Temporal-Spatial Discretization and Fractional Latency Techniques for Wave Propagation in Heterogeneous Media

by

Tom De Rybel

Ind. Ing., Hogeschool Gent, 2002
M.A.Sc., The University of British Columbia, 2005

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate Studies

(Electrical and Computer Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

February 2010

Abstract

This thesis presents the development of a novel, transient wave propagation simulator using time-decoupled transmission line models. The models are based on the electro-magnetic transient program (EMTP) power system transient analysis tools, extended to two dimensions. The new tool is targeted at acoustic wave propagation phenomena. The method, called TINA for transient insular nodal analysis, uses temporal interpolation and fractional latency to maintain synchronicity in heterogeneous media. The fractional latency method allows the model cells to operate at a local simulation time step which can be a non-integer ratio of the global simulation time step. This simplifies synchronicity and saves computation time and memory. Thévenin equivalents are used to interface the mesh cells and provide an abstraction of the cell content. Numerically, the method is of the transmission-line matrix (TLM) family. In the thesis, loss-less and distortion-less models are considered. The loss-less transmission line models are studied for their stability and numerical error, for which analytical expressions are derived based on the simulation parameters. A number of new relations were discovered and discussed. The TINA method is evaluated in 2D using acoustic experiments, and also a new method is proposed for obtaining impulse responses in time-domain simulation, based on a periodic, band-limited impulse signals.

Table of Contents

Appendices

List of Tables

List of Figures

Acknowledgements

> If I have seen farther than others, it is because I was standing on the shoulders of giants.
>
> *Isaac Newton*

When Newton made the famous statement that inevitably pops up in thesis acknowledgement sections world-wide, I believe that he had it only right in part. One must indeed stand on the shoulders of giants to look further, but one must also be able climb up there, understand the giants strange language, and find food and shelter. There is an undeniable need to live and work in an environment that allows one to spend so much time at lofty heights while the world turns on below.

Such environments, where curiosity is encouraged and the inevitable failures along the way are accepted as the cost of business, where advise and support is readily available, be it of a technical or personal nature, are rare and ultimately the creation of the efforts of many.

Over the years, I have been supported in so many ways and by so many people that I cannot, by any stretch of imagination, call this thesis my work alone. Neither can I acknowledge everyone in these pages, for fear that this section would exceed the length of the actual text. Nor do I even know the names of all people involved.

Thus, to those I missed: my deepest gratitude for your help. You are all heroes that made life more livable and the world a nicer place to be in. And I sure as hell couldn't have done it without you.

When I came to Canada almost eight years ago, in the last week of August, 2002, I barely knew the existence of three people and my understanding of the place was limited to rumors that the queen of England was somehow involved. And she only ended up getting tangled into the whole mess through ways that needs the likes of Terry Pratchett to make it sound reasonable, if only for a non-dictionary definition of the word. Oh, and Monty Python's "Lumber Jack" song taught me that British Columbia had many trees in it, and that the local woodsmen entertained rather strange vestimentary choices, given their line of work. South Park, the motion picture, finally explained all I had to know about the political relations between Canada and the United States. I had also consulted an atlas once, and noticed that large parts of the country looked like a sponge. Yes, one might say I definitely came prepared for the adventure.

Of those three people who helped me settle-in, using actual facts, one is no longer with

us. Eric Timothy "Tim" Baars, a friend who made me aware of the existence of the country in the first place, and taught me much about the English language, passed away on October 23rd, 2002, after a long fight with cancer. I still have his bicycle and helmet, a gift I use to explore the beautiful Pacific Spirit park around the university.

My first land lord, Andrew "Andy" Lai, a Malaysian ex-hippie who once toured South America on a motor cycle, and a descendant of a failed Chinese revolutionary (his great grand-father, I think), taught me how to cook, live on less money than most spend on a single meal per week, and otherwise kept the bunch of international rabble in his house in line. It was there that I also learned that Canada, and definitely Vancouver, tends to the multicultural side, with people from New Zealand, Poland, Germany, Belgium, and Malaysia under the same roof.

Getting accepted in the university itself, for what was then a masters degree (little did I know where that was going to go, although I had plenty of misguided hopes) was eased by the patience of Doris Metcalf, the department's graduate secretary, who managed to hold things off here, push things there, until I finally got all the paperwork in order. Also, when I discovered that the communications group I was accepted in wasn't my thing (they do theory there!), she also sent my file around again, I imagine with a fluorescent sticky that said: "Warning: severe allergy to theory".

And this is where things get going. My file was picked-up by Dr. José Martí who, at the time, was looking for someone who could hold a screwdriver and know which end goes forward (the pointy bit). In due course, the masters degree happened and, near the end of it, he offered me the opportunity to do a PhD with him. I, of course, said.... Let me think about it. After some consideration, I accepted, as this was by far the best offer I got, and stated, in capitals, that I Would Not Do Theory on this degree. It was to become a practical doctoral thesis. Upon hearing this, Dr. Marti gave me one of his patented little smiles. To put that smile in perspective, it is that smile that a fox makes when he knows he's got the chicken down and fried-up with a nice sauce, far before the chicken even realizes that now would be a good time to start getting worried and book a ticket to the Bahamas. As to who won this particular Epic Struggle Against Theory (also capitalized), well, see the remainder of this text. Suffice to say that my high school math teacher, who told my parents in no uncertain terms that I would, maybe, make a not altogether incompetent plumber, would have a fit. One could attribute this rather monumental change in perspective to the wisdom and guidance of an experienced and internationally acclaimed supervisor. Myself, I'd attribute it to being outfoxed in ways I'm only now starting to get a glimpse of. There are, indeed, reasons why Dr. Martí is the professor and I'm the student.

When it comes to the topic of this thesis, an acoustic transient simulator, it may seem incongruous with energy systems, the field I'm supposedly in. Like many long stories, it

started out with a question: how does the human ear work? To answer this ambitious question, a simulator was needed. This was to be based on power system analysis techniques, which use some interesting tricks that are unique to the field. As, later on, I found it nearly impossible to obtain good information on the human cochlea, and thus had no reliable way to actually verify any results, the topic shifted to the simulator itself. Of course, this acoustic simulator would also have to be verified somehow. Thus, I took a course in acoustics with Dr. Murray Hodgson. As seems to be a trend in my life, also this simple situation went completely out of hand. Allow me to explain.

When I started the whole "let's use acoustics to validate the simulator, as it'll be simpler to work with, with big wave lengths, inexpensive equipment, and easy to make a set-up because it can be big", I proved to be a total, complete, and utter fool. After completing the first couple experiments using equipment that amounted to the value of a nice Porsche, I still couldn't get any decent, let alone repeatable, results. And this is where Dr. Hodgson agreed to become my co-supervisor and teach me one or two thousand things about the realities of the field and how to get the universe to behave. This resulted in much gnashing of teeth, which in itself proved interesting. I was unaware that the universe had them.

Over the years, two degrees, and a countless number of "good ideas" aka "side projects" aka "scaring the bejezus out of the ethics committee, the health and safety board, as well as any bystander innocent or otherwise", quite a few crazy contraptions resulted. I will save myself the embarrassment of listing them here, but nearly all of them involved building something out of parts that were never meant to be used in quite that way.

In order to pull off such stunts, I frequently needed various mechanical parts made. Luckily, our department has a good machine shop staffed with the kind of people who could literally build the space shuttle from a lipstick drawing on the back of a napkin that was subsequently buried in the bottom cruft of a handbag for fifty years. And it would work. On the first try. In order of harassed most to slightly less, I owe a lot to Donald "Don" Dawson, Neil Jackson, and David "Fletch" Fletcher. Although, one could argue that I managed to harass the latter though various indirect ways and this order is skewed enough to be circular.

Then, there was the slight issue of getting parts in the first place. See, what one must know is that, with few glorious exceptions, suppliers are a lying, conniving, and utterly morally deprived bunch of puppy botherers. I write this in kindness, as the uncensored version would get me in trouble. Dealing with this lot of failed politicians requires a kind of calm perseverance, a particular ability to not reach down the phone line and strangle the troglodyte on the other end when said waste of skin uses its sincere voice to explain that the part was shipped last week and should arrive any day now. And of course they have it in stock. The person who handles these issues in our department is David "Chuch" Chu Chong, assisted by the now retired Lance

Jones and his current incarnation, Kristie Henrikson. David has perfected the art of sounding absolutely menacing and as unavoidable as a hovering mountain that is about to drop while being absolutely pleasant and relaxed at the same time. This is why having friendly and helpful assistants is necessary. Few cross the Purchaser and tell the tale. On the other hand, I'm happy to report that the lost art of getting parts yesterday is quite alive and well.

On more ethereal planes of existence, those of software and networks, expensive coasters, paper weights and door stops, there is the extensive IT support that our department houses. Out of my own experience, our system works much better than that in the computer science department, and people are far more flexible when it comes to strange, but not idiotic, requests. By the unwritten lore of IT, idiotic requests are, of course, met in much the same way starved saber tooth tigers meet bunnies. Anything else happens with an efficiency and effectiveness so unlike the general expectations from the industry that it becomes scary. My digital life has been made a lot easier by, in alphabetic order, since I don't know their MAC addresses: Andrey Vlassov, Chris Dumont, Ken Madore, Luca Filipozzi, and Rob Ross.

It might come as a surprise, but UBC has seen fit to allow me to mold the minds of the next generation. Instrumental in allowing me to do so, and getting paid well while at it, was Darla La Pierre, and her predecessor, Daria. I am very grateful for the continued teaching opportunities, as the extra income really helped keeping life fun. Being ever the scientist, I thus unleashed numerous didactic and otherwise diabolical experiments on what, by now, amounts to roughly one thousand two hundred students. During these activities, mostly on EECE375/474, I joined ranks with numerous professors as a teaching assistant and had a generally good time doing so. Of special note is Dr. Paul Davies, who guided me through the terms I was hired to be the instructor for the course, and prevented me from making many of the errors that result in paperwork avalanches and other occupational hazards. Over the terms, I've had the pleasure of working with Peter Vautour, the assigned technician for the course, who has a similar outlook on students as I do. It boils down to the likeness between the majority of them and the Egyptian mythological critter of choice that made the sun move through the heavens. I've been told the feeling is mutual. One thing I'm especially proud of is my earned nick name: Scary Tom.

Actually leaving the institute of higher learning that is UBC proved about as convoluted as leaving those other esteemed places where those deemed too unsettling to let run free in society are kept. The graduation process, which takes about nine months strangely enough, was made less painful by Cathleen Holtvogt. Using a good sense of humor, she managed to coral the committee into attending, battled the Powers that Require Much Paper, and finally got my files in order. It was all rather epic, if you ask me.

Of course, all work and no play will land you in the actual loony bin before you can say cuckoo. It is, in fact, of note how close the UBC hospital psychiatric ward is to the department

of electrical and computer engineering. Using the income from all the teaching assistant and instructor work, I managed to fulfill my life-long desire of making the country unsafe on land, on water, and in the air. It started by the discovery that I couldn't make it to the top floor of the building without being completely out of breath. This is when I picked-up this strange Vancouver fascination with hiking and joined the SFU hiking club, led by Andrew Haskell. Yes, that's the competing university. The UBC version of it was just far too serious, filled with the types who think of climbing Mount Everest as a day trip. Now, the fun in hiking is, I believe, based on oxygen deprivation in the brain. By the time you make it to the view point, anything and everything looks spectacular though the hallucinatory daze. But it is fun. With this group, I also hiked the wold famous West Coast Trail, as well as its little brother, the Juan de Fuca trail, both on Vancouver Island. The fame is well deserved. I also joined the UBC sailing club, got the White Sail I, II, and III certificates on dingies, and got roped into fixing the boats for a year. Fiber glass and gel coat were no longer strangers after that. People had a habit of crashing boats and sailing a submarine. Yes, this means exactly what you think it does: sailing the boat back in with the deck below the water line.

That leaves air. It was a long-standing dream of both my dad and me (and probably the majority of the human race) to learn how to fly. Unfortunately, when I was finally of age, the airport in my hometown was converted into a golf course. Three years ago, however, I decided to give it a go in Canada, at Pitt Meadows Airport (CYPK) on a Cessna 172. The grin went ear to ear and stayed there for at least two hours after that first flight. The people at Pitt told me about Boundary Bay (CZBB), since the distances were too much to make it reasonable to do my training there. Then, after discovering the Vancouver Cooperative Auto Network and becoming a member (I am now the proud owner of about half a door handle), and some shopping around the schools, I settled on Canadian Flight Center, and started off with Jill Smith. I tended to be a bit low on the natural skills (choice sentence after pulling the mixture to "off" for the second time in a row, delivered in a dead-pan, calm voice: "We need that to live. Carry on."). Eventually, when Jill got a real pilots job, I switched to Mike Stann, who also got a real job, then to Kwan Lam, who found a better paying instructor job in Malaysia, then to Amelia Kerr, who is looking for a real job. At this point, almost three years had gone by, the economy had tanked, and Jill was back again. With Amelia, I actually managed to finish the license and am now ready to point that old Piper Cherokee (C-GOUP, the orange and white rust bucket of love) into the sunset, with the landing light switched to "pulse" to make it my lighthouse in the sky. Kudos to all the good folks there who helped to make it happen!

The search for, and maintenance of, sanity is also continued during the day job. First, I'd like to tip my hat at the Lunch Team. This highly competitive and athletic endeavor at the Lunching Hour has been upheld over the years. Members of this team were: Khosro Kabiri

Chapter 1

Introduction

> There are more things in heaven and Earth, Horatio, than are dreamt of in your philosophy.
>
> *William Shakespeare, Hamlet*

In this chapter, we will discuss the philosophical background behind the idea of simulating a system, as well as the fundamental equivalences that exist between different fields of study. It serves to introduce the concepts in the thesis, as well as explain why power systems simulation techniques may be used to approach acoustical problems.

1.1   Scope of the Work

The objective of this thesis is the development of a new, one and two-dimensional, discrete-space, time-marching, transient wave propagation simulator based on time-decoupled transmission line models. These models are based on those used in transient power system analysis tools, such as the EMTP. The new tool is targeted at acoustic and electro-magnetic applications, although the approach is universal to most wave propagation phenomena.

The method, called TINA for transient insular nodal analysis, uses temporal interpolation and the newly developed fractional ratio latency technique to achieve synchronicity in heterogeneous media. Thévenin equivalents are used to interface the mesh cells and provide an abstraction of the cell content.

Numerically, the method is of the transmission-line matrix family. For the purpose of the thesis, only loss-less models are considered, which means their poles are all on, or close to, the unit circle in the z-domain. The used transmission-line models are studied for their stability and numerical error, for which an expression is derived based on the simulation parameters. A number of new relations were discovered and discussed. The method is evaluated in 2D using acoustic experiments and a new method for obtaining impulse response in time-domain simulations is shown and used.

1.2   Why Was This Research Done?

How would a new musical instrument sound? What would the acoustics of a new concert hall be like? Will this expansion chamber filter work to block the machine noise in the ventilation duct? How effective is this new radio antenna design? Will this microwave waveguide design work?

Traditionally, one had little choice but to go out and build it, or at least a scale model, hoping that prior experience was enough to achieve a reasonable result. One could improve the chances of success by using approximate calculations, by using calculus on simplified shapes, if the problem was to be tractable by hand or early machine calculation. In fact, the ability to try something out on paper, descriptive geometry as it was then called, was invented in the 1760s by Gaspar Monge and kept a closely-guarded French military secret until 1799, when Monge was allowed to publish his treatise on the subject [57]. Today, this art is referred to as technical drawing; the ability of representing a three-dimensional object in two dimensions.

Frequently, however, physical models of the system were not accurate enough, or the understanding of the system insufficient, to obtain a satisfactory result. Scale models often behave quite differently from the full-scale version. The latter is often the case with aircraft [82], where a design that was proven stable as a model ended up being dangerously uncontrollable when built full-scale. Another example can be found in the TRIUMF particle accelerator. The magnetic field was studied by means of a scale model. In this model, half-inch thick steel plates were used to construct the magnet geometry. Although great care was exercised, the ten inch thick steel plates used in the actual device had somewhat different metallurgical properties. As a result, the accelerator would not work until invasive modifications were made, involving grinding-off chunks of metal and welding bits to other places, so the required field distribution could be obtained. Acoustically, many a concert hall sounded disastrous until major retrofitting of acoustic panels, Helmholtz resonators, and other treatments took place.

With the advent of sufficiently powerful electronic computers in the sixties, numerical discretized-space methods took flight, such as finite-differences and finite-element methods. In general, these methods discretized space and geometry in small, elementary cells from which the physical system could then be built in much the same way one may build something out of Lego bricks. This allowed many of the simplifications required for hand calculations to be dropped. Conceptually, these methods caused a complete change in the way complex systems could be solved. One now only had to understand the behavior of the materials and how these materials interact with one another. The complexities of the geometry, or the thing the materials form, are no longer important.

Using these methods, it became possible to hear an instrument not even built, or test a shock

absorber that exists only in the engineer's mind. And all one had to do was define the geometry and the materials with their interactions. Thus, after someone figured-out the details of these elemental cells, anything could now be tried in the computer, with reasonable accuracy.

This is, of course, too good to be true. It was soon found that the ten thousands to millions of building blocks required for an accurate representation of a system present an enormous computational burden and thus take a long time to solve. The limitations on computational speed and memory meant that only coarse, highly simplified geometries could be studied. This was still a leap beyond the approximations required for traditional hand-solutions, but in the end, the discretized-time methods have their weakness in the sheer number of little building blocks required to represent a system in detail.

These issues are fundamental, but the applicability of the method has been steadily pushed forward as more powerful computers were developed. However, there are two ways of pushing-forward the usability of the methods: one can try to improve the computer, or one can try to improve the algorithm itself.

This thesis attempts the latter and focuses on a family of algorithms known as transmission-line modeling and finite-difference time-domain methods. Improvements are made in the sense that the presented methods allow one to trade some precision for significant numerical speed increases. Using these methods, larger and more detailed systems may thus be studied in less time, or complex systems may be solved where, without these methods, the available computers do not have enough memory to hold the problem or the computational time becomes too unwieldy to be practical. Thus, the work in this thesis serves to improve the usability of a fundamental tool used in many disciplines to give insight to complex systems and even allows one to study problems that cannot – yet – be built, or are simply so expensive to realize that one needs to be quite sure of the expected outcome before attempting construction. In the past, advancements to these methods have invariably lead to advances in many fields, and we hope that the work presented here may once more allow the state of the art in many fields to be pushed a little further.

1.3 Predicting the Future

The ultimate purpose of scientific theory is nothing short of a flawless prediction of the future. For a theory to hold, it must be falsifiable through experiment. In other words, there must be a way, through an experiment, by which it may be shown false, and thus the limits of its applicability discovered. And, most importantly, no theory can be Truth. They are approximations, their predictions limited to conditions that may take centuries to discover. And, when such a limiting case is found, the theory is to be rejected in favor of one that explains all that was

explained before, including the new discovery.

The purpose of a scientific theory to the engineer is an understanding of nature detailed enough that the behavior of a complete system may be predicted to a sufficient degree of accuracy for the purpose at hand. For engineering, the requirements are usually not as stringent as those in science. Newtons laws of motions are still perfectly adequate for most applications. Automobiles do not operate at a significant fraction of the speed of light. Thus, significant approximations may be made and a good enough result obtained in far less time than a complete solution of the standard model of physics [61] would have required. In fact, determining the interactions between the billions of atoms involved in a bicycle pump is, at the time of writing, not only impractical, but also impossible to calculate on any currently existing computer.

In the end, no matter if the objective is science or engineering, the predictions of the model must be tested to reality, thus compared with measurements of a physical entity or effect. Maxwell, in his famous treatise on electricity and magnetism [54, 55], in fact starts the discussion with an exposition on what it means, to measure.

Let us look at a classic high school example of such an engineering prediction. Suppose we take a ball, in vacuum, and we throw it. We are able to obtain the velocity of the ball as it leaves the thrower's hand, as well as the angle at which it leaves. We may weigh the ball beforehand, to ascertain its mass. From this, Newton's laws of motion [86] will give us an accurate prediction of where the ball will be at any time in the future. The more exact we know the initial conditions, thus, what the ball was doing when it left the thrower's hand, the more precise the future location and trajectory, if it were to interact with structures, will be known.

Of course, according to special relativity, the whole concept of time assumed by Newton is changed [22] and, according to quantum mechanics, there are strict limits on the observable, though Heisenberg's uncertainty principle [23] that imposes a trade-off between the accuracy of knowing either where something is and how fast it is going. Thus, the ultimate accuracy we may obtain in the prediction of this system is limited, even if we apply the relativistic corrections. However, for this day-to-day engineering problem on a human scale, the throwing of a ball against a wall, these considerations have so little impact as to be imperceptible to even the most accurate of instruments and the fact remains that we may compute an accurate future outcome for the system of the ball being thrown and who it will hit in a classroom.

1.4 Modeling the Outcome or the System?

There are two fundamentally different ways in which a system may be modeled. We may either model the outcome of the system, or we may model the system itself and see its outcome. The difference may be explained with an example.

Consider the current through a capacitor, given by the equation $i(t) = C\frac{\partial v(t)}{\partial t}$. For an ideal capacitor, this equation is a model for the behavior of the device, relating the rate of change of the voltage over the terminals at a moment t to the current at a moment t. This equation thus models the capacitor, and is used in cases where the internal workings of the capacitor are not important, just its behavior. We must thus know the behavior of the system, the capacitor, in order to derive a model of this type. Thus, this is a case of a model where the outcome of the system is modeled, but no insight into the actual mechanics in the box is shown, nor needed.

A different approach is one where we obtain extensive knowledge about the behavior of, e.g.: aluminum and air. Through experiments, we may find a model of their behavior, much like in the case for the whole capacitor above. However, what is done now is, using these constituent materials, the geometric structure of the capacitor is represented in the simulator. From the known, constituent materials, we may now see how the geometry of the capacitor gives rise to its function, how the fields and charges interact with the materials to store energy, how the behavior of the whole device comes to be. Thus, in this case, we model the system itself, and see the outcome. One could, theoretically, take this procedure to its logical extreme and model the individual constituent parts of the sub-atomic particles and the fundamental forces between them. Although such an extreme case in not practical for macroscopic objects, these methods do allow one to find the behavior of complex systems without any prior knowledge but that of the constituent parts.

The power in this method is thus that, if we know the materials well enough, that any geometry, be it a capacitor, resistor, inductor, antenna, etc... may be modeled and the behavior of the structure found. Even more, we gain insight in how the fields interact with the structure to give rise to the function. Compared with the other approach, we would only know how it behaves, but not why.

Both approaches have their place. When doing a typical circuit simulation, computing the inner workings of each capacitor is not only unpractical, but also unnecessary. On the other hand, when we have a system and we do not know its behavior, this may be found through the second approach and captured in laws of the first kind for simplified computation later on. Both methods tend to be complementary. In this thesis, however, we will be focusing on methods that model the system to find its outcome.

1.5   Unification in Nature

In this thesis, methods originating from electro-magnetics are used to compute acoustic phenomena. Although one could, at this point, simply give the one-dimensional, loss-less, acoustic wave equation and show it to have exactly the same form as the one-dimensional, loss-less,

electromagnetic wave equation, this does not truly show the deep equivalences on a fundamental level.

The great triumphs of the last century in physics were those of the unification of widely different fields [61, 81] and, with that unification, the emergence of deep insights into the nature of the universe. In fact, the standard model [61] incorporates all that is currently known in one, quantum mechanical equation, save for gravity. The unification of quantum mechanics and gravity is still unresolved, at the time of writing, and the last grand question of last century's physics.

The situation, in the nineteenth century before the unifications took place, was that physics was split in six different fields, namely [61]:

Dynamics The laws of motion which, when combined with Newton's laws of universal gravitation, describe the motion of celestial bodies, as well as all day-to-day objects on a human scale.

Thermodynamics The laws of temperature and heat energy, as well as the behavior of solids, liquids, and gases in bulk. It explains expansion, contraction, freezing, melting, and boiling.

Waves The study of oscillations of continuous media.

Optics The study of light.

Electricity The study of static and dynamic charges.

Magnetism The study of magnetic interactions.

By the beginning of the twentieth century, these branches had been reduced to two. The atomic hypothesis, strongly supported though the advances in chemistry, had shown that thermodynamics, wave mechanics, and dynamics were the same. The theories of the electromagnetic field had shown that optics, electricity, and magnetism were the same. Thus, all in nature could be explained by either particles, through the atomic theory, or waves, through the field theories.

As an aside, it is interesting, but disappointing, to note that these old territories are still very much alive. One of the issues that plagues the different traditional fields of science and engineering is that each method had its origins in some field, and the terminology and assumptions of that field go with it. The unifications had little effect on this, and thus these obstructions to interdisciplinary cross-fertilization remain to this day.

The particle-wave duality remained a fundamental problem. How was one to reconcile the relativity of motion with Maxwell's laws of electricity and magnetism? Or: how can we

reconcile the theories of light with the existence of atoms? This issue, and with it the final unification of classical physics, was performed by Einstein in 1905 [22], resulting in the well-known particle-wave duality.

At the same time, the quantum revolution was also in full swing. With the work of Planck, who discovered the quantized nature of black-body radiation, Schrödinger, and Heisenberg, the old Newtonian dream of the predictable, clock-work universe was shattered. The universe was made of countable, discrete entities with properties that could only be observed to degrees of precision. Statistics took over from certainty.

Eventually, also the strange quantum theories were reconciled with special relativity by the generation of Freeman Dyson and Richard Feynman [23, 81]. By 1980, all then known long-lived particles and forces could be explained by the standard model of elementary particle physics. All, except gravity. The quest for this final unification of what is known is still on.

Although, in this thesis, electro-magnetic and acoustic phenomena are treated as identical phenomena with different names for the governing parameters, from the standard model we learn that two different force carrying particles are involved. Electro-magnetic forces are transferred through photons while the mechanical forces for acoustics in matter are transferred by phonons. Essential is that phonons can only work in matter, thus do not have the infinite reach through vacuum that photons have. Put differently, for photons, absolute vacuum may serve as a medium through which they propagate, for phonons vacuum, and thus the absence of matter, becomes a barrier to propagation.

1.6 Equivalence Between the Electro-Magnetic and Acoustic Wave Equation

Although the underlying phenomena are different, both the propagation of electro-magnetic and acoustic energy, within certain approximations, are similar processes, when observed macroscopically. For both phenomena, we may observe an energy wave traveling through a medium. For both, there is a propagation speed set by the medium, and a wave length related to that speed and the frequency of oscillation. For both, different media give rise to similar effects of scattering, transmission, and reflection at the boundary between the media.

To illustrate these striking similarities, the one-dimensional wave equations for acoustics [13] and electro-magnetics [46] are shown in (1.1) and Table 1.1, for the loss-less case. Since the governing equations have the same form, it follows that a solution method devised for one will also be applicable to the other. This is why a solution method for electro-magnetics may also be used to solve for acoustic phenomena, within certain limitations, e.g.: [67, 68].

It should be noted here that these equations appear in many different fields of study, as many natural phenomena are, in essence, wave propagation problems. For example, the transmission line model used in this thesis as the basis of the actual numerical implementation of the solution to these equations, the Bergeron line model [6], finds its origins in hydraulics. More specifically, Bergeron studied the effects of both water hammer in pipes and lightning strikes on transmission lines and found that both could, to a good degree of approximation, be represented with the same model.

$$\frac{\partial^2 p}{\partial z^2} = \rho k \frac{\partial^2 p}{\partial t^2} \qquad \frac{\partial^2 v}{\partial z^2} = \mu \epsilon \frac{\partial^2 v}{\partial t^2}$$
$$\frac{\partial^2 u}{\partial z^2} = \rho k \frac{\partial^2 u}{\partial t^2} \qquad \frac{\partial^2 i}{\partial z^2} = \mu \epsilon \frac{\partial^2 i}{\partial t^2}$$

$$(1.1)$$

Acoustics	Electro-Magnetics
p = pressure	v = voltage
u = particle velocity	i = current
ρ = density	μ = permeability
k = compressibility	ϵ = permittivity
t = time	t = time
z = distance	z = distance

Table 1.1: The Acoustical and Electro-Magnetic 1D Wave Equations

1.7 Relativity and Uncertainty in Simulation

It is interesting to note that two of the underlying main concepts of modern physics, relativity and quantum-mechanical uncertainty appear naturally in the discretized-space time-domain methods this thesis talks about. Although, relativity is partly violated in the typical use of these methods, as the simulation is usually observed at all spatial locations instantaneously.

Let us first consider the uncertainty principle. Heisenberg's relation is not unique to physics. It is a fundamental property of any measurement in a discrete, thus quantized, system. For example, in information theory, the concept is known as the Nyquist limit. In a discretized-space simulation, where each constituent element in the simulation has a finite spatial extent, also a spatial Nyquist limit appears in addition to the temporal one.

Indeed, if a computation happens at discrete times, and each value is represented by a binary number, thus has a finite precision, and the simulation space is divided into finite cells,

what that really says is that the spatial, temporal, and quantity indicators are all quantized. The difference with the universe is that it operates at the Planck units [65] and a typical computation occurs at much coarser devisions. In fact, from the Planck units, one may compute the Nyquist frequency of the universe. Or: one may determine fixed limits to what may be observed in our universe, how small something may be, how fast something may move, so that we may still observe it. If something were to exceed those limits, one would get aliasing effects, thus the faster thing would appear slower, the smaller thing would appear larger, and so on. For example, when viewing a movie of a stationary wheel that accelerates, one first sees it turn forward, then become stationary again, then turn backward, become stationary, and eventually turn forward again, with this cycle repeating as the wheel turns faster and faster. By making a movie of the wheel, it is observed at twenty five frames per second, thus discretized in time, and aliasing effects occur as the wheel's rotational speed exceeds the Nyquist frequency of the camera.

From the discretization of the simulation, we may now express uncertainty relations, as is done in quantum mechanics. For example, if the spatial grid is discretized to one centimeter cells, we cannot locate an event with better precision than that. We may say it is happening in a particular cell, but we cannot be more precise. Same with time. We may say that something happened between two discrete time events, but we cannot say when precisely, as no detailed information is available between those discrete time events, just as there are no finer spatial devisions between the cells.

Thus now the subtlety: to know more precisely when something happens, we may reduce the simulation time step, and get more temporal precision. The problem is that, for an equal spatial discretization size, the acoustic wave still travels at the same speed. Thus, if before with the coarser time step, it would travel further between time steps, and thus encounter more cells, with the finer time step, it cannot travel as far, as there was less time to travel in. Thus, we know better when something happened, but the wave, having traveled only a little, may just as well still be in the same cell, and thus we will get less knowledge about the location of the wave between time steps, as it has less chance to interact with different cells.

As for relativity, these discrete-space methods, when operated in the time domain, are time-marching solutions where all phenomena may only propagate at a certain rate, set by the respective medium. Due to the Nyquist limits, any phenomena that attempts to move faster ends up being aliased, and thus appears slower again. Thus, the temporal and spatial discretizations place an upper limit to the highest velocity, and frequency, possible in the computation. More importantly, each medium will be configured for a finite propagation speed. Thus, since all phenomena may only travel at the speed of the medium, there is causality and we may also observe relativistic effects, as phenomena are habitually traveling at the maximum speed for

the phenomena in the medium and, unless the simulation is multi-physics, there are no other, potentially faster, means of conveying information. E.g.: in an acoustics simulation, electro-magnetic effects are usually not present, thus, for that simulation, the acoustic wave can be considered to have similar properties as a light wave would in the real world, and its speed is the absolute limit to how fast information may propagate in that medium.

However, this relativistic notion, and its complications, are usually circumvented, as the whole system is typically observed as a whole, not taking into account the time it takes the phenomena in the system to reach an observer, as it would in the real world. Thus, although the various waves operate at the maximum speeds in the medium, and thus relativistic effects would be expected, by using an observation technique that allows everything that happens to be seen at the same time, relativistic effects are mostly removed. By using this technique, there is a preferred observer who watches at a preferred, absolute time, which is coupled to the simulation time step. This is, however, not a requirement, and a non-privileged observer using relative time is easily possible.

Chapter 2

The TINA Method

> We do not try to model the outcome of the system. We try to model the system and see its outcome.
>
> *André LaMothe*

In this chapter, we will present a different formulation of the transmission-line modeling (TLM) method [31], called transient insular nodal analysis (TINA). The numerical solution used becomes a hybrid between TLM and circuit analysis approaches used in Power Systems, such as the electro-magnetic transients program (EMTP) [90]. We will show how the TLM method can be modified to enable time-decoupled, EMTP-style solutions and a temporal interpolation-based approach to synchronicity. The advantage of the TINA formulation lies in the extensive use of Thévenin equivalents, which enable the transparent incorporation of the various transmission line models developed for EMTP simulations to be used in a TLM-style method, potentially allowing complex and non-linear media to be modeled in a straight-forward way. No matrix formulations are used for numerical efficiency on normal PCs, as the basic linear algebra sub-programs (BLAS) performance/loading and unloading of small matrices in memory results in significant and detrimental performance penalties. Although the TINA formulation, in practice, is more computationally, and significantly more memory, intensive than the basic TLM formulation, the simplicity and flexibility of direct Thévenin equivalents offers a distinct practical advantage when combining different types of models in a simulation. Also, trivial parallelization of the method is possible due to the full time-decoupled nature of each mesh cell and latency techniques can readily be used to increase the numerical performance of the method, as is shown in Chapter 4. The ability to easily integrate widely dissimilar materials, with different properties and underlying models, in the same simulation environment and maintain transient-free synchronicity is another advantage. A link-resistor approach with dissimilar cell impedances is used in the TINA formulation, as opposed to the traditional link-resistor or link-line approaches in TLM, where all impedances are kept constant and stub lines are used to modify the base material parameters and achieve synchronism. Also, the TINA medium uses absolute parameters for all materials, as opposed to the relative ones used in TLM, where typically the base material is locally modified using stubs.

The contributions presented in this chapter lie in the introduction of the TINA approach, which presents a new way of incorporating and maintaining synchronism between dissimilar

materials in a simulation without the transients due to impedance discontinuities (and thus small internal reflections) that come from stub-based approaches. Additionally, the TINA formalism offers an automatic irregular grid in time that is still sampled in a regular pattern, due to the use of variable-length transmission lines (which could be used with stubs as well), which result in a Minkowsky space-time, which is Cartesian in space, but non-regular in time. The use of interpolation to affect sub-time step synchronicity adjustments in combination with variable line length for adjustments that exceed one time step, is another novelty in TLM-style simulations. The effect of interpolation on the transmission-line models will be studied in a further chapter.

2.1 The Transmission Line Modeling Method

Introduced by Peter Johns in 1971, based on an idea by Raymond Beurle [31], the transmission-line matrix method, also known as the transmission-line method, has seen continued interest in many disciplines as diverse as diffusion processes in food and heat flow in jet turbines to electro-magnetic wave propagation in wave guides.

The time-domain method is conceptually different from finite-element and finite-difference methods in that it does not attempt to discretize and solve the electro-magnetic field equations, or other equations that govern the flow of energy in the system, directly. Rather, it uses a circuit equivalent approach where the physical phenomenon of energy propagation and reflection through a mesh is simulated though an electrical network.

This approach, using circuit equivalents to model field problems, has in itself a long history, dating back to the pioneering work by Kirchhoff and Helmholtz. A particularly striking example was shown by Gabriel Kron in his papers on network equivalents for Maxwell's field equations [36, 37]. In fact, when observing Kron's work carefully, we may note the constituent cells of his equivalent networks are lumped π and T transmission line segments, which are electrically equivalent and may be converted from one form to the other. Such equivalent networks, as shown in Figure 2.1, are capable of representing elliptic (only R and G, Poisson's equation), parabolic (either R and C or G and L, diffusion equation), and hyperbolic (only L and C, loss-less wave equation (Helmholtz equation)) partial-differential equations [73].

In the end, the TLM method performs a field to circuit mapping. It allows the solution of field problems through the simpler problem of solving a representative electrical network. Such an approach is possible when the circuit components each govern a spatial extent much smaller than the length of the physical wave to be modeled in the medium. In other words, the coarseness of the mesh components must be much smaller than the wave being modeled, similarly to the grain in a photograph, which must be much smaller than the object shown

in the image for the grain to be no longer noticeable. This is an expression of a spatial discretization criterion, like a temporal discretization criterion is given by the Nyquist frequency [63]. Similarly, spatial aliasing effects and numerical errors occur when the wave is under sampled in space, as the continuous fields are necessarily represented by an equivalent circuit mesh with finite and discrete spatial resolution. These considerations equally apply to finite-element, boundary-element, and discretized-space methods in general.

What is different in TLM, compared to, for example, Kron's work, is that it does not rely on lumped circuit equivalents. It uses a mesh of continuous two-wire transmission lines. This makes it a more general formulation, and allows it to perform better at higher frequencies, where the transmission and reflection properties can no longer be considered lumped [47]. Second, the use of transmission lines allows for an elegant implementation in a computer, as, in the loss-less case, the system can be abstracted though a scattering process at the cell boundaries and propagation in the cells. The actual details of the line need not to be taken into account explicitly [16]. However, one may still represent the system as an electrical network of discrete transmission lines and solve it directly, which is what the TINA method will be based on.

Although the following discussion is centered on the TLM method, many of the proposed concepts also apply to finite-difference time-domain methods (FDTD) [96] which, in certain cases, can be shown to be equivalent formulations [9, 11].

2.1.1 TLM Circuit Equivalents in 1D and 2D

The derivation of TLM theory is usually done from a lumped circuit perspective. Then, it is shown that the lumped circuit equations, when the lumped elements are made infinitely small, is equal to the wave equation in one or more dimensions. Here, we will show how the T-circuit yields the Telegrapher's equation and how this circuit may be extended to a two-dimensional system. In the next section, the field to circuit mapping will be shown. In both cases, the derivation is based on [73].

One-Dimensional Wave Equation

Starting from the model in Figure 2.1, we have an elemental portion of transmission-line of length Δl, with R, G, L, and C the resistivity, conductance, inductance, and capacitance per unit length, respectively. The wave is assumed, without loss of generality, to propagate in the $+z$ direction, from generator to load.
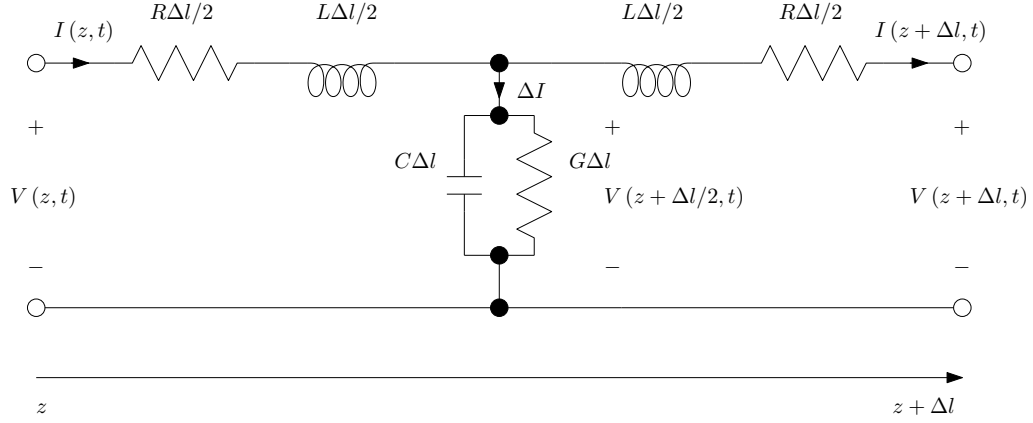
Figure 2.1: One-Dimensional T-Circuit of a Differential Length of Line

Applying Kirchoff's voltage law to the left loop, we obtain:

$$-\frac{V\left(z+\Delta l/2,t\right)-V\left(z,t\right)}{\Delta l/2} = RI\left(z,t\right)+L\frac{\partial I}{\partial t}\left(z,t\right) \tag{2.1}$$

Applying Kirchoff's current law to the main node, we obtain:

$$-\frac{I\left(z+\Delta l/2,t\right)-I\left(z,t\right)}{\Delta l} = GV\left(z+\Delta l/2,t\right)+C\frac{\partial V}{\partial t}\left(z+\Delta l/2,t\right) \tag{2.2}$$

Taking the limit of both above equations, as $\Delta l \to 0$, we can find continuous forms of the lumped circuit:

$$-\frac{\partial V}{\partial z} = RI + L\frac{\partial I}{\partial t} \tag{2.3a}$$

$$-\frac{\partial I}{\partial z} = GV + C\frac{\partial V}{\partial t} \tag{2.3b}$$

Differentiating the above equations with respect to z and t, respectively, we find:

$$-\frac{\partial^2 V}{\partial z^2} = R\frac{\partial I}{\partial z} + L\frac{\partial^2 I}{\partial z\partial t} \tag{2.4a}$$

$$-\frac{\partial^2 I}{\partial t\partial z} = G\frac{\partial V}{\partial t} + C\frac{\partial^2 V}{\partial t^2} \tag{2.4b}$$

Substituting the above equations into one another yields the wave equation, or more pre-

cisely, Telegrapher's equation in either voltage or current form [46]:

$$\frac{\partial^2 V}{\partial z^2} = LC\frac{\partial^2 V}{\partial t^2} + (RC + GL)\frac{\partial V}{\partial t} + RGV \tag{2.5a}$$

$$\frac{\partial^2 I}{\partial z^2} = LC\frac{\partial^2 I}{\partial t^2} + (RC + GL)\frac{\partial I}{\partial t} + RGI \tag{2.5b}$$

This equation is in fact very generic in form and appears in a wide variety of problems. In fact, by setting the appropriate parameters to zero, the equation reduces to the classic partial differential equations that govern Poisson, diffusion, and (loss-less) wave propagation problems.

