedances are related to the V_{ss} node, which is connected to the substrate through a via. The two methods differ in that the BEM / 3-layer FEM applies node contraction near the via, whereas the BEM / minimal hybrid FEM does not. As we observe, the 3D effects around the via are apparently pronounced enough to show the modeling differences. However, we also observe that the BEM / minimal hybrid FEM, which captures 3D effects only around via connections to the substrate, is not sufficient to approach the reference. The BEM / nominal hybrid FEM, however, which captures 3D effects around via connections, but also

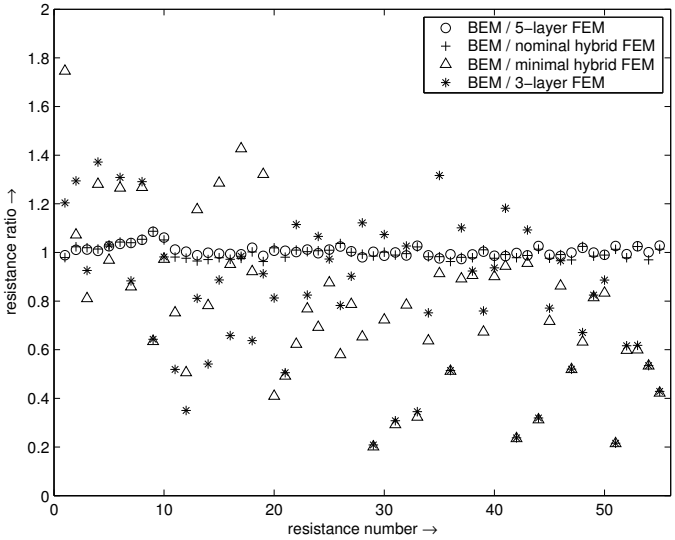

Figure 6.6: Resistance ratios for BEM/FEM versus 11-layer FEM. Only the BEM / 5-layer FEM and BEM / nominal hybrid FEM show an overall resistance ratio of approximately 1, indicating a good match with the 11-layer FEM.

underneath the polysilicon interconnect, matches the reference almost exactly.

Figure 6.8 shows the simulations on the sensor node indicated in the layout from Figure 6.1. The simulations show that the BEM / 5-layer FEM, and the BEM / hybrid FEM are close to each other, while the fully contracted situation is further away. This is consistent with the previous figures, except that the simulation with the BEM / minimal hybrid FEM is much better than what might be expected. This is most probably a coincidence, because the errors in the BEM / minimal hybrid FEM and BEM / 3-layer FEM resistance networks are in the same order (almost equivalent mean and standard deviation) and only a few port impedances calculated by the BEM / minimal hybrid FEM are slightly better than those calculated by the BEM / 3-layer FEM. This observation is confirmed by Figure 6.9, which shows simulations of the waveforms on the backgate of transistor 't4' from Figure 6.1. Here, the waveform for BEM / minimal hybrid FEM overlaps with BEM / 3-layer FEM waveform, while the other waveforms have a greater amplitude (though they are all similar in shape).

As a final observation, the simulations in Figure 6.8 all resemble each other, and the differences are only marginal. In particular, the accuracy of the BEM / 3-layer FEM may be the worst of the studied situations, but should not be considered 'poor', because it provides a consistent approximation of the BEM / 5-layer FEM method. Apparently, a *consistent* model is the most relevant factor in accurate simulations. In general, the behaviour of the methods is stable with respect to each other.

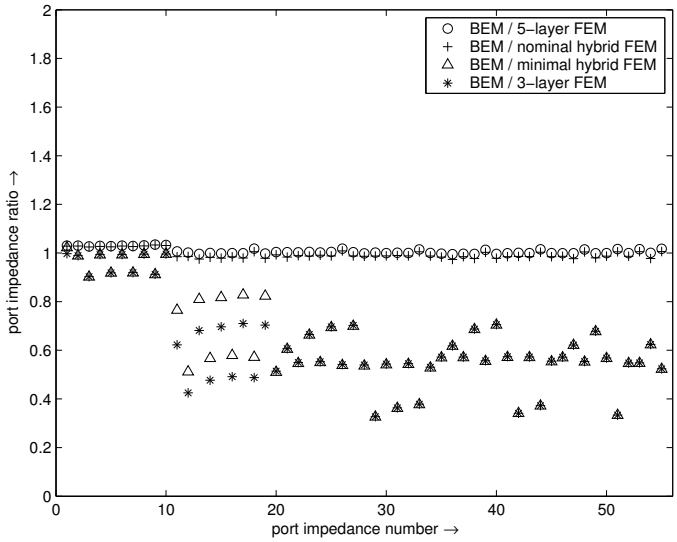

Figure 6.7: Port impedance ratios for BEM/FEM versus 11-layer FEM. Only the BEM / 5-layer FEM and BEM / nominal hybrid FEM show an overall port impedance ratio of approximately 1, indicating a good match with the 11-layer FEM. The BEM / minimal hybrid FEM shows only a small improvement over the BEM / 3-layer FEM.

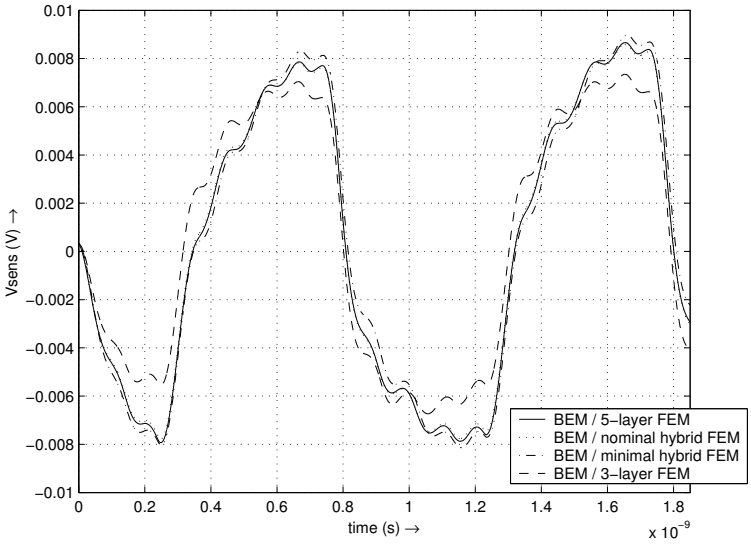

Figure 6.8: Simulations on the sensor node using networks generated with the BEM / 5-layer FEM, BEM / hybrid FEM and BEM / 3-layer FEM. All waveforms have the same shape, but the BEM / 3-layer FEM waveform has a smaller amplitude than the BEM / 5-layer FEM and BEM / hybrid FEM waveforms.

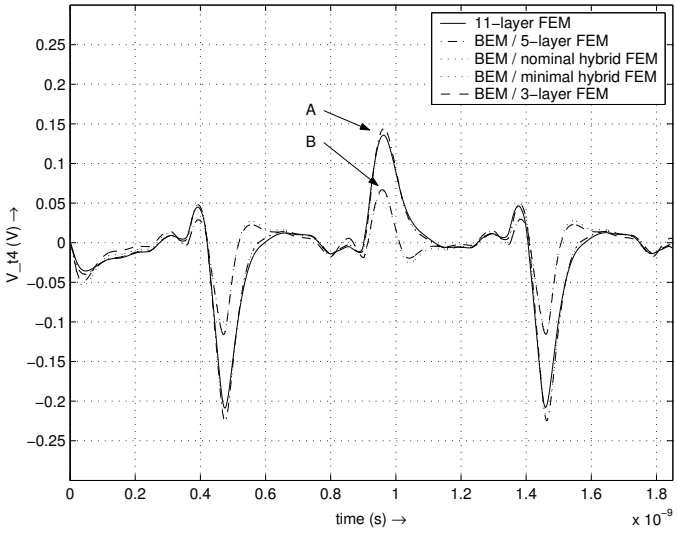

Figure 6.9: Simulations on the backgate of transistor 't4' from Figure 6.1 using the full FEM and BEM/FEM methods. All waveforms have the same shape, but the waveforms under Marker B have a smaller amplitude than those under Marker A. Marker A: full FEM, BEM / 5-layer FEM and BEM / nominal hybrid FEM. Marker B: BEM / minimal hybrid FEM and BEM / 3-layer FEM.

Table 6.2: Selected resistance values from the networks for an increasing number of layers in the layered FEM. Convergence in the resistance values is apparent.

# FEM layers	9	11	16	20
R(vss, sens) ($k\Omega$)	272.4	135.8	97.82	99.85
R(vss, t1) ($k\Omega$)	12.36	18.66	20.12	21.11
R(sens, t1) ($k\Omega$)	551.9	432.5	306.6	316.4

6.2.2 Layered FEM reference

As explained in Section 6.2.1, the BEM / 5-layer FEM method compares well to its equivalent 11-layer FEM method, while the extraction may be some 4–8 times faster. However, the layered FEM consisted only of 11 layers while the horizontal discretization was relatively coarse at maximum tile dimensions of $50\lambda\times50\lambda$, which is not expected to be a fully converged situation of the layered FEM. Therefore, this section will introduce a more fully converged situation of the layered FEM, which can subsequently be used as a reference for the BEM/FEM method in the accuracy study of Section 6.2.3.

From the vertical discretization from Table 5.1, we derived a coarser version consisting of 8 conductive layers (see Table B.1 in Appendix B) and two refined versions consisting of 15 and 19 conductive layers (see Tables B.3 and B.4 in Appendix B). These approaches result in a 9-layer, 16-layer and 20-layer FEM respectively. The 4 different vertical discretizations each use the same piecewise constant approximation of the doping profile shown in Figure 5.7. In this way, we may show convergence of the layered FEM method itself. It is also possible to use e.g. the 20-layer FEM for a more accurate piecewise constant approximation of the actual doping pofile. Such an approach would show convergence towards the physical situation, which is addressed briefly in Appendix C.

From Table 5.4, we previously observed that strong refinement in the horizontal discretization is necessary for convergence. However, due to the excessive computation times caused by the non-optimized solver, it was not feasible to choose the strongest refinement in the horizontal discretization in obtaining the layered FEM reference for the present example. Therefore, the horizontal discretization was chosen such that the reference is as close as possible to convergence, without using infeasible extraction times. Under these considerations, the horizontal discretization was chosen as $10\lambda\times10\lambda$ and the longest observed extraction time was 7 days, which occurred, as expected, for the 20-layer FEM.

Keeping the horizontal discretization fixed at $10\lambda\times10\lambda$, while increasing the vertical discretization from 9, through 11 and 16 to the 20-layer FEM, Table 6.2 and

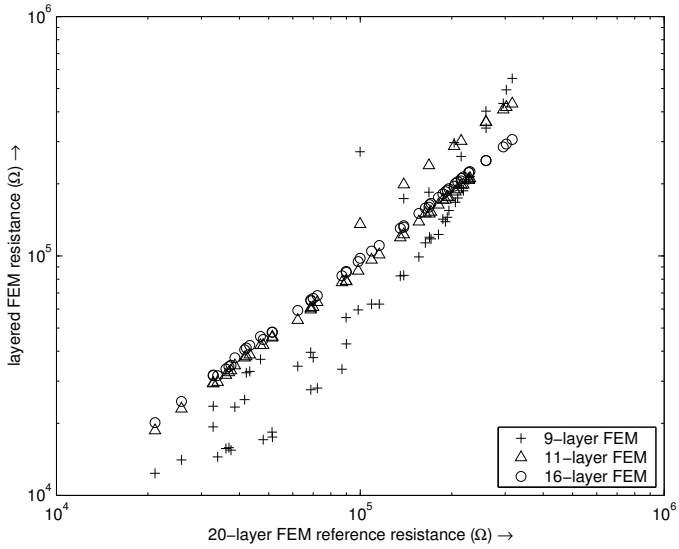

Figure 6.10: 9-, 11- and 16-layer FEM resistances versus 20-layer FEM resistances. Correlation improves as more FEM layers are introduced, indicating convergence.

Figures 6.10 and 6.11 confirm the converging behaviour in the resistance values under refinements in the vertical discretization. Also in simulations on the sensor node 'sens' and on the backgate of transistor 't4' we observe the convergence, as shown in Figures 6.12 and 6.13.

With respect to Figure 6.11, which uses the same numbering scheme for the resistances as Figures 6.6 and 6.7, it is interesting to note that the 11-layer FEM approximates the *local* resistances inside the ring oscillator very well, represented by resistances 11–55, while the *global* resistances to the 'sens' node contain significant errors, represented by resistances 1–10. It seems that local couplings mainly take place through the channel-stop layer, which is apparently taken into account with sufficient accuracy by the 2 conductive layers used in the 11-layer and 16-layer FEM, because a refinement to 3 conductive layers in the 20-layer FEM has no significant impact on the values of the local resistances. For the global resistances, however, the 16-layer FEM refines significantly with respect to the 11-layer FEM, which apparently takes the global resistances a step closer to convergence.

Having confirmed the converged behaviour of the 20-layer FEM at a horizontal discretization of $10\lambda \times 10\lambda$, we will now use it as a reference in a comparison with the BEM / 5-layer FEM.