Thus, from the above, in the limit where the spatial extent of the lumped components tends to zero, we find that the circuit equivalent of Figure 2.1 yields the lossy wave equation for one-dimensional waves. In TLM, if we construct a cascade of these circuit sections, we may model a one-dimensional problem where each section of line represents a small part of the physical system, e.g.: a small section of air in a narrow duct. If we take the spatial discretization Δl small enough with respect to the wave length and the geometry of the system, we find that we may model one-dimensional wave propagation using these electric circuits, as they implement the physical behavior of the electro-magnetic wave.

Later on in this chapter, the losses will be neglected in the further derivation of the TLM method, as well as the field to circuit comparison. This is done since in this thesis only loss-less wave propagation problems are studied and the inclusion of the losses significantly complicates the models and reduces clarity in the derivation of the methods.

In the TLM method, when applied to wave propagation problems, the losses are added back in later, as external to the wave propagation model itself. By doing so, the model remains simple and basic scattering techniques may be used to solve the propagation aspect, and modification to the cell central node are used to add non-ideal behavior [16, 18].

In the EMTP also, the losses are added to the ideal line model later on. Doing so keeps the core model, the Bergeron line model, computationally effective. The derivation of this model from the Telegrapher's equation is shown in Appendix A.

Keeping the losses in the model complicates the derivation. The line parameters of impedance and wave speed become frequency dependent [46], which precludes their direct solution in a time-domain computation, such as TLM or EMTP. One could solve them using a series expansion, convert to the frequency domain, or curve-fit the response using, for example, the frequency-dependent (FD) line model [51]. All of this adds computation. However, the EMTP line models, by keeping the Bergeron model at the core, are modular and allow different types of losses to be added around the basic line model. This approach gave rise to the

constant-parameter (CP, frequency-independent losses), distortion-less (Heaviside condition), and frequency-dependent line models (FD, based on curve-fitting the desired response to the model using LC sections). All of these are computationally effective and used extensively in EMTP power system simulations [20].

Thus, we can keep the derivations sufficiently general by assuming the loss-less case and treat the losses separate from the propagation model later on. In fact, the TINA method was especially designed to accommodate the EMTP line models in a TLM-style computation.

Two-Dimensional Wave Equation
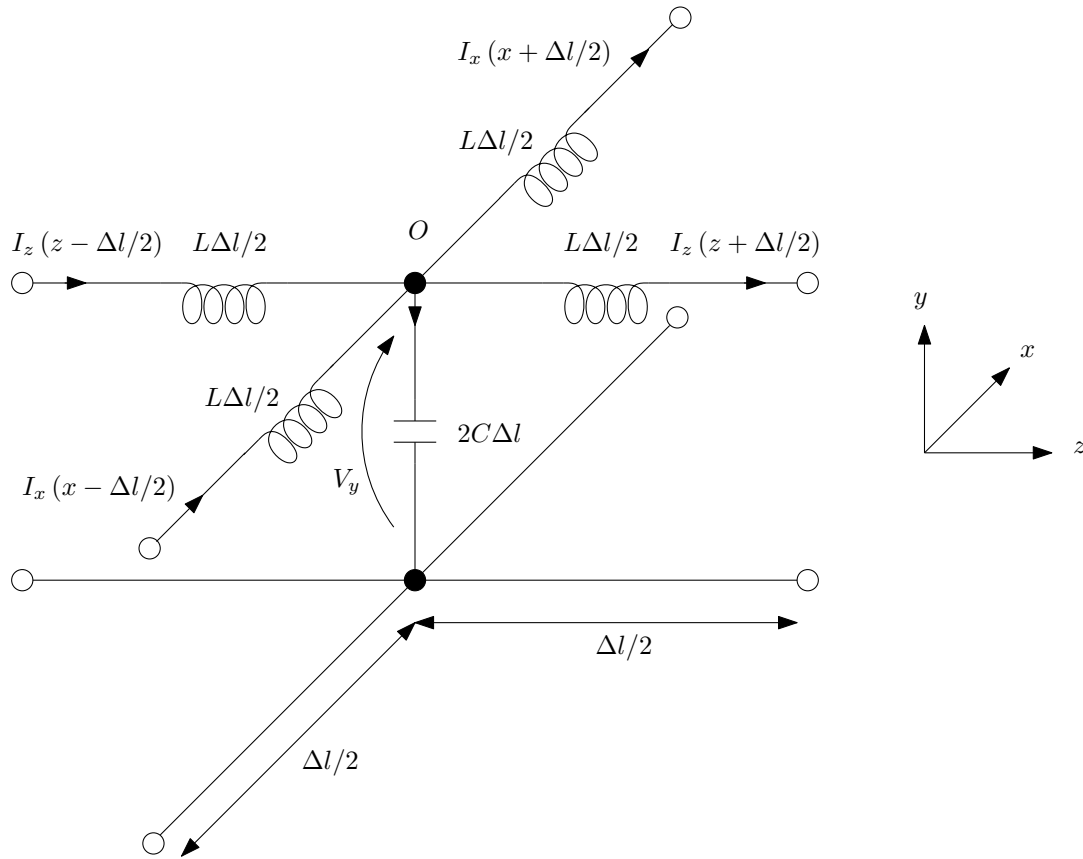
Figure 2.2: Two-Dimensional T-Circuit of a Differential, Loss-Less Area of Line

For the 2D case, the losses were neglected from the onset ($R = G = 0$). This was done to maintain clarity in the derivation, since only loss-less wave propagation is considered in this thesis. The result is the elemental area of 2D line, shown in Figure 2.2. The derivation follows along the same lines as for the 1D case above.

Note the presence of **double the capacitance** ($2C\Delta l$ **as opposed to** $C\Delta l$) in the 2D model.

This is due to an artifact of the 2D spatial discretization of space in Cartesian coordinates, where the diagonal propagation in a cell takes longer than the propagation along the main axis (diagonal of a square is longer than the sides of the square). The TLM method corrects for this in the model by including twice the permittivity ϵ in the medium, which increases the wave speed by $\frac{1}{\sqrt{2}}$ to make up for the apparent slow-down of the wave propagation in the model. This issue is discussed in detail in Section 2.1.3 of this chapter. The correction results in the difference between the circuit model and Maxwell's field equations, as is demonstrated later in Section 2.1.2. One could choose not to include this capacitance now and correct later, and have the circuit to field mapping have the same parameters. However, in TLM, it appears that the consensus [16, 18] is to modify the 2D cell model from the onset and build-in the correction. Thus, in this thesis, this convention was adopted.

Applying Kirchoff's voltage law around the loop in the x-y plane, we find:

$$\frac{V_y\left(x - \Delta l/2\right) - V_y\left(x + \Delta l/2\right)}{\Delta l} = \frac{L}{2}\frac{\partial I_x\left(x - \Delta l/2\right)}{\partial t} + \frac{L}{2}\frac{\partial I_x\left(x + \Delta l/2\right)}{\partial t} \tag{2.6}$$

Applying Kirchoff's voltage law around the loop in the y-z plane, we find:

$$\frac{V_y\left(x - \Delta l/2\right) - V_y\left(x + \Delta l/2\right)}{\Delta l} = \frac{L}{2}\frac{\partial I_z\left(z - \Delta l/2\right)}{\partial t} + \frac{L}{2}\frac{\partial I_z\left(z + \Delta l/2\right)}{\partial t} \tag{2.7}$$

Applying Kirchoff's current law at the central node, we find:

$$\frac{I_x\left(x - \Delta l/2\right) - I_x\left(x + \Delta l/2\right)}{\Delta l} + \frac{I_z\left(z - \Delta l/2\right) - I_z\left(z + \Delta l/2\right)}{\Delta l} = 2C\frac{\partial V_y}{\partial t} \tag{2.8}$$

Taking the limit $\Delta l \to 0$ of the above equations, we may find the continuous forms of the lumped circuit:

$$-\frac{\partial I_z}{\partial z} - \frac{\partial I_x}{\partial x} = 2C\frac{\partial V_y}{\partial t} \tag{2.9a}$$

$$\frac{\partial V_y}{\partial x} = -L\frac{\partial I_x}{\partial t} \tag{2.9b}$$

$$\frac{\partial V_y}{\partial z} = -L\frac{\partial I_z}{\partial t} \tag{2.9c}$$

Differentiating the above equations with respect to t, x, and z, respectively, we find:

$$-\frac{\partial^2 I_z}{\partial z \partial t} - \frac{\partial^2 I_x}{\partial x \partial t} = 2C\frac{\partial^2 V_y}{\partial t^2} \tag{2.10a}$$

$$\frac{\partial^2 V_y}{\partial x^2} = -L\frac{\partial^2 I_x}{\partial t \partial x} \tag{2.10b}$$

$$\frac{\partial^2 V_y}{\partial z^2} = -L\frac{\partial^2 I_z}{\partial t \partial z} \tag{2.10c}$$

Substituting the voltage equations in the first current equation above yields the 2D loss-less wave equation, or more precisely, the 2D Helmholtz equation:

$$\frac{\partial^2 V_y}{\partial x^2} + \frac{\partial^2 V_y}{\partial z^2} = 2LC\frac{\partial^2 V_y}{\partial t^2} \tag{2.11}$$

2.1.2 Equivalence Between Network and Field

A TLM simulation is made-up of a large number of building blocks shown in Figures 2.1 and 2.2, organized in a 1D or 2D regular, orthogonal mesh. Non-orthogonal meshes have, however, been studied. For example, cylindrical and spherical [1, 18], as well as un-structured meshes [16] have been successfully used with TLM.

In order to show the field to circuit mapping, we start from Maxwell's field equations [54, 55] for the electric and magnetic fields in differential form for uniform plane wave propagation in one direction for the time-varying case. We assume a simple medium with constant, scalar permittivity and permeability, and in which no free charges ($\rho = 0$) and currents ($\mathbf{J} = 0$) are present [71].

Maxwell's equations for such a case reduce to:

$$\bigtriangledown \cdot \mathbf{D} = 0 \tag{2.12a}$$

$$\bigtriangledown \cdot \mathbf{B} = 0 \tag{2.12b}$$

$$\bigtriangledown \times \mathbf{E} = -\mu\frac{\partial \mathbf{H}}{\partial t} \tag{2.12c}$$

$$\bigtriangledown \times \mathbf{H} = \epsilon\frac{\partial \mathbf{E}}{\partial t} \tag{2.12d}$$

Expansion of the curl equations for **E** and **H**, in a rectangular coordinate system, yields:

$$\frac{\partial E_z}{\partial y} - \frac{\partial E_y}{\partial z} = -\mu\frac{\partial H_x}{\partial t} \tag{2.13a}$$

$$\frac{\partial E_x}{\partial z} - \frac{\partial E_z}{\partial x} = -\mu\frac{\partial H_y}{\partial t} \tag{2.13b}$$

$$\frac{\partial E_y}{\partial x} - \frac{\partial E_x}{\partial y} = -\mu\frac{\partial H_z}{\partial t} \tag{2.13c}$$

$$\frac{\partial H_z}{\partial y} - \frac{\partial H_y}{\partial z} = \epsilon\frac{\partial E_x}{\partial t} \tag{2.13d}$$

$$\frac{\partial H_x}{\partial z} - \frac{\partial H_z}{\partial x} = \epsilon\frac{\partial E_y}{\partial t} \tag{2.13e}$$

$$\frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y} = \epsilon\frac{\partial E_z}{\partial t} \tag{2.13f}$$

These equations will now be adapted for 1D and 2D fields below.

Network and Field Equivalence in 1D

From [71], for uniform, one-dimensional plane waves, we assume variation in only one direction. We take this as the z-direction of a Cartesian coordinate system. Doing so sets the derivatives with respect to x and y to zero, thus $\frac{\partial}{\partial x} = 0$ and $\frac{\partial}{\partial x} = 0$. Applying this to (2.13) yields:

$$-\frac{\partial E_y}{\partial z} = -\mu\frac{\partial H_x}{\partial t} \tag{2.14a}$$

$$\frac{\partial E_x}{\partial z} = -\mu\frac{\partial H_y}{\partial t} \tag{2.14b}$$

$$0 = -\mu\frac{\partial H_z}{\partial t} \tag{2.14c}$$

$$-\frac{\partial H_y}{\partial z} = \epsilon\frac{\partial E_x}{\partial t} \tag{2.14d}$$

$$\frac{\partial H_x}{\partial z} = \epsilon\frac{\partial E_y}{\partial t} \tag{2.14e}$$

$$0 = \epsilon\frac{\partial E_z}{\partial t} \tag{2.14f}$$

Since the time-varying parts of H_z and E_z are zero, the fields of the wave are entirely transverse to the direction of propagation, as is expected of a plane wave in 1D. The remaining equations break into two independent sets, one relating E_y and H_x, and the second set relating E_x and H_y. These sets are merely an expression of the orthogonal electric and magnetic fields in the orthogonal coordinate system, perpendicular to the direction of propagation, z. Since both an electric or a magnetic field can exist in the x-y plane, e.g. in the y direction, its

corresponding magnetic or electric field, respectively, must necessarily exist in the x direction.

The wave propagation behavior may be illustrated with either set. We choose the pair E_y and H_x. Partially differentiating E_y with respect to z and H_x to t, we find:

$$\frac{\partial^2 E_y}{\partial z^2} = \mu\frac{\partial^2 H_x}{\partial t \partial z} \tag{2.15a}$$

$$\frac{\partial^2 H_x}{\partial t \partial z} = \epsilon\frac{\partial^2 E_y}{\partial t^2} \tag{2.15b}$$

Substituting the second equation into the first yields:

$$\frac{\partial^2 E_y}{\partial z^2} = \mu\epsilon\frac{\partial^2 E_y}{\partial t^2} \tag{2.16}$$

Thus, we find the one-dimensional loss-less wave equation, with the wave propagating in the z-direction. When we compare this result with the wave equation found from the circuit equivalent (2.5a), and set the losses to zero, thus $R = G = 0$, we find that Telegrapher's equation in voltage reduces to:

$$\frac{\partial^2 V}{\partial z^2} = LC\frac{\partial^2 V}{\partial t^2} \tag{2.17}$$

By comparing the equations (2.16) with (2.17), as well as (2.4) with (2.15), we find that the circuit and field representations of the wave propagation have the same form, in the loss-less case. We may now express the following equivalences:

$$E_y = V \tag{2.18a}$$

$$H_x = I \tag{2.18b}$$

$$\mu = L \tag{2.18c}$$

$$\epsilon = C \tag{2.18d}$$

Thus, in the equivalent circuit:

- The voltage at a node is E_y

- The current in the z-direction is H_x

- The inductance per unit length represents the permeability of the medium

- The capacitance per unit length represents the permittivity of the medium

For completeness, in Appendix A.1, a similar derivation is shown, including the losses.

Network and Field Equivalence in 2D

Returning to [73], the two-dimensional case is studied for transverse electric (TE) wave propagation. This implies that $E_x = E_z = H_y = 0$ and that all derivatives with respect to y are zero, thus $\frac{\partial}{\partial y} = 0$. As such, there is an electric field only in the y-direction, with its corresponding magnetic fields in the x and z directions. The magnetic field in the z-direction is now possible, as this is a 2D case, and thus both the x and z are perpendicular to the y direction, in which the electric field propagates. Thus, this is a plane wave in 2D, with the wave propagation possible in any direction in the x-z plane, where only those fields with an electric component in the y direction are considered for clarity.

Maxwell's equations of (2.13) thus reduce to:

$$\frac{\partial E_y}{\partial z} = \mu \frac{\partial H_x}{\partial t} \tag{2.19a}$$

$$\frac{\partial E_y}{\partial x} = -\mu \frac{\partial H_z}{\partial t} \tag{2.19b}$$

$$\frac{\partial H_x}{\partial z} - \frac{\partial H_z}{\partial x} = \epsilon \frac{\partial E_y}{\partial t} \tag{2.19c}$$

Partially differentiating E_y with respect to x and z, and H_x to t, we find:

$$\frac{\partial^2 E_y}{\partial z^2} = \mu \frac{\partial H_x}{\partial z \partial t} \tag{2.20a}$$

$$\frac{\partial^2 E_y}{\partial x^2} = -\mu \frac{\partial H_z}{\partial x \partial t} \tag{2.20b}$$

$$\frac{\partial H_x}{\partial z \partial t} - \frac{\partial H_z}{\partial x \partial t} = \epsilon \frac{\partial^2 E_y}{\partial t^2} \tag{2.20c}$$

Substituting the first and second equation in the third equation above yields:

$$\frac{\partial^2 E_y}{\partial x^2} + \frac{\partial^2 E_y}{\partial z^2} = \mu \epsilon \frac{\partial^2 E_y}{\partial t^2} \tag{2.21}$$

This is a Helmholtz, or loss-less wave equation, in two dimensions, for transverse-electric propagation in the x-z plane. We compare this result with the wave equation found from the 2D loss-less circuit equivalent (2.11), repeated below:

$$\frac{\partial^2 V_y}{\partial x^2} + \frac{\partial^2 V_y}{\partial z^2} = 2LC \frac{\partial^2 V_y}{\partial t^2} \tag{2.22}$$

By comparing the equations (2.21) with (2.22), as well as (2.10) with (2.20), we find that the circuit and field representations of the wave propagation have the same form, in the loss-less

case. We may now express the following equivalences:

$$E_y = V_y \tag{2.23a}$$

$$H_x = -I_z \tag{2.23b}$$

$$H_z = I_x \tag{2.23c}$$

$$\mu = L \tag{2.23d}$$

$$\epsilon = 2C \tag{2.23e}$$

Thus, in the equivalent circuit:

- The voltage at a node is E_y

- The current in the z-direction is $-H_x$

- The current in the x-direction is H_z

- The inductance per unit length represents the permeability of the medium

- Twice the capacitance per unit length represents the permittivity of the medium

It is important to note here that the circuit capacitance per unit length is **twice** the permittivity of the medium in the 2D case. This is critical when converting the input parameters for a medium in the simulation, typically given as the field parameters μ and ϵ, to the circuit parameters for the lines. Failing to adjust the capacitance will result in incorrect wave speed and line impedance in the computed results. The reason for this correction factor will now be discussed.

2.1.3 The Dispersion Relation of the Propagation Velocity in TLM

When propagating a wave on a TLM grid from a point source, it is noticed that the wave front does not travel in a perfect spherical front. In fact, as can be seen in Figure 2.3(a), in some directions it appears to go slower and there is significant fine detail in the wave where one would only expect perfect concentric symmetry. Also, the overall wave propagation seems slower than expected and for higher frequencies, the propagation does not seem to happen at all. These issues will now be discussed.

Wave Speed in the 2D TLM Mesh

To give an intuitive insight into this phenomenon, let us consider the grid in Figure 2.3(b). If a plane wave front were to propagate either perfectly left $\leftrightarrow$ right or up $\leftrightarrow$ down, the wave

(a) Point Source in a Grid   (b) Grid Distances
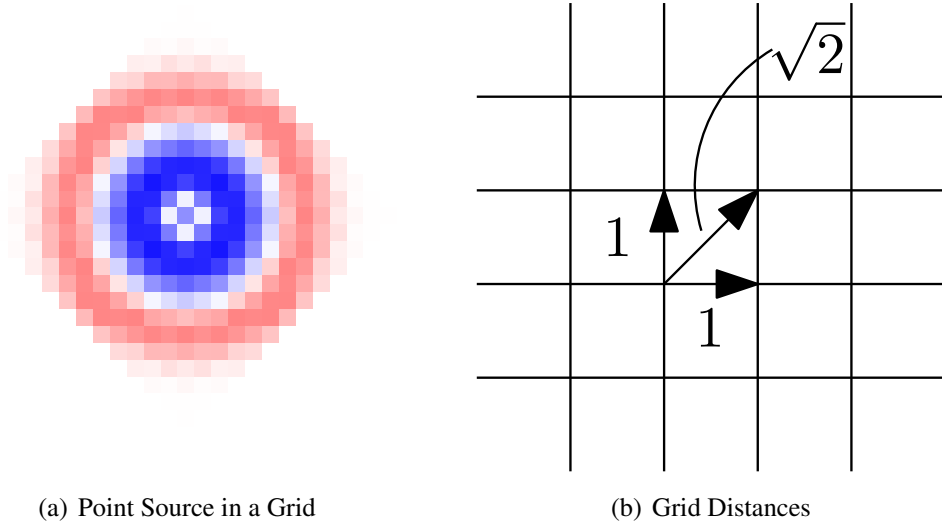
Figure 2.3: Velocity Dispersion in a 2D TLM Grid

will, as it propagates, always see a distance of one grid division between the cell center nodes, which are taken on the intersections of the grid lines. However, if a wave were to propagate at a 45 degree angle with the grid, the distance between grid nodes becomes $\sqrt{2}$. Thus, the perceived wave speed at these angles becomes less than that along the main grid lines. This is an issue common to most discretized-space methods

If the grid is made fine enough, this error eventually blends out and a clean wave front propagating at a coherent speed is observed some reasonable distance from the source. This effect can already be seen to happen in Figure 2.3(a). Near the source in the center, the wave front appears nearly diamond-shaped, and it quickly rounds-out the corners to appear more and more circular as the distance from the sources increases. It is generally taken that, if $\Delta l/\lambda \leq 1/10$, the variation in velocity is acceptably small [18]. This is demonstrated by the figure where, close to the sources, insufficient cells are available, compared to the size of the wave. As the wave propagates out, the wave becomes properly sampled.

It may be mentioned here that this $1/10$ rule for spatial discretization is a similar rule of thumb as used in interpolated EMTP line models [20]. This ratio is an expression of a spatial sampling criterion, much like the temporal sampling criterion given by the Nyquist frequency, which in practice must also be taken so that the sampling rate is five to ten times the highest frequency in the system.

The remaining problem, that the overall wave front travels too slow in 2D TLM, can be explained by looking at the equation for the wave speed on the transmission lines in TLM and the expression for the propagation constant in the 2D TLM mesh.

The wave speed in a transmission line is given by [16], where C_d and L_d are the capacitance and inductance per unit length, for a 1D transmission line, respectively:

$$a = \frac{1}{\sqrt{L_d C_d}} = \frac{1}{\sqrt{\mu_r \mu_0 \epsilon_r \epsilon_0}} \tag{2.24}$$

We may already observe the issue from the above equation, when comparing the field to circuit equivalences, given by (2.23): $C_d \neq C$. Thus, if we define the 2D medium by using the field parameters μ and ϵ, the resulting wave speed and line impedance will be wrong, as the circuit equivalent parameters do not have a 1:1 correspondence. In fact, the resulting mesh will have a propagation speed equal to $\frac{1}{\sqrt{2}}$ of the desired value. The solution is to change the ϵ parameter according to the field to circuit equivalences, thus divide the required ϵ by two internal to the model.

Dispersion in the 2D TLM Mesh

Dispersion means that the wave propagation is frequency dependent. For the loss-less 1D line, this is studied in detail in Chapter 3. Here, we will give the expression for 2D meshes, as shown in the literature. Intuitively, the issue may be thought of as follows: the TLM mesh consists of a large number of interconnected transmission lines. A transmission line functions as a resonator when the length equals specific fractions of the wave length, e.g.: $\lambda/2, \lambda/4$, and for these wave lengths functions as an open and short circuit, respectively. Thus, it is not unreasonable to expect that the TINA mesh will have similar cut-off effects at these wave lengths, due to the impedance discontinuities on the the mesh. These discontinuities arise from the nodes where four transmission lines connect. Thus, between two nodes, there is a transmission line with length equal to the spatial discretization Δl. The line is split in two half-lines, one half in each cell. When the impedances of these lines are equal, the result is a single line of length Δl, on which the resonator effects may happen.

From [73], the dispersion relation for a 2D TLM, uniform mesh is given by:

$$sin\left(\frac{\pi}{r_a}\frac{\Delta l}{\lambda}\right) = \sqrt{2}\, sin\left(\frac{\pi \Delta l}{\lambda}\right) \tag{2.25}$$

where $r_n = u_n/a$ is the normalized velocity, here the ratio of the velocity of the waves on the mesh relative to the desired wave speed. Solving this equation numerically for different values of $\Delta l/\lambda$ results in Figure 2.4. From this, we do indeed observe that the mesh blocks propagation at the quarter wavelength point. Thus, the TLM mesh can only represent Maxwell's equations over a range of frequencies from DC to the finest network cut-off frequency. This occurs at $\Delta l/\lambda = 1/4$. However, there is a significant error close to this frequency, and the

Dispersion of the Wave Velocity in the 2D TLM Mesh
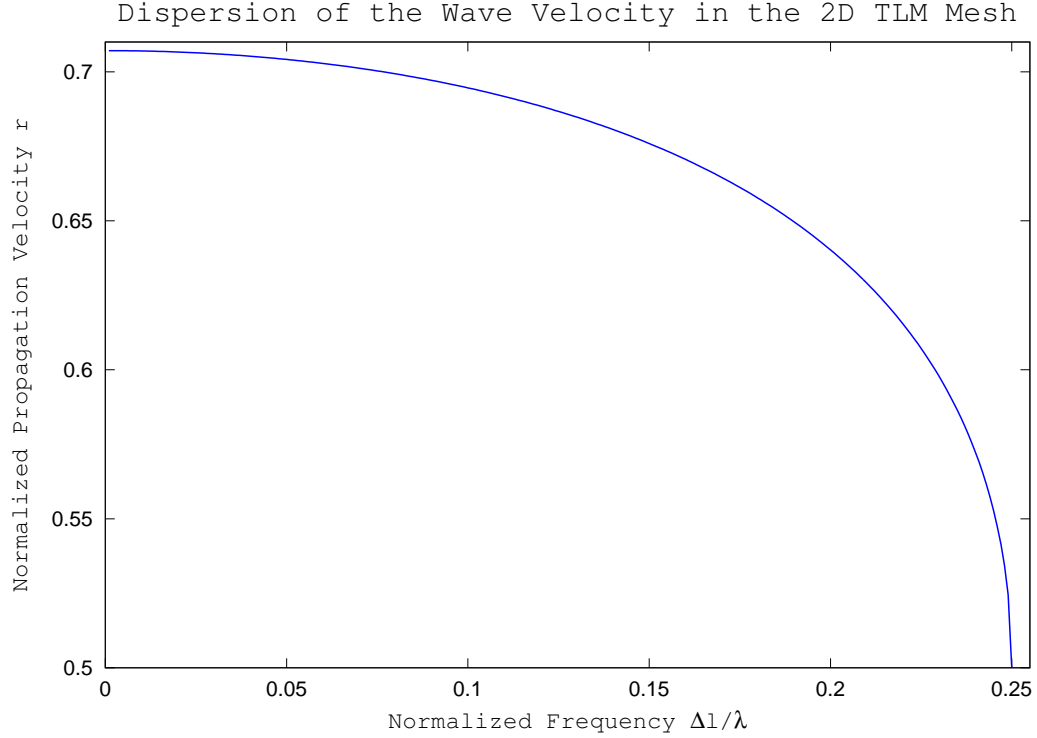
Figure 2.4: Dispersion of the Velocity of Waves in the 2D, Loss-Less TLM Mesh

usable range is thus lower. The rule of thumb, $\Delta l/\lambda \leq 1/10$ finds its origin here. Also note that the wave speed in the mesh indeed levels off at DC to $\sqrt{2}$ of the expected velocity.

2.1.4 The Use of Transmission Lines in TLM

The derivation of the loss-less TLM method has shown that the 1D TLM mesh is essentially a cascade of transmission lines, each obeying the loss-less wave equation. The loss-less 2D mesh consists of cells that are a combination of four loss-less 1D lines, joined at the central node [16] in either a series or shunt configuration. Thus, we may conclude that the 1D transmission line is the basic building block of 1D and 2D TLM (except for the correction factor applied to the capacitance per unit length or the permittivity ϵ, in the 2D case, depending on the particular formulation of the line model).

Now, in Appendix A.2, we showed that an exact solution to loss-less, one-dimensional wave equation is the Bergeron line model. In this model, no discretization is used, and the system is essentially two Thévenin equivalents, where the source is a time-delayed version of the currents and voltages on the other side of the model. This makes the model very easy and effective to implement in a time-domain solution, such as TLM.
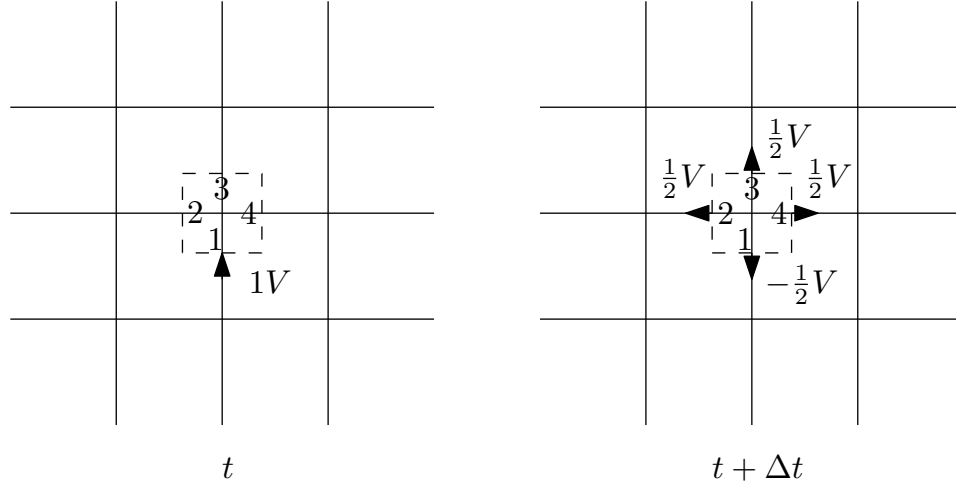
Figure 2.5: Dispersion of the Velocity of Waves in the 2D, Loss-Less TLM Mesh

However, the model is not directly implemented in TLM like this. Using the model directly would require a circuit solution at each cell's central node, as well as at each line connection between cells. Although these are small circuits which may be solved efficiently, in TLM the solution is based on scattering.

When looking at regular TLM, we may observe that all lines have the same propagation time, usually set to the base medium in use, and are loss-less. Losses and material variations are implemented on the cell's central node by means of stubs and resistive losses [16]. This means that a direct circuit solution is not required. Note that this will be a point where the new method presented later in this chapter diverges.

In TLM, the solution of the system is based on Huygens principle, thus scattering of the wave at impedance discontinuities. This solution is efficient when all lines have the same, typically one time step, time delay and materials are defined relative to a base material by means of matching stubs. In such a case, the system may be easily expressed in terms of a scattering matrix, which also includes any matching stubs. The scattering process describes the incident and reflected pulses, and forms the basis of the TLM algorithm, e.g.: [31, 73].

To illustrate this, let us evaluate the process when a wave enters a cell. The cell is connected to four identical neighbors, with normalized characteristic impedance $Z_c = 1$, though ports numbered 1-4, as shown in Figure 2.5. The traveling time of each line is normalized so that $\tau = \Delta t$. At time = t, a Dirac pulse of unit energy (S_i) enters the central cell from cell 1. One time step later, the pulse is scattered to the connecting cells.

Since line 1 has three other lines connected to it in the cell, its effective terminal impedance is 1/3. Since we know the relevant impedances, we also know the reflection coefficient for both

the reflected and transmitted voltage [71]. It is given by:

$$\Gamma = \frac{Z_L - Z_c}{Z_L + Z_c} = \frac{\frac{1}{3} - 1}{\frac{1}{3} + 1} = -\frac{1}{2} \tag{2.26}$$

Thus, the reflected and transmitted energy are, where the subscripts i, r, and t are incident, reflected, and transmitted, respectively:

$$S_r = S_i \Gamma^2 = \frac{1}{4} \tag{2.27a}$$

$$S_t = S_i(1 - \Gamma^2) = \frac{3}{4} \tag{2.27b}$$

Thus, a voltage impulse of $-\frac{1}{2}$ is reflected back in terminal 1, while a voltage pulse of $\frac{1}{2} = \sqrt{\frac{3/4}{3}}$ will be launched into each of the three other terminals.

Generalizing the above case for four impulses incident on a node, we may apply superposition by four single pulses, as the mesh is linear. Thus, if at a time t, voltage pulses $_tV_1^i$, $_tV_2^i$, $_tV_3^i$ and $_tV_4^i$, are incident on lines 1-4, at any junction node, the combined voltage reflected along line 1 at time $t + \Delta t$ will be given by:

$$_{t+\Delta t}V_1^r = \frac{1}{2}\langle {_tV_2^i} + {_tV_3^i} + {_tV_4^i} - {_tV_1^i}\rangle \tag{2.28}$$

In general, the total voltage impulse reflected along line n at time $t + \Delta t$ will be:

$$_{t+\Delta t}V_n^r = \frac{1}{2}\left[\sum_{m=1}^{4} {_tV_m^i}\right] - {_tV_n^i}, \quad n = 1, 2, 3, 4 \tag{2.29}$$

This may be described in matrix form, relating the reflected voltages at time $t + \Delta t$ to the incident voltages at the previous time step t:

$$\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \end{bmatrix}_{t+\Delta t}^{r} = \frac{1}{2}\begin{bmatrix} -1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & 1 \\ 1 & 1 & 1 & -1 \end{bmatrix}\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \end{bmatrix}_{t}^{i} \tag{2.30}$$

Also, an impulse leaving a node at position (z,x) in the mesh (thus, a reflected impulse),

becomes automatically an incident pulse at the neighboring node. Thus:

$$_{t+\Delta t}V_1^1\left(z, x + \Delta l\right) = {}_{t+\Delta t}V_3^r\left(z, x\right) \tag{2.31a}$$

$$_{t+\Delta t}V_2^1\left(z + \Delta l, x\right) = {}_{t+\Delta t}V_4^r\left(z, x\right) \tag{2.31b}$$

$$_{t+\Delta t}V_3^1\left(z, x - \Delta l\right) = {}_{t+\Delta t}V_1^r\left(z, x\right) \tag{2.31c}$$

$$_{t+\Delta t}V_4^1\left(z - \Delta l, x\right) = {}_{t+\Delta t}V_2^r\left(z, x\right) \tag{2.31d}$$

Thus, by applying the above equations (2.30) and (2.31), the magnitudes, positions, and directions of all impulses at time $t + \Delta t$ can be obtained at each node in the network, provided that the corresponding values at the previous time step were known. Thus, it is possible to solve the system through scattering without doing a formal network analysis at each node in the system. Also, note that the time steps are fully decoupled in time. Thus, it is possible to solve any part of the system independent from another, making parallelization trivial.

2.1.5   Synchronicity in TLM

Thus far, the discussion on TLM has focused on simulations with only one, homogeneous material in the system. To do any interesting work with the method, different material types must be integrated in the simulation. This may be easily done by locally changing the medium parameters μ and ϵ. However, this brings with it a new problem.

When changing the material parameters, it follows that both the impedance and wave speed of the material, and thus line segments, must change. The central node solution remains unchanged for the impedance changes; however, the connection solution now must take additional scattering at the impedance discontinuity between two neighboring areas into account.

There are two main methods used in TLM to address these issues. One is based on changing the impedance of the constituent transmission lines, and modifies the connection equations. This method, however, does not resolve the issues due to different wave speeds in the simulation. The second method uses loading stubs and addresses both. The latter method will now be discussed.

Connection Solution in the Presence of an Impedance Discontinuity

In general, a pulse moving from a source s to a neighboring node a experiences a reflection as it passes the impedance discontinuity between the area with impedance Z_s and the area with

impedance Z_a. These reflection coefficients are given by:

$$\rho_{s\to a} = \frac{Z_a - Z_s}{Z_a + Z_s} \tag{2.32a}$$

$$\rho_{a\to s} = \frac{Z_s - Z_a}{Z_s + Z_a} \tag{2.32b}$$

Now, the connection process, expressed in (2.31), must be changed accordingly. Each connection will now get a portion of the neighbor's contribution superimposed by a portion of its own signal reflected back from the boundary. The terms in the first equation, for example, must be substituted as follows, where ρ is the reflection and τ the transmission coefficient, respectively:

$$_{t+\Delta t}V_1^1\left(z, x + \Delta l\right) = \rho_{a\to s t+\Delta t}V_1^1\left(z, x + \Delta l\right) + \tau_{a\to s t+\Delta t}V_3^r\left(z, x\right) \tag{2.33a}$$

$$_{t+\Delta t}V_3^r\left(z, x\right) = \rho_{s\to a t+\Delta t}V_3^r\left(z, x\right) + \tau_{s\to a t+\Delta t}V_1^1\left(z, x + \Delta l\right) \tag{2.33b}$$

In essence, the connection process now completes a circuit solution based on incident and reflected waves. Also, it must be noted that this extended solution is only necessary at the boundary between two materials, thus in practice, the numerical cost of this is usually not very high, as the material boundaries tend to only make up a small percentage of the simulation.

What this method does not resolve is the issue caused by different traveling times in different media. When two, or more, media in a simulation have traveling times that are not an integer multiple of one another, only one of these can match the simulation time step. As a result, the other areas cannot operate normally as the discrete transmission line is a time delay and thus can only be a multiple of the simulation time step, as each past value corresponds to a discrete memory location. There is no simulation information for times in between these times that are multiples of the simulation time step.

This is the problem of synchronicity. One could change the size of the lines in the other media, to insure they match the simulation time step. However, that introduces the issue that the physical cells themselves now no longer match with one another and partial overlaps may happen. Thus, doing so moves the synchronicity issue from a temporal to a spatial expression, but does not solve the underlying issue. Later in this chapter, we will show an approach that overcomes these issues with this method.

The Use of Stubs

There is another way of handling the impedance differences, which also solves the issue of different traveling times in the computation. In this solution, all transmission lines in the mesh

(a) Scattering Connection   (b) Stub Matching
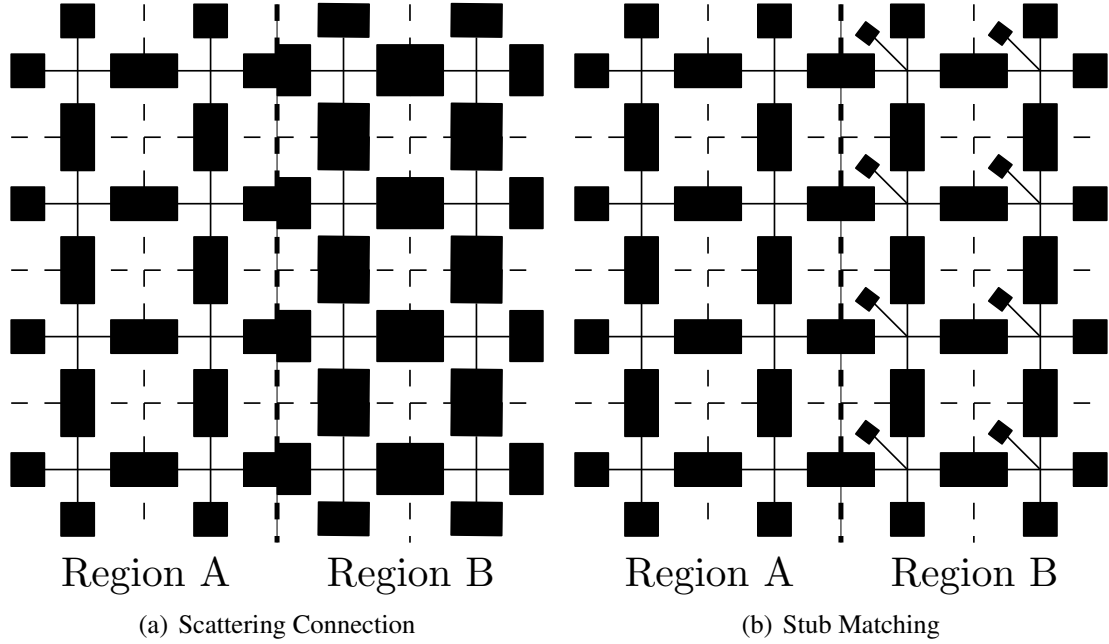
Figure 2.6: Methods to Incorporate Different Media in TLM

are kept at the same impedance and, at the cell nodes, a transmission line stub of length $\Delta l/2$ is connected to adjust the material parameters [16, 18]. The result of this approach is a modified scattering matrix, and the connection process can be left as it was, since the line impedance at the connections is now unchanged. The difference in concept between these methods is illustrated in Figure 2.6.

The concept of a stub comes from microwave engineering, where, due to the very short wave lengths involved, lumped components are unusable due to excessive parasitics. In this field, short lengths of transmission line are used to emulate capacitive and inductive behavior.

In TLM, we may use the concept to modify the behavior of a medium by adding such a stub to the central node in a cell. If we have two materials, A and B, we may write the following equivalence for the capacitive component of the line models in the mesh, where both materials are referred to vacuum, as is customary in TLM for electro-magnetics [18]:

$$C_A = \epsilon_A C_0 \quad C_B = \epsilon_B C_0 \tag{2.34}$$

Thus, we may say that:

$$\epsilon_B = \epsilon_A + \epsilon_S \tag{2.35}$$

There, S refers to the stub. Indeed, when looking at the circuit equivalent for a TLM node,

adding a line adds to the capacitance of the node. Thus, we may write:

$$C_B = C_A + C_S \tag{2.36}$$

The impedance within region B is then:

$$Z_B = \frac{\Delta t}{C_A + C_S} \tag{2.37}$$

Thus, the medium with characteristic impedance Z_B may be replaced by a parallel combination of the medium with impedance Z_A and a stub line. When $\lambda \gg \Delta l$, these stubs become part of the medium, in the same way the lumped line model components in the circuit equivalent did.

This stub should be a storage element, so that all energy that enters it is eventually reflected back into the mesh. Thus, open and closed stub lines must be used, depending on if we require to add capacitance or inductance to the system. We may, in fact describe the stub impedance as capacitive or inductive, depending upon the termination of the line. However, one can never have a line that has only capacitance without inductance, and vice versa. As a result, each of the above expressions is associated with a parasitic inductance and capacitance, respectively [16], for a standard stub with length $\Delta l/2$:

$$Z_{SC} = \frac{\Delta t}{2C_S} \qquad L_{error} = \frac{(\Delta t)^2}{4C_S} \tag{2.38a}$$

$$Z_{SL} = \frac{2L_S}{\Delta t} \qquad C_{error} = \frac{(\Delta t)^2}{4L_S} \tag{2.38b}$$

These errors may be reduced by making the simulation time step Δt smaller. This, however, will also make the simulation numerically much more expensive. Also, the error associated with the simulation depends on how much difference in material properties must be matched between different regions. As the required L and C goes up, so does the associated error.

In TLM, one will usually choose a base medium, and all other media will be defined relative to this by choosing appropriate stubs. Since the stubs allow us to influence both L and C, we may thus calculate the stub so the material matches the required μ and ϵ, and thus match both the desired impedance and wave speed, as both are related. When the errors associated with the matching stubs become too large, one may reduce the simulation time step to obtain sufficient accuracy. This change will increase the ratio of the wave length to the mesh size, and thus push the mesh closer to the limiting condition of $\Delta l \to 0$.

With the stub lines present in the system, the scattering matrix now becomes [18]:

$$
\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ V_S \end{bmatrix}^r_{t+\Delta t} = \frac{1}{4Z_S + Z} \mathbf{S} \begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ V_S \end{bmatrix}^i_t
\tag{2.39}
$$

with

$$
\mathbf{S} = \begin{bmatrix}
-(Z+2Z_S) & 2Z_S & 2Z_S & 2Z_S & 2Z \\
2Z_S & -(Z+2Z_S) & 2Z_S & 2Z_S & 2Z \\
2Z_S & 2Z_S & -(Z+2Z_S) & 2Z_S & 2Z \\
2Z_S & 2Z_S & 2Z_S & -(Z+2Z_S) & 2Z \\
2Z_S & 2Z_S & 2Z_S & 2Z_S & (Z-4Z_S)
\end{bmatrix}
\tag{2.40}
$$

In this solution, since all lines at the cell connections have the same impedance, the scattering process remains unchanged.

Stubs and Dispersion

The use of stubs results in an effective reduction of the wave propagation speed in a TLM mesh. Since the stub has parasitic behavior, it is expected that this speed reduction will be frequency dependent. The dispersion of a computation is directly related to the wave speed. Thus, one may expect that the dispersion relation for the 2D TLM mesh will change. From [18], the stub capacitance may be expressed as:

$$
C_S = \Delta x \Delta y \epsilon \left\{ 1 - \frac{\Delta t^2}{\mu \epsilon} \left[\frac{\Delta x^2 + \Delta y^2}{\Delta x^2 \Delta y^2} \right] \right\}
\tag{2.41}
$$

For a conventional, thus regular, TLM mesh with $\Delta x = \Delta y$, and with $\mu_r = \epsilon_r = 1$, the $\Delta t = \frac{1}{a\sqrt{2}}$, with a the wave speed. If ϵ_r and/or μ_r no longer equal to unity, we may find the following relation for the cut-off frequency of the mesh:

$$
\left(\frac{\Delta x}{\lambda} \right)_{cut-off} = \frac{1}{\pi} sin^{-1} \left[\frac{1}{\Delta x \sqrt{\mu_r \epsilon_r} \sqrt{F}} \right]
\tag{2.42}
$$

with:

$$
F = \frac{\Delta t^2}{\mu_r \epsilon_r} \left[\frac{\Delta x^2 + \Delta y^2}{\Delta x^2 \Delta y^2} \right]
\tag{2.43}
$$
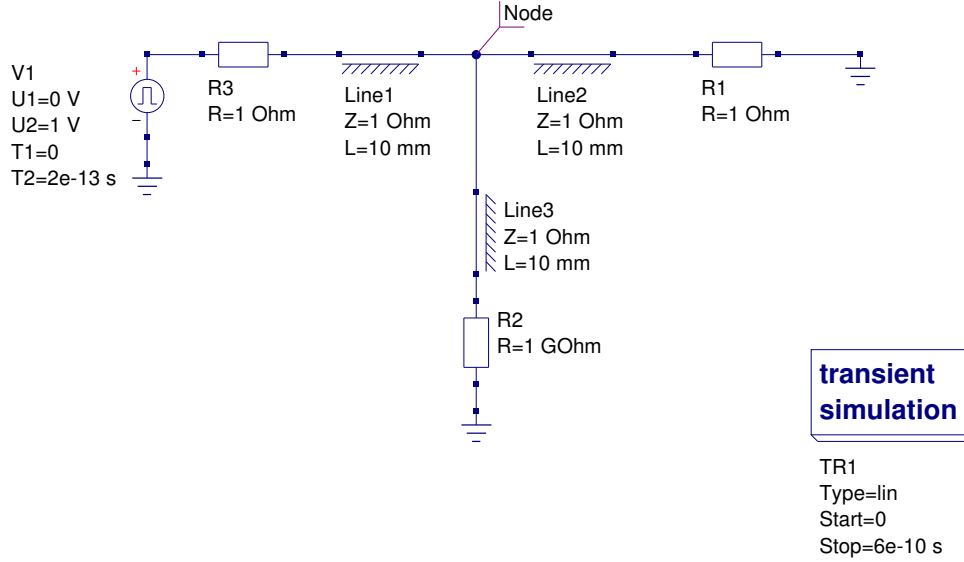
Figure 2.7: 1D TLM Node with Stub

The cut-off frequency is the highest when no dissimilar materials are present in the simulation. When F increases, the cut-off frequency of the mesh decreases. Thus, the more dissimilar the materials are, the larger the incurred error. The accuracy may be improved by reducing the size of the mesh cells, which is directly related to changing the sampling rate of the system, as the size of the cells is linked to the traveling time of the line segments, which is equal to an integer multiple, usually unity, of the simulation time step.

Stubs and Node Reflections

The use of stubs has a drawback in that the cell node exhibits an oscillatory transient response [18]. Indeed, the energy stored in the stub is reflected back to the node. Here, it encounters an impedance discontinuity, resulting in only a part of the energy in the line to be injected back into the mesh. Thus, these reflections cause a transient behavior.

To illustrate this, consider a 1D cell with an open circuit stub of length $\Delta l/2$, as illustrated in Figure 2.7. An incident step of magnitude 1 is injected. The stub line has the same impedance as the other lines, all set to unity. The lines are terminated in their characteristic impedance.

When this system is solved, the voltage behavior on the node is shown in Figure 2.8. The solution was computed using QUCS v0.0.14 [70], an open-source circuit simulator. The results show small spikes on the transitions. These are due to the use of an "infinite" rise-time, thus zero in one time step and unity in the next, pulse with the trapezoidal integration rule without critical damping adjustment (CDA) [48], and may be ignored.

The reflections due to the use of stubs are clearly visible. Although they die out after
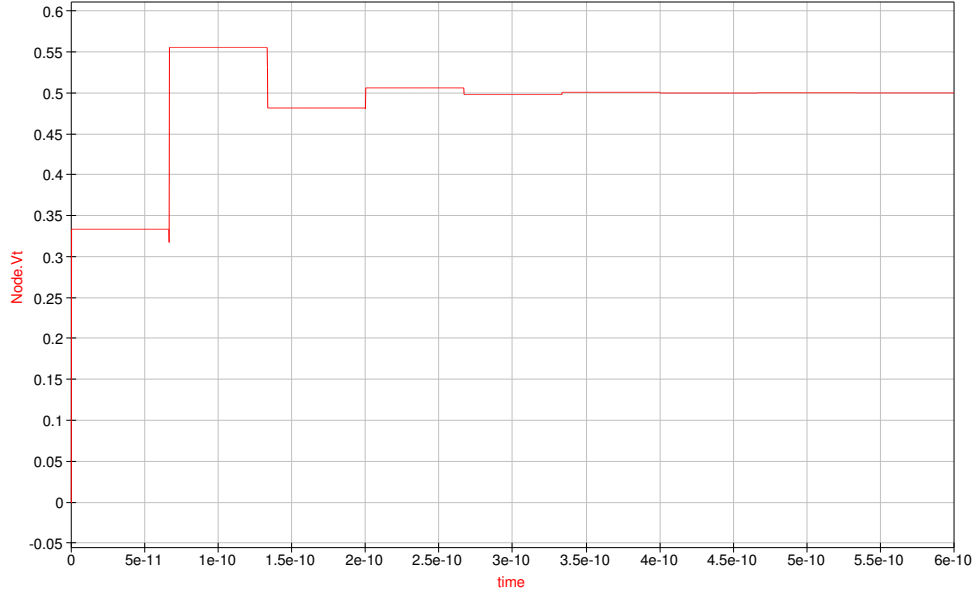
Figure 2.8: 1D Node Reflections

only a few time steps, their magnitude is significant and will influence the transient behavior of a computation. Thus, when using stubs, one must insure that the transient in the system is much slower than the numerical transients caused by the stubs. The transient response of the interpolation-based method introduced later in this chapter, Figure 2.12, has an inherently non-oscillatory response.

Note that, since the stub is half a time-step in length, the node would only be observed from the over-all simulation during the moments the wave-form is non-zero, as a reflection on this line would require it to be traveled twice, and thus synchronize with the overall simulation time step. The transient response is thus visible in the results.

2.1.6   3D in TLM

In this thesis, we limit ourselves to 1D and 2D computations. 3D solutions are possible, but the solution gains much complexity. For electro-magnetic fields, there is the phenomenon of polarization that does not exist in 1D and 2D fields. The acoustic equivalent would be the existence of compression, shear, and torsional modes in a solid. In order to capture these phenomena, the TLM cells must incorporate orthogonal field components on each connecting face. Where, in 2D, one only had to compute one voltage for one transmission line at each connection to the cell, here, two orthogonal transmission lines must connect, to form a plane in which the angle of the voltage vector may be described. Thus, the scattering matrix becomes 12 X 12, as there are six sides to the unit cell [16] for the series condensed node (SCN). This

is for a case where no stub lines are required.

However, for acoustics in gasses and many liquids, there is no polarization (shear and torsional forces with their associated mode conversion), and the 3D case is found simply by adding line segments in the third dimension and adjusting the various computations accordingly. Thus, the 3D acoustical node in these limited cases contains six line segments as opposed to four for the 2D case [66, 67], for a case where no stub lines are required.

When different media must be incorporated in a computation, the scattering matrix for solutions beyond 2D incorporates more that one stub line. For example, the 3D SCN for acoustics [66] has one shunt and three series stubs in the system. This brings the total to ten transmission line segments per cell.

The resulting circuit for the full model including polarization, incorporating stub lines, has no relation to anything physical. For 3D, it merely becomes a formal representation of the field solution at the centre of the node.

2.2   Transient Insular Nodal Analysis

In this thesis, a modified version of TLM was devised and used. The rationale behind this was the realization that EMTP power systems simulations have a host of very flexible and numerically efficient transmission line models available to them, varying from loss-less, delay-only, models to full-fledged frequency-dependent models with arbitrary response. These models thus showed promise to allow complex materials to be modeled in a TLM context.

Since the TLM method is, essentially, a circuit solution of a mesh of transmission lines, the EMTP models should be suitable for use in this context. In EMTP, temporal interpolation is used. Using this approach, stubs and their incurred node reflections, discussed above, may be prevented at a relatively modest increase in computation.

Functionally, the EMTP line models are based on an ideal transmission line, thus a time delay, with additional RGLC components to model losses and frequency-dependence, as required. Thus, one of the main advantages of TLM, time decoupling, remains present and, with it, the possibility of effective parallelization.

In order to make the TLM method suitable for direct insertion of EMTP transmission line models, we start from the observation that the majority of these models, in the time domain, may be represented by a Thévenin equivalent and interface more easily with the connecting circuit also represented by such an equivalent.

Thus, if we modify the TLM method so that its mesh cells present a two or four port, in 1D and 2D, respectively, and each port is a Thévenin equivalent, we may obtain building blocks from which a simulation case may be build. The simulator then only needs to provide the
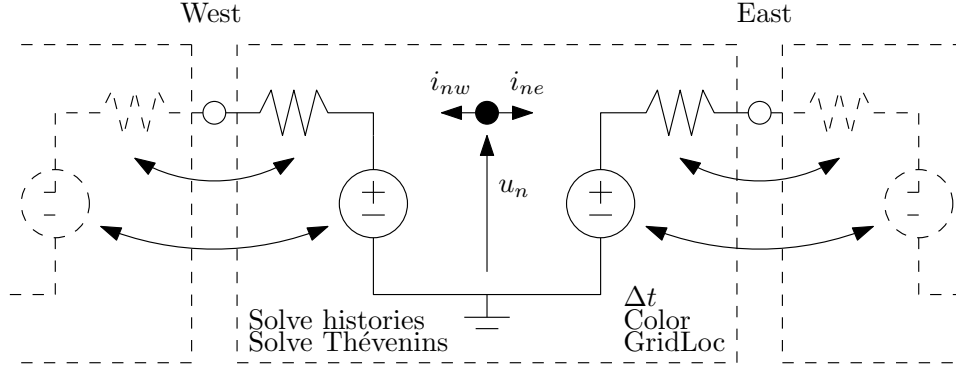
Figure 2.9: Loss-Less, 1D TINA Node Cell

connectivity between such cells, and manage the solution to instruct each cell to solve itself. Lastly, the simulator also manages the output processing, where the desired results from the cells are stored as required.

From this, the name of the new simulator, TINA, may be explained: Transient Insular Nodal Analysis. TINA is a transient simulator that exhibits full time-decoupling between its mesh cells allowing them to be solved independent from one another, thus makes them insular. The solution is based on a localized nodal analysis, only incorporating the components pertinent to the node under solution. Each insular cell is a fully self-contained model that exports a universal interface to the simulator: Thévenin equivalents for information exchange and a standardized API for the simulator to give instructions to the cell and obtain solution data from the cell. Such a cell, for one dimension, is shown in Figure 2.9. The internal details are not important to the TINA simulator, as long as the cell exposes suitable Thévenin equivalents, variables, operators, and central node state data.

From the figure, we may also observe that a link resistor formulation is used [18] to interface with the neighboring cells. In general, there are two choices: link-line and link-resistor. The first connects individual cells though a shared transmission line. The issue with this is when a material discontinuity is present. The line models used in this thesis cannot be readily split in such a way, where each half has different parameters, unless two separate line segments are used. In that case, one automatically has a link resistor formulation, where two neighboring cells are joined through a resistor. In TLM, however, a link-line formulation is possible with different materials when stub-matching is used at the nodes, and all the transmission lines in the mesh are thus identical. However, since TINA uses temporal interpolation, and thus line of different parameters for different media. Thus, a link resistor formulation is the more flexible and logical choice.

Also, TINA uses a single-wire formulation. This implies a ground, thus absolute reference,

throughout the simulation. This change from the two-wire formulation common in TLM was made to simplify the solution. The second wire is implied, as the ground serves as a mirror in transmission-line theory [7]. Since the TINA formulation uses a direct circuit formulation and solution, a two-wire formulation was not required and a solution to ground sufficient, as it is functionally equivalent.

In TINA, each cell is a complete material model that incorporates the transmission lines and all their associated components, as well as the node solution for that cell. This approach allows for model-specific optimizations, as well as potentially allow for any special functions to be integrated into the cells, such as non-linearities and various dependencies, such as temperature.

For the TINA implementation in this thesis, this segregation of the local and global solution was driven to its logical extreme, with each mesh cell a C++ object [44]. This implementation allowed a variety of cell models to be implemented and tested using the same solver code. Also, different solver synchronicity techniques, such as interpolation and fractional latency, could be used with the same, well-tested cell models. This proved convenient for development and presented a clean programming abstraction. However, the overhead incurred by literally hundred of thousands to millions of objects did result in a significant performance penalty. Future versions will likely be based on non-object approaches for the actual solution of the cells.

Another difference is the extensive use of interpolation in the TINA method for synchronicity, where TLM commonly uses stub lines. Although the TLM method can use either approach, its scattering matrix formulation favors stub lines as these can be implemented in a clean and efficient way, maintaining single time-step transmission lines.

TINA, with its self-contained cell solutions, can also be used with either, but since good transient response was a priority in the development of the method, the use of stubs was found unsatisfactory and temporal interpolation methods were chosen.

We will now show how the TINA model is obtained from TLM, and that the method is essentially identical, but for the formulation and solution approach. The changes allow for a more flexible simulation, allowing the direct use of EMTP line models, as well the new fractional latency technique, introduced later in Chapter 4.

2.2.1 The TINA Formulation

To obtain the TINA formulation, we must first return to the development of TLM. Although TLM uses a scattering solution, this is really just an implicit circuit solution of the cell. For TINA, we will use an explicit circuit solution for the cell and its boundaries.

The second observation is that the lumped element model used in TLM may be substituted
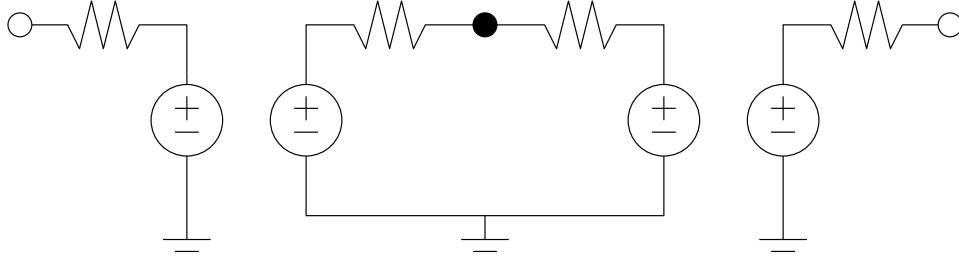
Figure 2.10: 1D TINA Node Cell Internals

by a transmission-line segment. This is, in fact, done in TLM though the scattering matrix formulation, which works with line parameters of impedance and the wave speed is implicit to one simulation time step. In TINA, the transmission line segments may be any temporal length and their actual models, which are derived in Appendix A, are solved explicitly. This choice allows the transmission line model to have any form, as long as it maintains time-step decoupling. It also allows the transmission line segments to have time histories that are deeper than one time step. This requirement is discussed in the next section of this chapter, where the temporal interpolation solution for synchronicity is introduced.

Thus, TINA is a time-decoupled, explicit circuit solution version of TLM that uses temporal interpolation for synchronicity. The latter is made directly possible due to the explicit use of the transmission-line models, which can have any size of history buffers. This is not as straight-forward in the common scattering TLM formulation, where a single time step history is expected.

By packaging all the lines and their associated components in cell models, each of which is time-decoupled from one-another, the basic TINA solution becomes trivial. All the solver has to do is tell the cells, in any order desired, to first solve their internal histories. When that is done, it tells the cells again, in any order desired, to update their external Thévenin equivalents. The solver will also query each cell for its internal node state information, which is the solution of the system. This sequence of events is identical to that used in EMTP circuit simulations [90]. Hence, one may think of TINA as a higher-dimensional extension to the EMTP power system simulator.

The actual circuit solutions in each cell are made through basic nodal analysis [79] of the circuit. Since the circuit is small and time-decoupled from the neighbors, this solution is trivial. The solution is further simplified, as each neighbor presents a Thévenin equivalent, which is fed through a length of time-decoupling transmission line that, in itself, also automatically presents a Thévenin equivalent, in the loss-less case. Thus, the node solution becomes a simple R-only circuit with two or four sources, for 1D and 2D, respectively. This situation is schematically

illustrated in Figure 2.10, for a 1D case. Each pair of sources and resistors forms a section of loss-less CP line model, as discussed in detail in the following chapters and Appendix A. The circuit shown in this figure is the contents of a typical TINA building block. When compared to the abstract cell of Figure 2.9, we may see how directly this fits in the general TINA concept of cell abstraction.

Since nearly all EMTP transmission line models have a similar circuit formulation [20], this elemental cell can be readily adapted to accommodate such models and add losses to the simulation. Also, other components may be integrated in the solution to add non-linearities to a simulation e.g.: [89]. None of this requires modification to the basic simulator algorithm itself, making the solve inherently modular.

A direct, algebraic solution of the circuit in the cell was chosen over a matrix formulation. The problem with the matrix formulation for such a small system was that constructing the matrix and solving it, even using machine specific, optimized mathematics libraries, such as GotoBLAS [24], proved to perform at roughly half the speed of the direct solution. The direct solution, when the equations of same form are grouped, and using a suitable, optimizing compiler, actually allows for some degree of parallelism in vector operations which modern CPUs offer in their mathematical units.

2.2.2 The Interpolated Line Model in TINA

The fundamental issue with the transmission line method, and also TINA, when different materials are incorporated in the simulation, is that the traveling time of the wave in each length of transmission line must be an integer multiple of the simulation time step to maintain synchronicity. Clearly, with different materials, this will not always be the case.

The reason why integer multiples are required has to do with the nature of the transmission-line models used. They are time-delays, thus the only values that exist in their memory, called history sources, are values at the simulation time steps. When the traveling time does not match with the simulation time step, data points in between the available information is required. This is discussed in Appendix A.

In TLM, stub-matching is usually employed to avoid this synchronicity issue. By choosing the base medium so that each length of line is an integer multiple, usually unity, of the simulation time step, and then adjusting the medium with a stub line where required, the desired local wave speed and medium impedance is obtained. In TINA, the other method available, temporal interpolation, is used, as this is the approach has a clean transient response, as is desired for a transient simulator. This is also the method used in the EMTP.

The concept behind the interpolation process and its requirements is detailed in Appendix

A.3, and the behavior of the interpolation process, with respect to stability and accuracy, is studied in Chapter 3. All of these studies are for the one-dimensional line, which is the fundamental building block of all TINA models.

From those chapters, it is found that the interpolation process, when applied to the lossless line is stable. It is even more stable than the ideal line model itself, as it adds some amount of dampening to the system due to the losses incurred by the low-pass filter effect of the interpolation. The low-pass filter effect is also present in the stub matching approach [18]. The TINA method, being nothing but a circuit solution, is in itself conserving and linear, as the circuit solutions are conserving through Kirchoff's laws. Thus, the whole system must also be conserving.

When stubs are not used, and different materials are present, the transmission line segments must now also have different lengths, or rather, history depths. This means that the temporal interpolation method to synchronicity will require more memory than the stub-matched approach.

To illustrate the issue, consider a material with wave speed 1.03 m/s, and one with wave speed 10 m/s. For a simulation time step, which would have to be linked to the faster material, there are two ways to incorporate the slower material.

Using stubs, the slower material would be composed out of the same transmission lines as the fast material, and suitable stub lines would be added to bring-down the speed of the material to the desired one. Thus, since all lines operate at the same time step, synchronicity is obtained and all lines may be chosen so that their traveling time is precisely equal to the simulation time step. In TLM is usually chosen so that the traveling time equals the simulation time step. Thus, the memory required to implement the lines is only one memory location for the one step deep history buffer.

Using temporal interpolation, the fast material would be chosen in the same way, but the slow material would be composed out of transmission lines that are much longer, temporally. Since the simulation operates at a rate roughly ten times faster than the traveling time of the line, it needs to store a data point for each value in the past in the buffer of the line. Thus, roughly simulation time steps. As such, the history buffers associated with the lines must be about ten steps deep. The temporal interpolation method thus requires about ten times more storage.

In addition, the interpolation of the history sources requires more computation than the stub line requires.

As such, temporal interpolation, while allowing more flexibility and direct compatibility with the EMTP transmission line models, does add a significant overhead to the computation, both in number of operations and memory use. That said, it does have a clean transient response

due to the absence of stubs.

Much of the lost performance may, however, be re-gained through the use of fractional latency, introduced in Chapter 4, which only requires interpolation at the borders of a material, as opposed to throughout its entire volume, and further allows a suitable local time step so that each line in the simulation may have history buffers one step deep. Using this, the numerical performance my be better than using stubs, as the overhead of calculating the stubs is removed, as they are no longer required.

2.2.3 Error Criterion

As with any discretized method, there are limits to how coarse a computation may be for a desired precision. Here, we discretize in both time and space. Both are intimately related, as will be seen below.

First, consider the spatial discretization. To prevent spatial aliasing, two cells per wavelength are a minimum. However, in practice, about five to ten cells per wavelength are typically used for accuracy [18]. From this, it follows that the shortest wavelength in the computation will determine the maximum allowed grid size in a material. This will thus be the highest occurring frequency. However, and this is the subtlety, it is the highest occurring frequency in the material with the slowest wave speed that will have the shortest wave length.

For the temporal discretization, this is given by the traveling time of the line segments in each grid cell. If each material has traveling times that are integer multiples of one another, and thus no interpolation is needed for synchronicity, the simulation time step simply becomes the shortest traveling time in the system. This will be the fastest medium in the system, where the size of that cell, and thus the line lengths in it, is set by the slowest medium.

If, however, interpolation is used for synchronicity, the simulation is no longer based on ideal line segments. The interpolation adds additional accuracy requirements, as discussed in Chapter 3. As is custom in EMTP [20], and shown in that chapter, the used simulation time step must be five to ten times smaller than the traveling time of the line segments.

Thus, two error criteria may be stated for the TINA method: one for when interpolation is used, and one for when it isn't. These are summarized in Table 2.1, where a_{max} is the wave speed of the fastest material and f_{max} is the highest frequency in the simulation. Δl_{max} and Δt_{max} are the maximum spatial and temporal discretization sizes allowed for these parameters, taking a ten times spatial oversampling into account. This automatically results in a ten times temporal oversampling. When interplation is used, an additional ten times temporal oversampling for the fastest line segments is required, resulting in a total hundred times temporal oversampling.

Discretization	Ideal Lines	Interpolated Lines
Spatial	$\Delta l_{max} = \frac{a_{min}}{10 f_{max}}$	$\Delta l_{max} = \frac{a_{min}}{10 f_{max}}$
Temporal	$\Delta t_{max} = \frac{1}{10 f_{max}}$	$\Delta t_{max} = \frac{1}{100 f_{max}}$

Table 2.1: TINA Error Criteria

Thus, the use of interpolation requires ten times more time steps than the non-interpolated computation. However, it must be noted that the use of stubs in TLM also requires a reduced time step to maintain accuracy [18], as it also exhibits a low-pass filtering effect.

2.2.4 Domain Termination

A problem common to nearly all discretized-space methods is that of domain termination. In the computer, only a finite expanse of space can be represented, meaning that for most cases, the outer bounds of the simulation space must somehow be terminated in a known way. Usually, one would prefer the medium to be unbounded, so that any wave incident on the edges of the simulation space continues as if the space were infinite. The two other common boundary conditions are that of the open circuit or short [16].

The termination into infinity for a 1D mesh is easy to obtain, and is simply the termination of the transmission line into a resistance of equal impedance to the line [46, 52].

Unfortunately, this termination into infinity is hard to obtain beyond 1D, and much has been written about this subject. The difficulty lies in the fact that one must terminate the medium, not the transmission lines that make-up the medium. Since it was already established in this chapter that the TLM, and by extention also TINA, mesh is anisotropic, with a different propagation speed depending on the direction of propagation. It follows that the effective medium characteristics are than also different in different directions of propagation. This poses an issue with the boundary, which has no direct information on the direction of wave propagation.

Lately, the consensus is that the split-field perfectly matched layer (PML), proposed by Berenger in 1994 [4], and its more recent generalization, stretched-coordinate PML [12, 85], perform best. This first technique is based on decomposing the electro-magnetic field into two un-physical fields in the terminating region and applying complex termination materials to these. In essence, the wave is de-composed in orthogonal components, which are now propagation direction in-sensitive and can thus be easily terminated in a suitable complex impedance to match the frequency-dependant (dispersion) behaviour of the mesh.

A newer formulation is based on a Cartesian to cylindrical coordinate system transforma-
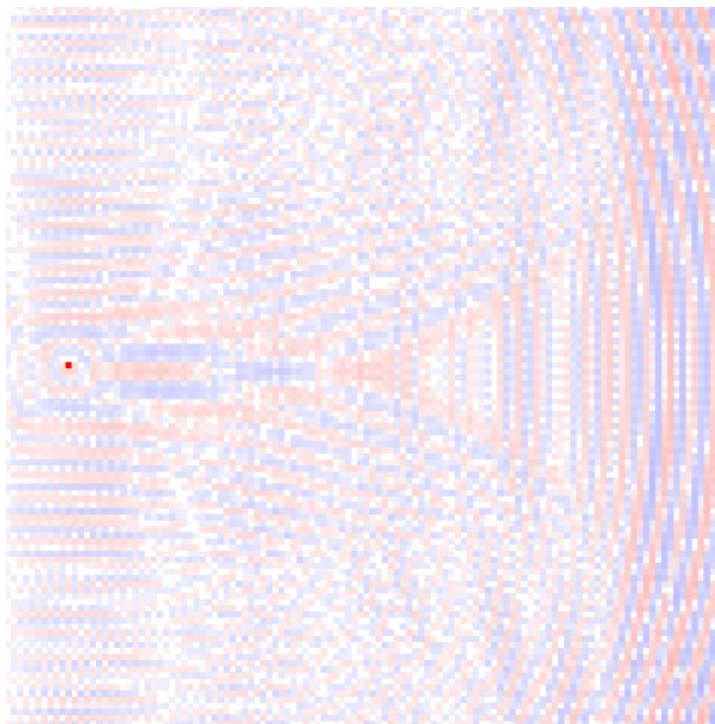
Figure 2.11: Boundary Reflections in a Resistively Terminated 2D TINA Mesh

tion, where one or more coordinates are mapped to complex numbers. It represents an analytic continuation of the wave equation into complex coordinates, replacing propagating, thus oscillating, waves by exponentially decaying waves.

In this thesis, a simple, resistive termination was used to terminate both the 1D and 2D grids [16]. However, this resulted in poor attenuation of the incident wave on the boundary for 2D, and significant reflections occurred. This is why the precision experiments in this thesis were done on closed ducts, so little energy would leave the system under investigation and have a chance to interact with this boundary. Also, since all systems measured were 2D, the termination of the 2D experimental setup into 3D free-space could not be modeled accurately, and thus some degree of inaccuracy would be incurred in those physical experiments. Chapter 5 shows those results.

Figure 2.11 show what happens when a 2D mesh is resistively terminated, using boundary conditions matched to the medium [16]. The errors were somewhat amplified in this figure to make the effect more pronounced. But, as will be clear in Chapter 5, for the open-plane experiments, the influence is quite significant.

It must, however, be noted that the matched boundary conditions work reasonable well if the wave is incident at right angles [16]. The incident at right angle limitation is reminiscent of the dispersion of the mesh, with different propagation velocity depending on the angle the

wave travels to the mesh coordinates. This is why a simple, resistive, termination does not work. The mesh has different properties depending on the angle of incidence, thus a simple, real resistance is insufficient. Spit-field PML works around this be decomposing the field in orthogonal components and presents suitable complex terminations to these.

For the purpose of the thesis, matched boundary conditions provided a sufficient, and computationally efficient, solution. Further work on TINA, however, would include the implementation of PML-style boundaries.

2.2.5   Validation of the TINA method

To evaluate both the correctness of the implementation, and the method itself, three types of validation were performed.

First, the TINA solver itself was compared with EMTP, for identical small-scale systems. This validates the numerical implementation of the solver and its models. Also an interpolated version of the models was tested this way.

Second, a time-domain evaluation of the step response for the interpolated model was performed to verify that the transient response is, indeed, clean and the use of interpolation justified when the transient response is important.

Third, the interpolated solution was compared with the non-interpolated version for larger simulations, with parameters chosen so the results should be nearly identical. This was to experiment with different simulation parameters and partially validate the error criteria.

Last, the TINA solutions were compared with measurements. No direct comparisons with a TLM solver were made, as the TINA method is a only a re-formulation of TLM. The experimental validation is detailed in Chapter 5.

Validation with EMTP

Using Mictrotran [56], the EMTP version developed at UBC, a number of 1D and 2D cases were constructed and compared to the TINA solutions. The 1D cases had three cells in them, the 2D cases were three by three grids. For each system, computations with and without interpolation were performed. Also, different materials were included in the tests.

It was found that all results were identical to the EMTP reference solutions, proving that the TINA solver and models are a correct implementation of both the circuit solution and the constituent loss-less transmission-line models.
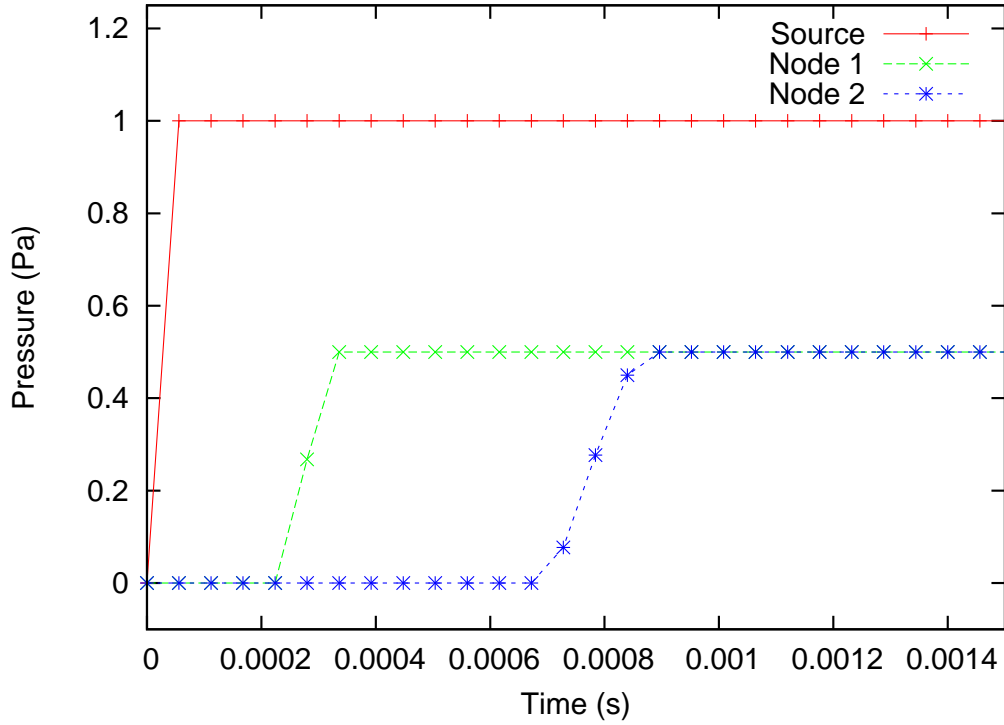
Figure 2.12: Step Response in a 1D TINA Mesh

Time-Domain Step Response of 1D TINA Mesh

Since one of the main reasons behind the use of interpolation for synchronicity was the clean transient response, compared to stubs in Figure 2.8, also this was evaluated. The experiment was performed on 1D grid of three cells in length. The simulation time step was 56 μs, the wave speed 2000 m/s, and the medium impedance 100 Ω. For this combination of parameters, interpolation is required. The mesh was excited with a step.

In Figure 2.12, the source data was taken at the ideal source node, before its matching resistor. Hence, the amplitude is twice that of the rest of the mesh. Also, the source was implemented so that it had no traveling time associated with it. Thus, the distance from the source to node 1 was $\Delta l/2$, while the distance between the grid nodes is Δl. Hence, the propagation time from the source node to node 1 is half that of the normal propagation time in the mesh, as would be the case between node 1 and node 2.

Two main effects are immediately apparent: the step response has no transient behavior, as would be the case with stub matching. Second, the low-pass filter effect of the interpolation removes some of the higher frequency components, causing the slope edges to roll-off, as higher frequency components are more and more attenuated as the signal passes through the
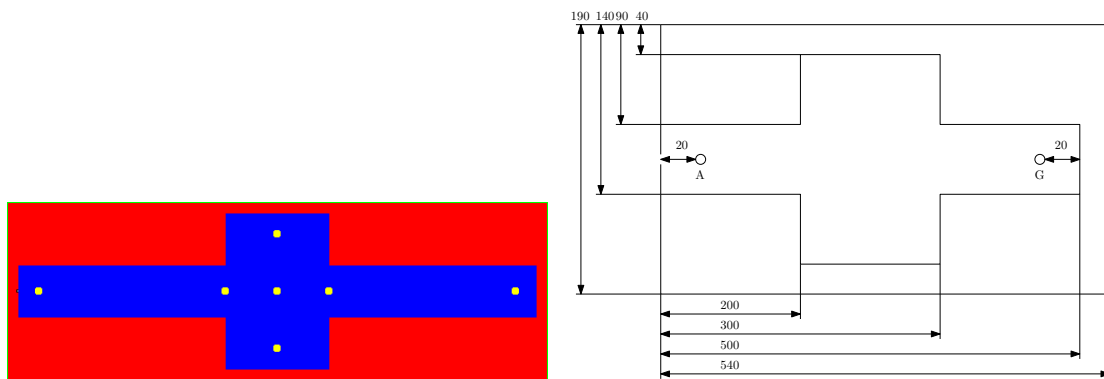
Figure 2.13: Expansion Duct Set-Up

mesh. Hence the need for increased oversampling when interpolation or stubs are used, so the higher-frequency performance may be retained, if required.