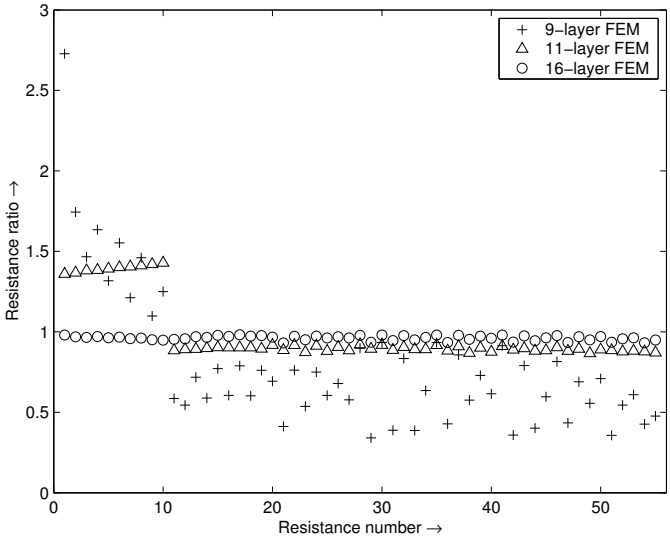

Figure 6.11: Resistance ratios for the 9-, 11- and 16-layer FEM versus the 20-layer FEM. Overall resistance ratio approaches 1 as more FEM layers are introduced, indicating convergence.

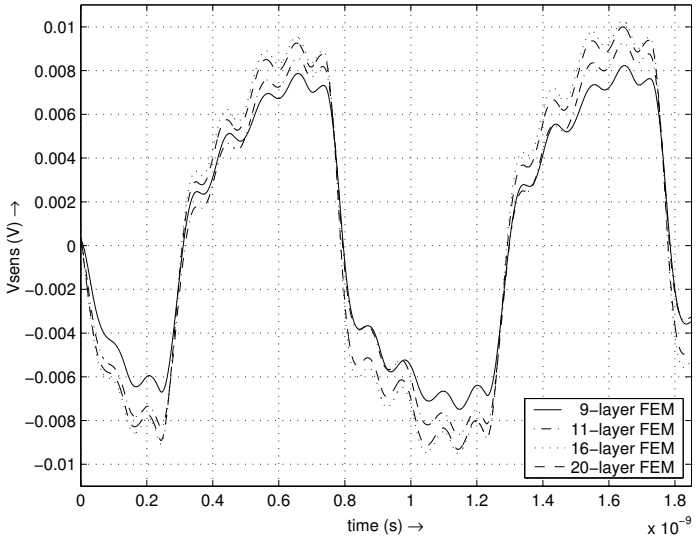

Figure 6.12: Simulations on the sensor node using networks generated with an increasing number of FEM layers in the layered FEM approach. All waveforms have the same shape, but slightly different amplitude. The waveforms approach the 20-layer FEM case as more FEM layers are introduced, indicating convergence.

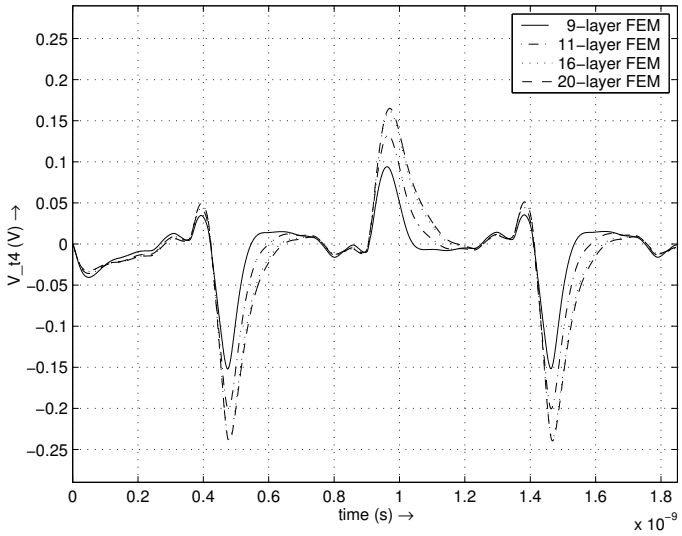

Figure 6.13: Simulations on the backgate of transistor 't4' using networks generated with an increasing number of FEM layers in the layered FEM approach. All waveforms have the same shape, but slightly different amplitude. The waveforms approach the 20-layer FEM case as more FEM layers are introduced, indicating convergence.

6.2.3 BEM / (contracted) FEM vs. layered FEM reference

This section will study the accuracy of the BEM / 5-layer FEM method compared to the 20-layer FEM reference from the previous section. In fact, the data for the BEM / 5-layer FEM method was already generated in Section 6.2.1 and compared to the 11-layer FEM there, such that the present section can suffice with providing BEM / 5-layer FEM versus layered FEM plots that are similar to those from Section 6.2.1, but now show comparisons to the 20-layer FEM reference from the previous section. Note that the comparison involves the BEM / 5-layer FEM with its *nominal* extraction settings, predominantly determined by the maximum tile dimensions of $50\lambda \times 50\lambda$, while the 20-layer FEM reference uses *refined* extraction settings determined by the maximum tile dimensions of $10\lambda \times 10\lambda$.

Figures 6.14 and 6.15 show the BEM/FEM networks compared to the 20-layer reference. Unfortunately, the BEM/FEM networks compare poorly to the layered FEM reference. Nevertheless, some clustering along the diagonal in Figure 6.14 and resistance ratios that neither closely approach zero nor reach very large values in Figure 6.15 confirm that correlation between the networks is not strong, but may still be considered to be present.

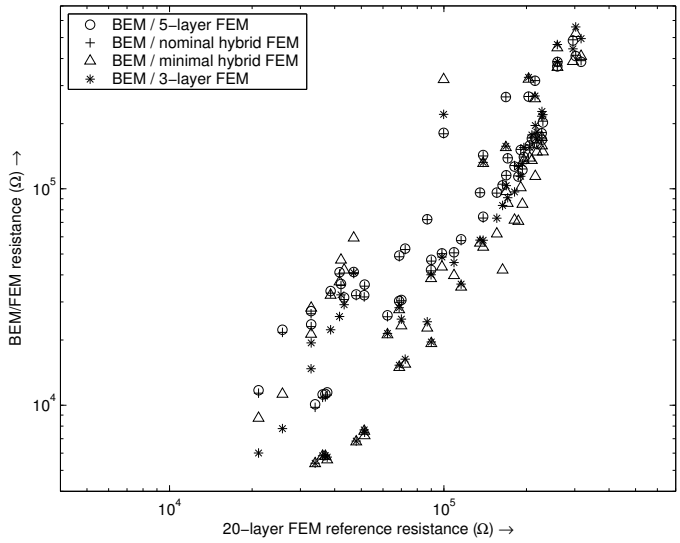

Figure 6.14: BEM/FEM resistances versus 20-layer FEM reference resistances. The differences are large, but correlation is still observed.

In the resulting simulations, however, the match between the resulting waveforms is much better than what initially might have been expected from the poort match in the resistance values. This is illustrated for BEM / 5-layer FEM versus the 20-layer FEM reference in Figures 6.16 and 6.17. The other BEM/FEM waveforms have been omitted from these two figures to avoid too many waveforms being plotted, but they closely relate to the BEM / 5-layer FEM waveform as previously shown in Figures 6.8 and 6.9 from Section 6.2.1.

The main common denominator in all these situations is the fact that all the discussed modeling methods provide *consistent* modeling of the physical situation, yet at different levels of accuracy. Despite sometimes large errors and weak correlation in the resulting networks, the simulations are all comparable, despite some marginal differences. From this, we may conclude that *consistency* of the model with the physics is significantly more important than the highest accuracy in the model itself. This also suggests that a relatively coarse, but consistent, substrate model is sufficient for accuracy in the simulations. As a result, the substrate model may quickly be obtained using coarse extraction settings.

6.2.4 Impact of Variations in BEM/FEM Extraction Parameters

As a final step in the practical analysis of the BEM/FEM method we check the impact of the BEM and FEM extraction parameters on the resulting BEM/FEM model. This can be achieved by a refinement of the extraction parameters with respect to the nomi-

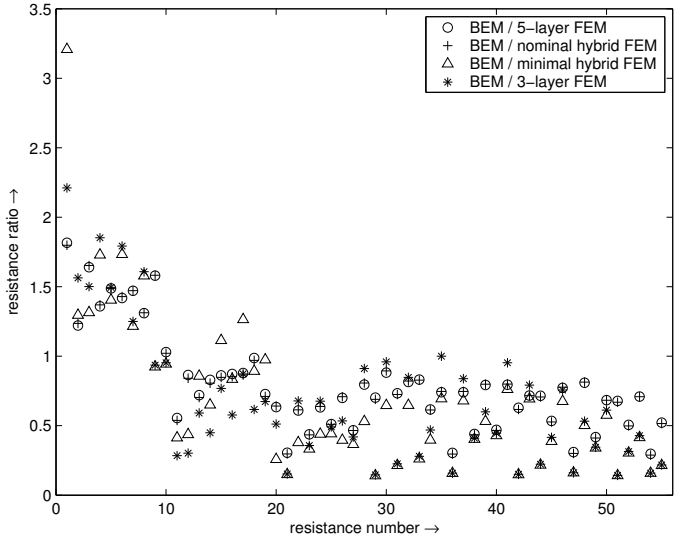

Figure 6.15: Resistances ratios for BEM/FEM versus the 20-layer FEM reference. Most resistance ratios do not approach one, indicating a poor match with the 20-layer FEM, but the ratios do not closely approach zero, nor do they reach very large values.

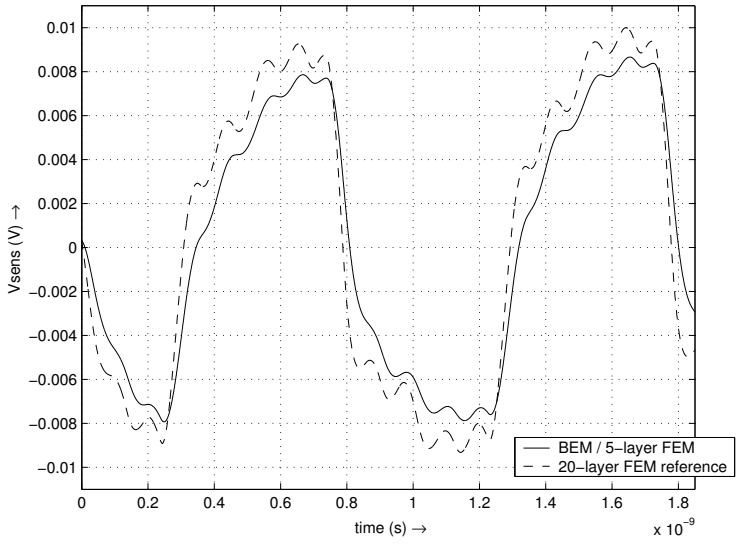

Figure 6.16: Simulations on the sensor node using networks generated by BEM / 5-layer FEM and the 20-layer FEM reference. The waveforms are only marginally different.

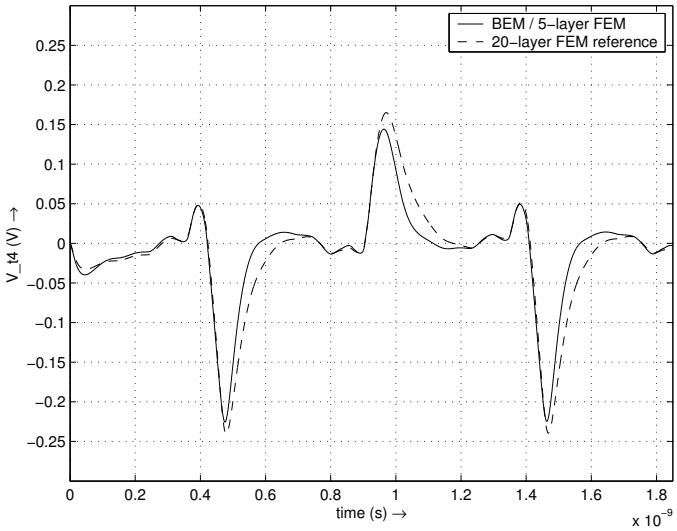

Figure 6.17: Simulations on the backgate of transistor 't4' using networks generated by BEM / 5-layer FEM and the 20-layer FEM reference. The waveforms are only marginally different.

nal situation presented at the beginning of Section 6.2. Here, we will particularly study the impact of the BEM window-size, the number of BEM panels per interface tile, and reduced maximum dimensions of the tiles in the horizontal discretization. The impact of the vertical discretization was already illustrated in Section 6.2.2, and, knowing that the BEM/FEM method can approximate its equivalent layered FEM situation closely, will not be considered here any further.

For this illustration, we chose 3 cases with particular refinements over the default parameters. The details, extraction statistics and a selection of resistances are shown in Table 6.3. Note that, in comparison to Table 6.1, the halos for the node-contraction are not present in this case, and therefore do not have influence on the tile division. Thus, the default BEM / FEM extraction shown in Table 6.3 is comparable to the BEM / 5-layer FEM extraction shown in Table 6.1, but uses a coarser tile division and is therefore significantly faster.

A comparison of the three refined situations to the default situation is shown in Figures 6.18 and 6.19. Despite differences that may be considered significant, the match between the networks can be considered (very) good, especially considering the much longer extraction times and much higher memory usage for the extractions with refined settings. Additionally, as shown in Figure 6.20, the simulations with the resulting networks seem to change very little with respect to the nominal BEM/FEM situation.