Performance Comparison of Interpolation in TINA

Three different spatial discretizations will be used for this comparison, namely 0.5 mm, 1 mm, and 2 mm square cells. The system under investigation is the expansion duct case, as described in Chapter 5. The normal simulation uses slightly adjusted parameters for the aluminum to insure no interpolation is used in the solution. The system is shown in Figure 2.13. The utmost left yellow dot is the "A" microphone position and the utmost right yellow dot is the "G" microphone position. Transfer functions between those two physical locations were calculated and compared.

The physical size of the microphone itself was also taken into account, which is a sizable portion of the wave length for the higher frequencies evaluated. Thus, in the simulation, the acoustic pressure over an area comparable to that of the physical microphone was averaged to obtain the result for that location.

The differences between the precise and adjusted materials are shown in Table 4.3, where ρ is the compressibility, k the compressibility, a the wave speed, and z the impedance. These parameters are acoustic ones, but relate directly to their electrical equivalents of permeability and permittivity.

From Figures 2.14 and 2.15, we see that the normal and interpolated methods are self-consistent and overlap. For the phase, the unwrapping algorithm made an number of incorrect unwraps due to small differences in the phase of the input signal. Thus, the error between the traces is not representative of the actual phase error incurred. The different spatial discretization, and thus accuracy by which the system may be represented, results in some deviations between the cases, as the physical size is somewhat different. The system has to be discretized

	Air	Al (correct)	Al (adjusted)
ρ	$1.1198 \times 10^3 \frac{g}{m^3}$	$2.6989 \times 10^6 \frac{g}{m^3}$	$2.6989 \times 10^6 \frac{g}{m^3}$
k	$6.99746 \times 10^{-9} \frac{ms^2}{g}$	$1.45138 \times 10^{-14} \frac{ms^2}{g}$	$1.584727 \times 10^{-14} \frac{ms^2}{g}$
a	$3.45383 \times 10^2 \frac{m}{s}$	$5.05261 \times 10^3 \frac{m}{s}$	$4.83537 \times 10^3 \frac{m}{s}$
z	$4.13769 \times 10^5 \frac{g}{sm^2}$	$1.36365 \times 10^{10} \frac{g}{sm^2}$	$1.30502 \times 10^{10} \frac{g}{sm^2}$

Table 2.2: Materials Used in the TINA Simulations
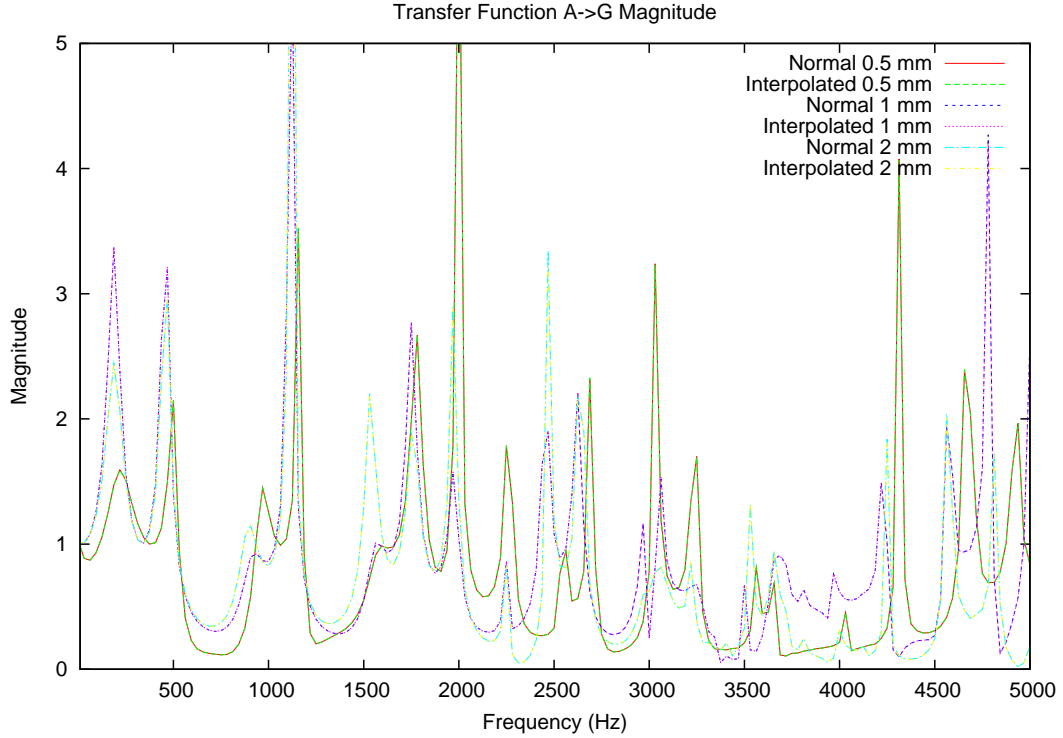
Figure 2.14: Magnitude Results for Various Mesh Sizes

to 1 mm to be able to accurately represent the dimensions. Also, for higher frequencies, the errors increase faster with coarser spatial discretization, mostly due to the spatial error becoming an increasingly more significant portion of the acoustic wavelength in the medium for those frequencies.

Attempts to cause insufficient temporal discretization or spatial oversampling proved problematic, as the bandwidth obtained from the required spatial discretization to represent the physical system and provide sufficiently accurate microphone positions was in the megahertz range, far beyond what could be verified through measurement with the available equipment. Thus, any conclusions drawn from such simulations would not be falsifiable through the available experiments. This also implies that, for the auditory acoustic range usually of interest,
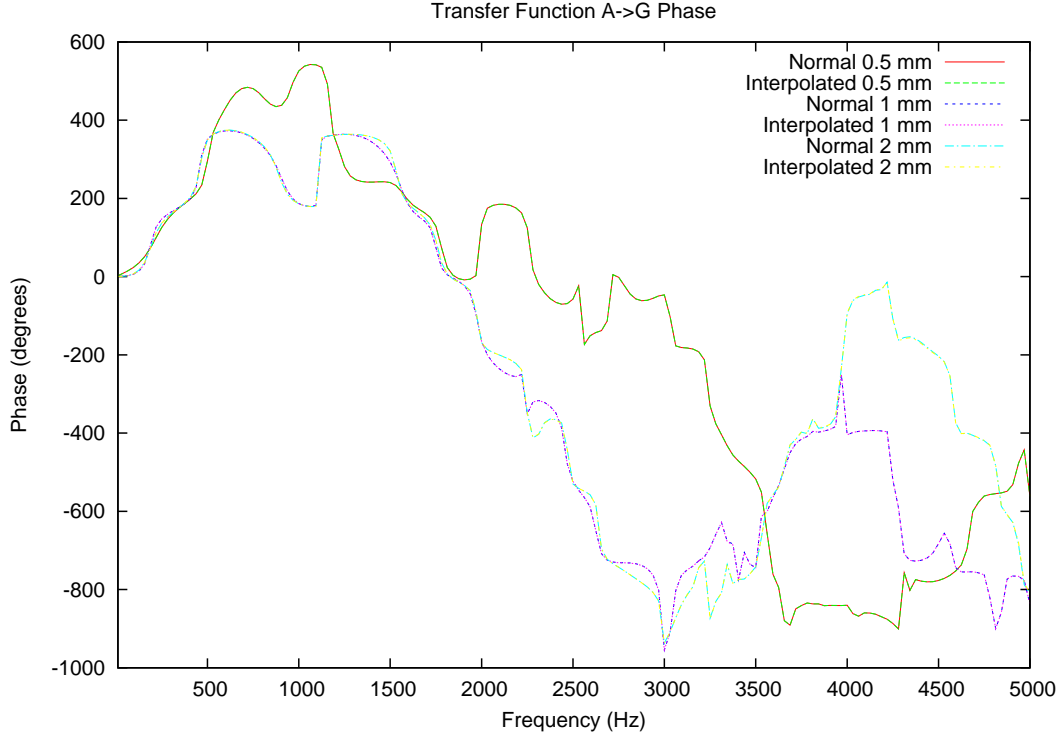
Figure 2.15: Phase Results for Various Mesh Sizes

when simulating small cases that require high, millimeter-scale precision, the available bandwidths are so high that a lot of computation is wasted. This is a problem of scale, and is discussed in Chapter 4.

More detailed comparisons between these cases may be found in Chapter 4, where the relative memory use and computation times for 1 mm cell size cases is discussed. These are long simulations, over seven hours compared to the less than half an hour required for the 2 mm cell sizes. Thus, any cumulative errors will be more likely to show there. The experimental validation of the method is discussed in Chapter 5.

2.3 Conclusions

In this section, we discussed the TLM method, and how TINA is a direct descendant of it. The TINA formulation, which is closely inspired by EMTP approaches, has the advantage of direct compatibility with EMTP transmission-line models, and allows complex material cells to be inserted into its mesh, as long as these cells obey the API and provide suitable Thévenin equivalents to the solver. The mesh cells are conceived as self-contained units that may be called upon to solve themselves, based on the EMTP solution algorithm.

The use of temporal interpolation allows for clean transient responses, at the cost of increased computation and memory consumption, when compared to stub matching techniques. However, the ability to have materials with arbitrary history buffer depths makes the method more flexible than the scattering matrix formulation, which is based on history buffers only one step deep. Also, using fractional latency, as discussed in Chapter 4 in this thesis, much of the lost performance may be re-gained.

The TINA method was evaluated for numerical correctness and consistency with itself and EMTP. The interpolation was shown to yield good results. Further experimental evaluation is performed in Chapter 5 and further numerical validation is performed in Chapter 4 of this thesis.

Future work would include the optimization of the TINA solver to avoid the use of objects in the actual network solution, as well as the implementation of PML mesh termination. Also, a detailed study of the distortions introduced by the mesh, inspired by prior work on TLM, would be of interest.

Chapter 3

Convergence, Dissipation, and Dispersion in 1D Loss-Less Line Models

> A small error in the former will produce an enormous error in the latter.
> *Henri Poincaré*

In this chapter, the properties of the one-dimensional line models themselves are studied in detail. Their numerical convergence, dissipation, and dispersion is evaluated analytically for the loss-less ideal and interpolated cases. Although the studied line models have been in wide use in EMTP power system simulations since the sixties, e.g.:[90], this analysis has not been performed before in the literature. Error expressions are derived that allow the incurred errors of an interpolated line under a given set of operating conditions to be determined. This forms a contribution of this thesis.

In previous chapters, we showed that, as in the TLM method, the mesh in the TINA method is constructed from a spatial combination of one-dimensional transmission lines. The derivation of these one-dimensional building-block line models is shown in Appendix A.

The computation of the TINA mesh itself is entirely based on the application of circuit analysis, and thus is stable and conserves energy. As such, the problem of evaluating the simulation stability is reduced to analyzing the stability of the one-dimensional line models used to construct the mesh. However, the numerical accuracy of the TINA mesh, in terms of dissipation and dispersion error, is determined by both the spatial discretization and the inherent accuracy of the line models used to construct the mesh.

From the analysis, a commonly used "rule of thumb", the necessity for choosing the simulation time step Δt at least ten times smaller than the line traveling time τ [20, 90], can be obtained. Also, the incurred error is quantified, making it possible to relax the requirement in certain cases.

3.1 Convergence of the Line Models

In order for the line models to be useful, the numerical solution must be stable. However, stability in itself is not a sufficient condition. Numerical methods can be stable, but still yield

the wrong results when a solution converges to a wrong value. Thus, a better criterion must be used.

What must be proven of a numerical scheme is convergence [83], which proves that the eventual solution is bound and yields the correct final answer. This is frequently non-trivial to show in a direct manner. The difficulties in showing convergence arise directly from the fact that the numerical method is defined by a recurrence relation, and the differential equation is defined by a differentiable function. However, convergence can be proven in an indirect fashion by showing both consistency and stability. Together, these two tests provide necessary and sufficient conditions for convergence, according to the Lax-Richtmyer equivalence theorem [40]. The theorem applies to TLM and TINA as well, as these methods can be show to be equivalent [9, 11].

> For finite-difference schemes for initial-value problems, the Lax-Richtmyer equivalence theorem states that a consistent finite-difference scheme for a partial differential equation, for which the initial value problem is well-posed, is convergent if and only if it is stable [83].

Thus, the problem of proving convergence is replaced by the equivalent and simpler to verify conditions of consistency and stability. The theorem can also be applied to higher-order schemes [72, 83].

Consistency

Consistency, when applied to a discretized model, is a test that insures that the solution tends to the correct answer. The method described by Strikwerda [83] is used. The test is done by establishing a point-wise convergence on each grid point or, in other words, by letting the spatial and temporal discretization intervals tend to zero. The discretized model should then approach the non-discretized one. Thus, the error in the numerical solution, when compared to the continuous model, should tend to zero.

Stability

Stability is established by testing if the numerical solution is bounded, or in other words, produces a finite output for a finite input. In this thesis, eigenvalue analysis was used to find the modes of the system. These modes give direct insight in the dynamics of the system, e.g.: [3, 38, 62, 83, 91]. This method is further detailed in Appendix B. It is then only a matter of insuring that each of these modes remains bounded to insure stability.

Well-posedness

Mathematical models of physical phenomena should have a unique solution that depends on the data in a continuous way, and is insensitive to small variations in the input data or model parameters. Such a model is referred to as well-posed. Or, in other words: does it even make sense to attempt to discretize the problem in the first place and can a well-behaving discrete model be obtained?

The study of well-posedness, pioneered by Hadamard in his seminal 1921 lecture [25], is a statement of the sensitivity of a problem to small changes in the input data. When we say that an equation is well-posed, it is meant that the initial value problem for the equation is well-posed.

In well-posed problems, small changes yield small differences in the solution. This helps the system to converge to a stable and correct value. A well-posed problem, when solved with a stable algorithm, has thus a good chance of reaching a correct and stable result. Ill-posed problems, by comparison, may yield wildly different results for small changes in the input. In the presence of noise, interpolation, modeling approximations, and numerical roundoff errors, such problems are typically difficult to solve. Thus, well-posedness is a test to ensure that an attempt to discretize a continuous-time system makes sense.

Another consideration is that a partial differential equation derived to model a physical system should be unaffected by the addition of, or changes in, the lower order terms, and by sufficiently small changes in the coefficients. This quality is referred to as robustness. It is important, as most models are derived by making assumptions and neglecting various effects deemed not important through an understanding of the system. Thus, these approximations change the model equations. The resulting model should be insensitive to the small perturbations caused by these omissions and still produce results that allow correct conclusions to be drawn from the computations.

To test for well-posedness, the system of governing equations must be evaluated. Such tests, based on finding the modes in frequency domain though eigenvalue analysis, were described by Kreiss [34]. Strikwerda, in his book [83], provides a treatment of the subject in the English language following the methods proposed by Kreiss.

The principle of the analysis, for equations of first order with one time derivative and constant coefficients, is to show that the real parts of the eigenvalues of the continuous-time system in frequency domain are always less or equal to a constant for all ω. In other words, the solutions, each of an exponential form, have exponents that ensure that each mode is bound for each frequency. This is a necessary and sufficient condition for well-posedness.

More exactly, the situation for the ideal line model, which is a system of two PDE of the general form:

$$\dot{\mathbf{u}}_t = \mathbf{Q}(\omega)\dot{\mathbf{u}} \tag{3.1}$$

In frequency domain, this system has a general solution of the form:

$$\dot{\mathbf{u}}(t,\omega) = e^{\mathbf{Q}(\omega)t}\dot{\mathbf{u}}_o(\omega) \tag{3.2}$$

The well-posedness condition for this system is now:

The necessary and sufficient condition for the system (3.2) to be well-posed [83] is that for each non-negative t, there is a constant C_t such that, for all $\omega \in \mathbf{R}^N$:

$$\|e^{\mathbf{Q}(\omega)t}\| \leq C_t \tag{3.3}$$

A necessary condition for the above to hold is that, for each eigenvalue $q(\omega)$ of $\mathbf{Q}(\omega)$, there is some constant q, such that, for all values of ω:

$$Re\ q(\omega) \leq q \tag{3.4}$$

It can also be seen that the eigenvalues of $\mathbf{Q}(\omega)$ govern the exponentials that make-up the solution of this model. Thus, the eigenvalues define the modes of the system, and their study will give the required information: whether or not the system remains bound to some constant value.

An exponential to the power of a general matrix, required for the first part of the well-posedness condition (3.3), is not trivial to compute. There is, however, a standard solution to this problem [62], outlined in Appendix B.3. The second part of the condition (3.4) is much easier, and necessary, but not sufficient.

The first part can be calculated if the $\mathbf{Q}(\omega)$ matrix is diagonalized. Hence the application of eigenvector and eigenvalue analysis. Aside from diagonalizing the system, the procedure also yields the individual eigenvalues that are needed to perform the second part of the well-posedness condition.

From the approach outlined in Appendix B.3, assuming that the eigenvalues are distinct, it is found that:

$$e^{\mathbf{Q}t} = \mathbf{V} \begin{bmatrix} e^{\lambda_1 t} & & 0 \\ & \ddots & \\ 0 & & e^{\lambda_n t} \end{bmatrix} \mathbf{V}^{-1} \tag{3.5}$$

Thus, the first part of the well-posedness condition (3.3) can be evaluated quite easily using

the above solution. Using the individual eigenvalues themselves, the second part of the well-posedness condition (3.4) can also be readily performed.

Now, applying all the above, starting from the transmission-line equations for the loss-less case [46], as derived in Appendix A.2:

$$\begin{cases} \frac{\partial v(x,t)}{\partial x} = -l\frac{\partial i(x,t)}{\partial t} \\ \frac{\partial i(x,t)}{\partial x} = -c\frac{\partial v(x,t)}{\partial t} \end{cases} \tag{3.6}$$

In order to find the modes from this system, needed for the well-posedness criteria, we need to solve it. This implies conversion to either Laplace or Fourier domain, as that reduces the problem from a system of PDE to a system of ODE. Diagonalization will then yield the required eigenvectors and eigenvalues.

Applying the Fourier transform to space, we thus find:

$$\begin{cases} j\omega V(\omega) = -l\frac{\partial I(\omega)}{\partial t} \\ j\omega I(\omega) = -c\frac{\partial V(\omega)}{\partial t} \end{cases} \tag{3.7}$$

Re-arranging:

$$\begin{cases} \frac{\partial I(\omega)}{\partial t} = -\frac{j\omega}{l}V(\omega) \\ \frac{\partial V(\omega)}{\partial t} = -\frac{j\omega}{c}I(\omega) \end{cases} \tag{3.8}$$

Writing in matrix form suitable for diagonalization though eigenvalue analysis:

$$\frac{\partial}{\partial t}\begin{bmatrix} V(\omega) \\ I(\omega) \end{bmatrix} = -\begin{bmatrix} 0 & \frac{j\omega}{c} \\ \frac{j\omega}{l} & 0 \end{bmatrix}\begin{bmatrix} V(\omega) \\ I(\omega) \end{bmatrix} \tag{3.9}$$

Computing the eigenvalues of (3.9):

$$det(\mathbf{Q} - \lambda\mathbf{I}) = 0 = \begin{vmatrix} -\lambda & -\frac{j\omega}{c} \\ -\frac{j\omega}{l} & -\lambda \end{vmatrix} = 0 \tag{3.10}$$

$$\lambda^2 - \frac{(j\omega)^2}{lc} = 0 \Leftrightarrow \lambda^2 = \frac{(j\omega)^2}{lc} \Leftrightarrow \lambda = \pm\frac{j\omega}{\sqrt{lc}} \tag{3.11}$$

Computing the eigenvectors, using these eigenvalues:

$$\mathbf{v_1} = \begin{bmatrix} 1 \\ -\sqrt{\frac{c}{l}} \end{bmatrix} \quad \mathbf{v_2} = \begin{bmatrix} 1 \\ \sqrt{\frac{c}{l}} \end{bmatrix} \tag{3.12}$$

Thus, $\mathbf{V}$, $\mathbf{\Lambda}$, and $\mathbf{V}^{-1}$ are:

$$\mathbf{V} = \begin{bmatrix} 1 & 1 \\ -\sqrt{\frac{c}{l}} & \sqrt{\frac{c}{l}} \end{bmatrix} \quad \mathbf{\Lambda} = \begin{bmatrix} \frac{j\omega}{\sqrt{lc}} & 0 \\ 0 & -\frac{j\omega}{\sqrt{lc}} \end{bmatrix} \quad \mathbf{V}^{-1} = \frac{1}{2} \begin{bmatrix} 1 & -\frac{1}{\sqrt{\frac{c}{l}}} \\ 1 & \frac{1}{\sqrt{\frac{c}{l}}} \end{bmatrix} \tag{3.13}$$

Where all the above matrices are in Fourier (ω) domain. We find that both the eigenvectors and eigenvalues have multiplicity one, and are thus distinct. The general solution for an exponential raised to the power of a matrix, outlined in Appendix B.3, can thus be used for the well-posedness analysis.

To calculate the norm of the matrix required for the first part of the well-posedness test (3.3), the expressions of (B.14) and (B.15) are required. Also, the following definitions regarding matrix norms are used [62]:

$$\|\mathbf{AB}\| \leq \|\mathbf{A}\|\|\mathbf{B}\| \tag{3.14a}$$

$$\|\mathbf{A}\|_\infty = \max_{1 \leq i \leq n} \sum_{j=1}^{m} |a_{ij}| \tag{3.14b}$$

Applying the above to the first well-posedness test, for the infinite norm [62], yields:

$$\|e^{\mathbf{Q}(\omega)t}\|_\infty \leq C_t \tag{3.15a}$$

$$\Leftrightarrow \|\mathbf{V}(\omega)e^{\mathbf{\Lambda}(\omega)t}\mathbf{V}(\omega)^{-1}\|_\infty \leq C_t \tag{3.15b}$$

$$\Leftrightarrow \|\mathbf{V}(\omega)\|_\infty \|e^{\mathbf{\Lambda}(\omega)t}\|_\infty \|\mathbf{V}(\omega)^{-1}\|_\infty \leq C_t \tag{3.15c}$$

$$\Leftrightarrow 2 \, |e^{\pm\frac{j\omega}{\sqrt{lc}}t}| \left(\frac{1}{2} + \frac{1}{2\sqrt{\frac{c}{l}}} \right) \leq C_t \tag{3.15d}$$

Since l and c are constant for a given case, only the exponential term requires consideration as it has variables. After simplification, the entire expression becomes a sum of two exponentials with purely imaginary exponents. It is thus bounded for any value of t and ω. Thus, the expression is changing, but always $\leq$ a constant value given by the line parameters, l and c.

From the eigenvalues, we can now apply the second part of the well-posedness test, as shown in (3.4), directly. Testing for the inequality, thus verifying if the real part of each eigenvalue is $\leq$ a constant:

$$Re\left(\pm\frac{j\omega}{\sqrt{lc}} \right) = 0 \leq C \tag{3.16}$$

Thus, we find that the equations governing the ideal line model are well-posed.
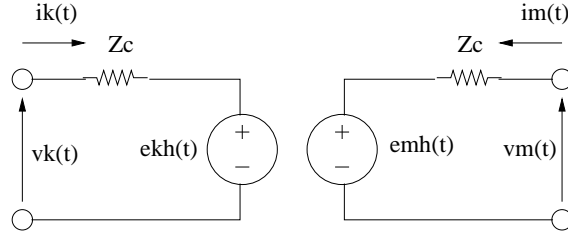
Figure 3.1: Lossless Line Model

3.1.1 The Ideal Line Model

Consistency

Derived in Appendix A.2, this line model is found as an exact solution to Maxwell's equations for the one-dimensional traverse electro-magnetic plane wave in the loss-less, non-dispersive, homogeneous case. It is thus not a discretized form of a continuous-time equation. As such, the consistency check for this model can be omitted, as the model is the solution itself, and the consistency test would thus be a test with itself.

Stability

The proof of stability requires a little more work. In order to do the analysis, eigenvalues are used to find the modes of the system, as outlined in Appendix B. It is the behavior of these modes that governs stability. We require all modes to be bounded if the solution is not to grow infinite for a finite input. To diagonalize the system, and thus find the eigenvalues, the model must first be written in terms of a single type of variables, e.g.. voltage. In case of the line models, we can write the equations in terms of only the voltage history sources e_{hk} and e_{hm} and applied external voltages v_k and v_m.

From Appendix A, (A.30), the ideal line model in voltage only is:

$$\begin{cases} e_{hk}(t) = -e_{hm}(t - \tau) + 2v_m(t - \tau) \\ e_{hm}(t) = -e_{hk}(t - \tau) + 2v_k(t - \tau) \end{cases} \tag{3.17}$$

Writing (3.17) in a matrix form suitable for eigenvalue analysis:

$$\begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^t = - \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^{t-\tau} + 2 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} v_k \\ v_m \end{bmatrix}^{t-\tau} \tag{3.18}$$

Now, to perform the eigenvalue analysis as outlined in Appendix B, we recognize that (3.18) has the form $\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t - \tau) + \mathbf{B}\mathbf{u}(t - \tau)$, where $\mathbf{u}$ is the forcing function, which we

can set to zero for the current analysis. Thus, the eigenvalues of **A** are the relevant modes, or the poles, of the system described by the model. As such, setting the external sources v_k and v_m to zero we find the eigenvalues as follows:

$$det\,(A - \lambda I) = 0 \Leftrightarrow \begin{vmatrix} -\lambda & -1 \\ -1 & -\lambda \end{vmatrix} = 0 \tag{3.19}$$

The eigenvector is then found as:

$$\lambda^2 - 1 = 0 \tag{3.20}$$

This yields the eigenvalues:

$$\lambda = \pm 1 \tag{3.21}$$

For stability, thus to insure a bounded output for a bounded input, all modes of the system, and thus eigenvalues, must be so that $|\lambda| \leq 1$. We find that they are equal to 1. This represents a special type of stability, where the solution is bound for a finite input, but does not increase nor decrease in value over time. Such a type of stability is described by Lyapunov [99]. In z-domain analysis, this would indicate that all poles are exactly on the unit circle. The ideal line model is thus marginally stable. This result is expected for the ideal model, as it has no losses of any kind and all energy in the model must thus be preserved. Further discussion of the stability conditions follows in Section 3.1.3, later in this chapter, where these results are compared to the interpolated line model.

Thus, since the ideal line model is both consistent and stable, as well as well-posed, we can conclude from the Lax-Richtmyer theorem that it is convergent.

3.1.2 The Interpolated Line Model

Derived in Appendix A.3, this line model is found from the ideal model by manually adding-in a linear interpolation function in time between two know history values in order to obtain the required history at a time that does not coincide with the simulation step. This addition of interpolation makes the solution no longer exact, and may thus introduce undesired behavior. The convergence aspect will be studied here.

Consistency

Consistency can be easily checked by comparison with the ideal model, by letting the temporal discretization (simulation time step) Δt and the spatial discretization Δx (grid size) tend to zero. The interpolated model should then become identical to the ideal model [83]. If that is

the case, the interpolated model will, for sufficiently small spatial and temporal discretization intervals, tend to the correct answer and thus be consistent.

To perform the proof, we first need the line equations in a single variable. Voltage is selected here, as it is convenient and the resulting expressions can be used for the following stability proof. Also, the line equations in terms of voltage for the ideal model were derived previously in Appendix A, (A.30). The same methodology was then used to obtain the interpolated version.

From Appendix A, (A.35), the interpolated line equations in voltage only are:

$$
\begin{cases}
\begin{aligned}
e_{hk}(t) &= 2v_m(t - \tau_{int}) - 2R\left[v_m(t - \tau_{int}) - v_m(t - \tau_{int} - \Delta t))\right] \\
&\quad - e_{hm}(t - \tau_{int}) + R\left[e_{hm}(t - \tau_{int}) - e_{hm}(t - \tau_{int} - \Delta t)\right] \\
e_{hm}(t) &= 2v_k(t - \tau_{int}) - 2R\left[v_k(t - \tau_{int}) - v_k(t - \tau_{int} - \Delta t))\right] \\
&\quad - e_{hk}(t - \tau_{int}) + R\left[e_{hk}(t - \tau_{int}) - e_{hk}(t - \tau_{int} - \Delta t)\right]
\end{aligned}
\end{cases}
\tag{3.22}
$$

Comparing (3.22) with the previous result from (3.17), it can be seen that both are similar in form, but in the interpolated case, there are extra terms that relate to an additional history, one Δt further in the past. The R term expresses where in between both known values the interpolated value is to be computed. The τ_{int} terms are the line length expressed as an integer multiple of simulation time steps Δt, rounded down.

To assess consistency, we must now let Δt to zero and subtract the interpolated equations from the ideal ones, which are an exact solution for the loss-less model under investigation. The result should be zero, thus no error in the limit for initially small temporal discretization, for the interpolated model to be consistent. From the equations, we also find that the spatial discretization is of no consequence to consistency in the 1D model. This is due to the fact that the interpolated model is based on an exact solution, modified with an interpolation that operates only in the temporal domain.

Thus, with $\Delta t \to 0$ it can be found though inspection of (A.15) and (A.16) that $\tau_{int} \to \tau$ and $R \to 0$.

Substituting these values in (3.22), it reverts to the expressions for the ideal case, as given in (3.17). Subtracting this limit result for the interpolated model from the exact expressions for the ideal line model (3.17), the remainder is found to be zero. The interpolated line model is thus consistent.

Stability

The stability analysis for the model will be performed along the same lines as in Section 3.1.1, at least, initially. Due to the multiple history depths, the equations cannot be written in a form

that is directly amenable to standard eigenvalue analysis to evaluate the modes as described in Appendix B.

Conversion to the z domain is used from that point forward. The stability is then assessed by computing the system transfer function matrix $\mathbf{H}(z)$. The eigenvalues (poles) of this expression, without a forcing function, are used to assess stability.

However, the resulting closed-form polynomial expressions in z have variables as exponents, making it impossible to find closed-form roots in general. Specific cases can be evaluated, though. Numerical examples are given to illustrate the stability trends, and show that the interpolated model tends to the ideal case as the number of time steps per line increases, and the interpolation error becomes less and less.

A less rigorous stability proof is then shown, based on verification of the long-term outcome of the system (homogeneous part) as $t \to \infty$. This proof is not as rigorous, as it assumes a long-term outcome is possible.

Thus, writing the interpolated line model equation in voltage only (3.22) in a matrix form suitable for eigenvalue analysis:

$$
\begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^t = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \left\{ -(1-R) \begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^{t-\tau_{int}} - R \begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^{t-\tau_{int}-\Delta t} \right.
$$
$$
\left. +2(1-R) \begin{bmatrix} v_k \\ v_m \end{bmatrix}^{t-\tau_{int}} + 2R \begin{bmatrix} v_k \\ v_m \end{bmatrix}^{t-\tau_{int}-\Delta t} \right\} \tag{3.23}
$$

We see that (3.23) has the form:

$$
\mathbf{x}(t) = \mathbf{A}_1 \mathbf{x}(t - \tau_{int}) + \mathbf{A}_2 \mathbf{x}(t - \tau_{int} - \Delta t) + \mathbf{B}_1 \mathbf{u}(t - \tau_{int}) + \mathbf{B}_2 \mathbf{u}(t - \tau_{int} - \Delta t) \tag{3.24}
$$

The various $\mathbf{A}$ and $\mathbf{B}$ terms can be readily obtained from (3.23) as:

$$
\mathbf{A}_1 = \begin{bmatrix} 0 & -(1-R) \\ -(1-R) & 0 \end{bmatrix} \quad \mathbf{A}_2 = \begin{bmatrix} 0 & -R \\ -R & 0 \end{bmatrix}
$$
$$
\mathbf{B}_1 = \begin{bmatrix} 0 & 2(1-R) \\ 2(1-R) & 0 \end{bmatrix} \quad \mathbf{B}_2 = \begin{bmatrix} 0 & 2R \\ 2R & 0 \end{bmatrix} \tag{3.25}
$$

$\mathbf{u}$ is the forcing function, which is not required for the current analysis. This brings the relevant part of the equation entirely in the same terms of $\mathbf{x}$.

The issue is that the various RHS terms of $\mathbf{x}$ are in different time scales, making it impossible to find the system modes through regular eigenvalue analysis as outlined in Appendix

B. This is because the **A** matrices for each different time step cannot be combined in a single expression that can be diagonalized to directly evaluate the system behavior.

To overcome this problem, the system is converted to z-domain using the following conversions of the indices to make them consistent with the implied concept of time in this discrete domain:

- $\tau_{int} \rightarrow N$ Where N is the length of the line as an integer number of simulation time steps, rounded down, per definition of (A.15).

- $\Delta t \rightarrow 1$ One time step in z-domain is a single delay.

- $t \rightarrow n$ The total simulation time is implicit in z-domain. It is an integer number of simulation time steps.

Using this notation, and setting the excitation to unity, the skeletal equation of (3.24) becomes:

$$\mathbf{x}(n) = \mathbf{A}_1\mathbf{x}(n - N) + \mathbf{A}_2\mathbf{x}(n - (N + 1)) + \mathbf{u}(n) \tag{3.26}$$

Converting (3.26) to z-domain:

$$\mathbf{x}(z) = \mathbf{A}_1 z^{-N}\mathbf{x}(z) + \mathbf{A}_2 z^{-(N+1)}\mathbf{x}(z) + \mathbf{u}(z) \tag{3.27}$$

Grouping terms we find:

$$\mathbf{x}(z) = \left[\mathbf{I} - \mathbf{A}_1 z^{-N} - \mathbf{A}_2 z^{-(N+1)}\right]^{-1}\mathbf{u}(z) \tag{3.28}$$

The system behavior is now described by a single matrix in Z-domain. This form is thus suitable for eigenvalue analysis. Since the system is in z-domain, the resulting eigenvalues are the poles of the transfer function in z-domain, and must be within the unit circle for stability [63].

Substituting the **A** terms back into the above equation we find:

$$\mathbf{x}(z) = \begin{bmatrix} 1 & Rz^{-(N+1)} + (1 - R)z^{-N} \\ Rz^{-(N+1)} + (1 - R)z^{-N} & 1 \end{bmatrix}^{-1}\mathbf{u}(z) \tag{3.29}$$

Computing the eigenvalues using Maxima [53], we find:

$$\lambda_1 = \frac{1}{Rz^{-(N+1)} + (1 - R)z^{-N} + 1} \quad \lambda_2 = \frac{1}{-Rz^{-(N+1)} - (1 - R)z^{-N} + 1} \tag{3.30}$$

To evaluate stability, we inspect the above expressions. It becomes clear that, to ensure both eigenvalues are ≤ 1, we need to find the poles (roots) of polynomial-like expressions

with exponents in N, where N is a positive integer and a variable. This form is not a standard polynomial, and the roots cannot be found directly. Thus, a closed-form expression cannot be found in the generic case. However, since N is constant for any given case, evaluation in the specific is possible. It is thus possible to exhaustively work through each value of N, as practical values will be limited to a couple hundred.

When evaluating (3.30) numerically, the poles of these expressions in z-domain show a number of definite trends, as illustrated graphically below for a number of cases for varying R and N.

From the plots (Figures 3.2 and 3.3), it can be seen that the poles follow a number of distinct trends. As the line becomes longer (N increases), the poles tend closer and closer to the unit circle. Looking at (3.30), it can indeed be seen that, as $N \to \infty$, the expressions for the eigenvalues tends to:

$$\lim_{N\to\infty} \lambda_1 = \lim_{N\to\infty} \frac{1}{z^{-N} + 1} \quad \lim_{N\to\infty} \lambda_2 = \lim_{N\to\infty} \frac{1}{-z^{-N} + 1} \tag{3.31}$$

Thus, all poles, as $N \to \infty$, tend toward the unit circle, but never exceed it. Also, the influence of the interpolation factor R decreases. This is due to the fact that the interpolation only acts on one segment of the line. The effect of interpolation is thus reduced with increasing length.

A second observation is that the interpolation factor R affects the symmetry of the pole placement in the z plane. For both small and large values of R (recalling that $0 \le R < 1$) the effect of the interpolation is small. This is explained by the fact that, for such values of R, the point to be interpolated is very close to a known value. Thus, the interpolation error is significantly reduced. For $R = 0.5$, however, the interpolated point is in the middle of the interval, and thus at the highest uncertainty. The error is thus larger.

The effects of N and R will be studied in more detail in Section 3.2, later in this chapter.

From the trends exhibited by the poles for various parameters, it can thus be inferred that the interpolated line model is stable. As the model is also consistent, and the governing continuous-time equations well-posed, from the Lax-Richtmyer theorem, the interpolated line model is convergent.

An additional observation is that the model tends toward Lunyapov stability [99] as it tends closer to the ideal line model. When the interpolation is of significant influence, there are clear errors associated with the model and the poles are further away from the unit circle. Thus, the model may have dissipative properties which make it more stable. This dissipation and dispersion are studied in Section 3.2.