Table 6.3: The Impact of the BEM and FEM settings. The resistance values change only marginally, but computation times and memory usage increase significantly, or even drastically.

	fine FEM / large BEM win	def. FEM / fine BEM	fine FEM / def. BEM	default BEM/FEM
max. tile dimensions	$30\lambda \times 30\lambda$	$50\lambda \times 50\lambda$	$30\lambda \times 30\lambda$	$50\lambda \times 50\lambda$
BEM win (μm)	12.5	12.5	1.5625	1.5625
# BEM panels per tile	1	8	1	1
time (min:sec)	790:36	45:33	49:12	25:32
memory (Mb)	811	340	84.8	57.7
R(vss, sens) ($k\Omega$)	172.2	156.1	167.4	159.0
R(vss, t1) ($k\Omega$)	10.47	9.500	11.08	9.606
R(sens, t1) ($k\Omega$)	431.0	425.6	458.0	450.0

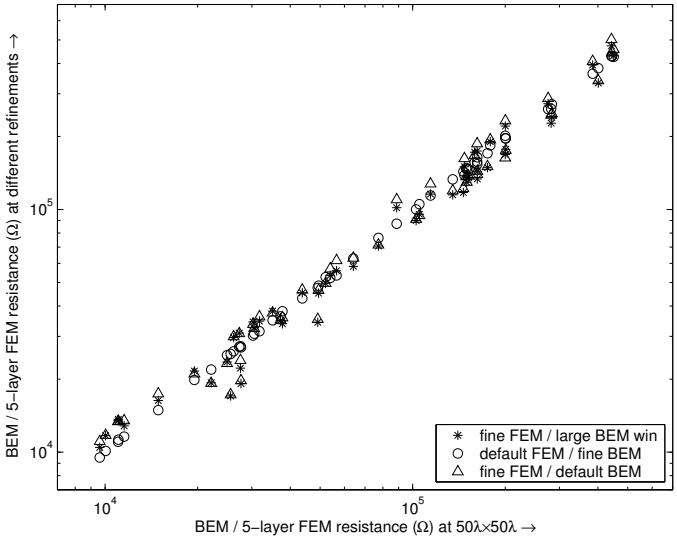

Figure 6.18: BEM/FEM resistances generated with refined extraction settings versus BEM/FEM resistances generated with the default extraction settings. The differences between the networks are small, a strong correlation is observed.

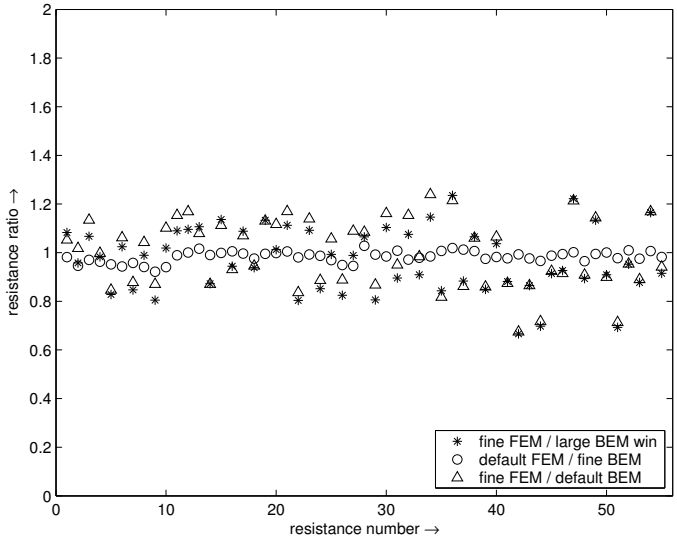

Figure 6.19: Resistance ratios of BEM/FEM resistances generated with refined extraction settings versus BEM/FEM resistances generated with the default extraction settings. All resistance ratios are near 1, indicating a good match between the networks.

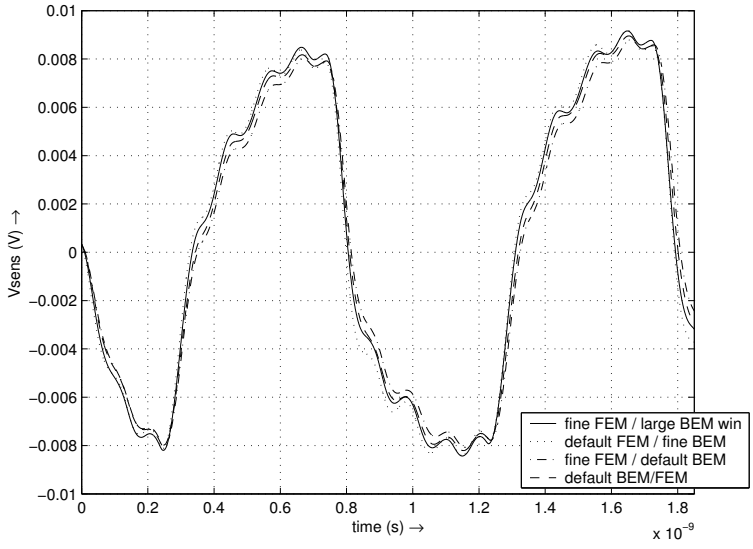

Figure 6.20: Simulations on the sensor node using networks generated by the BEM/FEM method with different refinements in the extraction settings. The waveforms are almost identical.

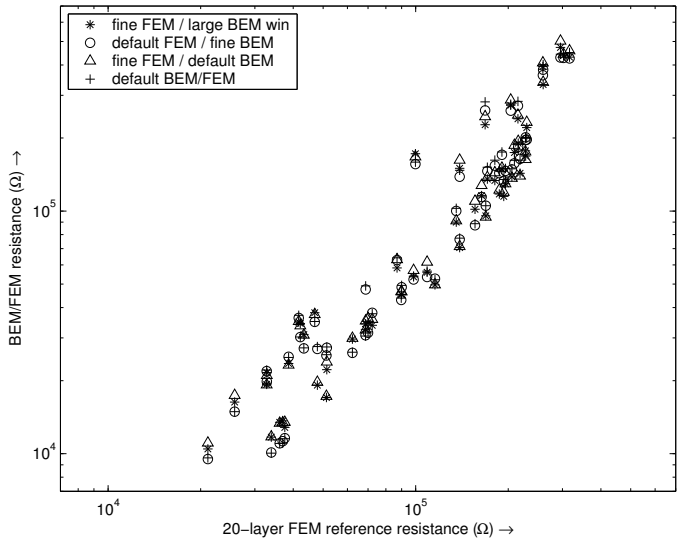

Figure 6.21: BEM/FEM resistances generated with different refinements in the extraction settings versus 20-layer FEM resistances. The differences are significant, but correlation is observed.

Furthermore, knowing that the 'default FEM / fine BEM' extraction only refines the BEM part with respect to the default settings from the previous sections, we observe that the BEM settings significantly increase the extraction time and memory usage of the overall BEM/FEM method, but contribute relatively little to the overall accuracy of the model and the simulations. This confirms the previous observation that the deep substrate plays only a limited role in the overall model and can therefore be captured with a coarse and sparse BEM. Under these circumstances, the BEM/FEM method will be the most efficient, and will combine the good characteristics of the BEM and the FEM.

A comparison to the 20-layer FEM reference from Section 6.2.2 is shown in Figures 6.21 and 6.22. We observe that, compared to default BEM/FEM, refinements in the BEM/FEM extraction settings do not significantly improve the accuracy of the networks with respect to the reference. The same holds for the simulations in Figure 6.23.

From these results we can conclude that the simulations change only marginally, almost irrespective of the accuracy of the substrate model itself. Apparently, the extraction settings for the BEM/FEM, or layered FEM, methods have influence on the accuracy of the resulting networks, as well as on the extraction times and memory usage, but do not significantly change the level of consistency between the physical situation and the resulting models. As it seems, the consistency between the physical situation and the models predominantly determines the behaviour of the simulations.

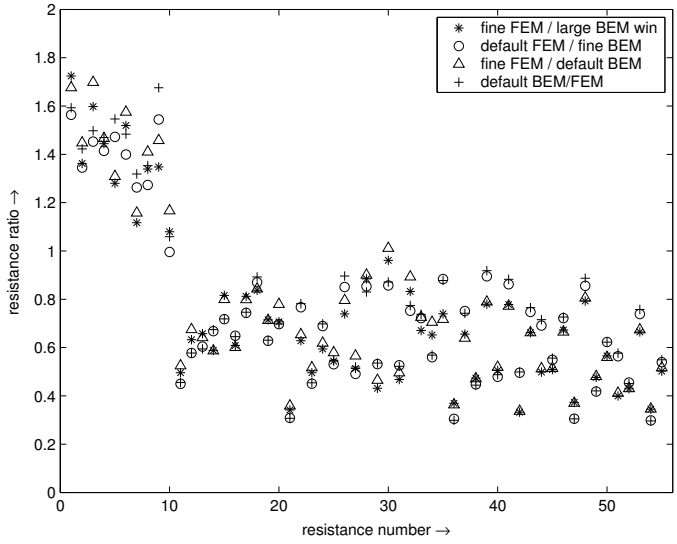

Figure 6.22: Resistance ratios for BEM/FEM generated with different refinements in the extraction settings versus the 20-layer FEM. Most resistance ratios do not approach one, indicating a poor match with the 20-layer FEM, but the ratios do not closely approach zero, nor do they reach very large values.

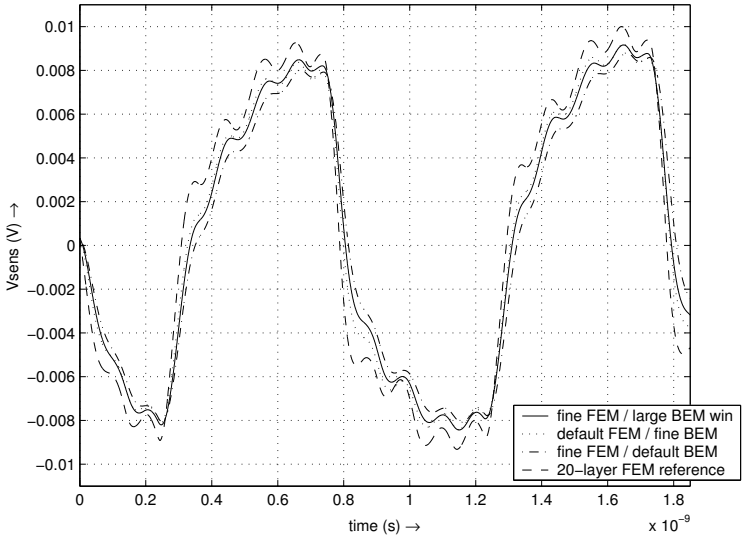

Figure 6.23: Simulations on the sensor node using networks generated by the 20-layer FEM reference and the BEM/FEM method with different refinements in the extraction settings. The waveforms differ only marginally.

From this we may conclude that relatively coarse settings in the BEM/FEM extraction are already sufficient for accuracy in the simulations, while the BEM/FEM method works as efficiently as possible. A similar observation is apparently also exploited by the FDM-based substrate modeling tool SubstrateStorm, as mentioned previously.

6.3 Consistency in the Substrate Model

As observed in the previous sections, *consistency* between the substrate model and the physics plays a very important role in the simulations, even more than the highest accuracy. In this respect, the statement from Section 5.2 that the n-well would be assumed an ideally conducting domain in the extractions and simulations that followed, may already have introduced inconsistencies in the resulting models, possibly even invalidating the results from the previous sections. To what extent these inconsistencies have impact on the results will be studied in Section 6.3.1.

In the same context of model consistency, the inconsistencies introduced by a purely BEM-based substrate model (as already shown in Section 5.6.4) will be studied in Section 6.3.2.

6.3.1 The Impact of the Well-Resistance

As announced at the end of Section 5.2, the n-well is assumed an ideally conducting domain. In this section, we will identify whether this assumption is valid in this example, and whether it will affect the consistency of the overall model.

The n-type doping concentration in the n-well is approximately $4 \cdot 10^{15} cm^{-3}$ (recall Figure 2.3). This is approximately equivalent to a conductivity of 100 S/m [Nea92]. Together with the known thickness of the well, 5.4 μm, a layered FEM may be constructed according to this data. The extractions and simulations below used a 2-layer FEM for the n-well, which was added to the BEM / 5-layer FEM approach from Section 5.3.2 for the remainder of the substrate. The default extraction settings were used. The extraction took almost exactly 3 hours, which is well over 3 times longer than the corresponding BEM / 5-layer FEM extraction from Table 6.1. This is caused by the relatively fine FEM mesh and BEM contact division in, and underneath, the n-well due to the presence of many layout features like the transistors, contacts and interconnect. Note, however, that the resistances shown in Table 6.1 do not change by taking the n-well into account with a layered approach, because the n-well model is completely separated from the remainder of the substrate model by the junction capacitance.

The simulation results are shown in Figure 6.24. The waveforms are almost identical, which confirms that the n-well may be considered an ideally conducting domain

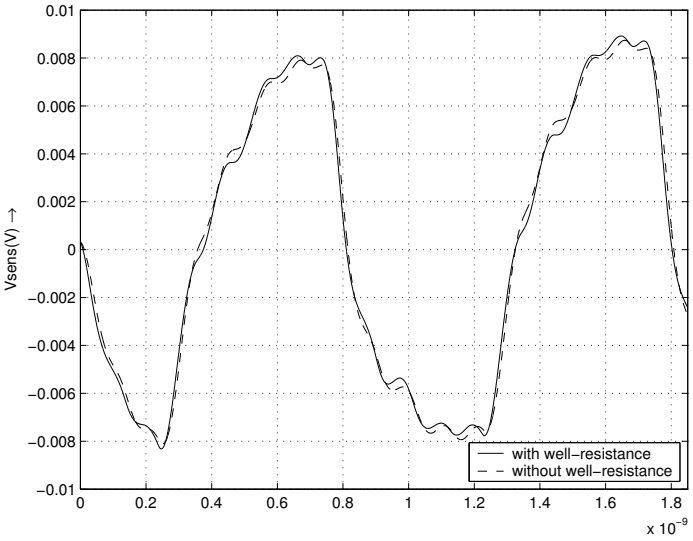

Figure 6.24: Simulations on the sensor node with and without taking the well-resistance into account. The waveforms are almost identical.

in this example, without introducing significant inconsistencies. However, this might have been otherwise if the n-well would have covered much larger areas in a larger circuit.

6.3.2 BEM/FEM Approach vs. Plain BEM Approach

Having established the behaviour of the BEM/FEM method with respect to (refinements in) its parameters, and its behaviour with respect to a full FEM method for the ring oscillator with guarded sensor node, the behaviour of the BEM/FEM method with respect to the plain BEM may also be established. Even though this comparison was already done for a simplified situation in Table 5.11 from Section 5.6.4, the comparison previously limited to resistance values may presently be expanded with a comparison of the simulated waveforms on the sensor node.

Recall that the plain BEM approach would lift the n-well out of the substrate, and either model the substrate with a single-layer BEM (i.e. ignore the channel-stop pattern) or a double-layer BEM (i.e. assume the channel-stop pattern to be a uniform layer). Clearly, the plain BEM approach would yield a *consistent* model in only a limited set of situations. The BEM/FEM method, though computationally typically more intensive, stays close to the actual physics of the substrate, and therefore provides a consistent model which is inherently more accurate.

Table 6.4: BEM/FEM vs. BEM. The differences between the BEM/FEM and plain BEM resistances are large, and consistency between the networks is not observed.

substrate model	BEM / 3-layer FEM	BEM / contracted 3-layer FEM	1-layer BEM	2-layer BEM
time (min:sec)	8:25	5:19	5:13	5:12
memory (Mb)	34.9	25.1	87.9	65.2
R(vss, sens) ($k\Omega$)	205.6	288.3	913.5	23.67
R(vss, t1) ($k\Omega$)	7.171	7.006	1791	14.82
R(sens, t1) ($k\Omega$)	567.8	528.7	266.4	9.864

In this simulation, we will compare the BEM/FEM method to the plain BEM. For a valid comparison, we will choose a situation where the BEM/FEM method performs well, and match the extraction times of the plain BEM to this situation. In this case, we have chosen for the minimum FEM settings in the BEM/FEM method: the 0.5 μm thick channel-stop layer (1000 S/m) will be handled with a 2-layer FEM, and the underlying 5 μm (10 S/m) will also be handled with a 2-layer FEM. The vertical FEM discretization in the top 5.5 μm then consists of three layers. This situation is actually a BEM/FEM version of the 9-layer FEM approach presented in Section 5.5 (see also Table B.1 in Appendix B).

Table 6.4 shows the extraction times and a selection of resistance values. Having observed that the BEM / full 3-layer FEM extraction took some 7 minutes, which would require a very strong refinement in the BEM mesh for a comparable extraction time, the extraction time was further reduced by applying node contractions in the channel-stop layer. The contraction halo was chosen according to the minimal settings introduced earlier, such that the 3-layer FEM is only applied near the contacts (V_{ss} and sens) and that it is contracted into a 2-layer FEM elsewhere in the channel-stop layer. We observe that the extraction time is then reduced to approximately 3.5 minutes, while the consistency between the models is preserved (also in comparison to Table 6.1). The settings for the plain BEM extractions are then chosen such that the extraction times are comparable to these 3.5 minutes. Under these circumstances, the 1-layer (10 S/m) and 2-layer BEM (1000 S/m top-layer, 10 S/m bulk) show a significant difference in the memory usage. This is because the 2-layer BEM uses a coarser mesh (smaller matrix) due to the slower convergence in the Green's function. We observe that the obtained resistance values are not consistent with the BEM/FEM networks and that the differences between the simulations in Figure 6.25 are large. Note also that the amplitude and shape of the BEM/FEM simulation in Figure 6.25 are comparable to the earlier BEM/FEM simulations from Section 6.2.1.

With respect to the performance of the BEM / 3-layer FEM method, we observe that the extraction is slightly faster than the BEM / minimal hybrid FEM extraction from Table 6.1, and that the obtained resistance values are consistent with each other.

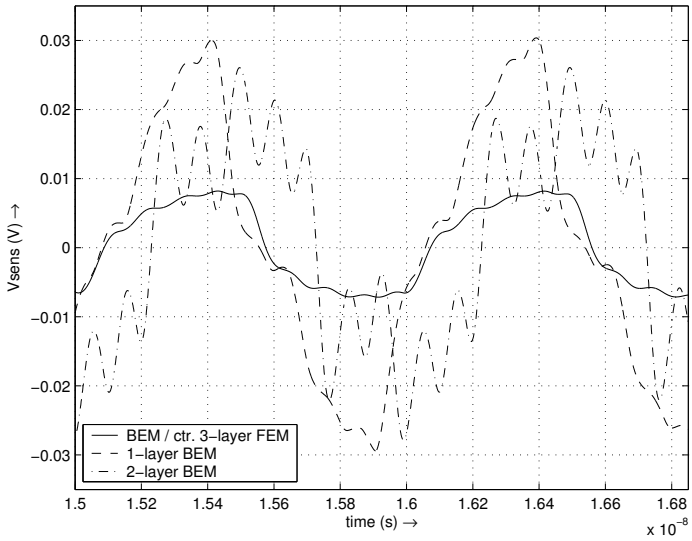

Figure 6.25: Simulations on the sensor node using networks generated by the BEM / contracted 3-layer FEM, the 1-layer BEM and the 2-layer BEM. The waveforms have significantly different shapes.

Knowing that the BEM/FEM method typically stays closer to the physics than the plain BEM approach, we can conclude that the BEM/FEM method reaches more accuracy in the same amount of time. Furthermore, high accuracy in a model that is not entirely consistent is not sufficient for accurate simulations.

6.4 Overall Consistency in the Partitioned Problem

As mentioned in Chapter 1, Figure 1.3, the original modeling problem can be partitioned into an interconnect, transistor and substrate part. Having shown in the previous sections that the BEM/FEM method yields a consistent model for the substrate, the consistency of the overall model with respect to the other partitions may also be determined. In particular, we will study the relevance of the interconnect in the overall model; the details of the transistor models are outside the scope of this thesis.

In [Gen01] and [Sch03] the modeling and the importance of the interconnect capacitances in the substrate model were already discussed. The observations from those papers are also confirmed for the ring-oscillator with guarded sensor node. In summary, [Sch03] studied the relevance of coherence between the interconnect and the substrate models to the consistency of the overall model by identifying the separate and simultaneous noise contributions from the interconnect and the transistors to the

noise waveforms observed in the substrate. A similar study may be performed for the present example of the ring oscillator with guarded sensor node. In particular, 3 cases will be studied:

- Substrate noise from transistors. That is, separate noise injection from the transistors into the substrate; capacitive coupling from the interconnect to the substrate is connected to a fictive ideal ground.
- Substrate noise from interconnect. That is, separate noise injection from the interconnect into the substrate through capacitive coupling; the backgates of the transistors are connected to an ideal ground or voltage source, depending on the transistor type.
- Substrate noise from both. That is, simultaneous noise injection from interconnect and transistors into the substrate; the default case for the simulations in this thesis.

The simulation results from these 3 cases are shown in Figure 6.26. We observe that the capacitive coupling from the interconnect to the substrate is a strong contributor to the noise in the substrate. In the context of consistency in the model, the capacitive coupling between interconnect and substrate is apparently very important to take into account.

Apart from the interconnect capacitance, the interconnect resistance may also have influence on the results, as shown in Figure 6.27. Even though the waveforms are comparable in shape, omitting the interconnect resistance increases the oscillation frequency of the oscillator by some 7%. The speed-increase is almost entirely caused by neglecting the resistance of the polysilicon interconnect which carries the actual oscillating signal. This is because polysilicon has a relatively high resistance, which, when neglected, significantly reduces the RC-delay along the signal path.

Again, these simulations confirm that any model should be *consistent* with the physical situation, otherwise the resulting simulations may be significantly different.

6.5 Summary and Concluding Remarks

Based on the same objectives defined in Chapter 5, this chapter applied the BEM/FEM method to a more realistic example consisting of a ring oscillator with a sensor node at some distance away that was surrounded by an n-well guard ring.

The layered FEM, BEM/FEM and plain BEM modeling methods were applied to this layout. For comparable extraction settings, the BEM/FEM technique, using a coarse and sparse BEM, proved to be significantly faster than the layered FEM, particularly when using node contractions. The observed acceleration was about a

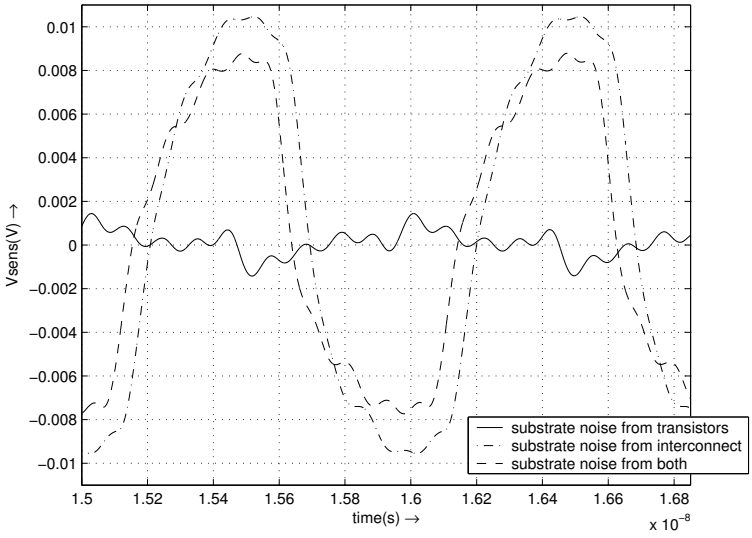

Figure 6.26: Simulations on the sensor node when substrate noise from transistors, interconnect, or both is taken into account. The waveforms have significantly different shapes.

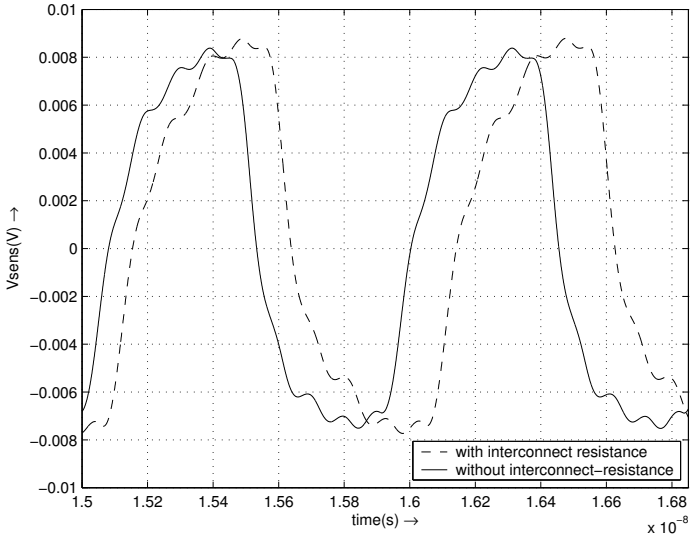

Figure 6.27: Simulations on the sensor node with and without taking the interconnect resistance into account. The waveforms have the same shape and approximately the same amplitude, but slightly different wavelength.

factor of 4 without loss of accuracy, to a factor of 8 with some loss of accuracy, while in all cases the memory usage was approximately equal and remained in the order of 40 – 100 Mb.

Regarding accuracy, the resistances in the layered FEM and BEM / (contracted) FEM networks sometimes showed significant differences, but, nonetheless, correlation between the networks was observed, indicating *consistency* between the networks. Simulations with these networks typically proved to be closer to each other than what initially might have been expected according to the differences in the networks. Indeed, the main common denominator between the layered FEM and BEM / (contracted) FEM methods is that they produce models that are consistent with the physical situation, yet at different levels of accuracy. Apparently, *consistency* of the models with the physical situation is the main factor that determines the behaviour of the simulations.

In an additional set of extractions, the BEM/FEM model showed stable behaviour for various settings in the BEM and FEM parts, and clearly reaches better accuracy and consistency in the resulting models than the plain BEM approach. Particularly, the previous observation that a coarse and sparse BEM is sufficient for accuracy in the overall BEM/FEM model, is confirmed. With respect to consistency in the overall model, capacitive coupling from interconnect to substrate was identified as having significant impact on the noise levels in the substrate.