(a) $R = 0.1$, $N = 1$ (b) $R = 0.1$, $N = 10$ (c) $R = 0.1$, $N = 100$

(d) $R = 0.5$, $N = 1$ (e) $R = 0.5$, $N = 10$ (f) $R = 0.5$, $N = 100$

(g) $R = 0.9$, $N = 1$ (h) $R = 0.9$, $N = 10$ (i) $R = 0.9$, $N = 100$
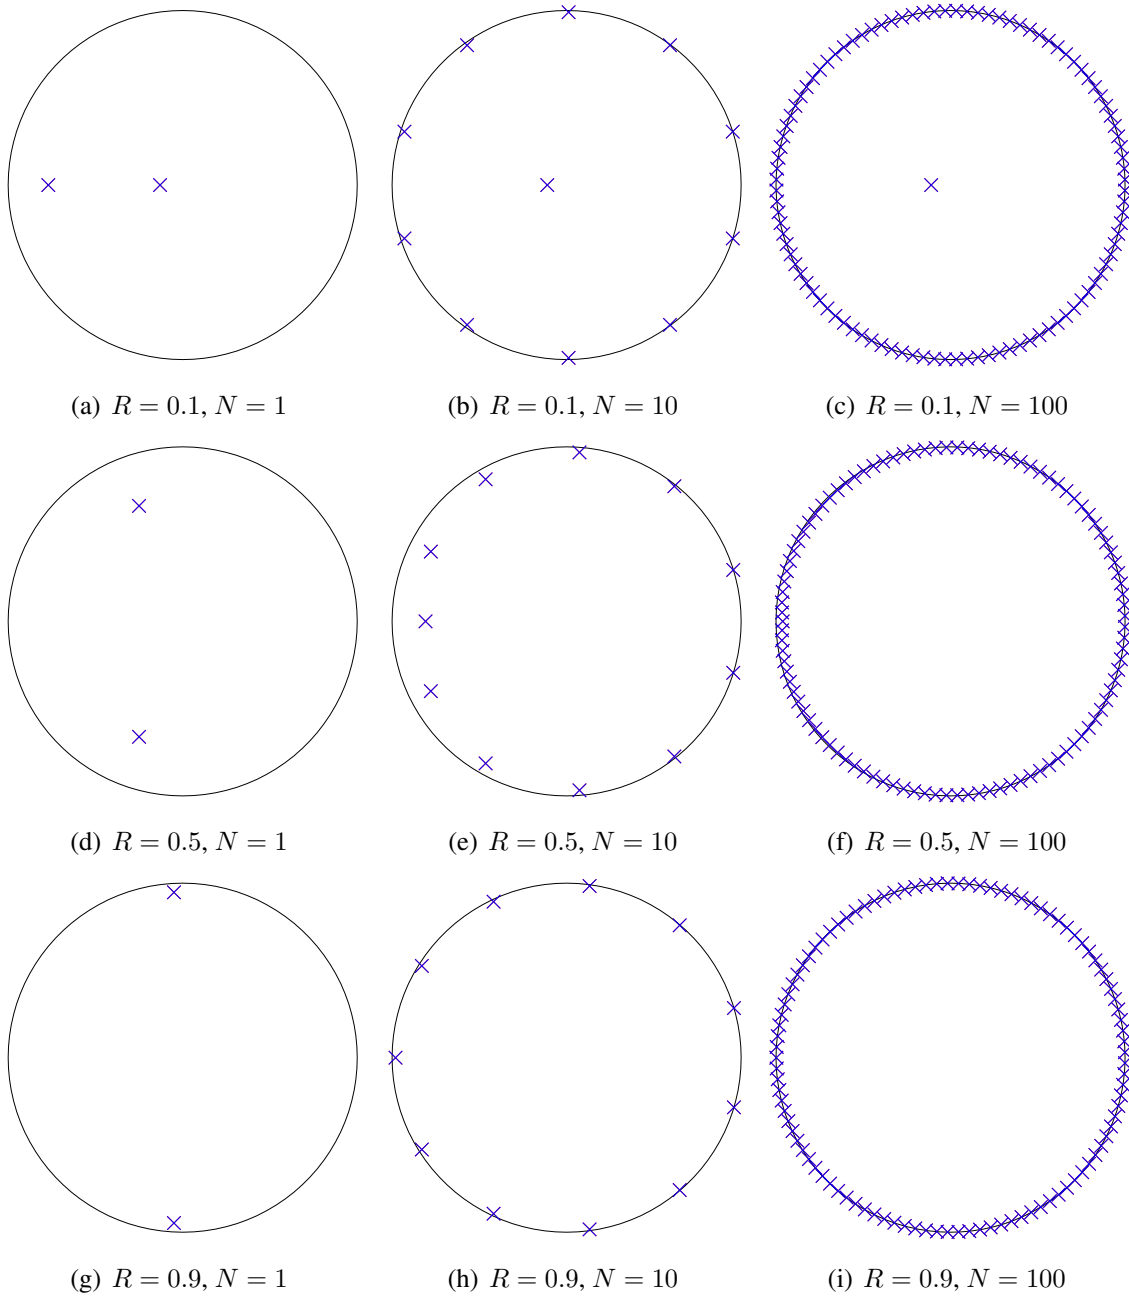
Figure 3.2: Pole Plots for Various Combinations of R and N for λ_1

(a) $R = 0.1, N = 1$ (b) $R = 0.1, N = 10$ (c) $R = 0.1, N = 100$

(d) $R = 0.5, N = 1$ (e) $R = 0.5, N = 10$ (f) $R = 0.5, N = 100$

(g) $R = 0.9, N = 1$ (h) $R = 0.9, N = 10$ (i) $R = 0.9, N = 100$
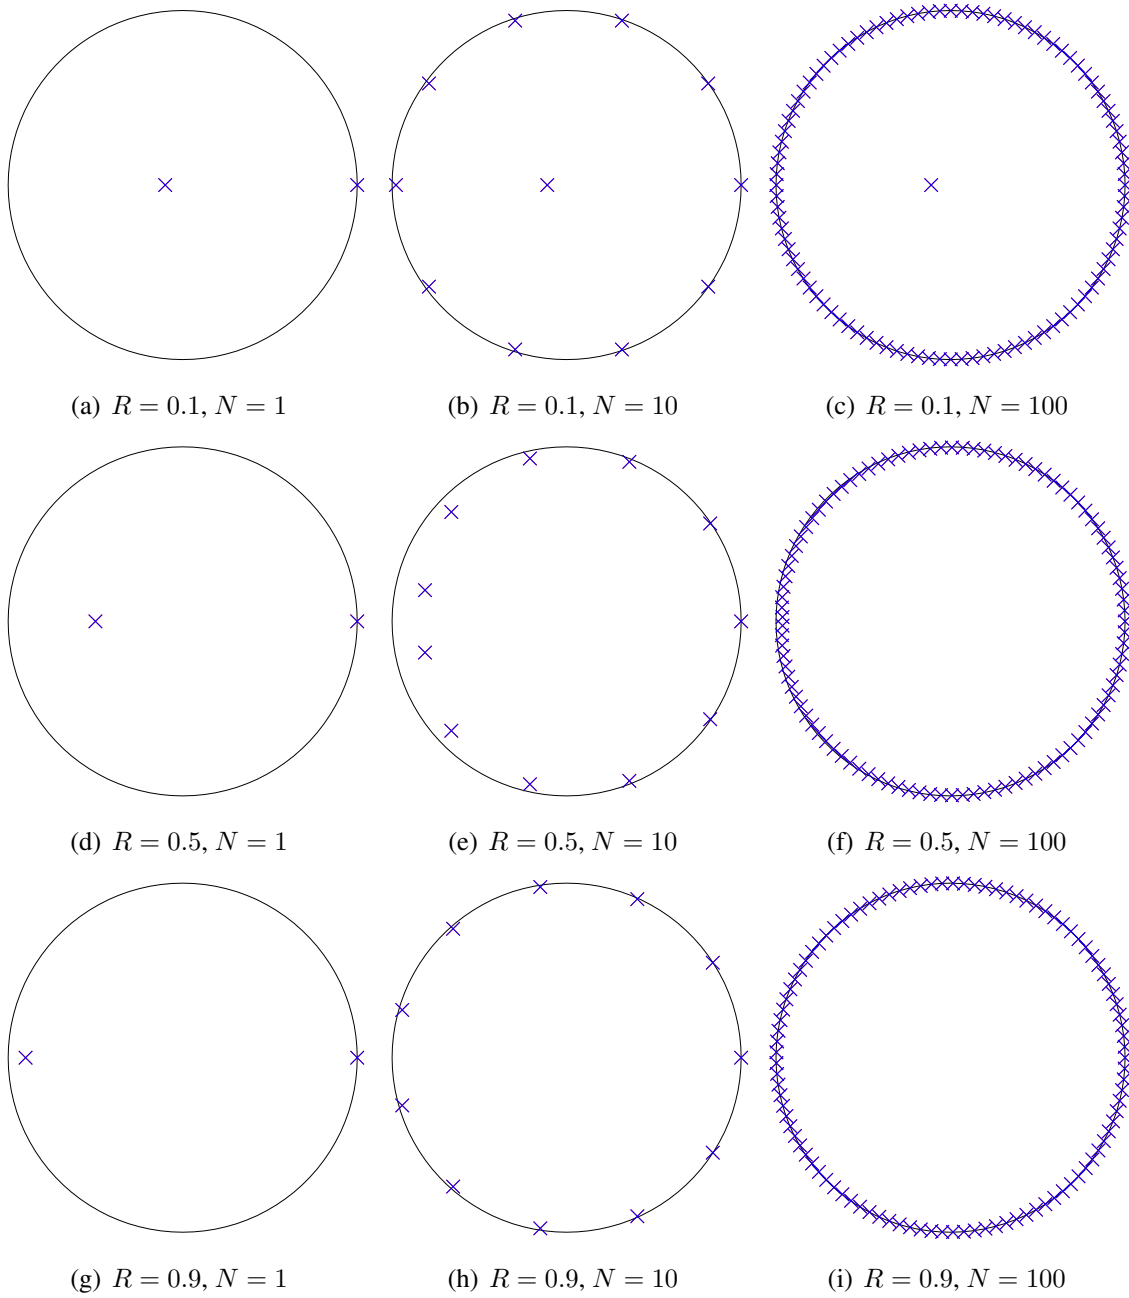
Figure 3.3: Pole Plots for Various Combinations of R and N for λ_2

A different approach to stability

Although the above proof is mathematically rigorous, the resulting eigenvalues cannot be evaluated in the general case due to the variable in the exponent of the polynomials. Evaluation in the specific case is, however, straight-forward.

To strengthen the inferred conclusion of the above section regarding the stability of the interpolated line model, a different approach will now be used to show stability. The proof is based on finding the long-term solution ($t \to \infty$) of the interpolated line model. This particular solution must be zero or bound, in case there are no losses and no input or a bound input. This proof is not as rigorous, as a long-term solution is assumed to exist.

Thus, starting from (3.24):

$$\mathbf{x}(t) = \mathbf{A}_1\mathbf{x}(t - \tau_{int}) + \mathbf{A}_2\mathbf{x}(t - \tau_{int} - \Delta t) + \mathbf{B}_1\mathbf{u}(t - \tau_{int}) + \mathbf{B}_2\mathbf{u}(t - \tau_{int} - \Delta t) \quad (3.32)$$

In the long run, when $t \to \infty$, $(t - \tau_{int}) \approx (t - \tau_{int} - \Delta t)$. Using this simplification, we find:

$$\mathbf{x}(t) \approx (\mathbf{A}_1 + \mathbf{A}_2)\mathbf{x}(t - \tau_{int}) + (\mathbf{B}_1 + \mathbf{B}_2)\mathbf{u}(t - \tau_{int}) \quad (3.33)$$

Writing in matrix form, and substituting:

$$\begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^t \approx - \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} e_{hk} \\ e_{hm} \end{bmatrix}^{t-\tau_{int}} + 2 \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} v_k \\ v_m \end{bmatrix}^{t-\tau_{int}} \quad (3.34)$$

This expression is the same as the ideal line model (3.18). Setting the excitation to zero, it was shown to be Lyapunov stable. Thus, the particular solution is bound. From this, we can conclude that the interpolated line model is also Lyapunov stable, when the simulation time is long enough for the particular solution to be reached.

3.1.3 A Physical Interpretation of Marginal Stability in the Line Models

From the preceding analysis, it is found that both the ideal and interpolated line models have eigenvalues/poles in the z-domain that are unity/on the z-domain unit circle. Thus, the line models are/tend to marginally stable. This is a curious stability condition where the outcome may be oscillatory, but is still bounded. Or, in different words, the energy in the model is conserved, but the response may be infinite in duration.

For ideal transmission lines, this is the expected behavior. If we consider an ideal line terminated in a short on both ends and we inject a pulse, this pulse will keep bouncing back and forward forever, as there are no losses and no dispersion. This means that the model, for a

finite input, produces an infinite response (in duration) which is still bound (in the total energy in the system). Thus, although the response is oscillatory, the system is stable in the sense that it does not gain energy. The interpolated model has some losses, to be investigated in the following section, and thus only tends to this infinite response behavior, as some energy is lost over time. This can be confirmed from the z-domain analysis above for the interpolated model, where the poles are in fact only tending to the unit circle. The interpolated model is thus more stable than the ideal model, due to these losses.

These types of stability are describe by Lyapunov [99]. It can be summarized as that for a dynamic system that starts out near an equilibrium point x_e will stay near this equilibrium point forever. If that condition is true, then the system is Lyapunov stable. If a solution starts out near x_e and converges closer toward it, the solution is asymptotically stable. Exponential stability guarantees a minimal rate of decay, and thus an estimate of how fast the solution will converge.

3.2 Dissipation and Dispersion

In a kinder world, the performance of the loss-less, interpolated line model would be identical to the ideal one. In practice, this is not the case and some frequency-dependent behavior occurs due to the interpolation process. In fact, the interpolation results in a low-pass filter effect, as will become apparent in this section.

The errors incurred through interpolation are frequency-dependent losses which, in frequency domain, result in a magnitude and phase error. The magnitude error is the dissipation error, while the phase error gives the dispersion error. Or, in different words, the magnitude loss and phase shift for each frequency as it passes through the model.

As opposed to the previous section, where only the internal behavior of the models was considered, here the boundary conditions play a role as well. For example, reflections on the terminals due to a load mis-match will cause part of the energy to be reflected back into the model, where it will be affected by the non-ideal behavior once more. Thus, a direct comparison of the ideal and interpolated model by only evaluating their transfer functions will not provide all the required information.

Since, in this section, we must evaluate the models in the presence of boundary conditions to measure their response, we must first define both the used boundary conditions, as well as the performance metric used in the evaluation.

To facilitate the analysis, a version of the line models converted to frequency domain was used, since it is the magnitude and phase response of the model with respect to frequency that is of interest. The conversion from time to frequency domain is illustrated in Appendix A,

Sections A.2.3 and A.3.3.
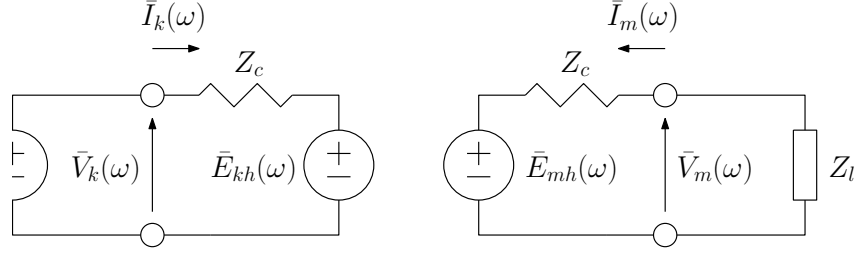
Figure 3.4: Loss-Less Line Model with Boundary Conditions in Frequency Domain

As illustrated in Figure 3.4, the boundary conditions are, on the k-port of the model, an ideal voltage source, and on the m-port, a variable impedance. This impedance will be adjusted to highlight both typical and border cases.

The performance metric is a normalised, relative measure that compares the performance of the interpolated model with to the ideal one. Since the latter is an exact solution, it is chosen as the benchmark. As the used models are expressed in frequency domain, the ratio is one of phasors:

$$\overline{\text{Error}} = \frac{\overline{\text{Interpolated model}}}{\overline{\text{Ideal model}}} - 1 \tag{3.35}$$

Thus, the evaluation is simply a calculation of the line model in frequency domain, with identical boundary conditions for both the ideal and interpolated model, and then dividing the respective output phasors to find the error ratio (3.35) for each frequency. Plotting the results in magnitude and phase will yield the dissipation and dispersion error for the case under investigation.

Let us now derive a general error expression for a line terminated in a complex load, and driven from an ideal source. Two such expressions must be derived, one in voltage and one in current, to accommodate shorted and open-circuit boundary conditions. From Appendix A, the ideal and interpolated line model in frequency domain are:

$$\text{Ideal} \qquad \begin{cases} \overline{V}_k - Z_c\overline{I}_k = \left(\overline{V}_m + Z_c\overline{I}_m\right)e^{-j\omega\tau} \\ \overline{V}_m - Z_c\overline{I}_m = \left(\overline{V}_k + Z_c\overline{I}_k\right)e^{-j\omega\tau} \end{cases} \tag{3.36a}$$

$$\text{Interpolated} \quad \begin{cases} \overline{V}_k - Z_c\overline{I}_k = \left(\overline{V}_m + Z_c\overline{I}_m\right)e^{-j\omega\tau_{int}}\left[1 + R\left(e^{-j\omega\Delta t} - 1\right)\right] \\ \overline{V}_m - Z_c\overline{I}_m = \left(\overline{V}_k + Z_c\overline{I}_k\right)e^{-j\omega\tau_{int}}\left[1 + R\left(e^{-j\omega\Delta t} - 1\right)\right] \end{cases} \tag{3.36b}$$

We observe that both are of the same general form:

$$\begin{cases} \overline{V}_k - Z_c\overline{I}_k = \left(\overline{V}_m + Z_c\overline{I}_m\right)A \\ \overline{V}_m - Z_c\overline{I}_m = \left(\overline{V}_k + Z_c\overline{I}_k\right)A \end{cases} \tag{3.37}$$

which is a general expression of an input subjected to a complex phase shift A. This expression is now used to find a general solution for the model, including boundary conditions as shown in Figure 3.4. The appropriate terms will then be substituted for A after the solution is found.

The boundary conditions for voltage are thus:

$$\overline{V}_k = \overline{V}_s \ \overline{I}_m = -\frac{\overline{V}_m}{Z_l} \tag{3.38}$$

Using these boundary conditions, the system reverts to two equations in two unknowns, and can thus be solved. Substituting:

$$\begin{cases} \overline{V}_s - Z_c\overline{I}_k = \left(\overline{V}_m - Z_c\frac{\overline{V}_m}{Z_l}\right) A \\ \overline{V}_m - Z_c\frac{\overline{V}_m}{Z_l} = \left(\overline{V}_s + Z_c\overline{I}_k\right) A \end{cases} \tag{3.39}$$

Solving for voltage, we find:

$$\frac{\overline{V}_m}{\overline{V}_s} = \frac{2Z_l A}{\left(Z_l - Z_c\right) A^2 + Z_l + Z_c} \tag{3.40}$$

Repeating the above for current, using a current injection $\overline{I}_s$ instead of a voltage injection $\overline{V}_s$, the following boundary conditions are applied:

$$\overline{I}_k = \overline{I}_s \ \overline{V}_m = -\overline{I}_m Z_l \tag{3.41}$$

Using these boundary conditions, the system reverts to two equations in two unknowns, and can thus be solved. Substituting:

$$\begin{cases} \overline{V}_k - Z_c\overline{I}_s = \left(Z_c\overline{I}_m - Z_l\overline{I}_m\right) A \\ -Z_l\overline{I}_m - Z_c\overline{I}_m = \left(\overline{V}_k + Z_c\overline{I}_s\right) A \end{cases} \tag{3.42}$$

Solving for current, we find:

$$\frac{\overline{I}_m}{\overline{I}_s} = \frac{2Z_c A}{\left(Z_l - Z_c\right) A^2 - Z_l - Z_c} \tag{3.43}$$

Using the expressions for voltage (3.40) and current (3.43), the appropriate dissipation and dispersion errors, found from the ratio between the interpolated and ideal models, as defined in (3.35), can be computed. In order to obtain these voltage and current error ratios between the interpolated and ideal models, we recognize that, if both models use the same injection,

the respective $\overline{V}_s$ and $\overline{I}_s$ terms are the same and the ratios are found by simple division of the interpolated by the ideal voltage or current expression, respectively:

$$\overline{error}_V = \frac{\overline{V}_{m\,interpolated}\overline{V}_{s\,ideal}}{\overline{V}_{m\,ideal}\overline{V}_{s\,interpolated}} - 1 = \frac{A_{interpolated}\left[A_{ideal}^2 + \frac{K+1}{K-1}\right]}{A_{ideal}\left[A_{interpolated}^2 + \frac{K+1}{K-1}\right]} - 1 \qquad (3.44a)$$

$$\overline{error}_I = \frac{\overline{I}_{m\,interpolated}\overline{I}_{s\,ideal}}{\overline{I}_{m\,ideal}\overline{I}_{s\,interpolated}} - 1 = \frac{A_{interpolated}\left[A_{ideal}^2 - \frac{K+1}{K-1}\right]}{A_{ideal}\left[A_{interpolated}^2 - \frac{K+1}{K-1}\right]} - 1 \qquad (3.44b)$$

where $K = \frac{Z_l}{Z_c}$.

The A terms for the models are given by:

$$A_{\text{Ideal}} = e^{-j\omega\tau} \qquad (3.45a)$$

$$A_{\text{Interpolated}} = e^{-j\omega\tau_{int}} + Re^{-j\omega\tau_{int}}\left(e^{-j\omega\Delta t} - 1\right) \qquad (3.45b)$$

where the τ_{int} factor can be substituted as $\tau_{int} = \tau - R\Delta t$ (Appendix A.3.2) to write the expressions in a single variable of τ:

$$A_{\text{Ideal}} = e^{-j\omega\tau} \qquad (3.46a)$$

$$A_{\text{Interpolated}} = e^{-j\omega\tau}\left[e^{j\omega R\Delta t} + Re^{j\omega R\Delta t}\left(e^{-j\omega\Delta t} - 1\right)\right] \qquad (3.46b)$$

Substituting these A back with the appropriate terms for the ideal and interpolated line model:

$$\overline{error}_V = \frac{1 + e^{2j\omega\tau}\frac{K+1}{K-1}}{e^{j\omega R\Delta t}\left[1 + R\left(e^{-j\omega\Delta t} - 1\right) + \frac{e^{2j\omega(\tau - R\Delta t)}}{1 + R(e^{-j\omega\Delta t} - 1)}\frac{K+1}{K-1}\right]} - 1 \qquad (3.47a)$$

$$\overline{error}_I = \frac{1 - e^{2j\omega\tau}\frac{K+1}{K-1}}{e^{j\omega R\Delta t}\left[1 + R\left(e^{-j\omega\Delta t} - 1\right) - \frac{e^{2j\omega(\tau - R\Delta t)}}{1 + R(e^{-j\omega\Delta t} - 1)}\frac{K+1}{K-1}\right]} - 1 \qquad (3.47b)$$

The above are the error expressions, in voltage and current, for the interpolated line model, compared to the ideal line model, with boundary conditions applied.

In order to analyze the behavior of these expressions, the variables are now expressed in dimensionless ratios to permit a more general evaluation of the results.

R Interpolation factor $R = \left(\frac{\tau}{\Delta t} - \lfloor\frac{\tau}{\Delta t}\rfloor\right)$. Relative to one Δt and $0 \leq R < 1$

ρ Reflection coefficient $\rho = \left(\frac{K-1}{K+1}\right)$, incorporates both Z_l and Z_c through $K = \frac{Z_l}{Z_c}$

L Length of the line in expressed as # of simulation time steps $L = \left(\frac{\tau}{\Delta t}\right)$ (can be non-integer)

Using the dimensionless variables R, ρ, and L, we will now analyze the expressions (3.47) in terms of the above parameters for various cases.

3.2.1 Influence of the Interpolation Factor

The interpolation factor R describes where in the last history interval the unknown value has to be computed. Intuitively, the closer this value is to the known values on either side of the interval, or the less difference between these known boundary values, the better the accuracy of the estimate will be. Half-way in the interval, $R = 0.5$, should then be the worst. Thus, the location of the interpolated value in the interval, as well as the simulation step size compared to the highest frequency of interest, will have an influence on the accuracy of the simulation.

In order to isolate the effect of the interpolation factor from the line length, and avoid influences due to reflections from the model terminals, matched conditions, $\rho = 0$, were chosen. Using this condition, the line length τ drops out of the equations as the terms that incorporate it are forced to zero. This is discussed in more detail in the next section, where the influence of the reflection coefficient ρ is studied. Finally, the simulation time step was evaluated at Δt of 1s, 0.5, and 0.1s.

Using these conditions, the error expressions (3.47) are reduced to:

$$\overline{error}_V = \frac{1}{e^{j\omega R \Delta t}\left[1 + R\left(e^{-j\omega \Delta t} - 1\right)\right]} - 1 \qquad (3.48a)$$

$$\overline{error}_I = \frac{1}{e^{j\omega R \Delta t}\left[1 + R\left(e^{-j\omega \Delta t} - 1\right)\right]} - 1 \qquad (3.48b)$$

Plotting these equations for varying R, Δt and ω, in voltage error only as both voltage error and current error expressions are identical, Figures 3.5-3.10 result. It has to be noted that by varying R, the line length is also changing, but this had no visible effect in the error plots as the line length does not feature in the governing equations. Also, in the plots, only values of the interpolation factor up the $R = 0.5$ were shown due to this complete symmetry of the results around this value, e.g.: the trace for $R = 0.4$ was identical to the trace for $R = 0.6$. This omission reduced visual clutter. In addition, this symmetry in the error also coincides with the intuitive assessment of the influence of the location in the interpolation interval on the interpolation error, given earlier in this section.

From the plots and error equations, we can conclude, under matched conditions ($\rho = 0$):

- The position in the interpolated interval has an influence on the interpolation error.

- The worst-case interpolation error occurs at $R = 0.5$, half-way in the interpolated interval.

(a) $\overline{error}_V$ and $\overline{error}_I$ Magnitude

(b) $\overline{error}_V$ and $\overline{error}_I$ Magnitude (detail)

Figure 3.5: Voltage and Current Error Magnitude Plots for Various values of R, $\Delta t = 1$s

(a) $\overline{error}_V$ and $\overline{error}_I$ Phase

(b) $\overline{error}_V$ and $\overline{error}_I$ Phase (detail)

Figure 3.6: Voltage and Current Error Phase Plots for Various values of R, $\Delta t = 1$s
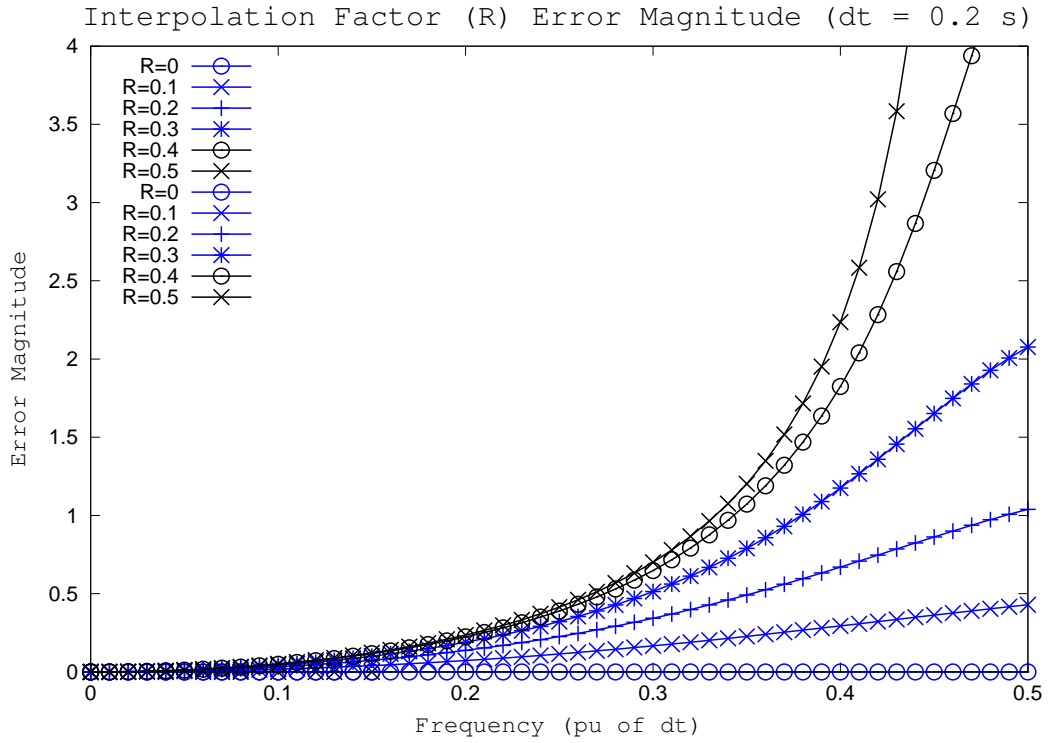
(a) $\overline{error}_V$ and $\overline{error}_I$ Magnitude
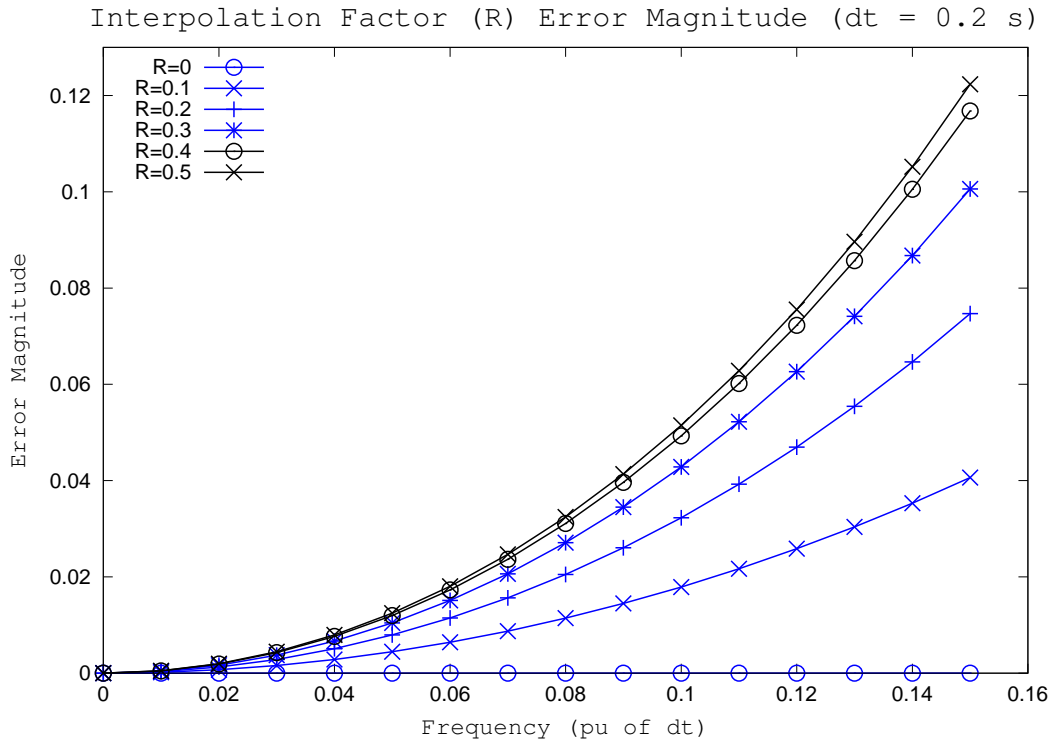
(b) $\overline{error}_V$ and $\overline{error}_I$ Magnitude (detail)

Figure 3.7: Voltage and Current Error Magnitude Plots for Various values of R, $\Delta t = 0.2$s
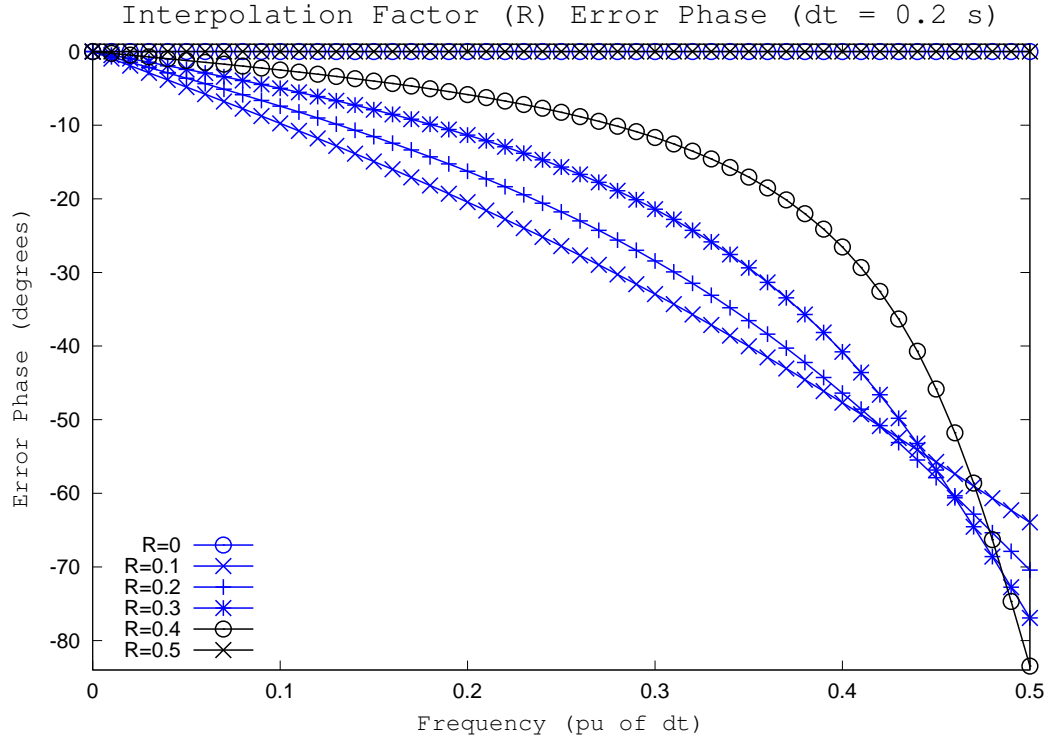
(a) $\overline{error}_V$ and $\overline{error}_I$ Phase
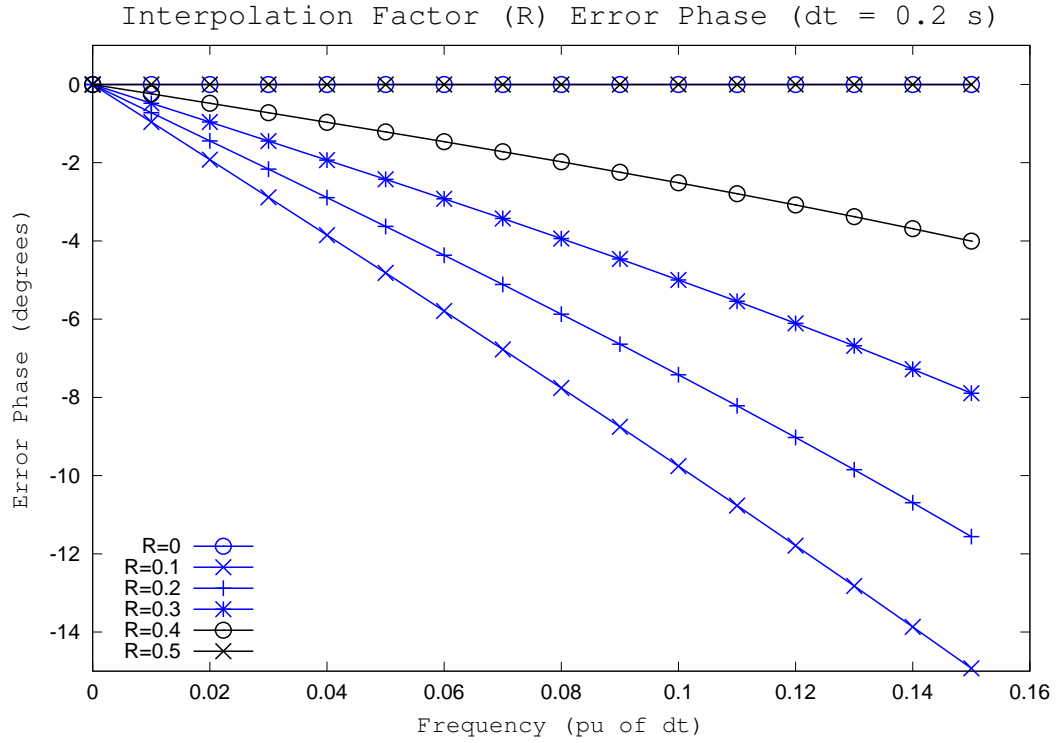
(b) $\overline{error}_V$ and $\overline{error}_I$ Phase (detail)

Figure 3.8: Voltage and Current Error Phase Plots for Various values of R, $\Delta t = 0.2$s
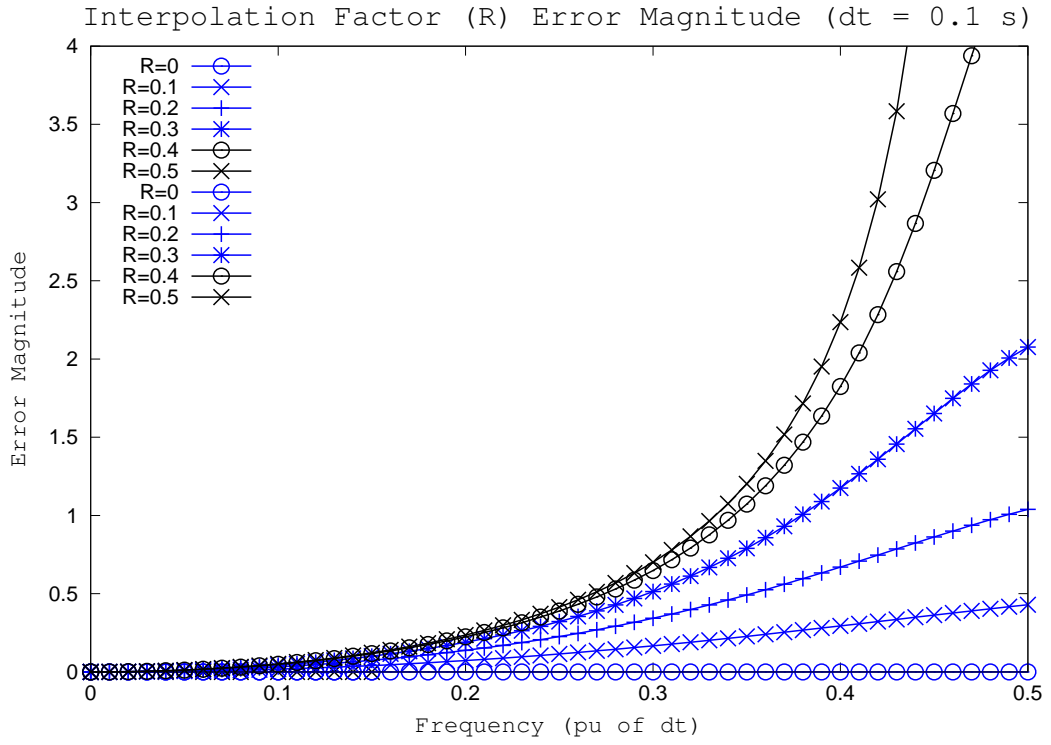
(a) $\overline{error}_V$ and $\overline{error}_I$ Magnitude
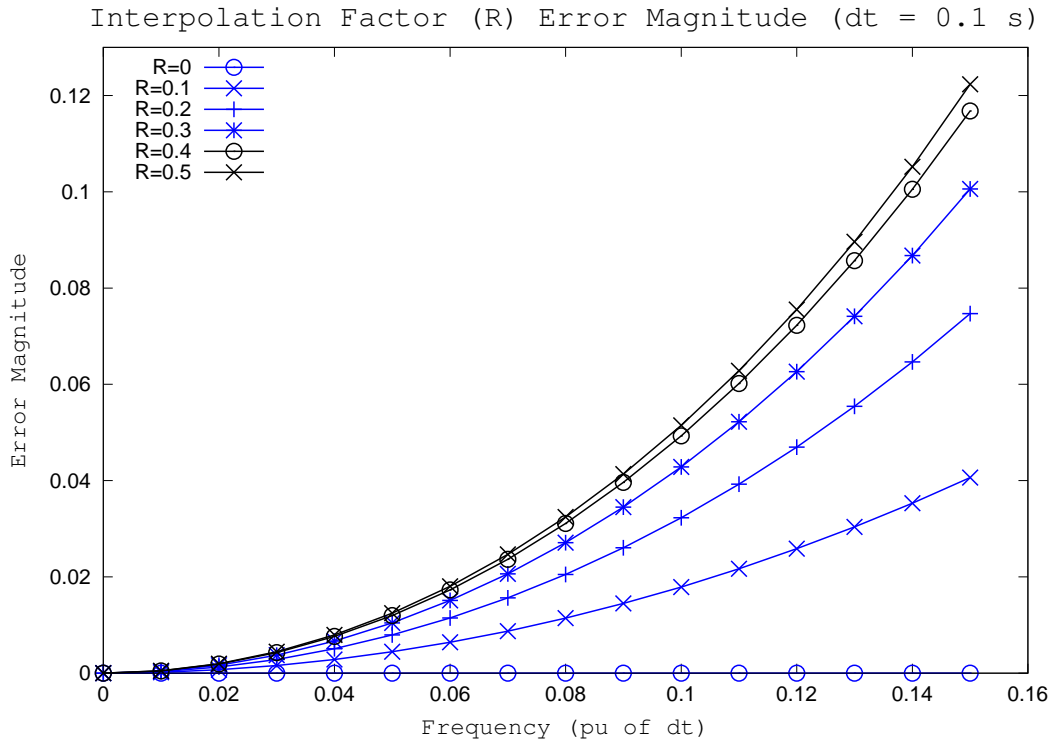
(b) $\overline{error}_V$ and $\overline{error}_I$ Magnitude (detail)