In summary, the conceptual comparison between the FEM and BEM/FEM techniques presented in Section 4.5 is confirmed by the simulations in this chapter. For a layered FEM approach to the substrate, the FEM layers applied in the deep substrate are replaced by a coarse and sparse BEM in the BEM/FEM approach. The BEM part in the BEM/FEM extraction is then very fast and contributes relatively little time to the overall BEM/FEM extraction. Furthermore, the BEM/FEM model uses fewer variables than the layered FEM, but retains sparsity. As a result, the BEM/FEM method is faster than the layered FEM, but reaches equivalent accuracy. Knowing that a plain BEM can be even faster, but can typically not reach equivalent accuracy, the BEM/FEM method provides a new speed-accuracy trade-off that is a crossover between the individual BEM and FEM methods.

7

Conclusions

In this thesis, we have addressed a combined BEM/FEM method for substrate resistance modeling. The method consistently partitions the modeling problem, such that it exploits the structure of the substrate modeling problem and the properties of the BEM and FEM modeling techniques. In particular, the modeling problem is partitioned such that the FEM is applied in the top few microns of the substrate, while the BEM is applied for the deep substrate.

With respect to the combination, the most relevant property of the FEM is as follows:

- For a FEM with a piecewise linear basis function, the FEM discretization can be interpreted as a sparse resistance network.

Using this property, and exploiting the possibilities for contracting equipotential nodes in the FEM network, the local couplings in the substrate can be captured with a sparse, reduced-order resistance network.

The most relevant properties of the BEM are now as follows:

- The BEM reference node at infinity is a valid representation of the deep substrate in finite domains of relatively large dimensions.
- The horizontal resistances between the contacts in the BEM model quickly become large for increasing distance between the contacts, whereas the vertical resistances towards the reference node become constant. This implies that the

vertical resistances towards the reference node at infinity are relatively more important.

- A windowing technique applied to the BEM significantly improves sparsity in the BEM model and computational efficiency of the BEM technique at the cost of only little accuracy.

Using these properties, and the observation that the current density in the deep substrate is relatively small, the global couplings through the deep substrate can now be captured with a coarse and sparse BEM.

In combination, the overall result is then a sparse, reduced-order BEM/FEM method which uses inherently fewer variables than a full FEM, while retaining sparsity and accuracy. This is also reflected in the extraction times: using a standard solver based on Gaussian elimination ordered by minimum degree, BEM/FEM is 4–8 times faster than full FEM, while using an equivalent amount of memory. Ultimately we conclude that the combined BEM/FEM method inherits the good characteristics from its constituting parts. That is, the combined BEM/FEM method operates in a new trade-off between the speed of the BEM and the accuracy of the FEM.

Additional Observations

The first two BEM properties from the list above are actually related to a more fundamental observation: the infinite-domain BEM is a valid approximation of a finite-domain substrate with relatively large dimensions, as long as the contact areas are not nearer to the edges of the finite domain than a few percent of its lateral dimensions. Whereas this observation is easily verified for heavily doped substrates, this thesis showed that this observation also holds for lightly doped substrates.

With respect to the accuracy of the BEM/FEM method we conclude that we are apparently operating in an area where, on the one hand, accuracy in the networks can well be obtained, but where, on the other hand, high accuracy in the networks is not crucial for accuracy in the simulations. The networks from the previous chapter sometimes showed poor accuracy in individual resistances, but nevertheless the networks proved to be globally consistent with each other. Simulations with these networks then showed only marginal differences in the resulting waveforms. The main common denominator between these networks is the physical situation, for which they are all consistent approximations, but at different levels of accuracy. From this, we conclude that *consistency* of the model with the physics is the most relevant factor in accurate simulations. This observation is apparently already applied in some commercial substrate modeling tools, which tend to operate at some distance from the convergence point and, with that, apparently aim for consistency, rather than the highest accuracy. Therefore, this observation may also be valuable in future optimizations to the efficiency of the BEM/FEM method, or in the development of new modeling methods.

Future Work

Even though the BEM/FEM method has shown to be faster than the FEM and more accurate than the BEM, its performance may still be optimized further.

For example, the extractions in this thesis all used a straightforward solver based on Gaussian elimination ordered by minimum degree. Even though this solver allowed a fair comparison between the different methods, other solvers may have a better computational efficiency. Therefore, future research into dedicated solvers may find a way to improve the overall efficiency of the layered FEM and BEM/FEM methods. However, the BEM/FEM method may require a different dedicated solver than the FEM method, making a fair comparison between the methods less straightforward.

Other optimizations to the solver may be aimed at the sparsity of the resulting model. For example, the plain BEM approach does not require internal nodes, therefore does not suffer from fill-in during the windowed solution procedure and results in a sparse model between the terminal nodes, where the sparsity is guided by the window size used in the extraction. The BEM/FEM modeling approach, however, despite using a typically small BEM window size, does use internal nodes which cause fill-in during the Gaussian solution procedure, resulting in an inherently full model between the terminal nodes. Future research on global solution methods for combined BEM/FEM problems may find a way to avoid this fill-in and provide an overall sparse model between the terminal nodes.

Apart from improvements to the solver, the BEM/FEM method may also be optimized in other respects. Currently, the FEM mesh, as wel as the BEM/FEM interface mesh, is generated through a straightforward meshing procedure based on a tile-wise division of the layout. The meshing procedure may be optimized (possibly including adaptive meshing) for better performance of the overall BEM/FEM method. In this context, it may be viable to apply a Delaunay triangulation instead of the tile-wise meshing approach. This might also suggest applying a Delaunay-based interpolation method in the top few microns of the substrate, instead of a FEM. Combined with the coarse and sparse BEM for the deep substrate, the resulting BEM / interpolation method may be more consistent with the physical situation than traditional interpolation methods, while taking relatively little extra computation time. Future research into this approach could be worthwile.

Additionally, the problem partitioning may also be applied in a different way. Currently, the whole top few microns of the substrate are handled with the FEM, such that the interface towards the BEM is a plane flat on top of the deep substrate; the current implementation of the BEM is optimized for flat contacts on top of a resistive domain. If the BEM capabilities would be extended to accommodate 3D structures rather than flat contacts, the FEM domain can be reduced to a minimum and a performance gain might be expected. However, such an approach to the BEM requires a more advanced Green's function, and may require heavier BEM-settings to reach appropriate accuracy. In particular, the BEM part of the BEM/FEM model will then

play a more prominent role capturing local couplings, rather than just coarsely capturing the global couplings, so the BEM will most probably require finer mesh settings and a larger BEM window to reach the required accuracy. It requires further study to determine to which extent such an approach would actually improve (or decrease) the performance trade-off of the BEM/FEM method.

A

Technology Descriptions

Technology information may be passed to the SPACE layout to circuit extractor through a technology file. The general syntax of the technology data can be found in [MSp]. Below, two ways of passing technology data to the BEM/FEM and layered FEM methods are discussed. Additionally, the way of passing technology data to the traditional, plain BEM approach is also presented for comparison.

The units applied in the technology descriptions below are as follows:

- distance: μm
- conductivity: S/m
- sheet resistance: Ω/sq
- vertical resistance: $\Omega \cdot \mu m^2$
- area capacitance ('acap'): $aF/\mu m^2$
- edge capacitance ('ecap'): $aF/\mu m$

A.1 High-Level Description

With the application of the BEM/FEM method in the SPACE layout-to-circuit extractor, and the introduction of the layered FEM, the technology language of the SPACE

layout-to circuit extractor was expanded with some appropriate technology statements. These statements may be called as follows. First, the thickness of the FEM domain may be defined by a depth offset in the top of the BEM domain

```
set bem_depth <depth>
```

Then, the characteristics of the FEM area may be defined through a number of 'wafer' statements which have the following general syntax:

```
wafer : condition : <conductivity> <thickness> <#layers> : <switches>
```

Basically, this statement divides a layer of given thickness and conductivity into the given number of layers, according to the principle from Figure 4.16. The switches provide additional information and may be any combination of the following three statements:

- restype: defines whether the doping-type of the domain is p or n (default: restype=p).
- subconn: defines whether the bottom of the layer-stack should be connected to the BEM (default: subconn=on).
- viamask: defines from which mask the halo for the node contractions is derived (no default).

Stacked wafer statements (i.e. defined by the same condition) will be connected along their common interface, unless the 'restype' is different (in that case, a junction capacitance should be defined elsewhere in the technology description). The same holds for adjacent wafer statements. Note, that each (stacked) wafer statement should match the bem_depth.

As an example, the channel-stop layer and n-well may be captured in the technology language as follows.

```
set bem_depth 5.5

# channel-stop
  wafer : cs  : 1000   0.5   3 :
  wafer : cs  :   10   5.0   3 :

# n-well
  wafer : cwn :  100   5.5   1 : restype=n subconn=off
```

The channel-stop area uses a stack of two wafer statements, of which the total thickness matches the bem_depth. The first wafer statement results in 3 FEM layers (w1_1,

w1_2 and w1_3, from top to bottom), and the second wafer statement also results in 3 FEM layers (w2_1, w2_2 and w2_3, from top to bottom). Because the two wafer statements are stacked, layers w1_3 and w2_1 are adjacent, and w2_1 is merged into the w1_3 layer. The n-well wafer statement results in only a single FEM-layer: w3_1.

Junction capacitances between the n-well and the remainder of the substrate may be defined as follows. Note that '@sub' in the 'acap' statement refers to the top of the BEM domain.

```
junction capacitances nwell :
  # name      :  condition     : mask1 mask2 : capacitivity

  # bottom
    acap_cw   : cwn            : w3_1 @sub   : 100

  # sidewall
    ecap_cw1 : cs -cwn         : w1_1 -w3_1 :  12.5
    ecap_cw2 : cs -cwn         : w1_2 -w3_1 :  25
    ecap_cw3 : cs -cwn         : w1_3 -w3_1 : 137.5
    ecap_cw5 : cs -cwn         : w2_2 -w3_1 : 250
    ecap_cw6 : cs -cwn         : w2_3 -w3_1 : 125
```

Finally, the deep substrate may be defined as follows

```
sublayers :
  # name        conductivity  top
    substrate       10        0.0
```

Note that the top of the deep substrate is now relative to the bem_depth.

A.2 Low-Level Description

If the high-level technology description is not sufficient to properly capture the technology situation, also a low-level technology description is possible. This basically requires to define the individual FEM layers in the layered FEM manually.

For example, the conductive layers in the vertical substrate discretization from Table 5.1 may be mapped to a layered FEM, and BEM/FEM description as follows.

A.2.1 layered FEM

The layered FEM is generated according to Table 5.1. The description defines the horizontal and vertical resistivities for the FEM discretization, and the capacitance

statements for the n-well. Notice the distributed approach to the sidewall capacitances of the n-well.

```
# layered FEM sheet resistances
conductors :
  # name      : condition : mask : resistivity
    cond_cs1  :    cs1    : cs1  :       25316
    cond_cs2  :    cs2    : cs2  :        4804
    cond_cs3  :    cs3    : cs3  :        3997
    cond_cs4  :    cs4    : cs4  :       11117
    cond_cs5  :    cs5    : cs5  :       95511
    cond_cs6  :    cs6    : cs6  :      286944
    cond_cs7  :    cs7    : cs7  :       34288
    cond_cs8  :    cs8    : cs8  :        7770
    cond_cs9  :    cs9    : cs9  :        2130
    cond_cs10 :    cs10   : cs10 :         583.6
    cond_cs11 :    cs11   : cs11 :         743.5

# layered FEM vertical resistances
contacts :
  # name      : condition : lay1  lay2 : resistivity
    cont_1    : cs1  cs2  : cs1   cs2  :         126.6
    cont_2    : cs2  cs3  : cs2   cs3  :         354.9
    cont_3    : cs3  cs4  : cs3   cs4  :       65200
    cont_4    : cs4  cs5  : cs4   cs5  :      169000
    cont_5    : cs5  cs6  : cs5   cs6  :       40400
    cont_6    : cs6  cs7  : cs6   cs7  :       29300
    cont_7    : cs7  cs8  : cs7   cs8  :      554000
    cont_8    : cs8  cs9  : cs8   cs9  :     2020000
    cont_9    : cs9  cs10 : cs9   cs10 :     7370000
    cont_10   : cs10 cs11 : cs10  cs11 :    26900000

junction capacitances nwell :
  # name     :  condition    : mask1 mask2 : capacitivity

  # bottom
    acap_cw  :  cwn      cs6 :  cs6  cwn   : 100

  # sidewalls
    ecap_cw1 : !cwn -cwn cs1 :  cs1 -cwn   :   5.0
    ecap_cw2 : !cwn -cwn cs2 :  cs2 -cwn   :  22.3
    ecap_cw3 : !cwn -cwn cs3 :  cs3 -cwn   : 180.2
    ecap_cw4 : !cwn -cwn cs4 :  cs4 -cwn   : 247.3
    ecap_cw5 : !cwn -cwn cs5 :  cs5 -cwn   :  84.3
```

A.2.2 BEM/FEM

The BEM/FEM method substitutes layers cs6 – cs11 from the layered FEM with the BEM. This is implemented by the connection of layer cs5 to the BEM interface (@sub) in the contacts list, and the definition of the deep substrate through the 'sublayers' segment at the end. Notice that the distributed approach to the sidewall capacitances of the n-well is also applied.

```
# layered FEM sheet resistances
conductors :
  # name     : condition : mask : resistivity
    cond_cs1 :    cs1    :  cs1 :       25316
    cond_cs2 :    cs2    :  cs2 :        4804
    cond_cs3 :    cs3    :  cs3 :        3997
    cond_cs4 :    cs4    :  cs4 :       11117
    cond_cs5 :    cs5    :  cs5 :      118343

# layered FEM vertical resistances
contacts :
  # name     : condition : lay1 lay2 : resistivity
    cont_1   :  cs1 cs2  :  cs1  cs2 :         126.6
    cont_2   :  cs2 cs3  :  cs2  cs3 :         354.9
    cont_3   :  cs3 cs4  :  cs3  cs4 :       65200
    cont_4   :  cs4 cs5  :  cs4  cs5 :      169000
    cont_5   :  cs5      :  cs5 @sub :           0

junction capacitances nwell :
  # name     :  condition    : mask1 mask2 : capacitivity

  # bottom
    acap_cw  :  cwn          :  @sub cwn   : 100

  # sidewalls
    ecap_cw1 : !cwn -cwn cs1 :  cs1 -cwn   :   5.0
    ecap_cw2 : !cwn -cwn cs2 :  cs2 -cwn   :  22.3
    ecap_cw3 : !cwn -cwn cs3 :  cs3 -cwn   : 180.2
    ecap_cw4 : !cwn -cwn cs4 :  cs4 -cwn   : 247.3
    ecap_cw5 : !cwn -cwn cs5 :  cs5 -cwn   :  84.3

sublayers :
  # name       conductivity  top
    substrate     10.0       0.0
```

A.2.3 Node contractions (BEM/ hybrid FEM)

The application of node contractions in the BEM/FEM method requires a halo within which the 3DFEM is applied. In this technology description the halo mask is arbitrarily called 'rbc', and may be derived from any mask that is physically present in the layout. Typically this may be a via-mask that defines a connection between a supply-line and the substrate, or this may be a mask that defines interconnect which is close to the substrate. The contraction is then guided by using the rbc-mask in the definitions for the vertical resistances of the FEM. Notice (again) that the distributed approach to the sidewall capacitances of the n-well is also applied.

```
# grow FEM halo 'rbc' over 0.25 microns
    resize : rbc : rbc : 0.25e-6

# layered FEM sheet resistances
conductors :
  # name     : condition : mask : resistivity
    cond_cs1 :    cs1    :  cs1 :       25316
    cond_cs2 :    cs2    :  cs2 :        4804
    cond_cs3 :    cs3    :  cs3 :        3997
    cond_cs4 :    cs4    :  cs4 :       11117
    cond_cs5 :    cs5    :  cs5 :      118343

# layered FEM vertical resistances
contacts :
  # name     :   condition   : lay1 lay2 : resistivity
    cont_1   :  cs1 cs2  rbc :  cs1  cs2 :         126.6
    cont_c1  :  cs1 cs2 !rbc :  cs1  cs2 :           0
    cont_2   :  cs2 cs3  rbc :  cs2  cs3 :         354.9
    cont_c2  :  cs2 cs3 !rbc :  cs2  cs3 :           0
    cont_3   :  cs3 cs4      :  cs3  cs4 :       65200
    cont_4   :  cs4 cs5      :  cs4  cs5 :      169000
    cont_5   :  cs5          :  cs5 @sub :           0

junction capacitances nwell :
  # name     :  condition    : mask1 mask2 : capacitivity

  # bottom
    acap_cw  :  cwn          :  @sub cwn   : 100

  # sidewalls
    ecap_cw1 : !cwn -cwn cs1 :  cs1 -cwn   :   5.0
    ecap_cw2 : !cwn -cwn cs2 :  cs2 -cwn   :  22.3
    ecap_cw3 : !cwn -cwn cs3 :  cs3 -cwn   : 180.2
    ecap_cw4 : !cwn -cwn cs4 :  cs4 -cwn   : 247.3
    ecap_cw5 : !cwn -cwn cs5 :  cs5 -cwn   :  84.3

sublayers :
```

```
  # name       conductivity  top
    substrate     10.0       0.0
```

A.3 Traditional Technology Description

Using the traditional, plain BEM modeling approach, the technology data would have been passed to SPACE as follows.

A.3.1 Single-layer BEM

The technology description for the plain BEM is very straightforward, but notice that the distributed approach to the sidewall capacitances of the n-well can not be applied in the plain BEM approach. Furthermore, notice that the channel-stop layer is not taken into account in any way.

```
junction capacitances nwell :
  # name    :  condition    : mask1 mask2 : capacitivity

  # bottom
    acap_cw :  cwn          : @sub cwn    :          100

  # sidewall
    ecap_cw : !cwn -cwn     : @sub -cwn   :          539.1

sublayers :
  # name       conductivity  top
    substrate     10.0       0.0
```

A.3.2 Double-layer BEM

This technology description is equivalent to the previous one, but now introduces the channel-stop layer as a BEM layer.

```
junction capacitances nwell :
  # name    :  condition    : mask1 mask2 : capacitivity

  # bottom
    acap_cw :  cwn          : @sub cwn    :          100

  # sidewall
    ecap_cw : !cwn -cwn     : @sub -cwn   :          539.1
```

```
sublayers :
  # name          conductivity  top
    channel-stop     1000       0.0
    substrate          10.0    -0.5
```

B

FEM Layer Divisions

This appendix presents the conductive layer divisions for the 9-layer, 11-layer, 16-layer and 20-layer FEM used in some of the simulations of Chapters 5 and 6. The 11-layer FEM is explained with more detail in Chapter 5 itself. With respect to the layer divisions presented in this appendix, the layer division in Table B.2 for the 11-layer FEM (data copied from Table 5.1 for convenient comparison here) is assumed the basic division, which, by further division of the given conductive layers, is refined to the layer divisions in Tables B.3 and B.4 for the 16-layer FEM and 20-layer FEM, respectively. The layer division in Table B.1 for the 9-layer FEM is a heuristically chosen simplification of the layer division in Table B.2. Based on a conductive layer division, the corresponding layered FEM may be constructed according to the principle from Figure 4.16.

Table B.1: Conductive layer division for 9-layer FEM

layer	thickness (μm)	approximate conductivity (S/m)
1	0.5	1000
2	5.0	10
3	0.404	10
4	0.293	10
5	5.54	10
6	20.2	10
7	73.7	10
8	269	10

Table B.2: Conductive layer division for 11-layer FEM.

layer	thickness (μm)	approximate conductivity (S/m)
1	0.100	790
2	0.346	975
3	3.26	50
4	1.69	10
5	0.404	10
6	0.293	10
7	5.54	10
8	20.2	10
9	73.7	10
10	269	10

Table B.3: Conductive layer division for 16-layer FEM

layer	thickness (μm)	approximate conductivity (S/m)
1	0.1	790
2	0.346	975
3	1.09	50
4	1.09	50
5	1.09	50
6	1.0	10
7	0.69	10
8	0.404	10
9	0.293	10
10	3.08	10
11	12.2	10
12	29.1	10
13	58.1	10
14	89.7	10
15	179.3	10

Table B.4: Conductive layer division for 20-layer FEM

layer	thickness (μm)	approximate conductivity (S/m)
1	0.1	790
2	0.146	975
3	0.2	975
4	0.5	50
5	0.59	50
6	1.09	50
7	1.09	50
8	1.0	10
9	0.69	10
10	0.404	10
11	0.293	10
12	1.0	10
13	2.08	10
14	6.0	10
15	12.0	10
16	24.0	10
17	48.0	10
18	96.0	10
19	182.37	10

C

Closer Approximation of the Doping Profile using Layered FEM

Section 6.2.2 showed convergence of the layered FEM method for the fixed (coarse) piecewise constant approximation of the doping profile shown in Figure 5.7. It was shown that convergence was reached using a 20-layer FEM. However, even though more FEM layers were introduced (see Appendix B), and the convergence of the method itself was shown, the additional FEM layers were introduced in the context of the same coarse approximation of the doping profile. In other words, the convergence experiments from Chapter 5 did not use the layered FEM to improve the approximation of the doping profile.

Therefore, in this appendix, we will use the 20-layer FEM for a better approximation of the doping profile as shown in Figure C.1. The accompanying conductive layer division is shown in Table C.2.

Applying the default 20-layer FEM and the improved 20-layer FEM to the ring-oscillator with guarded sensor node from Figure 6.1, error margins of some 30 % become apparent. A selection of resulting resistances is shown in Table C.1. We observe significant differences, but Figures C.2 and C.3 show that the differences stay within approximately 30 %, and may therefore be considered only marginal, which is also confirmed by the simulations in Figures C.4 and C.5.

From these results we may conclude that improved approximation of the doping profile has only marginal influence on the resulting resistances and relatively small

Table C.1: Selected resistance values for the ring oscillator with guarded sensor node using 20-layer FEM with the default (Table B.4, Figure 5.7) and improved (Table C.2, Figure C.1) approximation of the doping profile. Both extractions used maximum tile dimensions of $20\lambda \times 20\lambda$ in the horizontal discretization and took about 18.5 hours, using 108 Mb of memory. The maximum error in the resistance values is approximately 30 %.

doping profile approximation	default	improved
R(vss, sens) ($k\Omega$)	112.6	160.9
R(vss, t1) ($k\Omega$)	20.58	18.80
R(sens, t1) ($k\Omega$)	322.7	475.4

influence on the simulation results. This confirms a previous conclusion from this thesis: the overall consistency between the model and the physics is more relevant to the behaviour of the simulations than the highest accuracy in the model.

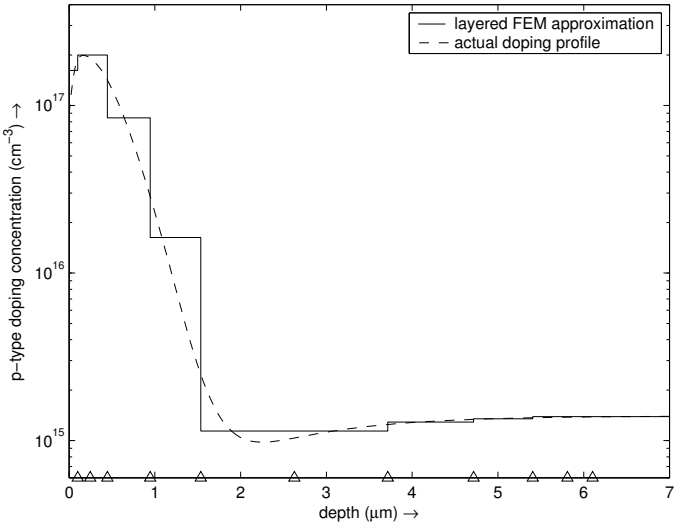

Figure C.1: A 20-layer FEM approximation of the channel-stop doping profile with improved accuracy over the approximation from Figure 5.7. The triangular tick marks along the horizontal axis are transition points between the layers defined in Table C.2.

Table C.2: Conductive layer division for Figure C.1

layer	thickness (μm)	approximate conductivity (S/m)
1	0.1	790
2	0.146	975
3	0.2	975
4	0.5	400
5	0.59	100
6	1.09	8.2
7	1.09	8.2
8	1.0	9.3
9	0.69	9.7
10	0.404	10
11	0.293	10
12	1.0	10
13	2.08	10
14	6.0	10
15	12.0	10
16	24.0	10
17	48.0	10
18	96.0	10
19	182.37	10

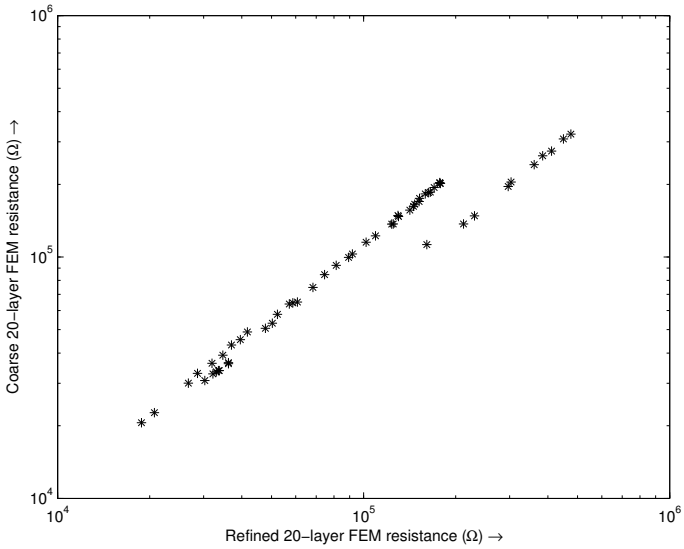

Figure C.2: Default 20-layer FEM resistances versus improved 20-layer FEM resistances. Despite some differences, correlation between the networks is clear.

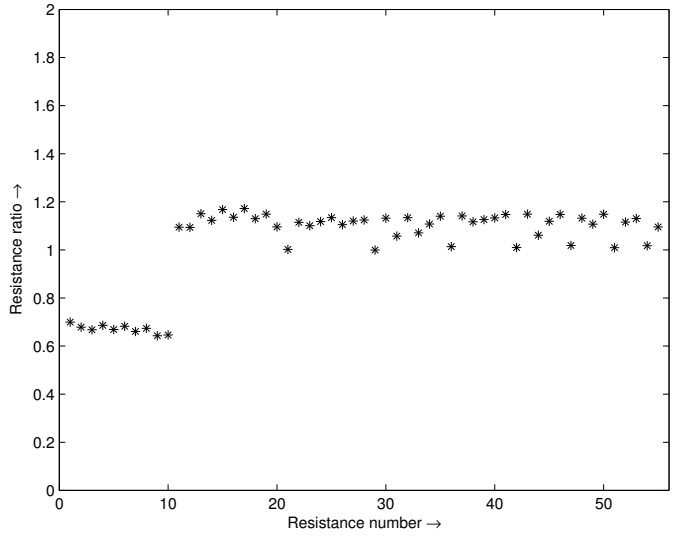

Figure C.3: Resistance ratios of the default 20-layer FEM versus the improved 20-layer FEM. All resistance ratios are near 1, except those related to the sensor node (see page 111 for details about the numbering scheme).