Figure 3.9: Voltage and Current Error Magnitude Plots for Various values of R, $\Delta t = 0.1$s
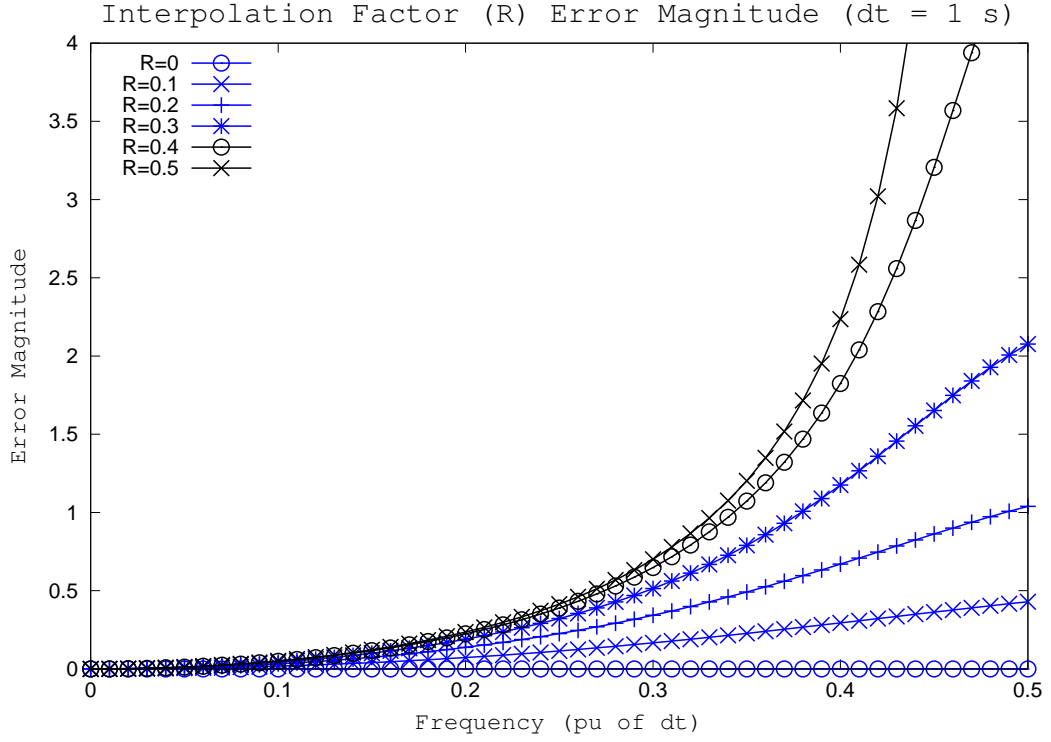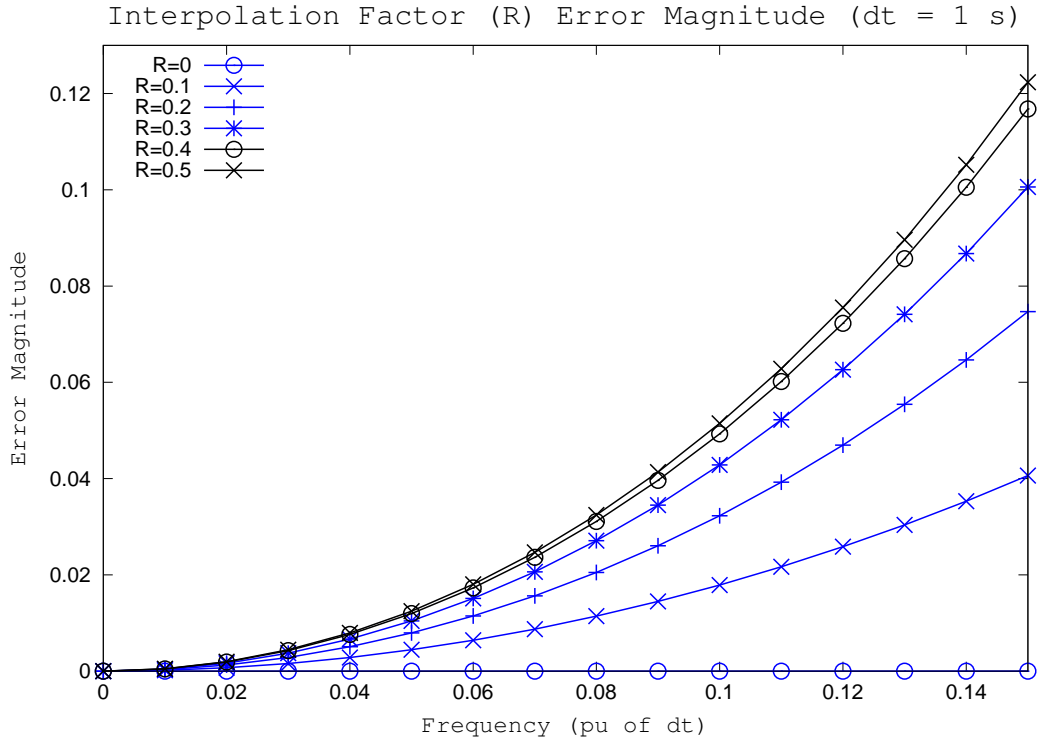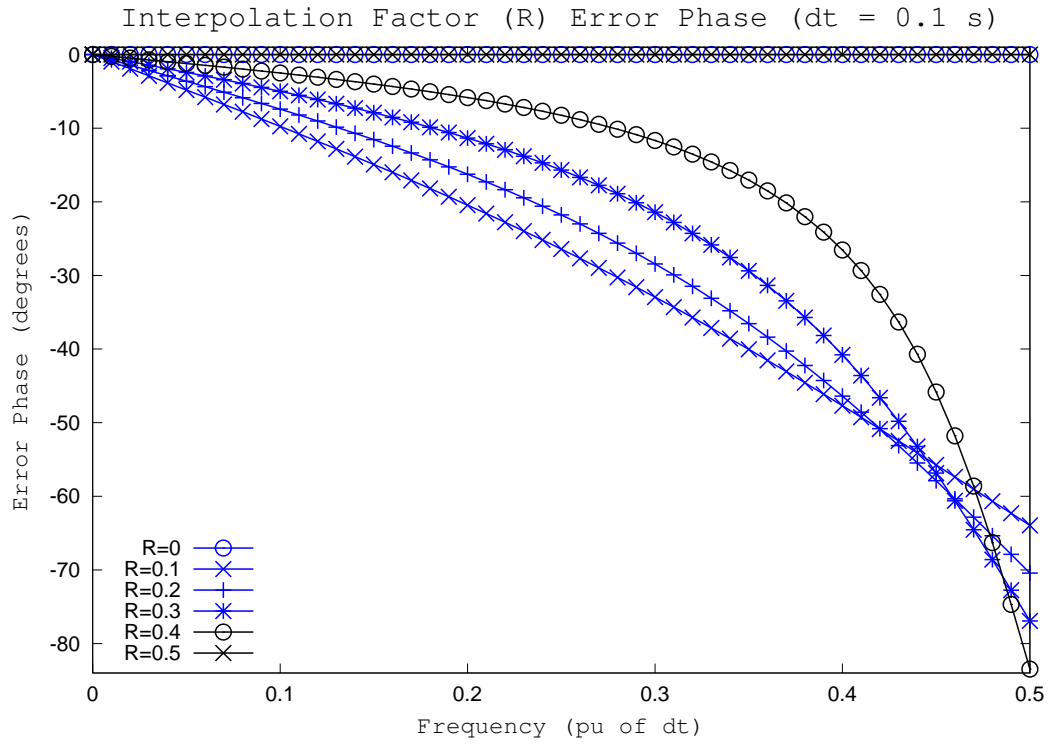
(a) $\overline{error}_V$ and $\overline{error}_I$ Phase
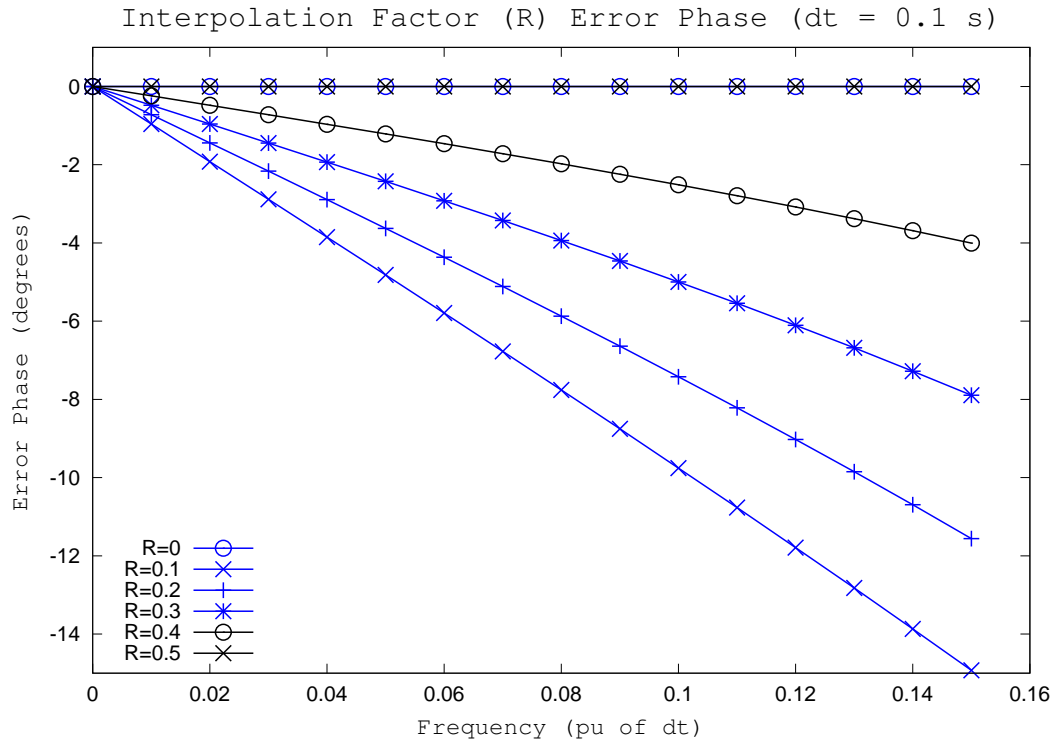
(b) $\overline{error}_V$ and $\overline{error}_I$ Phase (detail)

Figure 3.10: Voltage and Current Error Phase Plots for Various values of R, $\Delta t = 0.1$s

- The phase error is zero for $R = 0$, no interpolation, and $R = 0.5$.

- The length of the line, under matched conditions, has no impact on the interpolation error. The error does depend on the line length under non-matched conditions.

- Interpolation results in a reduced usable bandwidth due to rapidly increasing errors.

- The shape of the error graphs remains unchanged with different simulation time steps Δt.

- There is a dependency on the simulation time step Δt. The smaller the time step, the higher the acceptable error cutoff frequency, thus the higher the usable bandwidth in absolute terms. It remains, however, the same relative fraction of the total simulation bandwidth.

3.2.2 Influence of the Reflection Coefficient and Line Length

The matching conditions, given by the reflection coefficient ρ, have a significant influence on the model error due to a cumulative effect as reflected energy passes through the model multiple times. Also the line length, especially for boundary conditions close to shorted and open circuit conditions, results in difference in the error behavior, as the line becomes a resonator. In order to evaluate these effects, the worst-case condition for the interpolation factor $R = 0.5$ was chosen. Again, for ease of interpretation of the plots, the simulation time step was chosen $\Delta t = 1s$.

Choosing a fixed value for R does lock the possible line lengths τ to multiples of half the simulation time step Δt, as the expressions $R = \frac{\tau}{\Delta t} - \lfloor \frac{\tau}{\Delta t} \rfloor$ and $L = \frac{\tau}{\Delta t}$ must be satisfied. Thus, the allowable line lengths L for the plots are $L = 1.5, 2.5, 3.5, ...$

The reflection coefficient ρ is bound between $-1 \leq \rho \leq 1$, thus between shorted circuit ($\rho = -1$), over matched conditions ($\rho = 0$), to open circuit conditions ($\rho = 1$) [46, 71]. As such, a limit case will be encountered while evaluating the error expressions for matched conditions, as the expressions require the evaluation of $\frac{1}{\rho}$. However, this case is easily resolved by inspecting (3.40) and (3.43), and observing that, under matched conditions where $Z_l = Z_c$, only the respective A terms remain, which are independent of the line and load impedances. Computing the error ratios from these terms, and comparing with (3.47), it is thus found that, under matched conditions: $\frac{1}{\rho} = \frac{K+1}{K-1} = 0$. Thus, the error equations reduce to (3.48) in this case.

In general, however, the equations (3.47) cannot be further reduced. Thus, a mutual dependency of the line length, given by τ, and the reflection coefficient, given by ρ exists, even with

with a fixed value for R and Δt. Plotting for varying line lengths L and matching conditions ρ, Figures 3.11 - 3.18 result.

The plots for voltage and current are nearly identical. The difference is the location of the various resonances, e.g.: the dip for $\rho = 1$ for the voltage error is identical to the dip for $\rho = -1$ for the current error. This is expected behavior for transmission lines [7, 46].

Also, the various resonances in the error plots correspond with the resonance frequencies of the line sections. Indeed, a transmission line terminated in an open or short circuit (or conditions tending to these cases) will behave as a resonator at the nominal $\frac{\lambda}{4}$ and $\frac{\lambda}{2}$ frequencies. At such a resonance, the magnitude of the input versus output tends to limit conditions in the lossless case, and small discrepancies, e.g.: due to model losses, quickly result in large relative errors. Still, the center frequencies at which the resonances occur is not affected by the models inaccuracies, as will be demonstrated in Section 3.3.

From these figures and the error equations, under worst-case interpolation error conditions $R = 0.5$, we can conclude that:

- The error due to the matching conditions, within a usable band sufficiently far from the Nyquist frequency, is bound, as the reflection coefficient itself is also bound to the interval $-1 \leq \rho \leq 1$.

- For matched conditions, the error experiences a minimum.

- The line length has an influence on the error, especially for greater cases of mismatch, as resonances appear on the line. The line becomes a resonator near open and close-circuit boundary conditions and the many reflections accumulate error.

- Line length has no influence under matched boundary conditions.

- Interpolation results in a reduced usable bandwidth, which is further reduced due to cumulative error on reflected energy in case of mismatched boundary conditions.

3.3 Direct Comparison Between the Ideal and Interpolated Models

In the previous sections, the normalised error ratios for the interpolated line model, compared to the ideal model, were computed. During worst case conditions, with the interpolation factor halfway in the time interval ($R = 0.5$), a simulation time step Δt close to the traveling time of the line, and limit cases for the boundary conditions (open and shorted), the errors grew rapidly into usability.

(a) $\overline{error}_V$ Magnitude

(b) $\overline{error}_V$ Magnitude (detail)

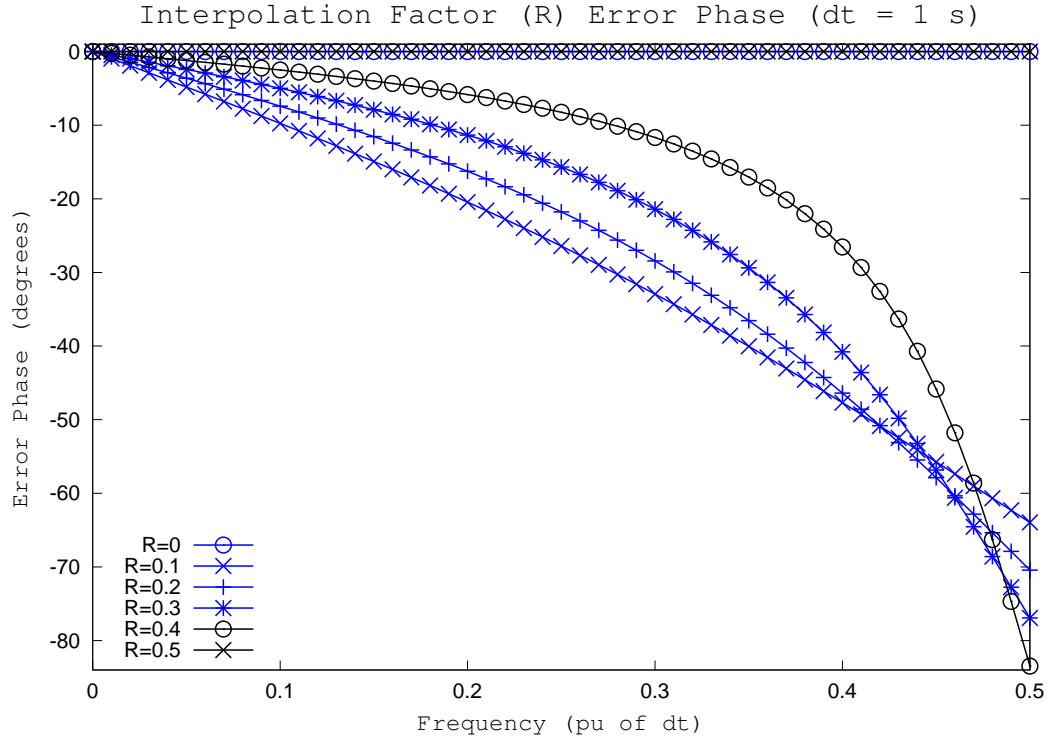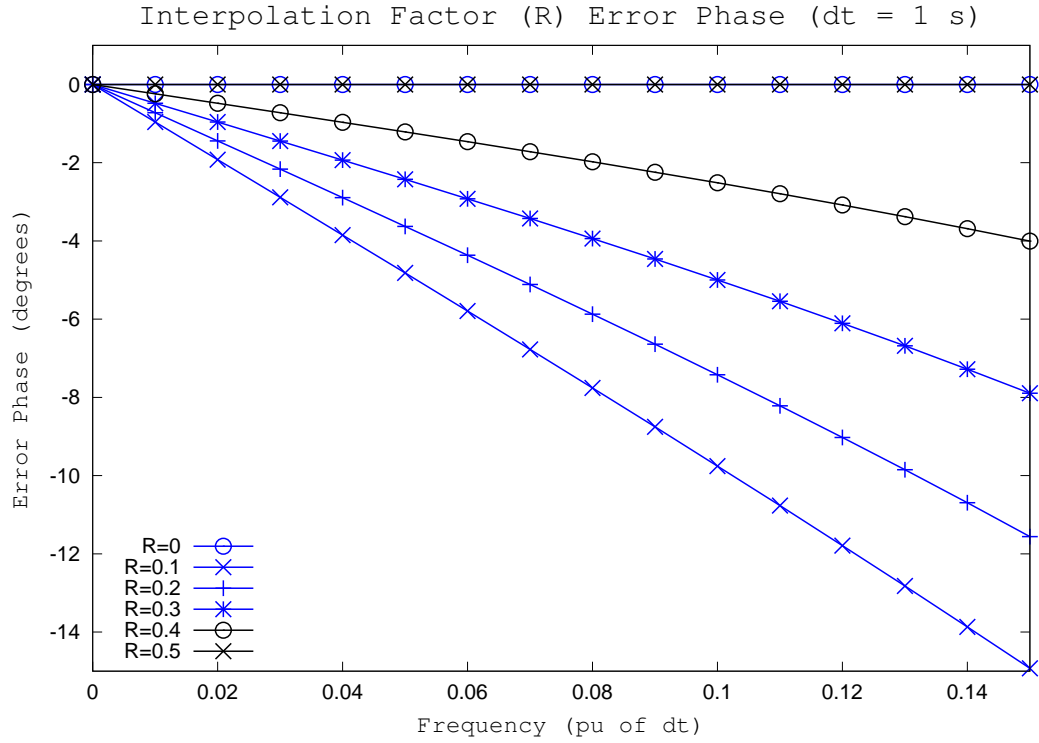Figure 3.11: Voltage Error Magnitude Plots for Various values of ρ and $L = 1.5$

(a) $\overline{error}_V$ Magnitude

(b) $\overline{error}_V$ Magnitude (detail)

Figure 3.12: Voltage Error Magnitude Plots for Various values of ρ and $L = 2.5$

(a) $\overline{error}_V$ Phase

(b) $\overline{error}_V$ Phase (detail)

Figure 3.13: Voltage Error Phase Plots for Various values of ρ and $L = 1.5$

(a) $\overline{error}_V$ Phase

(b) $\overline{error}_V$ Phase (detail)

Figure 3.14: Voltage Error Phase Plots for Various values of ρ and $L = 2.5$

(a) $\overline{error}_I$ Magnitude

(b) $\overline{error}_I$ Magnitude (detail)

Figure 3.15: Current Error Magnitude Plots for Various values of ρ and $L = 1.5$

(a) $\overline{error}_I$ Magnitude

(b) $\overline{error}_I$ Magnitude (detail)

Figure 3.16: Current Error Magnitude Plots for Various values of ρ and $L = 2.5$

(a) $\overline{error}_I$ Phase

(b) $\overline{error}_I$ Phase (detail)

Figure 3.17: Current Error Phase Plots for Various values of ρ and $L = 1.5$

(a) $\overline{error}_I$ Phase

(b) $\overline{error}_I$ Phase (detail)

Figure 3.18: Current Error Phase Plots for Various values of ρ and $L = 2.5$

It may now be of interest to compare the actual model responses under these adverse conditions directly in a graphical way as opposed to the numerical ratio of the responses [19]. This way, the nature and practical limitations of the inaccuracies can become more clear. In fact, if the peak value of the resonances is of no interest, the accuracy remains usable far longer than the previous results would let one to believe. Also, more suitable values for the simulation time step Δt will be shown for comparison.

To perform the comparison in frequency domain, to obtain the frequency response of a section of line under limit conditions of open and shorted circuit, we first must derive the appropriate equations. We will derive input-output relations for the models. Here, these are true transfer functions, not weighted errors as used in the previous analysis.

For the open line, it is more natural to compute the input to output voltage ratio, as this avoid limit conditions due to the zero current flow at the open end of the line. Starting from general line equations in frequency domain (3.37) and applying boundary conditions $\overline{I}_m = 0$:

$$\frac{\overline{V}_m}{\overline{V}_s} = \frac{2A}{A^2 + 1} \tag{3.49}$$

For the shorted line, the input to output current ratio is used, as this avoids the limit condition due to the zero voltage at the shorted end of the line. Starting from general line equations in frequency domain (3.37) and applying boundary conditions $\overline{V}_m = 0$:

$$\frac{\overline{I}_m}{\overline{I}_s} = -\frac{2A}{A^2 + 1} \tag{3.50}$$

To find the specific equations for the ideal and interpolated line, we now substitute the appropriate A terms from (3.46) and, after some manipulation, find for the ideal line model, for open and shorted cases, respectively:

$$\text{Open} \quad \frac{\overline{V}_m}{\overline{V}_s} = \frac{2e^{-j\omega\tau}}{e^{-2j\omega\tau} + 1} \tag{3.51a}$$

$$\text{Shorted} \quad \frac{\overline{I}_m}{\overline{I}_s} = -\frac{2e^{-j\omega\tau}}{e^{-2j\omega\tau} + 1} \tag{3.51b}$$

For the interpolated line model the corresponding equations, for open and shorted cases, respectively, are:

$$\text{Open} \quad \frac{\overline{V}_m}{\overline{V}_s} = \frac{2e^{-j\omega\tau}\left[e^{j\omega R\Delta t} + Re^{j\omega R\Delta t}\left(e^{-j\omega\Delta t} - 1\right)\right]}{e^{-2j\omega\tau}\left[e^{j\omega R\Delta t} + Re^{j\omega R\Delta t}\left(e^{-J\omega\Delta t} - 1\right)\right]^2 + 1} \tag{3.52a}$$

$$\text{Shorted} \quad \frac{\overline{I}_m}{\overline{I}_s} = -\frac{2e^{-j\omega\tau}\left[e^{j\omega R\Delta t} + Re^{j\omega R\Delta t}\left(e^{-j\omega\Delta t} - 1\right)\right]}{e^{-2j\omega\tau}\left[e^{j\omega R\Delta t} + Re^{j\omega R\Delta t}\left(e^{-J\omega\Delta t} - 1\right)\right]^2 + 1} \tag{3.52b}$$

These equations are now used to evaluate the frequency domain behavior of the terminated line models. We will only evaluate the voltage equations, as the current equations have the same response, with only a sign change.

The influence of the interpolation factor R and the simulation time step Δt is studied in Figures 3.19-3.21. The traveling time of the line in all cases was kept constant, within the temporal resolution afforded by the simulation time step and the variation added to the line length due to the interpolated sections.

It can be seen that, when the simulation time step Δt is close to the traveling time of the line τ ($L = 1$, thus the line is only one history deep), the worst-case condition evaluated in the previous section result in significant errors. However, in the plots, it also becomes clear that the phase change at the resonance of the open line section does happen at the right frequency, and so does the resonant peak, when compared with the ideal line, for all values of R. Thus, the errors are magnitude errors.

It must also be noted that the various peaks for different R are shifted with respect to each other. This is correct behavior. The influence of the interpolated interval is large, as in Figure 3.19, where the line itself is only one time step deep and, when $R = 0.5$, the interpolated interval adds 50% to the line length of $L = 1$, resulting in a line that is in fact $\tau = 1.5s$, with $\Delta t = 1s$, as opposed to the nominal $\tau = 1s$. Longer line lengths with increasing R cause the $\frac{\lambda}{4}$ resonance to shift to lower frequencies.

When the simulation time step becomes smaller, as in Figures 3.20 and 3.21, the accuracy of the model quickly increases, as the simulation bandwidth goes up by a factor 5 and 10, respectively. From the plots of previous section, it was seen that the usable bandwidth was roughly 1/5th to 1/10th of the bandwidth. By increasing the bandwidth, the model now produces good results. Also, since the interpolated interval is now much smaller compared to the length of the line, the total traveling times for the various cases of R are now closer to each other and all resonances are thus closer to each other as well.

3.4 EMTP 1/10 Rule of Thumb for Interpolated Lines

When interpolated transmission lines are used in EMTP simulations, it is customary to ensure that the simulation time step Δt is five to ten times smaller than the traveling time τ of the line [20, 52]. This is in addition to the usual requirements on bandwidth, as given by the Nyquist theorem [64]. In practice, it is advisable to use only one fifth to one tenth of the available bandwidth for accuracy, as opposed to half. Applying both considerations results in a generally acceptable error.

From the preceding study, similar conclusions are reached. In fact, using the error expres-

(a) Voltage Magnitude

(b) Voltage Phase

Figure 3.19: Frequency Response of Open Line for Various values of R and $L = 1$, $\Delta t = 1s$

(a) Voltage Magnitude

(b) Voltage Phase

Figure 3.20: Frequency Response of Open Line for Various values of R and $L = 5$, $\Delta t = 0.2s$

(a) Voltage Magnitude

(b) Voltage Phase

Figure 3.21: Frequency Response of Open Line for Various values of R and $L = 10$, $\Delta t = 0.1s$

sions derived above, it is now possible (in the cases for which the expressions are valid), to compute the error for a given case and adjust the simulation parameters to achieve an acceptable result. Also, depending on the value of the interpolation factor R, the ratio of Δt to τ may be relaxed. This can be seen in Figures 3.5-3.10 and 3.19-3.21 of the preceding analysis.

However, also the matching conditions play an important role, as they amplify the incurred errors due to the cumulative effect of multiple reflections, and thus passes through the model. This is shown in Figures 3.11-3.18.

Thus, the selection of the simulation time step, especially during severe mis-matched conditions and simulation parameters where the interpolation factor R is close to 0.5, requires great care, as the bandwidth over which the simulation gives acceptable results may be severely restricted. Conversely, in simulations where the lines experience relatively matched conditions and R is close to zero or unity, the computation may be sped-up as the larger allowable time step Δt results in less steps to be computed for a given simulation time, required bandwidth permitting. Equations (3.47) can help in the selection process, at least for cases where the section of line is driven from an ideal source. The same procedure of derivation can, however, be used to derive a more generic equation for a source with an impedance.

3.5 Conclusions

In this chapter, the behavior of the transmission line models that make-up TINA and the EMTP was studied. It was found that both the ideal and interpolated line models are stable. The interpolated model was shown to be dissipative, and the requirement for the EMTP $\frac{1}{10}$ rule was shown to maintain accuracy when interpolation is to be used.

The resulting error expressions can, under applicable conditions, be used to estimate the error incurred, and thus choose the simulation parameters for a suitable balance between performance and accuracy. It was found that the location of the interpolation in the interval, as well as the ratio of the time step to the traveling time of the line, and the termination conditions have a significant influence on the performance of the interpolated model. These relations, as applicable to the interpolated line model, have not yet been discussed in the literature, and may allow the simulation error constraints to be relaxed in certain cases.

Future work would be the extension of the discussion of the interpolated line to generic termination conditions, so a source with an impedance can be considered, as well as a further generalization of the stability proof for this model.

Chapter 4

Fractional Sub-Area Latency

> The law of the constant velocity of light in empty space, which has been confirmed by the development of electro-dynamics and optics, and the equal legitimacy of all inertial systems (special principle of relativity), which was proved in a particularly incisive manner by Michelson's famous experiment, between them made it necessary, to begin with, that the concept of time should be made relative, each inertial system being given its own special time. ...According to the special theory of relativity, spatial co-ordinates and time still have an absolute character in so far as they are directly measurable by stationary clocks and bodies. But they are relative in so far as they depend on the state of motion of the selected inertial system.
> *Albert Einstein*

In this chapter, the concept of fractional, sub-area latency is introduced. It is an extension and generalization of previous work on latency and temporal multi-grid techniques, allowing the ratio between the simulation time step and that of the area where latency is applied to be a non-integer multiple. Also, current use of latency and temporal multi-grid techniques in TLM and finite-difference methods has been limited to a single, integer, latency ratio for the whole system, or integer ratios for the sub-grids. Here, the use of multiple different, non-integer ratios for different parts of the simulation is developed, which allow both speed improvements in the solution and offer a new way of maintaining synchronicity in the simulation. The issues with non-integer multi grids, the asynchronous nature of the simulation and the need for extrapolation, and thus prediction, has been solved within the framework of the TINA method.

The contributions are thus the development of fractional latency, which requires a different approach compared to traditional latency and temporal multi-grid techniques, and the use of this technique, in compatible, time-decoupled solution methods such as TINA and TLM, to allow different parts of a simulation to operate at different local simulation time steps that need not be integer multiples of each other. As a result, the various sub-areas operate asynchronous from one another. This yields important performance gains, depending on the particular simulation case. It is demonstrated how the asynchronous simulation can be re-synchronized and, in time-decoupled methods, how the extrapolation necessary for fractional latency can be avoided. The fractional latency technique also offers a new way of maintaining synchronicity in a simu-

lation between the various sub-area in TLM-style methods, such as TINA, each with their own simulation time step, through its use of interpolation in time. This removes the need for local synchronicity methods, such as stub matching and temporal interpolation in the transmission lines themselves.

4.1 The Need for Latency Techniques in TINA

Latency techniques allow different parts of a simulation to operate at a different, local simulation time step. This ability is important, as in many cases, portions of a computation may be run slower without sacrificing accuracy, and thus important computational gains can be achieved.

The requirement to have multiple time steps in a simulation follows naturally when different cell sizes, and different materials, are used in a solution mesh. In electro-magnetics, most materials have wave propagation speeds that are relatively close to each other, typically within the same order of magnitude. However, when the FDTD and TLM methods were being used for non-EM problems, order of magnitude variations in wave speed readily occur. Acoustics is one such field. Another issue is the ratio between the wave length and the structure in which the energy must propagate.

Architectural acoustics in particular presents great difficulties to discretized-space methods, due to both the wave propagation speed variation between air and building materials, which can easily span over an order of magnitude, and the large solution space, compared to the wave length, resulting in lengthy computations.

Other methods are then commonly used, such as ray-tracing. This is a method that originated in computer graphics to compute three-dimensional scenes where virtual light beams are reflected and scattered off obstacles and finally observed at a receiver. Although computationally efficient for these applications, ray-tracing methods do not give phase information. However, recent work on a ray-tracing method that preserves phase information is under early development [10]. Discretized-space methods based on wave propagation naturally yield phase information and are thus of interest, and are actively developed, by the acoustics community, e.g.: [66, 67].

To clarify the situation, the following example may be considered: suppose we wish to study an acoustical system composed out of air and metal. The wave speed for air is roughly 340 m/s, while the wave speed for the metal can easily be 4000 m/s. Thus, an order of magnitude difference. It is clear that, given a minimum requirement of ten mesh cells per wavelength for accuracy, the air portion of the simulation, which has a short wavelength compared to the metal, will require a much finer spatial discretization than the metal.

If the same spatial discretization grid is used throughout the simulation, as is typically the

case in less developed solvers, it becomes clear that the mesh cell size is determined by the area requiring the finest discretization. This means that large areas of the simulation may be oversampled, and thus significant computational resources squandered for unneeded accuracy [15, 16, 98]. The situation becomes worse when the temporal discretization requirements of over-sampled areas are considered. The result is a simulation with unreasonably high bandwidth for most areas, just so sufficient accuracy may be maintained in one part. Sections 2.2.3 and 4.4.6 treat this in more detail as error criteria for the TINA method are discussed.

However, if one could operate these computationally unnecessary expensive, over-sampled areas at a more suitable, slower simulation time step, a significant portion of the lost computation can be recovered while sufficient accuracy can be maintained.

Typically, multi-grid techniques [15, 39, 98] would then be used in FDTD and TLM methods to allow for a suitable different spatial discretization within each simulation area and maintain accuracy. Spatial multi-grid techniques, however, are complex to implement due to the grid mis-match on the interface between the various different grid size areas. This is especially the case when the grids sizes are not integer multiples of each other, and partial cell overlaps exist. Hence, most multi-grid methods are restricted to integer multiples so that no partial overlaps exist, or otherwise unstructured grids are used, such as typically done in finite element methods, which adds significant computation to TLM-style methods [16]. Temporal multi-grid techniques are typically easier to implement and mis-matches are treated using interpolation [98], but have not been extended to TLM, or fractional ratios due to synchronization and extrapolation requirements, and thus stability and accuracy issues.

In this chapter, we will treat fractional sub-area latency as a multi-grid technique in time to be used within the time-decoupled mesh that makes-up the TINA method. By allowing fractional time-step ratios between sub-areas, each area in the TINA mesh can be operated at a local time step which is an integer multiple of the local transmission line τ, and thus not require internal interpolation, as was shown earlier in this thesis to maintain synchronicity. The use of fractional time steps, in itself, results in significant performance gains as the interpolation will only be on the interface between areas to maintain synchronicity, which will be explained further in this chapter, and no longer in each cell in the latency area. Accuracy of the computation will increase due to this, as each interpolated cell no longer forms a parasitic, cascaded low-pass filter. Using fractional latency, the interpolation errors are only incurred within the latency cells on an area boundary, thus only one instance of interpolation per boundary cell as opposed to thousands within an area, allowing more relaxed simulation parameters, and thus further reduced computational time.

In a sense, multi-grid methods in space and time are an expression of the space/time duality principle [22, 45]. We can either use an irregular grid in space, or an irregular "grid" in time.

Multi-grid techniques are an expression of the first, and latency is an expression of the latter. In the TINA computation, irregularity in time is used throughout and synchronicity is maintained though temporal interpolation. Temporal techniques are easier to handle in a digital computer, as this can be achieved directly by changing the depth of the history storage. In TINA, the spatial parameter is kept constant while the temporal one is allowed to be irregular. In most TLM and finite-difference methods, the converse applies. Ideally, a combination of both irregularity in space and time should be used to match a material area to the simulation with the desires spatio-temporal discretization.

The literature describing previous uses of multi-grid and latency methods will now be discussed, followed by the development of fractional latency as an extension and generalization of the method. It will be shown that, for time-decoupled meshes, the prediction required in fractional latency may be avoided and replaced with a look-ahead in the transmission line histories that connect the various TINA cells. This approach effectively removes the prediction from the technique and replaces it by a know value. Also, a method will be introduced to re-synchronize the now asynchronous, event-driven simulation when output must be generated.

Finally, the implementation of the fractional latency solution in EMTP and TINA is discussed and the simulation results compared to non-latency techniques.

4.2 Background

When discussing the latency technique, it must first be noted that different fields of study have developed different terminologies for what is, essentially, the same method. In electrical circuit simulation, incorporated in programs such as EMTP for power systems, where the systems are one-dimensional electrical circuits, or more correctly, graphs, the technique is know as latency, and applies to either elements of the electrical circuit being modeled, or entire sub-areas. Latency is the term we will use throughout this thesis.

In TLM and finite-difference time-domain methods, the technique is known as multi-grid in time, although sometimes the term sub-gridding is used as well, depending on the context. Throughout the discussion it must also be noted that TLM and FDTD formulations of wave propagation are closely related, and thus have significant overlap. In fact, both methods have been proven to be equivalent formulations under certain conditions [9, 11].

Due to the parallel, but mainly independent development, of the methods, the discussion will be split in two halves. First, we will discuss the TLM and FDTD methods, which have a longer history regarding these methods. Then, EMTP latency will be discussed, as this formed the basis for the development of the fractional sub-area latency method.

4.2.1 Multi-Grid Methods in TLM and FDTD

Soon after the development of the finite-difference time-domain method by Yee [96] for electro-magnetics in 1966, it became clear that in cases where small features in large structures had to be studied, such as a wire in a larger cavity, the computational requirements would quickly become excessive. This is due to the large difference of scale between the structures, and thus large, finely sub-divided simulation mesh over a large simulation space. Similar issues were found in the transmission-line modeling method proposed by Johns [31] in 1971, and is in fact an issue common to all regular-grid discretized-space methods.

In FDTD and TLM, methods were consequently proposed to embed thin wires and narrow slots right into the cell description, or other means of obtaining local solutions for such features smaller than the spatial discretization size. The concept was first proposed in [28], then later extended to arbitrarily oriented wires [21, 41] and eventually to multi-wire looms [5, 92]. In parallel, methods to model the thin wires as an equivalent circuit, and take into account the interaction with the field in the mesh allowed the wires to be placed between the cells [93], while another description placed them in the center of the cells [68].

Although computationally efficient, as the mesh can be kept at the desired coarseness for free-space, such techniques have many drawbacks, as the detail in the local features of the structure, and the corresponding local field, is lost. Allowable shapes and sizes of the features are restricted. Also, the spatial location of the wire is restricted to the center of the cell, or to the boundary between the cell. This places restrictions on the spatial accuracy of the resulting simulation. [15, 68, 98].

Consequently, work was undertaken to allow the simulation mesh to have different cell sizes so it could be locally matched to the required feature size. For example, a thin 1 mm wire interacting with an electro-magnetic wave of several cm length could be locally surrounded by a fine, sub-mm mesh while free-space could be modeled with much larger cells. Such methods are known as multi-grid, or variable mesh algorithms. In 1981, one of the earliest methods, a two-stage solution to allow part of the mesh to operate at a local, finer discretization was developed by Kunz and Simpson [39]. The method would first calculate the fields in the coarse region, and would then use spatial and temporal interpolation to obtain the tangential fields on the boundary with the finer grid. Using these results, the finer part of the grid could then be computed. This was later extended to allow for variable simulation step sizes [98], allowing for each region in the mesh to have a different, and appropriate simulation step. Both the space and time discretization in the mesh were integer multiples of each other.