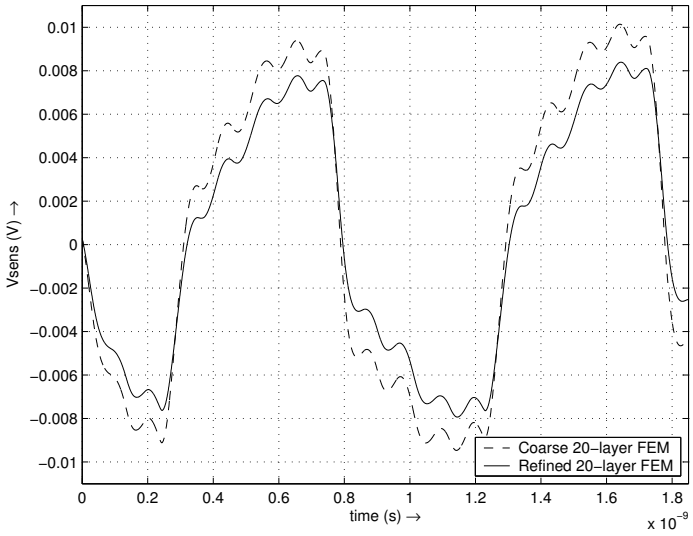

Figure C.4: Simulations on the sensor node using networks generated with the default 20-layer FEM and the improved 20-layer FEM. The waveforms have the same shape, but different amplitude.

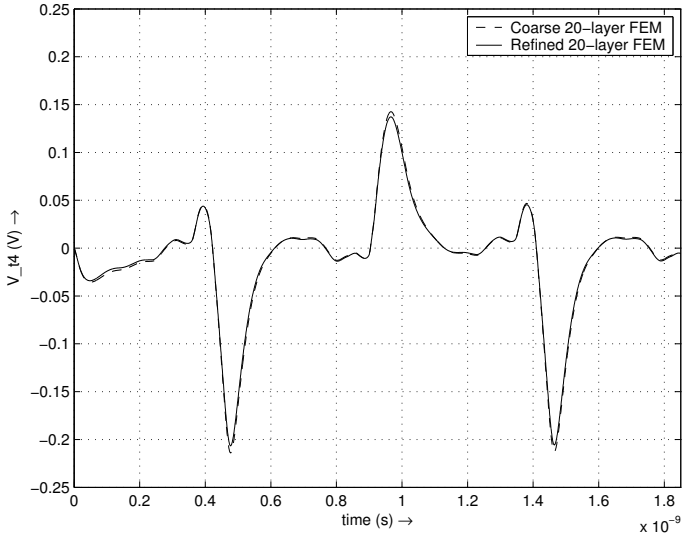

Figure C.5: Simulations on the backgate of transistor 't4' using networks generated with the default 20-layer FEM and the improved 20-layer FEM. The waveforms are nearly equal.

References

[Ara99] X. Aragonès, J.L. González and A. Rubio. *Analysis and Solutions for Switching Noise Coupling in Mixed-Signal ICs*. Kluwer, 1999.

[Bee98] F. Beeftink, A.J. van Genderen, N.P. van der Meijs and J. Poltz. "Deep-Submicron ULSI Parasitics Extraction Using SPACE". In *Proc. DATE, Designer Track*, pages 81–86, 1998.

[Bra04] T. Brandtner and R. Weigel. "Hierarchical Simulation of Substrate Coupling in BiCMOS Structures Using the Boundary Element Method". In *W.H.A. Schilders, E.J.W. ter Maten and S.H.M.J Houben (Eds.), Scientific Computing in Electrical Engineering, Proceedings of SCEE 2002*, pages 137–145. Springer, 2004.

[Bre78] C.A. Brebbia. *The Boundary Element Method for Engineers*. Pentech Press, 1978.

[Chi94] E. Chiprout and M.S. Nakhla. *Asymptotic Waveform Evaluation and Moment Matching for Interconnect Analysis*. Kluwer, 1994.

[Cle94] F.J.R. Clement, E. Zysman, M. Kayal and M. Declercq. "LAYIN: Toward a Global Solution for Parasitic Coupling Modeling and Visualization". In *Proc. CICC '94*, pages 537–540, 1994.

[Dew87] P.M. Dewilde and E.F. Deprettere. "Approximate Inversion of Positive Matrices with Applications to Modelling". *Modeling, Robustness and Sensitivity Reduction in Control Systems, Springer, NATO ASI Series*, pages 212 – 238, 1987.

[Don03] S. Donnay and G. Gielen (Eds.). *Substrate Noise Coupling in Mixed-Signal ASICs*. Kluwer Academic Publishers, 2003.

[Gen88] A.J. van Genderen and N.P. van der Meijs. "Extracting Simple but Accurate RC Models for VLSI Interconnect". In *Proc. ISCAS*, pages 2351 – 2354, 1988.

[Gen91] A.J. van Genderen. *Reduced Models for the Behavior of VLSI Circuits*. PhD thesis, Delft University of Technology, 1991.

[Gen96a] A.J. van Genderen and N.P. van der Meijs. "Using Articulation Nodes to Improve the Efficiency of Finite-Element based Resistance Extraction". In *Proc. 33rd DAC*, pages 758–763, 1996.

[Gen96b] A.J. van Genderen, N.P. van der Meijs and T. Smedes. "Fast Computation of Substrate Resistances in Large Circuits". In *Proc. European Design and Test Conference*, pages 560–565, Paris, France, March 1996.

[Gen01] A.J. van Genderen, N.P. van der Meijs and E. Schrik. "Modeling Capacitive Effects via the Substrate". In *ProRISC IEEE 12th Annual Workshop on Circuits, Systems and Signal Processing*, pages 366–370, Veldhoven, The Netherlands, November 2001.

[Gra90] H.C. de Graaf and F.M. Klaassen. *Compact Modeling for Circuit Design*. Springer Verlag, 1990.

[Hal87] J.E. Hall, D.E. Hocevar, P. Yang and M.M. McGraw. "SPIDER - A CAD System for Modeling VLSI Metallization Patterns". *IEEE Transactions on CAD*, 6(6):1023 – 1031, November 1987.

[Har68] R.F. Harrington. *Field Computation by Moment Methods*. Krieger, reprint edition, 1982 (1968).

[Hoo75] A.T. de Hoop. *Theorie van het Elektromagnetische Veld*. Delft University Press, Netherlands, 1975.

[Joa94] K. Joardar. "A Simple Approach to Modeling Cross-Talk in Integrated Circuits". *IEEE Journal of Solid-State Circuits*, 29(10):1212 – 1219, October 1994.

[Joh84] T.A. Johnson, R.W. Knepper, V. Marcello and W. Wang. "Chip Substrate Resistance Modeling Technique for Integrated Circuit Design". *IEEE Transactions on CAD*, 3(2):126 – 134, April 1984.

[Kol04] K.-J. van der Kolk. "The Unigreen Module". Internal Report, 2004.

[Kre93] E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, seventh edition, 1993.

[Li 02] H. Li , J. Carballido, H.H. Yu, V.I. Okhmatovski, E. Rosenbaum and A.C. Cangellaris. "Comprehensive Frequency-Dependent Substrate Noise Analysis Using Boundary Element Methods". In *Proc. ICCAD*, pages 2 – 9, 2002.

[MB3] *W. Liu et al., Manual: BSIM3v3.3 MOSFET Model – User's Manual, UC Berkeley, Dept. EECS, www-device.eecs.berkeley.edu/ bsim3*.

[MDa] *Manual: Davinci, 3DFEM-based device simulator, part of the Taurus environment by Synopsys, www.synopsys.com.*

[Mei92] N.P. van der Meijs. *Accurate and Efficient Layout Extraction*. PhD thesis, Delft University of Technology, 1992.

[Mei95] N.P. van der Meijs and A.J. van Genderen. "Delayed Frontal Solution for Finite-Element based Resistance Extraction". In *Proc. 32nd DAC*, pages 273 – 278, 1995.

[Mei01] N.P. van der Meijs. "Model Reduction for VLSI Physical Verification", Presentation at MACSI-net workshop on model reduction, www.macsinet.org. October 2001.

[MFe] *Manual: FEMLAB, Multiphysics Modeling Based on Finite Element Technology, www.comsol.com*.

[Mil99] J.J.H. Miller, W.H.A. Schilders and S. Wang. "Application of Finite Element Methods to the Simulation of Semiconductor Devices". *Reports on Progress in Physics*, 62(3):277 – 353, 1999.

[Mit94] A.R. Mitchell and D.F. Griffiths. *The Finite Difference Method in Partial Differential Equations*. Wiley, 1994.

[MRa] *Manual: Raphael, Field Solver with FDM and BEM support, part of the Taurus environment by Synopsys, www.synopsys.com.*

[MSp] *Manual: SPACE Layout-to-Circuit Extractor, Delft University of Technology, www.space.tudelft.nl*.

[MSu] *Manual: SubstrateStorm, dedicated substrate modeling tool, currently withdrawn from the market, but previously part of the Cadence design environment, www.cadence.com.*

[Nea92] D.A. Neamen. *Semiconductor Physics and Devices – Basic Principles*. Richard D. Irwin, Inc., 1992.

[Nel88] H. Nelis, E. Deprettere and P. Dewilde. "Approximate Inversion of Positive Definite Matrices, Specified on a Multiple Band". In *Proc. SPIE*, San Diego, 1988.

[Nel89a] H. Nelis. *Sparse Approximations of Inverse Matrices*. PhD thesis, Delft University of Technology, 1989.

[Nel89b] H. Nelis, P. Dewilde and E. Deprettere. "Efficient Modeling of Interconnections in a VLSI Circuit". In *Proc. ISCAS*, pages 872 – 876, Vol. 2, 1989.

[Nor73] D.H. Norrie and G. de Vries. *The Finite Element Method*. Academic Press, New York, 1973.

[Now96] E.B. Nowacka, N.P. van der Meijs and P.M. Dewilde. "The Hybrid Element Method for EMC Problems in VLSI Circuits". In *Third International Conference on Computation in Electromagnetics*, pages 72–77, Bath, United Kingdom, April 1996.

[Now97] E.B. Nowacka. *Hybrid Models for Parasitic Capacitances in Advanced VLSI Circuits*. PhD thesis, Delft University of Technology, 1997.

[Pfo96] M. Pfost, H.-M. Rein and T. Holzwarth. "Modeling Substrate Effects in the Design of High-Speed Si-Bipolar IC's". *IEEE Journal of Solid-State Circuits*, 31(10):1493 – 1501, October 1996.

[Pol87] S.J. Polak, C. den Heijer, W.H.A. Schilders and P. Markowich. "Semiconductor Device Modelling from the Numerical Point of View". *International Journal for Numerical Methods in Engineering*, 24:763 – 838, 1987.

[Pol05] D. Poljak and C.A. Brebbia (Eds.). *Boundary Element Methods for Electrical Engineers*. WIT Press, 2005.

[Pre85] F.P. Preparata and M.I. Shamos. *Computational Geometry: An Introduction*. Springer Verlag, 1985.

[Ram94] S. Ramo, J.R. Whinnery and T. Van Duzer. *Fields and Waves in Communication Electronics*. Wiley, 3rd edition, 1994.

[Saa03] Y. Saad. *Iterative Methods for Sparse Linear Systems*. 2nd ed., Philadelphia : Society for Industrial and Applied Mathematics, 2003.

[Sch02a] E. Schrik and N.P. van der Meijs. "Combined BEM/FEM Substrate Resistance Modeling". In *Proceedings of the 39th Design Automation Conference*, pages 771–776, New Orleans, LA, August 2002.

[Sch02b] E. Schrik, P.M. Dewilde and N.P. van der Meijs. "Theoretical and Practical Validation of Combined BEM/FEM Substrate Resistance Modeling". In *Proceedings of the International Conference on Computer-Aided Design*, pages 10–15, November 2002.

[Sch03] E. Schrik, A.J. van Genderen and N.P. van der Meijs. "Coherent Interconnect/Substrate Modeling Using SPACE - An Experimental Study". In *Proceedings of the 33rd European Solid-State Device Research Conference*, pages 585–588, Estoril, Portugal, September 2003.

[Sch04a] E. Schrik and N.P. van der Meijs. "Substrate Resistance Modeling by Combination of BEM and FEM Methodologies". In *W.H.A. Schilders, E.J.W. ter Maten and S.H.M.J Houben (Eds.), Scientific Computing in Electrical Engineering, Proceedings of SCEE 2002*, pages 364–372. Springer, 2004.

[Sch04b] E. Schrik and Q. Wang. "Apparent Accuracy Problems in SubstrateStorm". Internal Report, 2004.

[Sha03] Y. Shapira. *Matrix-Based Multigrid: Theory and Applications*. Kluwer, 2003.

[Sil96] P.P. Silvester and R.L. Ferrari. *Finite Elements for Electrical Engineers*. Cambridge University Press, 3rd ed., 1996.

[Sil99] L.M. Silveira and N. Vargas. "Multilevel Finite Difference Methods for the Characterization of Substrate Coupling in Deep Sub-Micron Designs". In *Proc. Symposium on Integrated Circuits and Systems Design '99*, pages 26 – 29, 1999.

[Sil04] J.M.S. Silva and L.M. Silveira. "Multigrid-Based Substrate Coupling Model Extraction". In *Proc. ISCAS '04, Volume 5*, pages 169 – 173, 2004.

[Sme95a] T. Smedes, N.P. van der Meijs, A.J. van Genderen, P.J.H. Elias and R.R.J. Vanoppen. "Layout Extraction of 3D Models for Interconnect and Substrate Parasitics". In *Proc. ESSDERC*, The Hague, The Netherlands, September 1995.

[Sme95b] T. Smedes, N.P. van der Meijs and A.J. van Genderen. "Extraction of Circuit Models for Substrate Cross-talk". In *Proc. ICCAD*, pages 199–206, 1995.

[Sta94] B.R. Stanisic, N.K. Verghese, R.A. Rutenbar, L.R. Carley and D.J. Allstot. "Addressing Substrate Coupling in Mixed-Mode IC's: Simulation and Power Distribution Synthesis". *IEEE Journal of Solid-State Circuits*, 29(3):226 – 238, March 1994.

[Str73] G. Strang and G.J. Fix. *An analysis of the Finite Element Method*. Prentice Hall, 1973.

[Ver93] N.K. Verghese, D.J. Allstot and S. Masui. "Rapid Simulation of Substrate Coupling Effects in Mixed-Mode IC's". In *Proc. CICC '93*, pages 18.3.1 – 18.3.4, 1993.

[Ver95a] N.K. Verghese, D.J. Allstot and M.A. Wolfe. "Fast parasitic extraction for substrate coupling in mixed-signal ICs". In *Proc. CICC '95*, pages 121 – 124, 1995.

[Ver95b] N.K. Verghese, T.J. Schmerbeck and D.J. Allstot. *Simulation Techniques and Solutions for Mixed-Signal Coupling in Integrated Circuits*. Kluwer Academic Publishers, 1995.

[Web50] E. Weber. *Electromagnetic Fields, Theory and Applications – Volume 1 Mapping of Fields*. John Wiley & Sons, 1950.

[Wem95] I.L. Wemple and A.T. Yang. "Integrated Circuit Substrate Coupling Models Based on Voronoi Tessellation". *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 14(12):1459 – 1469, December 1995.

[Wil84] D.R. Wilton, S.M. Rao, A.W. Glisson, D.H. Schaubert, O.M. Al-Bundak and C.M. Butler. "Potential Integrals for Uniform and Linear Source Distributions on Polygonal and Polyhedral Domains". *IEEE Transactions on Antennas and Propagation*, AP-32(3), March 1984.

[Yen64] W.C. Yengst. *Procedures of Modern Network Synthesis*. Macmillan, 1964.

[Zha98] J. Zhao, W.W.M. Dai, S. Kapur and D.E. Long. "Efficient Three-Dimensional Extraction Based on Static and Full-Wave Layered Green's Functions". In *Proc. 35th DAC*, pages 224 – 229, 1998.

[Zie83] O.C. Zienkiewicz and K. Morgan. *Finite Elements and Approximations*. John Wiley & Sons, 1983.

Summary

"A combined BEM/FEM Method for IC Substrate Modeling"

The research presented in this thesis was done in the context of physical verification of IC designs. In general, this encompasses taking an IC-design (i.e. a layout) as input and extracting a circuit model from it, including all relevant physical effects (e.g. field couplings and crosstalk) which would also be present on the actual chip. Simulations with the model then provide insight into the behaviour of the actual chip, before its fabrication. As such, physical verification is an integral part of the design flow and is used as a pre-fabrication step to do performance analysis or to locate design errors.

With the advances in IC-technology, the 'relevant physical effects' mentioned above become more complex. One of these effects is crosstalk through the substrate. The substrate is the actual semiconducting domain in which the transistors are embedded, and through which the transistors, or larger segments of the chip, may influence each other's behaviour. This last effect is known as substrate crosstalk. The research presented in this thesis has aimed at developing modeling techniques that consistently model the substrate as a crosstalk medium. Apart from accuracy of the modeling techniques, the focus is on computational efficiency and reduced-order modeling due to the ever-increasing complexity of chips.

The substrate modeling problem is identified as a potential problem in a conductive domain, which can mathematically be captured by the Laplace equation. The Laplace equation can conveniently be solved by modeling techniques like the Finite Element Method (FEM), or the Boundary Element Method (BEM). The FEM and the BEM both have their advantages and limitations. In short, the FEM is accurate and flexible, but typically slow, whereas the BEM is typically faster, but less flexible and only accurate in more restricted situations. When applied to modern substrate technologies, however, the accuracy limitations in the BEM and the speed limitations in the FEM are typically emphasized.

In this context, much research from literature has aimed at increasing the speed of the FEM through efficient meshing and solution techniques. Our approach, however, aims at making the BEM more flexible by combining it with a FEM. In particular, the FEM can be applied in specific subdomains that require improved flexibility or accuracy, while the BEM can be applied in the remaining subdomain that does not

have these requirements. The combination of the models then takes place on the boundary between the BEM and FEM subdomains.

Using these observations, we have developed a combined BEM/FEM method for the substrate modeling problem. The method consistently partitions the modeling problem, such that it exploits the structure of the substrate modeling problem and the properties of the BEM and FEM modeling techniques. In particular, the global couplings through the deep substrate are captured with a coarse and sparse BEM, while the local couplings are captured with a reduced-order FEM that contracts equipotential nodes. The resulting sparse, reduced-order BEM/FEM method uses inherently fewer variables than a full FEM, while retaining sparsity and accuracy. This claim may be illustrated in a straightforward way: where a full FEM would define nodes in the deep substrate, the BEM/FEM method handles the deep substrate with a sparse BEM, which requires no nodes there. The improved behaviour of the BEM/FEM method as opposed to the full FEM method is also reflected in the performance: using a standard solver based on Gaussian elimination ordered by minimum degree, BEM/FEM is significantly faster (about 4–8 times) than full FEM, while using an equivalent amount of memory.

The BEM/FEM method has shown to be a reliable method with proven convergence and stable behaviour. Furthermore, the accuracy of the method compares well to an independent 3DFEM modeling tool. However, during the course of this work, it was observed that the highest accuracy is not necessarily the most important characteristic of a good model. In particular, the studied network models sometimes showed poor accuracy in individual resistances, but nevertheless proved to be globally consistent with each other. Simulations with these networks then showed only marginal differences. The main common denominator between these models is the physical situation, for which they are all consistent approximations, but at different levels of accuracy. Thus, *consistency* between the model and the physical situation in the substrate proves to be more important than the highest accuracy in the model itself. This observation may be valuable in future optimizations to the efficiency of the BEM/FEM method, or in the development of new modeling methods.

Ultimately this thesis concludes that the combined BEM/FEM method inherits the good characteristics from its constituting parts. That is, the combined BEM/FEM method operates in a new trade-off between the speed of the BEM and the accuracy of the FEM. With this, the BEM/FEM method has become a new method in the roster of available substrate modeling methods.

Eelco Schrik
Delft, 2006

Samenvatting

”Een Gecombineerde BEM/FEM Methode voor IC Substraat Modellering”

Het onderzoek beschreven in dit proefschrift werd gedaan in de context van fysische verificatie van IC-ontwerpen. In het algemeen behelst dat het extraheren van een netwerkmodel uit een IC-ontwerp (een zogenaamde layout) waarbij het model alle relevante fysische effecten bevat (bijvoorbeeld veldkoppelingen en overspraak) die ook aanwezig zouden zijn op de daadwerkelijke chip. Simulaties met dit model geven dan inzicht in het gedrag van de daadwerkelijke chip, vòòr zijn fabricage. Als zodanig is fysische verificatie een integraal onderdeel van het ontwerpproces van IC's en wordt voorafgaand aan de fabricage toegepast voor prestatie-analyses van de chip of voor het localiseren van ontwerpfouten.

Met de vooruitgang in IC-technologie worden de bovengenoemde 'relevante fysische effecten' steeds complexer. Eén van deze effecten is overspraak via het substraat, waarbij individuele transistoren, of grotere segmenten van de chip, elkaars gedrag beïnvloeden via het halfgeleidende domein (het substraat) waarin ze liggen ingebed. Het onderzoek beschreven in dit proefschrift had tot doel het ontwikkelen van modelleringstechnieken die op consistente wijze het substraat als overspraak-medium kunnen modelleren. Naast de nauwkeurigheid van de modelleringstechnieken, lag de nadruk ook op de rekenkundige efficiëntie en gereduceerde-orde modellering wegens de immer toenemende complexiteit van chips.

Het substraatmodelleringsprobleem wordt geïdentificeerd als een potentiaalprobleem in een geleidend domein en kan wiskundig worden beschreven met een Laplace verglijking. De Laplace vergelijking kan worden opgelost door modelleringstechnieken zoals de Eindige Elementen Methode (Finite Element Method, FEM) en de Randelementen Methode (Boundary Element Method, BEM). De FEM en de BEM kennen beide hun voor- en nadelen. Samengevat is de FEM nauwkeurig en flexibel, maar typisch langzaam, terwijl de BEM typisch sneller is, maar minder flexibel en alleen nauwkeurig voor een meer beperkte set situaties. Wanneer deze methoden echter worden toegepast op moderne substraattechnologieën, worden de nauwkeurigheidsbeperkingen van de BEM en de snelheidsbeperkingen van de FEM typisch benadrukt.

In dit kader heeft veel onderzoek uit de literatuur zich gericht op het versnellen van de FEM door efficiënte vermazings- en rekentechnieken. Onze aanpak richt zich echter op het verbeteren van de flexibiliteit van de BEM door hem te combineren met een FEM. In het bijzonder kan de FEM worden toegepast in specifieke subdomeinen die meer flexibiliteit of hogere nauwkeurigheid vereisen, terwijl de BEM kan worden toegepast in het resterende subdomein waar deze eisen niet van toepassing zijn. De combinatie tussen de modellen vindt dan plaats op de grens tussen de BEM en FEM subdomeinen.

Gebruikmakend van deze observaties hebben wij een gecombineerde BEM/FEM methode ontwikkeld voor het substraatmodelleringsprobleem. De methode partitioneert het modelleringsprobleem op consistente wijze, zodanig dat de structuur van het substraatmodelleringsprobleem en de eigenschappen van de BEM en de FEM worden geëxploiteerd. In het bijzonder worden de globale koppelingen door het diepe substraat berekend met een grove, ijle BEM, terwijl de lokale koppelingen worden berekend met een gereduceerde-orde FEM die equipotentiale knooppunten contraheert. De resulterende, ijle, gereduceerde-orde BEM/FEM methode gebruikt inherent minder variabelen dan een volledige FEM, terwijl ijlheid en nauwkeurigheid behouden blijven. Deze bewering kan eenvoudig worden geïllustreerd: waar een volledige FEM knooppunten in het diepe substraat zou definiëren, behandelt de BEM/FEM methode het diepe substraat met een ijle BEM waarvoor er geen knooppunten in het diepe substraat zijn vereist. Het verbeterde gedrag van de BEM/FEM methode tegenover de volledige FEM methode wordt ook gereflecteerd in de prestaties: bij gebruikmaking van een standaard rekentechniek gebaseerd op Gaussische eliminatie geordend naar minimale graad, is de BEM/FEM methode significant sneller (ongeveer 4–8 keer) dan de volledige FEM, terwijl beide methoden ongeveer evenveel geheugen gebruiken.

De BEM/FEM methode heeft aangetoond een betrouwbare methode te zijn met bewezen convergentie en stabiel gedrag. Voorts is de nauwkeurigheid van de methode goed vergelijkbaar met die van een onafhankelijk 3DFEM softwarepakket. Met de voortgang van het onderzoek is echter geobserveerd dat de hoogste nauwkeurigheid niet noodzakelijkerwijs de meest belangrijke eigenschap is van een goed model. In het bijzonder toonden de bestudeerde netwerkmodellen soms een slechte nauwkeurigheid in de individuele weerstanden, maar bleken desalniettemin toch globaal consistent met elkaar te zijn. Simulaties met deze netwerken toonden dan slechts marginale verschillen. De gemene deler tussen deze modellen is de fysische situatie waarvoor al deze netwerken consistente benaderingen zijn, maar met verschillende niveaus van nauwkeurigheid. Aldus is de *consistentie* tussen het model en de fysische situatie in het substraat belangrijker dan de hoogste nauwkeurigheid in het model zelf. Deze observatie kan waardevol zijn bij toekomstige optimalisaties aan de efficiëntie van de BEM/FEM methode, of bij de ontwikkeling van nieuwe modelleringstechnieken.

Uiteindelijk concludeert dit proefschrift dat de gecombineerde BEM/FEM methode de goede eigenschappen van zijn twee bouwstenen erft. Dat wil zeggen, de gecombineerde BEM/FEM methode werkt in een nieuwe wisselwerking tussen de snelheid van de BEM en de nauwkeurigheid van de FEM. Hiermee is de BEM/FEM methode nieuw toegevoegd aan de lijst van beschikbare substraatmodelleringsmethoden.

Eelco Schrik
Delft, 2006