In TLM, similar developments took place and mesh refinement techniques, both local (multi-grid mesh) and global (variable or hybrid mesh) were developed to operate in space and/or time [69, 74, 94, 95]. Of special mention is the technique pioneered by Saguet and Pic

(a) Multi Grid Mesh (b) Variable, or Hybrid, Mesh
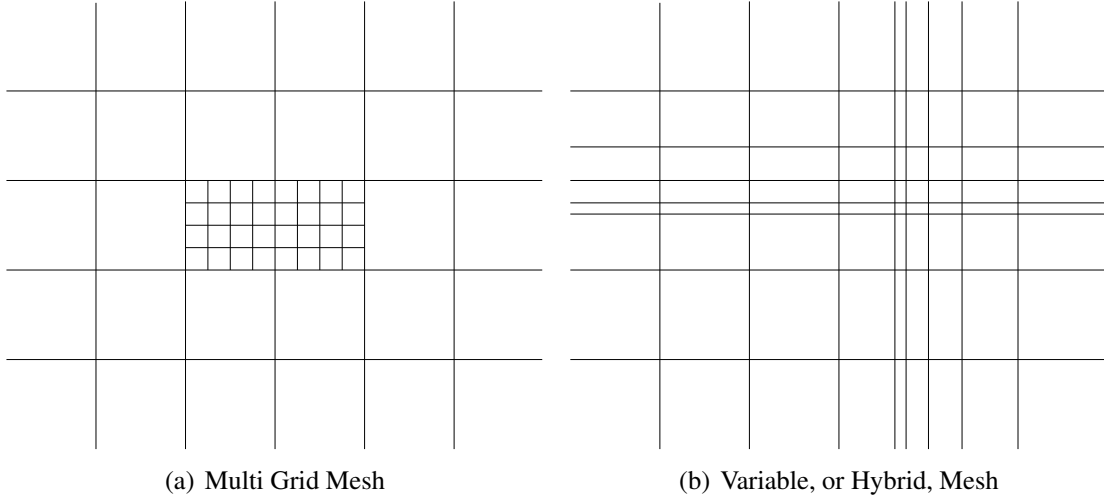
Figure 4.1: Mesh Types

[74], called variable of hybrid mesh. Illustrated in Figure 4.1(b), it does not use a purely local mesh refinement. The advantage is that a one-to-one spatial relationship is established between the cells, avoiding much of the complexity of the multi grid methods where the energy distribution on the boundaries between different mesh sizes must be handled [15, 26, 27], illustrated in Figure 4.1(a). Again, integer multiples for both spatial and temporal grid ratios are maintained to aid synchronicity between the various regions [14].

From the literature, we find that the shown spatial and temporal multi-grid techniques have only been achieved with integer multiples between time steps. In the development of the fractional latency method, which is a multi-grid technique in time, it will become clear that the difficulty with non-integer temporal multiples is that the simulation become asynchronous and that they require the use of extrapolation. This makes the method potentially unstable. Also, the potential need for iterative methods to improve the accuracy of the extrapolation prediction will increase the numerical burden. However, due to the formulation of the TINA method, the extrapolation issues can be readily overcome by the proposed fractional latency technique, as will be demonstrated. It will also be shown how the asynchronous nature of the simulation may be maintained and synchronized to an observer clock so synchronized output can be generated.

4.2.2 Latency in EMTP

Within the field of power systems analysis, there is no spatial discretization, as the computations are circuit simulations. They take the form of graphs, not discretized space. Thus, multi-grid methods are necessarily limited to the time dimension. The use of different time
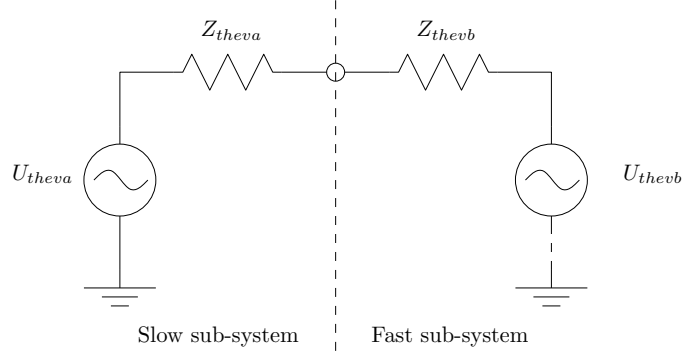
Figure 4.2: Two Connected Thévenin Equivalents in Different Sub-Systems

$$k\Delta t_a + \Delta t_b \quad k\Delta t_a + 2\Delta t_b$$

$$k\Delta t_a \qquad\qquad\qquad (k+1)\Delta t_a$$
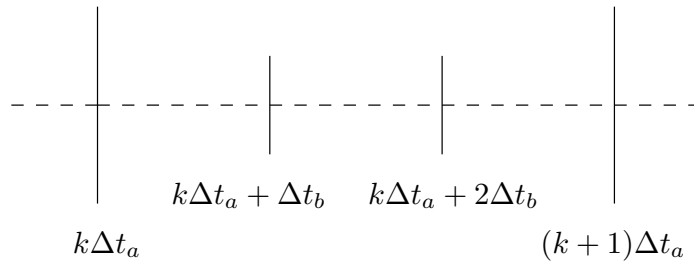
Figure 4.3: Integer Latency Solution Time-Line

steps for different parts of a simulation is referred to as latency.

The technique is based on earlier work on multi-rate simulation programs for very large-scale integrated circuits (VLSI), using the waveform relaxation method [42, 75, 97]. Later, the use of latency in power systems simulations, using EMTP, was studied [78] and initial work on numerical stability and accuracy of the method was shown in [43]. The method was further developed in 2002 by Moreira [59, 60] in the context of the Multi-Area Thévenin Equivalents (MATE) technique [49] that allows for efficient parallelization [87, 88] of power system simulations.

Within the context of EMTP, the latency technique shows the interesting property that, for integer multiples of the involved time steps, only interpolation of the slower time step needs to be used to find the intermediate fast ones. This is because we know the next step needed to interpolate with, as in EMTP the history sources compute their next result using only past and current information [20, 60].

To illustrate the concept, let us represent two parts of a system, each by their own Thévenin equivalent. This representation has the advantage that it is a universal description for any system at a given simulation time and used as the core sub-system interconnection concept in MATE and TINA, as well in the transmission line models used in this thesis. In Figure 4.2, the

U_{theva}

interpolated   exact

exact

Time

$k\Delta t_a \quad k\Delta t_a + \Delta t_b \quad k\Delta t_a + 2\Delta t_b \quad (k+1)\Delta t_a$
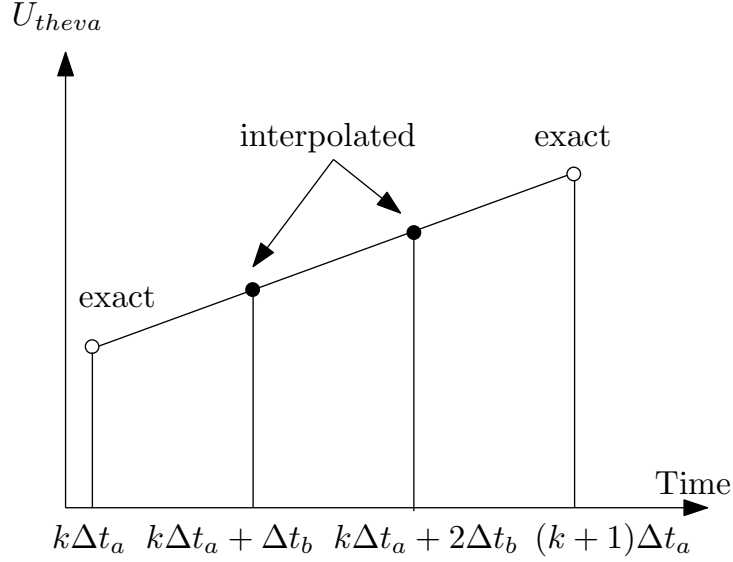
Figure 4.4: Integer Latency History Interpolation

situation is illustrated. The left-hand side of the figure is the slow sub-system, while the right-hand side is the fast sub-system. Thus, each sub-system can be allowed to run at a different simulation time step, where the slow sub-system runs at a larger local time step than the fast one, and maintain local accuracy. When we take the fast system three times faster than the slow system, the corresponding simulation time line is shown in Figure 4.3.

Thus, as time advances through the simulation, we observe that the sub-systems are solved in sequence. Both are solved at time $k\Delta t_a$, then the fast sub-system is solved twice, followed by another joint solution. Thus, for the fast system, there are no corresponding values from the slow system to compute its internal results with and vice-versa. However, this is not so. If the simulation is arranged so that the slow system is solved first, either from zero or known initial conditions, an exact value for its solution can be found, as EMTP solutions use current and present data to compute the next step. Assuming that the slow system varies little within each step, which is a key requirement for latency techniques to work, the next value for the slow system can be found. Then, using interpolation, the intermediate values that are needed by the fast system can be estimated. This means that the interpolation happens between two exact values, and no prediction is used. Also, the system remains in synchronism at the rate of the slowest sub-system. The situation is outlined in Figure 4.4. In EMTP latency techniques, linear interpolation is typically used.

The temporal interpolation used in EMTP latency techniques is typical of other forms of temporal interpolation for discrete-time, time-domain solutions discussed in the literature concerning TLM and FDTD techniques. It is based on integer multiples of the respective time

steps, which allow the use of interpolation between known values. Since the time steps are integer multiples of each other, the time steps always align, meaning the fast system, after a fixed number of solutions, coincides with the slow one. We will now show why only integer multiples have been used so far in the development of latency techniques.

4.3 Development of Fractional Latency

Taking a step back from latency, in a practical simulation, the different materials do not have simulation time steps that are clean integer multiples of one another. For this reason, the TINA method makes extensive use of temporal interpolation to achieve synchronicity to a common time step throughout the simulation. That implies that entire areas of a certain medium require interpolation in all constituent cells to match them to this time step. Thus, a material area becomes a system of cascaded filters, due to the low-pass effect of the interpolation in each cell, and the distortion increases. This happens in the stub-matching method used in TLM as well [18], which adds a matching transmission-line stub to each cell.

If it were possible to allow each material area to operate at a local time step that is an integer multiple of the transmission time τ of the individual lines that make-up the material cells, then the interpolation in the material cells could be omitted. Aside from speed-gains, this also removes an important and cumulative source of error.

Returning to latency techniques, the classical, integer multiple latency cannot help to completely remove the interpolation in the cells, but it can reduce the memory use and computational burden as the time steps in the local areas can be chosen closer to optimum values. This is especially the case when a constant grid size is used and there are large variations in wave speed in the simulation. If it were possible to use non-integer ratios, then also the interpolation within each cell could be removed as the local time step could be exactly matched to the local transmission time of the lines.

This is what fractional sub-area latency allows: each sub-area, or material, can operate at a convenient, local simulation time step that is a precise, integer multiple of the local transmission time τ. There is still need for interpolation to implement the fractional latency itself, but only on the boundaries between the different material regions. Thus, fractional latency maintains the advantages of integer latency while also allowing synchronicity to be maintained at a much reduced numerical and memory cost, compared to interpolation in every cell. This may result in better accuracy, as the interpolation is now reduced to only the boundary cells as opposed to the entire material area. Thus, fractional latency combines the advantages of integer latency and temporal interpolation as a means to synchronicity, while also mitigating some of the biggest drawbacks of interpolation for synchronicity.
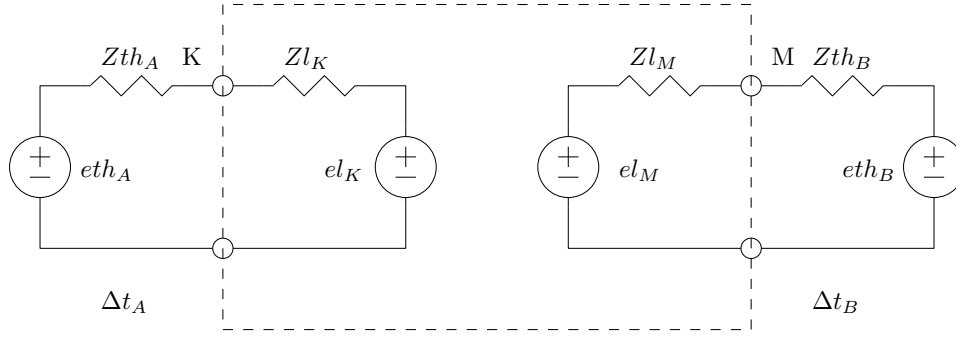
Figure 4.5: Fractional Latency Bi-Directional Cell

4.3.1 Equations and Parameters for Fractional Latency

Fractional latency, as it is considered here, is conceived as a special cell that is introduced between two parts of a system. Since connectivity in TINA is entirely based on Thévenin equivalents, this is also the concept that will be used for the latency cell, to allow for easy insertion into the simulation grid.

The advantage of using Thévenin equivalents is that they offer a universal abstraction for any circuit in a time-domain, time-marching solution scheme. Thus, the bi-directional latency cell can be described, and connected, as shown in Figure 4.5. It is essentially two Thévenin equivalents, where each represents a time-interpolated version of the circuit connected to the other side of the cell.

The governing equations, which are temporal interpolations of the Thévenin equivalent circuit components are now given, for the current time step t_A or t_B as:

$$
\begin{align}
el_K(t_A) &= eth_B(t_B + \Delta t_B) + R_K\left[eth_B(t_B) - eth_B(t_B + \Delta t_B)\right] \tag{4.1a} \\
el_M(t_B) &= eth_A(t_A + \Delta t_A) + R_M\left[eth_A(t_A) - eth_A(t_A + \Delta t_A)\right] \tag{4.1b} \\
Zl_K(t_A) &= Zth_B(t_B + \Delta t_B) + R_K\left[Zth_B(t_B) - Zth_B(t_B + \Delta t_B)\right] \tag{4.1c} \\
Zl_M(t_B) &= Zth_A(t_A + \Delta t_A) + R_M\left[Zth_A(t_A) - Zth_A(t_A + \Delta t_A)\right] \tag{4.1d}
\end{align}
$$

The interpolation factors, illustrated in Figure 4.6, are computed as:

$$
\begin{align}
R_K &= \frac{(t_B + \Delta t_B) - t_A}{\Delta t_B} \tag{4.2a} \\
R_M &= \frac{(t_A + \Delta t_A) - t_B}{\Delta t_A} \tag{4.2b}
\end{align}
$$

In Figure 4.6, to compute $t_A + \Delta t_A$ in a sub-area operating at a local time step of Δt_A, we

101

require a solution at time t_A. We have the solution at $t_A - \Delta t_A$, but to compute the solution at t_A, we require information from the connecting sub-area, which operates at Δt_B. Since this sub-area does not have this result available, the neighbor values needs to be interpolated so they can be used in the sub-area operating at Δt_A at time step t_A to find $t_A + \Delta t_A$. Using linear interpolation, the interpolation factor R_K is found. The gray dot in the figure represents the interpolated value for the neighboring sub-area operating at Δt_B, interpolated in time to match the time step of the area operating at Δt_A.

As opposed to the interpolated transmission lines, which were studied in Chapter 3 and Ap-
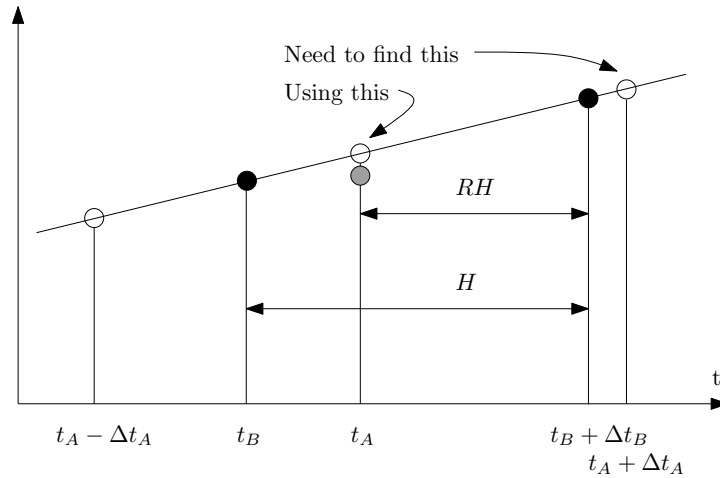
Figure 4.6: Fractional Latency Interpolation Interval
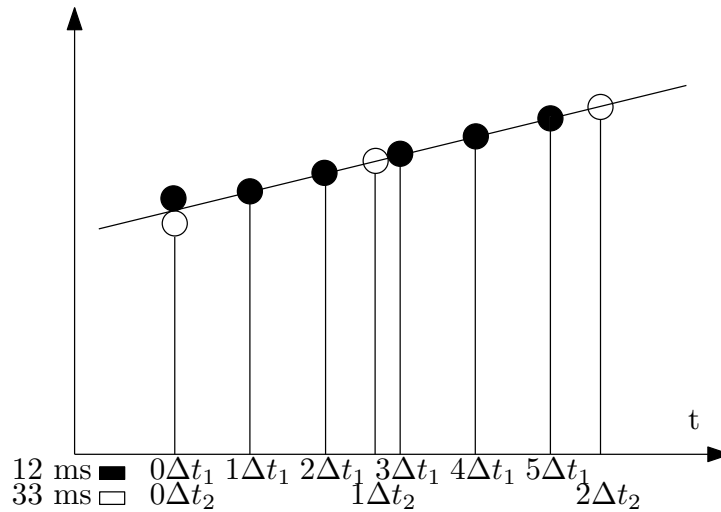
Figure 4.7: Fractional Latency Solution Issue

Event	Δt_1 (12 ms)	Δt_2 (33 ms)	Δt_3 (44 ms)	Δt_4 (100 ms)
0	0	0	0	0
1	12			
2	24			
3		33		
4	36			
5			44	
6	48			
7	60			
8		66		
9	72			
10	84			
11			88	
12	96			
13		99		
14				100
15	108			
16	120			
17	132	132	132	
18	144			
19	156			
20		165		
21	168			
22			176	
23	180			
24	192			
25		198		
26				
27				200
28	204			

Table 4.1: Numerical Fractional Latency Simulation Event Table

pendix A.3, the interpolation factor R is no longer a constant factor, but varies. The reason for this is illustrated in Figure 4.7 and Table 4.1, which shows a numerical example of simulation events for two and four latency areas, respectively. It can be seen that the location of the value to be interpolated in an interval changes with each time step. Thus, different interpolation factors must be used for each time step. The whole simulation, using fractional latency, operates asynchronous. Data is only synchronized on the boundaries between two areas using the latency cells, and for output processing, as discussed in Section 4.3.6. Temporal interpolation is used each time to achieve synchronicity.

4.3.2   The Need for Extrapolation

From the above equations, Table 4.1, and Figure 4.7, the main issues with fractional latency, at least in a generic EMTP context, will become clear: there is a chicken and egg problem. Throughout the discussion, keep in mind that EMTP solutions are found in two steps. First, using current and past data, the next, future step value for the history sources is computed [20]. During this step, the components in the system are not changed. When all the internal histories are solved, the system itself is computed using these new histories to obtain the complete solution for the next time step. Thus, there is a difference between knowing the values of the history sources at the next time step and knowing the system solution at the next time step. The cases discussed below are based on a simplified version of the numerical example given in Table 4.1.

In the following discussion, a simple case consisting of two sub-areas is considered. The local time steps are $\Delta t_1 = 12ms$ and $\Delta t_2 = 22ms$. The ratio between the time steps is thus: $\frac{\Delta t_1}{\Delta t_2} = \frac{12ms}{33ms} = 0.6363....$ Fractional latency applies. Each sub-system has a number of internal components, with their associated history sources. However, that internal complexity is abstracted to a Thévenin equivalent on each connection port [30]. Thus, the whole circuit is similar to that shown in Figure 4.5, where two sub-areas are joined using a universal, bi-directional fractional latency cell. As such, the overall circuit must still be solved together in order to reach a solution, as there is no decoupling in the solution. This last point is at the core of the issue with fractional latency ratios.

Starting at $t = 0$, values for both the $\Delta t_1 = 12ms$ and $\Delta t_2 = 22ms$ sub-systems can be computed through, for example, zero initial conditions. Since the entire system is defined, the EMTP solution algorithm allows us the compute the updated history sources for the next time step based on the known boundary conditions presented to each sub-area through the Thévenin equivalents. Using the updated histories, the whole system may be solved and new values for the Thévenin equivalents for the next time step found. Thus, at $t = 0$, we can compute system

solutions for $1\Delta t_1$ and $1\Delta t_2$.

Now, since the sub-system operating at Δt_2 hardly changes relative to the sub-system operating at Δt_1, which is the main condition for latency, we may use latency to compute the next few values for the faster sub-system using interpolation of the Thévenin equivalent presented by the slower sub-system. Thus, a local system solution for $2\Delta t_1$ can be found using the interpolated Thévenin equivalent presented by the slower sub-system. Also, the Thévenin equivalents of the faster sub-system can be updated, as the interpolated Thévenins for the slower sub-system are available. Thus, computation is saved as during the computations the slower sub-system needs not be solved for, only its Thévenin equivalents must be interpolated.

The solution will now fail as the slow sub-system needs to be computed at $t = 1\Delta t_2$. To compute the new history values for this sub-area and solve the system, we require the Thévenin equivalent obtained from the fast sub-system at this time step. However, the solution of this faster sub-system has only progressed until $t = 2\Delta t_1$. Although we could compute the next value of the internal history sources of this faster sub-system, its Thévenin equivalent cannot be found, as this requires a full system solution including the slow sub-system. However, since the internal histories of the slower system could not be computed for want of the boundary conditions presented through the Thévenin equivalent associated with the faster sub-system, the solution stalls. To find the solution for the slow sub-system, we need the solution of the fast sub-system. To find the solution of the fast sub-system, we require the solution of the slow sub-system. There is thus a dead-lock, a chicken and egg problem.

In classical latency, as implemented in EMTP [60], and all other temporal interpolation methods we found in the literature, the dead-lock problem is avoided by only allowing integer ratios between the time steps. Indeed, if the time steps at $3\Delta t_1$ and $1\Delta t_2$ overlap, the same situation as at the zero time step is encountered: the whole system is defined and thus solvable. This is likely the reason why we could find no examples of fractional latency in the literature.

At the core of the problem lies the fact that both sub-systems are not decoupled, and can thus not be solved wholly independent from each other. Such a solution is possible when transmission lines are present on the connections between sub-systems, which offer such decoupling in time, e.g.: [29]. This type of solution is discussed in Section 4.4.3. However, since in a generic EMTP computation one cannot rely on the presence of transmission lines, we may continue the solution if extrapolation is used. This allows us to guess one of the missing solutions and break the dead-lock. This solution is discussed in Section 4.4.1. The linear extrapolation used has the same form as the interpolation equations, only the interpolation factor R is allowed to grow larger than unity. Again, the extrapolation will be avoided in the TINA context, where each cell is time-decoupled from the other.

4.3.3 Fractional Latency for Simulation Synchronization

The main advantage of fractional latency over standard latency is that the ratios between sub-areas no longer need to be an integer ratio. Stepping back from this chapter, in other parts of this thesis, the use of interpolation within transmission line models was extensively studied, Chapter 3 and Appendix A, and used to construct the TINA mesh when dissimilar materials were present and maintain synchronicity. Also in EMTP simulations, temporal interpolation in the line models is used to achieve synchronicity in the computations. The reason temporal interpolation is used was in these cases is match the traveling time of the line τ with the simulation time step.

Returning to the latency concept, if only integer ratios are allowed, interpolation in the line models will frequently be required. This introduces a double error: the error caused by the interpolation at the border of a sub-area to match the different sub-area time steps, and the error caused by the interpolation within the transmission lines that make up the sub-area to match them to the local sub-area time step.

Using fractional latency, it is now possible to choose the time step of the sub-area as a clean, integer multiple of the transmission time τ of the line. Doing so removes the need to interpolate the transmission line, and the exact line model can be used. Aside from accuracy improvements, this also removes the computational burden of having to interpolate every transmission line segment in a sub-area. In a TINA simulation, with easily ten to hundreds of thousands of transmission line segments in an area, this results in significant speed gains.

Thus, fractional latency allows us to take the interpolation out of the line models and place it at the boundaries of the sub-area. As a result, the number of cells performing interpolation is reduced to only the sub-area border. The effect of cascading thousands of filters, as the interpolation has a distinct and important frequency response, as studied in Chapter 3, is removed. In TLM methods, the need for stub-lines at the center of each mesh cell used to achieve synchronicity [16, 18] is thus also removed.

Through the use of fractional latency, the errors and computational load usually associated with synchronicity methods could thus be significantly reduced. Using the technique, all the interpolation needed for the volume associated with a sub-are can be moved to only the cells associated with the bounding shell of the sub-area.

4.3.4 Choice of the Latency Master Time Step

Initially, there are two ways fractional latency could be implemented. One could choose to make the master time step, the base reference for the whole simulation to which all areas are eventually synchronized, the slowest or the fastest in the simulation. This decision has

implications on the details of the solution algorithm discussed in the next section, and thus must be addressed first.

Exploring the implications of this choice, questions of a philosophical nature appear and, in the end, the decision becomes one of personal preference, unless special technical considerations force the decision.

When choosing the master time step the fastest one in the system (thus, the sub-areas are "local slow"), the slower areas have a reduced bandwidth. Therefore, any signal in the system must be limited by the bandwidth of the slower sections or temporal aliasing effects will be present. This adds the difficulty of insuring that no signals exceeding any of the local band widths exist in the simulation to comply with all the different Nyquist criteria. Thus, there is a requirement for band-limited signal sources and care must be taken when non-linearities are introduced in the system. More important, the system output data, which is produced at each instance of the master time step, contains many redundant time steps due to excess bandwidth from over-sampled slow ares that cannot be used, as it is not valid for the whole simulation domain. For 2D and 3D simulations, this excess bandwidth results in wasted storage and processing time to write the output. This is especially the case with fractional latency, where each output cell that is not an integer multiple of the simulation master time step must be interpolated in time to obtain the result at the desired time step.

When choosing the master time step the slowest one in the system (thus, the sub-areas are "local fast"), the faster areas have an increased bandwidth. Again, the issue with non-linearities occurs, but if the sources are embedded in space that runs at the master time step, they are automatically band-limited, as the source is locked to the bandwidth afforded by the master time step. One could use a similar approach in the previous case and insure the sources are all in an area that runs at the slowest time step in the simulation. The advantage of the local fast approach is in the amount of output data that needs to be processed, and the natural limitation on the output bandwidth, leaving no doubt as to the Nyquist limit of the output data. The disadvantage is that the spatial wave propagation in fast sub-areas may be temporally under-sampled and spatial aliasing in the output data may exist. The data is still correct at each instance, but strobe-like visual effects do appear when a simulation is played-back as a movie. Compare, for example, the effect where in movies car wheels appear to suddenly turn backwards as the car slows down in a scene. The local slow approach does not have this phenomenon, as the output is saved at the rate of the fastest time step.

In the TINA implementation of fractional latency, the local fast approach was chosen to reduce the amount of data to be processed. Since output interpolation must be done at each stored time point, not having to do this at the rate of the faster system resulted in orders of magnitude less steps to be output-processed, interpolated, and stored.

107

(a) Sequencing with Two Time Steps   (b) Sequencing with Three Time Steps
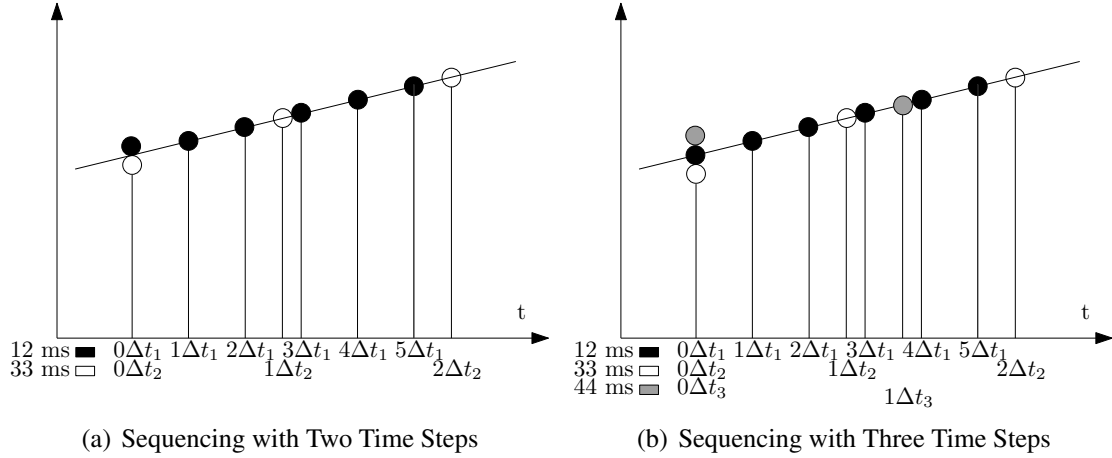
Figure 4.8: Fractional Latency Solution Sequencing Examples

The spatial aliasing effects were visible in all our acoustic simulations involving metal boundaries, thus fast materials compared to the slow air. But if this wave behavior within the fast material must be studied, the master time step (by means of configuring the simulation bandwidth) may be reduced to a sufficiently short interval to remove those spatial aliasing effects. Since the simulations experiments in this thesis were all studying the acoustic field in air, and not the bounding materials, the area of interest, being air and the slowest, did not suffer from spatial aliasing in the output data, and thus the computation and storage advantages of the local fast method could be obtained.

It must be noted that the spatial aliasing only happens in the output processing of the simulation, not in the simulation itself, as the fast areas are solved at an appropriate, local simulation time step within the computation. Due to the event-driven nature, bordering fast sub-areas exchange data at an appropriate rate. Thus no spatial nor temporal aliasing occurs in the simulation itself. It is only in the output processing that the effect occurs.

4.3.5 Solution Algorithm

Using the local fast approach, the main simulation and output processing is thus operated at the slowest rate in the system. The second consideration is the order in which the sub-areas are solved. This sequence is important, as otherwise the solution can dead-lock. In this thesis, an event table-based solution sequencing system was implemented.

First, consider the two simplified cases shown in Figure 4.8 for two and three sub-areas. The examples are based on Table 4.1. When looking at the figures and table, it is clear that the various sub-areas require solution at different absolute simulation times.

Let us first establish the concept of absolute simulation time, which is key to the algorithm

developed to manage the event table. At the core of the problem is that, using the fractional latency concept, the simulation is essentially asynchronous. With normal latency, the various sub-areas run at different rates, but still synchronize every so many steps. This is not the case with fractional latency. Thus, the standard idea of time in EMTP, TLM, and finite-difference computations, which is simply that of time-marching and wholly implicit, cannot be used as there is no clear relationship between the temporal conditions between sub-areas.

Using fractional latency, we thus have to synchronize an asynchronous system to a regular time step in order to generate valid output. Since there is no direct relation between the individual sub-areas, we still need to establish a relation of some type. There is an entire field of study that investigates such synchronicity problems. Called the science of synchrony, it brings mathematics, physics, and biology together to study how spontaneous order occurs naturally in the universe. Steven H. Strogatz's introduction to the topic [84] suggest a solution to the issue at hand.

In essence, it comes down to the realization that by observing the system as a whole at regular intervals, like using a strobe light to find the rotational speed of machinery, the asynchronous system may be synchronized at the observational times.

In order to define the very notion of "at the same moment" in a fractional latency simulation, the concept of absolute time must be introduced. Observing at regular intervals thus becomes a matter of obtaining values at regularly spaced absolute times and interpolating between absolute times if no solution for the required observational moment exists, an issue which will be investigated in Section 4.3.6.

The concept of absolute time is readily implemented in a time-marching solution. Since each sub-area is aware that it is being solved, by virtue of the solution code for that area being called by the simulator, and it is aware of its local simulation time step, it is a matter of counting the number of these steps and multiplying the count by the local simulation time step. The result is the absolute time of that sub-area. By means of this time, each solution for that sub-area can be supplied with an absolute time stamp, and stored in the correct order with respect to the other sub-areas. The result is a table of time-stamped simulation events, each associated with solution data, similar as Table 4.1.

To establish the solution order of the simulation using these absolute time stamps we return to Figure 4.8 and Table 4.1. From these, we can observe how the solution sequencing algorithm must work.

First, at the zero time, all sub-systems must be solved. The order in which is important. When multiple areas are to be solved at the same absolute time, we must always solve the sub-area with the largest simulation time step first. Doing so insures the future values needed to interpolate boundary conditions for the faster sub-areas are available. The solution for the

system will now yield the future values, each at the local time step of $1\Delta t_x$.

The second rule in the algorithm is to always solve the area which has the smallest absolute time stamp. In the example being studied, that would mean that all areas will be solved at $t = 0$, the slowest sub-area first. When all sub-areas whose local absolute time was set to zero are solved, the next are to be solved will be the sub-area operating at Δt_1. It has, at that time, the smallest absolute time and only that area needs to be solved at that absolute time. It needs interpolated boundary conditions from the connecting sub-areas for solution. In order for this interpolation to be possible, a solution for that neighboring area must exist at a later absolute time than the absolute time at which the solution for the sub-area is sought. Due to the described algorithm, this is always the case.

When the above two rules are applied, an event table may be calculated to control the solution sequencing in the simulation. In practice, the event table is computed on the fly and has as many entries as there are sub-areas in the system. An example solution sequencing with absolute time stamps, using the above rules is shown in Table 4.1.

Thus, the algorithm to determine the solution order may be captured in a flow chart. Figure 4.9 illustrates the process, including the output processing step, as discussed in the section below. The box labeled "Solve the latency area" refers to the standard EMTP-style solution used throughout this thesis, where the latency cells provide the information from what normally would be the boundary Theévenin equivalents from other latency areas so the local area may be solved at its local time step.

4.3.6 Generating Output

Regardless of the choice of the master time step, output generation will be required each time the areas operating at the master time step, which is directly related to the required simulation bandwidth, are called. One could generate output at any time, if required, but by synchronizing to a latency area, significant computation can be saved, as the area the output generation is synchronized to does not require interpolation. In TINA and the EMTP example later in the chapter, output generation is synchronized to the slowest time step in the system, which serves as the master time step. The required step for output processing is reflected in the flowchart of Figure 4.9, where the output processors are called when the master time step areas are computed.

Output generation is then a matter of obtaining readings from the asynchronous simulation that corresponds to the absolute time the output is requested for. Thus, for the master time step areas and those areas that, at the time have the same local absolute time, this is simply the results computed at the current time step.
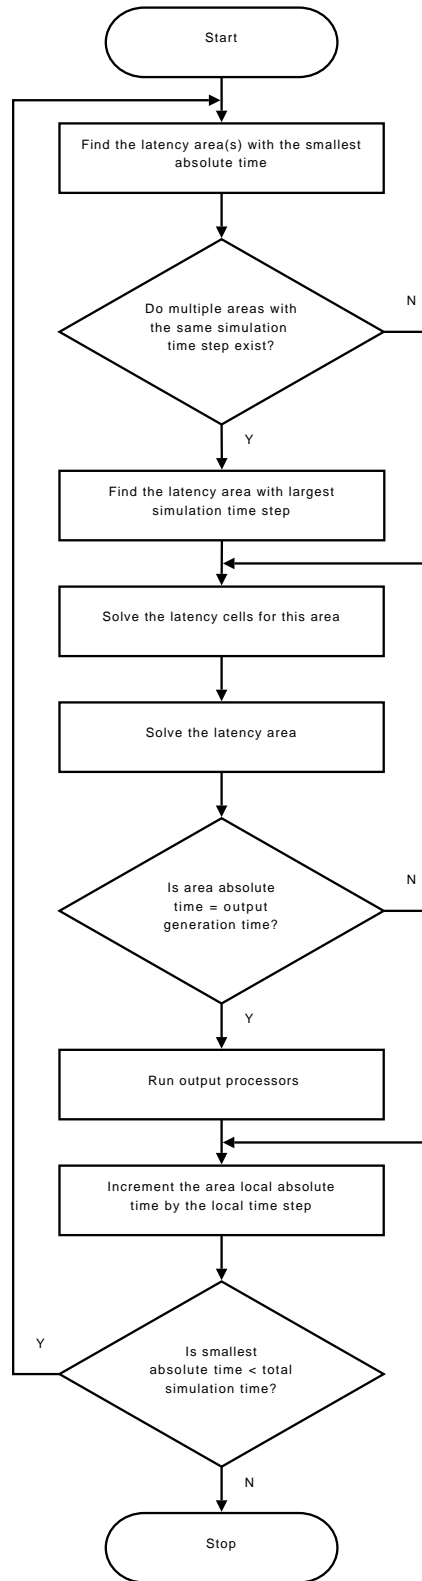
Figure 4.9: Fractional Latency Solution Sequencing Algorithm

For the other areas, where no simulation data exists for the required time step (thus, where fractional latency is applicable), interpolation over the future and past solution will be used. Due to the solution sequencing algorithm described in the previous section, it is ensured that either the area is at the correct absolute time for output, or a solution for a future time, relative to the current absolute time, exists and interpolation can be used to find the output for the required absolute time.

The process of generating the output data itself, at least in our implementation, is done using the same linear interpolation (4.1) used for the fractional latency process. It has to be noted that the errors associated with output interpolation do not influence the simulation itself. The output results do not participate in the actual computation, they are only stored to disk for post-processing of the data.

4.3.7   The Accuracy and Stability of Fractional Latency

The fractional latency interpolation algorithm and the resulting latency cell has a non-flat frequency response due to the interpolation. One important difference, compared to the interpolated line model studied in Chapter 3 is that the interpolation factor R, is not constant and potentially varies over all values in the normal interval $0 \leq R < 1$. Thus, the interpolation error is not constant either. By assuming worst-case conditions, $R = 0.5$, the error criteria derived for the interpolated line model may be indicated initially, but will be explored further later in this chapter, when the simulation parameters are discussed in Section 4.4.6. However, no formal assessment of the stability and accuracy of the fractional latency cells will be given.

Qualitative comparisons between latency and standard cases in the TINA context indicate that accuracy is comparable to what is to be expected from the interpolated transmission lines, with a much reduced computation time. Given the close relationship between the used equations and methodology, this is not unexpected.

In a more general EMTP context, however, the extrapolation brings with it a host of stability issues, as can be seen from the simulation results. A formal stability analysis would be quite similar to the methods used in Chapter 3 and a Courant condition [83], which places limits on the acceptable time step and spatial discretization size, should be derived. We leave this analysis for future work.

4.4   Implementations of Fractional Latency

In this thesis, two implementations of the universal fractional latency cell, as described earlier in this chapter, were made. One was an EMTP-style computation using extrapolation. This
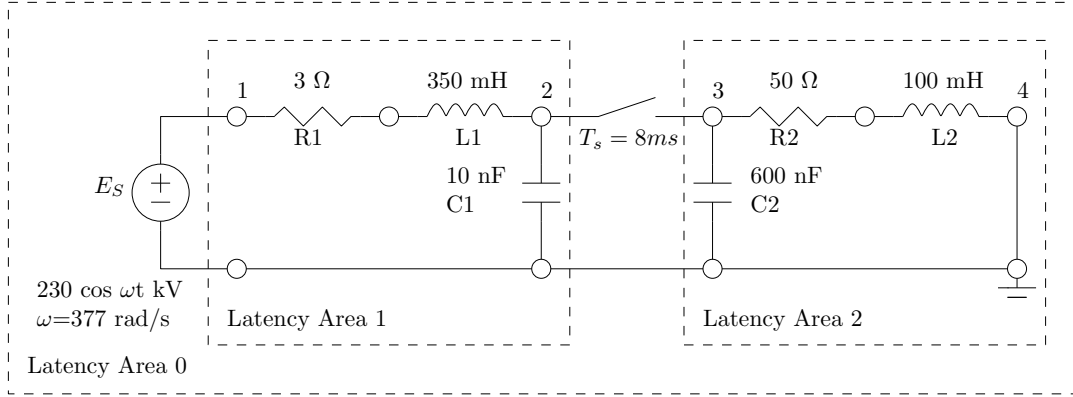
Figure 4.10: Fractional Latency Domains in EMTP Solution Case

test case was used for the original latency work, and repeated here to compare the fractional latency results, which require extrapolation in this context, to the normal EMTP solution. The second implementation is in the TINA framework, where the extrapolation can be omitted as the TINA solution is fully time-decoupled. Performance and accuracy comparisons between the ideal solution, synchronicity through interpolation, and fractional latency techniques will be made.

4.4.1 Fractional Latency in EMTP

The example used for the latency implementation is that shown in Figure 4.10, as obtained from [52]. It is a lumped system where each side, before and after the switch, has a different time constant. Thus, each side should be computable with a different local time step. As an additional complexity, in our simulation both the source and the switch operate in the base latency area, which is the one operating at the slowest rate. Thus, there are three latency areas in the system.

To implement this system using the fractional latency cells described in this chapter, the various sub-areas need to be abstracted as Thévenin equivalents and connected to the fractional latency cells from Figure 4.5 at each area boundary. This means, since this computation consists of only a few lumped components and is thus not time-decoupled, that the equivalent system for the entire system behind a given Thévenin equivalent must be found. This is of course computationally not effective. Hence, the latency formulation commonly used only operates on the history sources internal to a latency area themselves, e.g.: [43, 60]. However, the fractional latency cell was only applied here to investigate the extrapolation effects.

A more practical use in an EMTP context of the fractional latency cell method could likely be found through the use of Multi-Area Thévenin Equivalents (MATE) [49, 87]. In this method,

Figure	EMTP Δt	Area 0 Δt	Area 1 Δt	Area 2 Δt	EMTP t_s	Frac. Lat. t_s
4.11	35 μs	35 μs	3.51 μs	3.54 μs	0.015995 s	0.01589 s
4.12	35 μs	35 μs	35.1 μs	35.4 μs	0.015995 s	0.01547 s
4.13	35 μs	100 μs	100.1 μs	35.4 μs	0.015995 s	0.0154 s
4.14	35 μs	200 μs	100.1 μs	35.4 μs	0.015995 s	0.0088 s

Table 4.2: EMTP Case Simulation Parameters and Switch Times

sub-areas may be abstracted, and connected through Thévenin equivalents. It was shown [87] that this method is effective at parallelizing EMTP simulations and it could likely be expanded with the fractional latency technique to allow for different, non-integer local time steps between sub-areas, further increasing the performance.

The circuit of Figure 4.10 was implemented in Octave (an open-source Matlab derivative) both using a standard EMTP solution and the fractional latency cells. Since the use of fractional latency for this small system is inherently inefficient due to the overhead associated with finding the various Thévenin equivalents, compared to the circuit size, no timing results will be given.

4.4.2 EMTP Results

The circuit, in the EMTP comparison case, was computed with a time step of 35 μs while the fractional latency case was computed with various time steps, as given Table 4.2. The switching event was set at 8 ms, however, the switch will only open at the fist zero crossing of the current through the switch. This causes the switching moment to vary between simulation as a phase shift may cause the switch event to be delayed by a period of the natural frequency of the system, in addition to the limitations of the time step of the master area in which the switch resides. Since the magnitude of the resulting oscillation is strongly dependent on the voltage at the switching moment, this difference in switching time resulted in significant differences in amplitude.

A problem with the system being studied is that the required time steps of all the sub-areas are relatively close to one another. For the fractional latency to work, a sufficient ratio in time steps must exist. This is due to the interpolation used in the method, where it is assumed that one system varies much slower than another, and is thus essentially a stiff system during the interpolated interval. In Chapter 3, it was shown that for the interpolated line model, a ratio of about ten times is required.

This is the case in Figure 4.11. The simulation operates both sub-areas at a much smaller simulation time step ration of about $\frac{1}{10}$. The results are rather good, except of course the
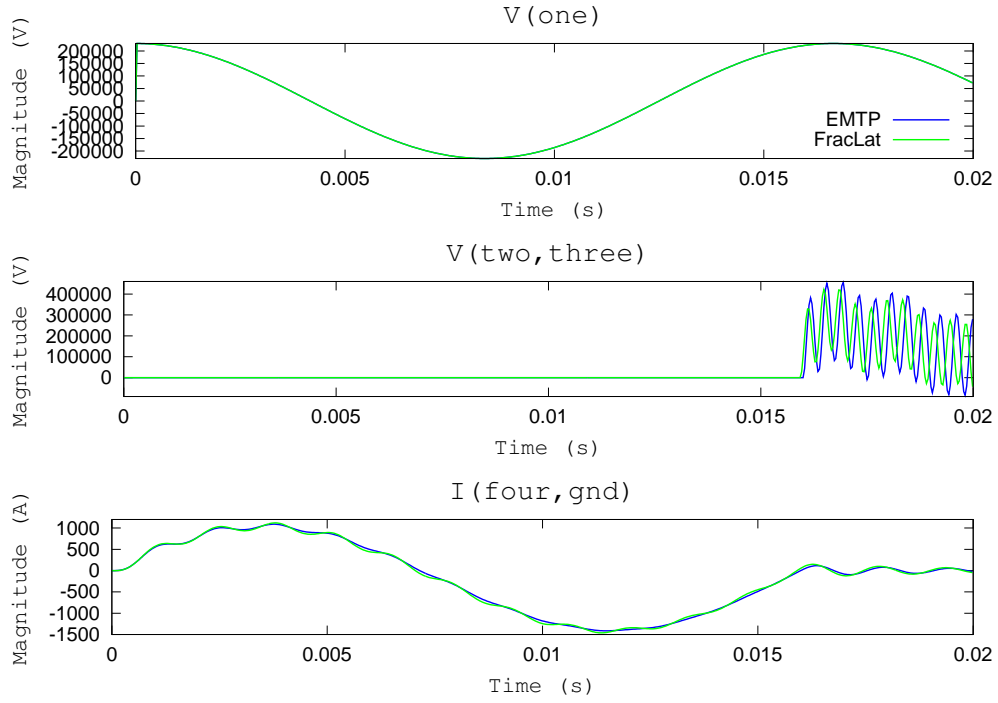
Figure 4.11: Fractional Latency EMTP Solution Stable Case 1
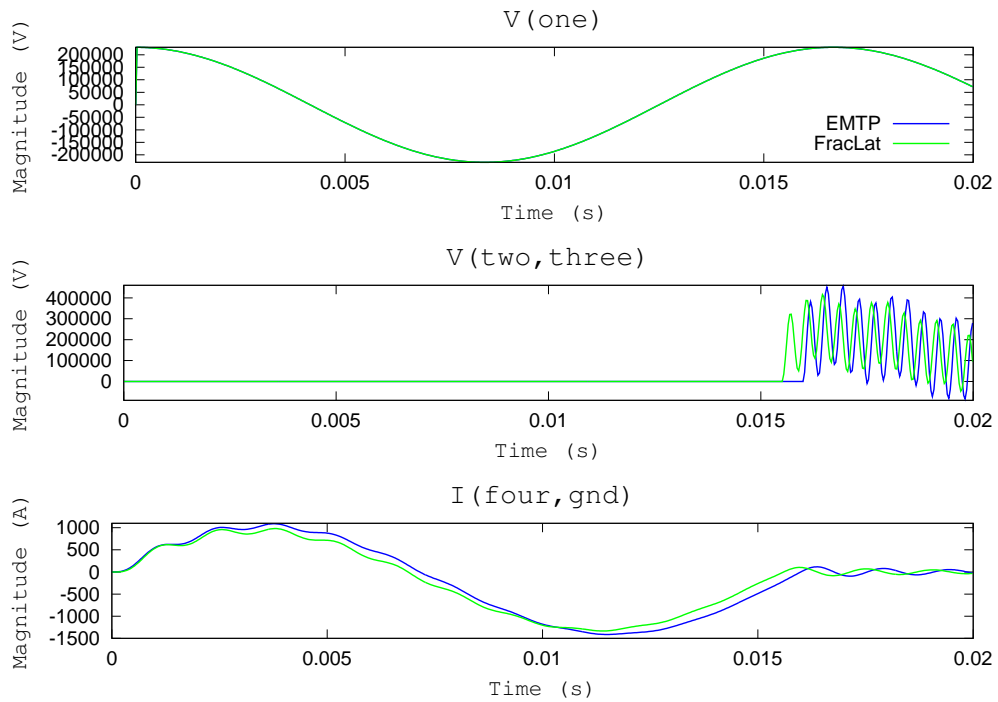
Figure 4.12: Fractional Latency EMTP Solution Stable Case 2
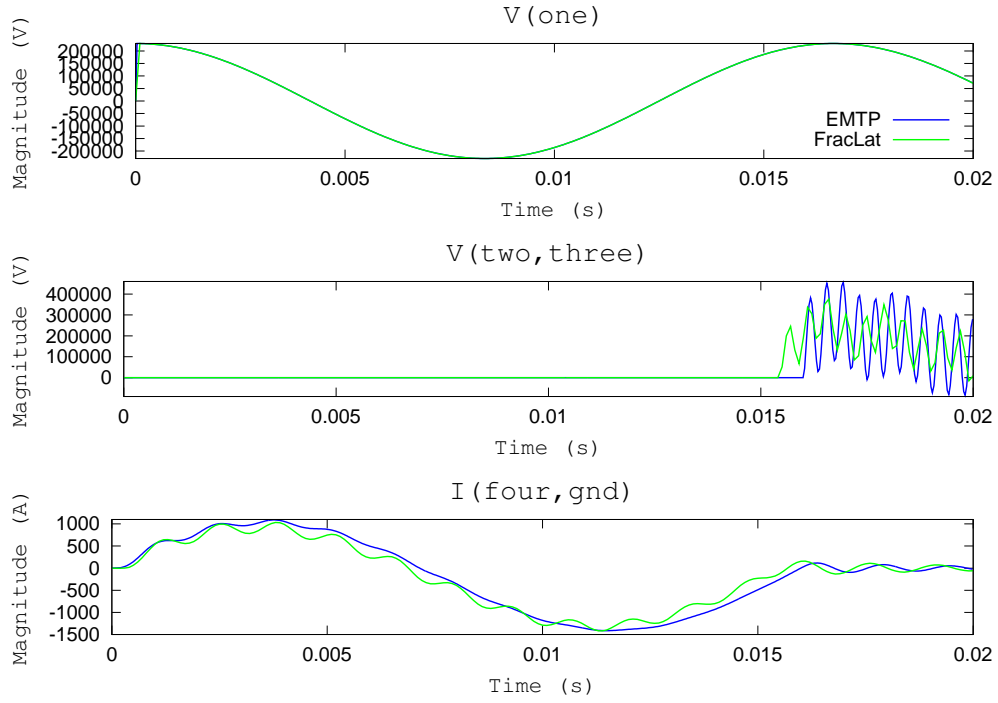
Figure 4.13: Fractional Latency EMTP Solution Stable Case 3
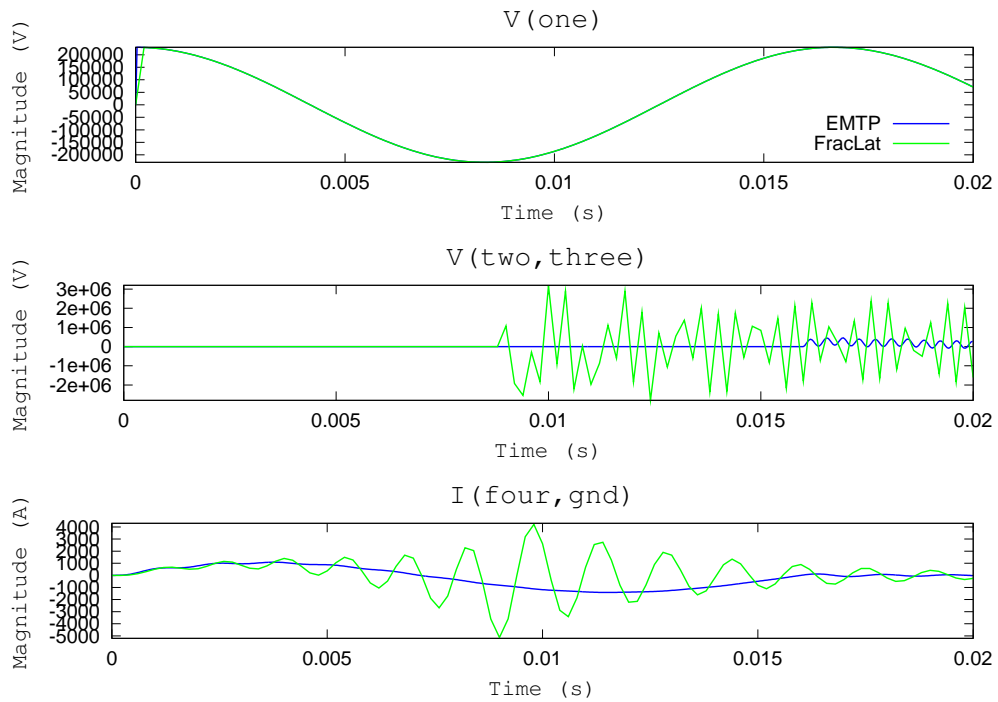
Figure 4.14: Fractional Latency EMTP Solution Unstable Case

increased computational time. The studied case did not have enough natural difference in the natural frequencies of each area, and thus resulting simulation time step, to make fractional latency work in a reasonable way. The ratio was thus artificially increased.

In the next cases, Figure 4.12, the base latency area and sub-ares were operated at about the same rate as the EMTP solution, thus about 35 μs. The results are not as good, but still follow the general shape of the expected result. The time step ratio between the base and sub-areas was thus almost unity.

In Figure 4.13, the base area and latency area 1 were slowed-down to roughly $\frac{1}{3}$ of the EMTP case rate, while the fast oscillating area was left to operate at a faster rate. The results degrade further, mainly due to the shifting of the switching event, but are otherwise comparable to the previous cases. Here, the fractional latency is being used as intended, having one part run at a faster rate than the others.

The last case, in Figure 4.14 operated with latency area and main area time steps that were too low for the system being studied and not only are the results wrong, they are also unstable as the extrapolation is pushed too far to maintain stability. With the used sampling parameters, the Nyquist bandwidth for the slowest area, the base area, is 2.5 kHz. This is sufficient for the 60 Hz source embedded in it. The fast oscillation has a frequency of about 2.2 kHz, and takes place in a sub-area with a Nyquist frequency of 14 kHz. Thus, the local time steps are according to the EMTP guidelines.

The instability is likely due to the system being too close to the stability limit imposed by the extrapolation. Since extrapolated solutions usually have stability limits, the Courant limits, much lower than the Nyquist limits [83], this is not unexpected. In practice, one would choose the areas guided by an analysis of the eigenvalues, to insure that the local simulation time steps are sufficient for the simulation [60], and then choose the time steps so the Courant condition for the method is met.

Table 4.2 summarizes the used simulation parameters for all the figures, as well as the actual times at which the switch opens for each simulation.

All in all, the extrapolation performs reasonably well, but the effects of the latency matching cells, due to the lack of sufficient oversampling between sub-areas, shows as a significant phase shift in the signals. In order to use the fractional latency cell in EMTP-style solutions, as described in this chapter, a stability analysis of the method would need to be performed and the Courant condition determined.

However, from the above, it is clear that the fractional latency cell concept is not a good one for EMTP simulations. The concept was not developed for it and the application to EMTP was only done to explore the concept and illustrate the issues with extrapolation that happen with the method in non-decoupled methods. Fractional latency could likely be used in EMTP,

but it should be applied to the history sources themselves, by modifying the traditional latency formulation. Alternatively, it could be integrated in the MATE concept [49, 87], where a system is split in sub-areas, communicated through Thévenin equivalents. We leave these issues as future work, as they are not the focus of this thesis.

4.4.3   Fractional Latency in TINA

The fractional latency method was developed for use in TINA, or other time-decoupled methods, not a regular EMTP. Hence, the choice of formulation as a universal, stand-alone cell that can be readily and transparently introduced in a discretized-space solution grid.

The issues with extrapolation in the EMTP implementation of fractional latency can be readily overcome within the TINA framework. As was stated earlier in this chapter, when transmission lines are present, the ability to "look ahead" becomes available without extrapolation, and fractional latency can be implemented directly. Since each and every cell in TINA is connected through transmission lines, full time-decoupling is in effect and transmission line "look ahead" can be readily implemented. In the next section, the equations are derived for the 1D case, which is the basic building block for the higher-dimensional meshes used in TINA.

4.4.4   Avoiding Extrapolation with Transmission-Line Decoupling

To illustrate how the prediction in fractional latency may be avoided, we will now show how to add fractional latency to an ideal line segment. In TINA, every part of the system is coupled though such transmission lines, which has the benefit of both time decoupling as well as a readily available Thévenin equivalent to work with. The resulting expressions are similar to the interpolated transmission line, except that they now feature prediction, which will be overcome through transmission-line look-ahead.

Since any number of ideal line segments can be cascaded, resulting in a single segment of transmission time equal to the whole, the following argument also holds for a one-dimensional section of TINA mesh connected to two neighboring areas, where both neighboring domains have equal, but different local time step than the sub-area the lines are in.

Conceptually, the ideal line consists of two Thévenin equivalents, with the sources given by the the model equations and the resistance equal to the line impedance. The fractional latency cell is essentially a Thévenin conversion system, that takes the equivalent on one end and presents an interpolated equivalent on the other. The resemblance between the circuits is coincidental. Combining both systems yields the interpolated line. We will apply the fractional latency conversion to the model, resulting in the circuit of Figure 4.15.
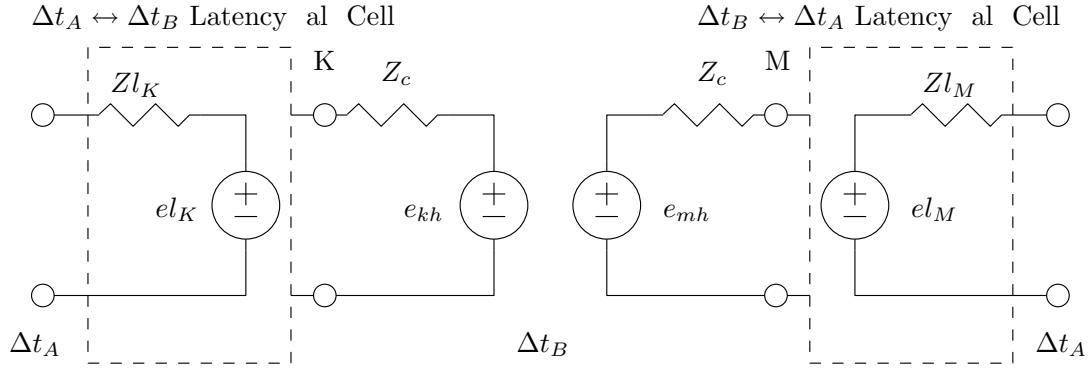
Figure 4.15: Fractional Latency Applied to a Transmission Line Section

Note that in this figure, only fractional latency half cells are shown. This was done to clarify the derivations. In the shown figure and derivation, only the matching of the line models Thévenins is performed. The second part of the system, matching the external system Thévenins to the line model latency area is not shown. These second halfs are required to find the various voltages and currents at the terminals of the line model and solve the system. In Figure 4.16, the whole system, comprising both half-cells for each size is shown.

Note that there is no coupling between both Thévenin equivalents in the full fractional latency cell (Figure 4.5). Rather, each Thévenin uses the Thévenin presented by the system neighboring on one side to find the values for the Thévenin it presents on the other side. The fact that the full latency cell circuit looks similar to the transmission line model used in this thesis is coincidence and not functional.

Starting from the ideal line equations (A.9) for the Thévenin history sources, from Appendix A:

$$e_{kh}(t) = v_m(t - \tau) + Z_c i_m(t - \tau) \tag{4.3a}$$

$$e_{mh}(t) = v_k(t - \tau) + Z_c i_k(t - \tau) \tag{4.3b}$$

The relevant fractional latency equations (4.1) are:

$$el_K(t_A) = eth_B(t_B + \Delta t_B) + R_K \left[eth_B(t_B) - eth_B(t_B + \Delta t_B) \right] \tag{4.4a}$$

$$el_M(t_B) = eth_A(t_A + \Delta t_A) + R_M \left[eth_A(t_A) - eth_A(t_A + \Delta t_A) \right] \tag{4.4b}$$

Now, applying the fractional latency equations (4.1) to both end of the line, with constant

impedances and using the A as the external area time step and B as the internal area time step:

$$el_K(t_A) = e_{kh}(t_B + \Delta t_B) + R_K [e_{kh}(t_B) - e_{kh}(t_B + \Delta t_B)] \tag{4.5a}$$

$$el_M(t_A) = e_{mh}(t_B + \Delta t_B) + R_K [e_{mh}(t_B) - e_{mh}(t_B + \Delta t_B)] \tag{4.5b}$$

Substituting the line equations and simplifying, with the local absolute time given by t_B:

$$el_K(t_A) = v_m(t_B - \tau + \Delta t_B) \tag{4.6a}$$
$$+Z_c i_m(t_B - \tau + \Delta t_B)$$
$$+R_K [v_m(t_B - \tau) - v_m(t_B - \tau + \Delta t_B)]$$
$$+Z_c R_K [i_m(t_B - \tau) - i_m(t_B - \tau + \Delta t_b)]$$
$$el_M(t_A) = v_k(t_B - \tau + \Delta t_B) \tag{4.6b}$$
$$+Z_c i_k(t_B - \tau + \Delta t_B)$$
$$+R_K [v_k(t_B - \tau) - v_k(t_B - \tau + \Delta t_B)]$$
$$+Z_c R_K [i_k(t_B - \tau) - i_k(t_B - \tau + \Delta t_b)]$$
$$\tag{4.6c}$$

The resulting equations have the same form as the interpolated line model, except that the LHS part is shifted one step ahead compared to the interpolated line model. This is where the extrapolation comes in. Also, the τ value used here is the actual length of the line, not a rounded-up integer multiple of the local time step as with the interpolated lines. However, in practice when using fractional latency, the local time step will usually be an integer multiple of the transmission time of the line. Thus, we can think of the transmission line with fractional latency as an interpolated line with the required history values shifted one step in the future, thus requiring the current and future values to compute a solution, as opposed to current and past values, as in the interpolated line model.

This is where the extrapolation, the requirement for the future value, may be avoided. The transmission line history buffers, as the model is implemented in this thesis, Appendix A, already store the $(t_B - \tau + \Delta t_B)$ values, as these are the values to be used at the next time step, and are thus simply waiting for use in the buffer. These are exact values, computed by the ideal line model. Since they are in transit, nothing else will happen to them, except the time delay in the model. We thus have access to exact, future values already in memory which may be used in the latency equations to match the different Δt domains. This is the advantage of time-decoupled methods. Since each area is decoupled, they may be solved independently from one another and we can obtain the area solution, not just the history sources, for the next step.

This solution may then be used with fractional latency to remove the need for extrapolation. Transmission-line decoupling is a classic example of a time-decoupled method, and at the basis of both the TLM and TINA methods. TINA has the advantage that the formulation is entirely based on Thévenin equivalents, which makes interfacing the line models and the fractional latency cells trivial.

4.4.5 Integration of the Fractional Latency Cells in the TINA Mesh

The latency cell, as described in the EMTP example, is used again and split in two uni-directional half-cells. This time, however, they will gain the ability to request the next value from the connecting line segments. In practice, this construction is shown in Figure 4.16, for a 2D example where the central cell is surrounded by four different latency domains, and must thus use the uni-directional fractional latency cells to interface with those neighboring regions. Using the latency cells, which are modeled like any other cell in TINA and present Thévenin equivalents to the outside world, they can be inserted into the grid at will and allow easy separation of a simulation into sub-areas.

The choice of using half-cells was done for purely practical reasons regarding the way the individual cells in TINA were implemented. Each cell in the simulation uses a pointer to the neighboring cell's Thévenin data. It uses this pointer to access that Thévenin equivalent from each neighboring cell in order to solve itself. Since this neighbor sits in a different latency area, its results must be converted to the local required time. This is done through the latency cell. It is inserted in the simulation as a 1D block (in the drawing indicated by the W-E and N-S blocks) and connected in the mesh.

The insertion process consists of initializing a uni-directional latency cell, obtaining the address the cell to be solved was pointing to, connecting the latency cell to that address, and configuring the cell to be solved to point to a new address, that of the output side of the fractional latency cell. The half-cell for the other data flow direction will be inserted when the neighbor cell is processed.

For a given latency area, the latency half-cells are solved before the mesh cells, and thus may obtain properly interpolated Thévenin equivalents from the neighbors that the connected mesh cell may then use to compute itself. By splitting in half-cells, the uni-directional nature of the Thévenin information exchanges is exploited and a symmetrical distribution of the half cells across an area barrier can be used. This simplifies the programming of the system. The latency cell requests its required data from the neighbor, interpolates it, and presents a new Thévenin to the internal cell. Later, when the actual mesh cells are solved, it will pull its data from the latency cell it is connected to. Thus, inserting the latency cells merely becomes a matter of
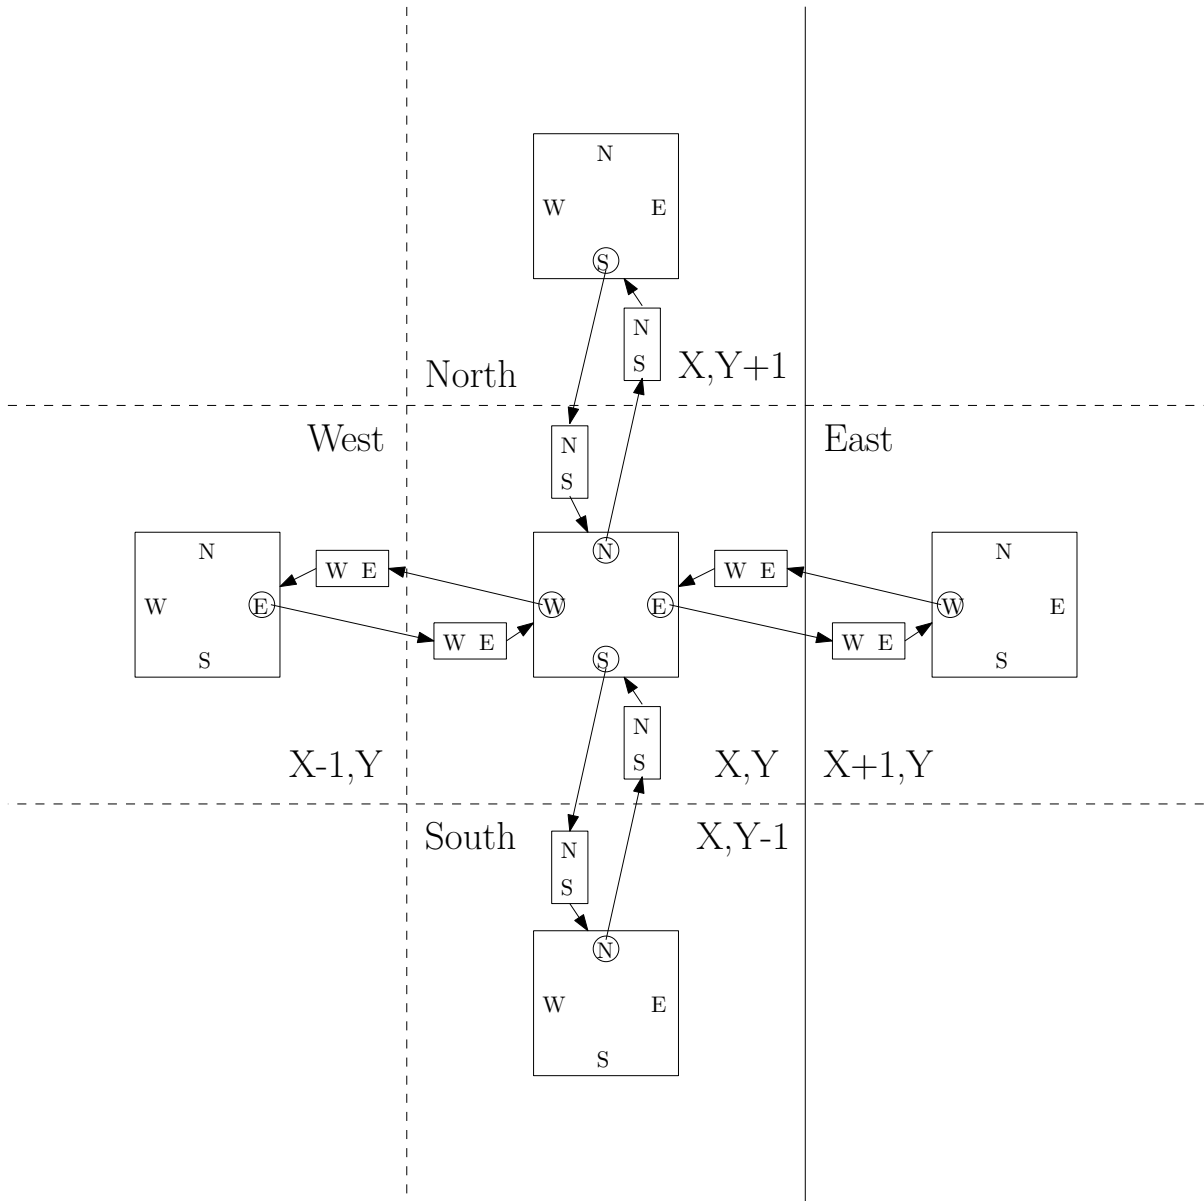
Figure 4.16: Fractional Latency Cells in a TINA Mesh

moving data connectivity pointers around and connecting the latency cell appropriately.

4.4.6   Simulation Parameters in TINA for Fractional Latency

To evaluate the error incurred by fractional latency, it would be of interest to lower the ratio of the time step between the latency areas and the master time step base area. This would increase the error in the simulation due to the interpolation used in the latency, likely in a similar way as studied in Chapter 3. Doing so, we could establish a qualitative limit for the interpolation time step ratios between the base area running at the master time step and the other areas in the system.

However, this proves impossible in the TINA implementation. We can increase the latency ratio, making the already huge sub-area bandwidths even larger compared to the usable simulation bandwidth given by the slower areas. But, we cannot reduce it below the spatial discretization limit. Or: if the system is discretized with, e.g.: $10\frac{cells}{\lambda}$, the interpolation and fractional latency ratios cannot be lower than 10 without also adjusting the spatial discretization factor.

The only way to reduce the fractional latency time step ratio is to also reduce the spatial discretization. This, however, also influences the accuracy of the simulation and makes it impossible to study the parameters independent from one another.

In the end, the situation may be explained from the fact that, when using fractional latency, we enforce that the length of each transmission-line segment must be an integer number of the traveling time τ and that the local sub-area time step is an integer multiple N of τ, thus: $\Delta t = \frac{\tau}{N}$. Doing so avoids the use of interpolation for that line segment and increases both precision and execution speed. This is the main reason for using the more complex fractional latency versus the regular latency. Thus, the shortest a line segment may be is one Δt, thus one memory location ($N = 1$). This puts a hard limit on the lowest simulation time step possible in a latency area, which is thus so that $\Delta t_{max} = \tau$.

The traveling time τ itself is given by the line parameters, which yield the wave speed a, and the spatial discretization of the problem, which determine the size of the cell, and thus the physical length of the line segments used therein. Thus: $\tau = \frac{\Delta l}{2a}$, noting that TINA uses half-line segments with a node in the cell center.

The spatial discretization Δl is now chosen so that the physical extent of the wave is sufficiently sampled. It is a spatial Nyquist criterion, and given by the shortest wave length in the latency sub-area. Typically, one samples so that $\Delta l = \frac{\lambda}{SD}$, where SD is the spatial discretization and typically chosen as 10 cells per wavelength. The shortest wave length λ_{min} itself is given by the wave speed in the local medium a, as well as the highest frequency in

the simulation, BW as: $\lambda_{min} = \frac{a}{BW}$. In TINA, the bandwidth BW is defined by the user for this purpose, and is equal to the desired Nyquist frequency of the computation, and the highest occurring frequency.

Thus, combining the above considerations together, the maximum time step usable in a latency area, Δt_{max}, may be expressed as:

$$\Delta t_{max} = \frac{\Delta l}{2a} \tag{4.7}$$

The spatial discretization used in the simulation $\Delta l_{simulation}$, then, is a single global parameter for a regular grid in space that must be less than the following limitation, where the wave speed is that of the area with the slowest wave speed in the simulation a_{min}:

$$\Delta l_{simulation} = \frac{a_{min}}{SD\ BW} \tag{4.8}$$

The spatial discretization is a constant for the whole simulation, as a regular grid in space is used in the TINA implementation discussed in this thesis. This means that this global parameter is determined by that sub-area with the slowest wave speed, thus shortest wave length for the given bandwidth. This will be the master time step area, or base area, since it must operate at the lowest simulation time step as the output generation is synchronized to it. From the above equations, it follows that the area with the slowest wave speed is also that with the largest simulation time step, thus is solved at the slowest rate, and has the longest traveling times for its line segments.

In a TINA simulation, SD, BW, and Δl are global parameters chosen by the user and verified by the program to be in compliance with all areas in the simulation. The program will also use these parameters to put limitations on the simulation time steps of each latency area. Thus, to return to the original argument in this section, of why the latency sub-area ratios could not be reduced below the spatial discretization, we may now discover why.

First, the simulation base area operates at the lowest simulation time step in the system. This was discussed in Section 4.3.4. From the above, it also follows that it is linked to the material with the slowest wave speed. When now regarding a latency sub-area, this must by definition operate at a faster rate and have a higher wave speed. The rate between the local time steps in these areas is the parameter of interest.

In the TINA implementation in the thesis, with the base area time step $\Delta t_{base} = \frac{1}{2BW}$, the base area cannot have any transmission lines in it, as these will automatically be spatially under sampled. We find, in fact, that if we were to define a medium in this area composed of transmission lines with $\tau = \Delta t_{base}$, the spatial sampling ratio is unity.

Thus, the time step ratio between the base area and a sub-area which has lines must be

at least the spatial discretization factor, further compounded by the requirements posed by the traveling time of the lines. Mathematically, the criterion for the minimum time step ratio between the base area and a sub-area becomes:

$$\frac{\Delta t_{latency\ area}}{\Delta t_{base}} = SD\frac{a_{latency\ area}}{BW\Delta l} \tag{4.9}$$

It must be noted that, in the TINA implementation, each spatial cell is split in two half-lines. Thus, the Δl parameter used in this discussion has to be halved to find the values used in the simulation. Also, keep in mind that the 2D and 3D grid require an adjustment factor for the line parameters, which is the same as for TLM [18], and a result of the field to circuit mapping, as discussed in Chapter 2.

In the above equation, the wave speed of the area under investigation $a_{latency\ area}$ may not be slower than the slowest wave speed in the simulation a_{min}. Otherwise, the spatial discretization needs to be adjusted, and thus the independence of these variables is lost, making the evaluation of the time step ratio between the simulation base area and the latency areas difficult.

When both wave speeds are equal, the above expression may be reduced to the following limiting condition:

$$\frac{\Delta t_{latency\ area}}{\Delta t_{base}} = SD^2 \tag{4.10}$$

We may also see that the minimum time step ratio between the base area and the various sub areas is similar, in practice, to that found in the error criterion when interpolation is used in TINA, as discussed in Section 2.2.3. With an oversampling of 10 for both the spatial and temporal factors, we found a time step ratio of 100. Here, a spatial discretization of 10 squared also yields 100 as a minimum ratio between the time steps as an accuracy condition.

4.4.7 TINA Results

To evaluate the use of fractional latency in the TINA method, three cases will be studied for both accuracy and numerical performance. These cases will be computed using highly over-sampled non-interpolated and interpolated lines, to serve as a base reference, as well as different time step ratios of fractional latency. The non-interpolated case is the reference, but has to operate with slightly adjusted parameters for some of the used materials, to insure the various time steps match.

The physical size of the microphone itself was also taken into account, which is a sizable portion of the wave length for the higher frequencies evaluated. Thus, in the simulation, the acoustic pressure over an area comparable to that of the physical microphone was averaged to obtain the result for that location.
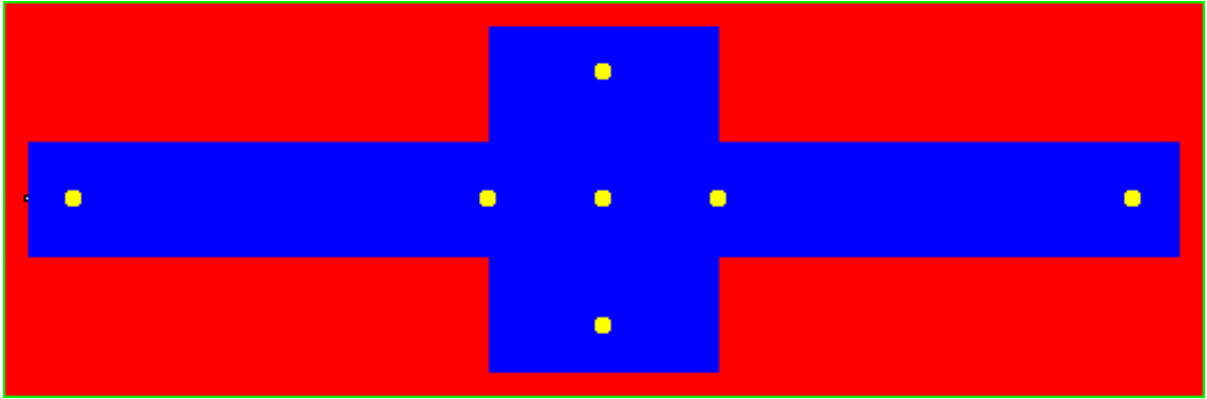
(a) Full Boundary
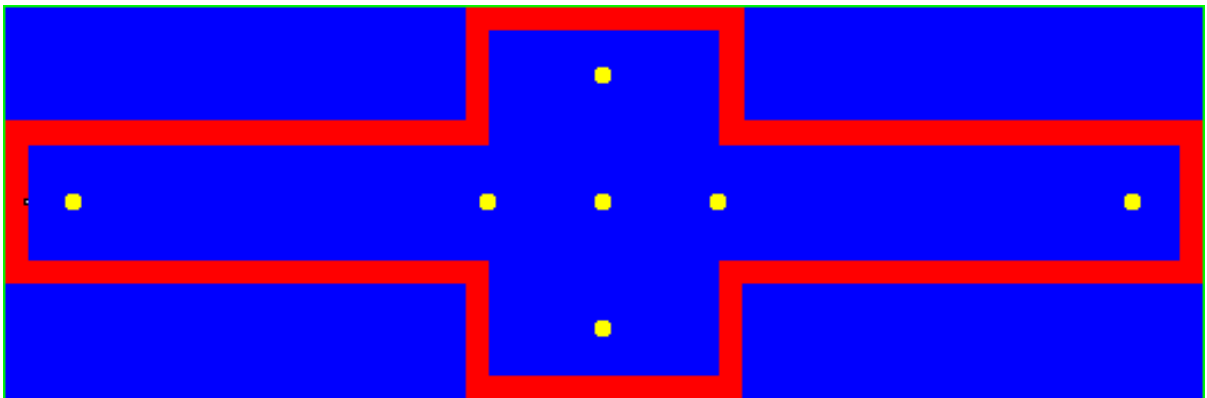
(b) Medium Boundary
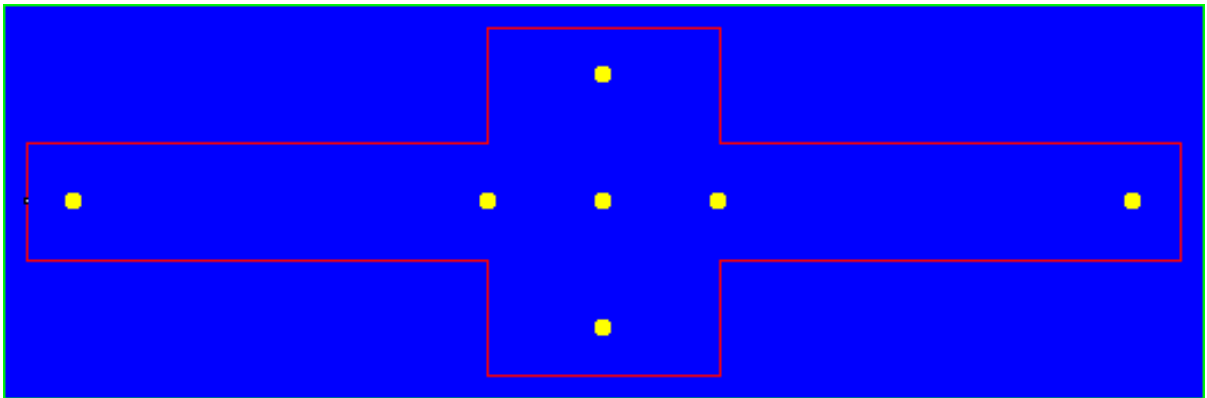
(c) Thin Boundary

Figure 4.17: Simulation Cases
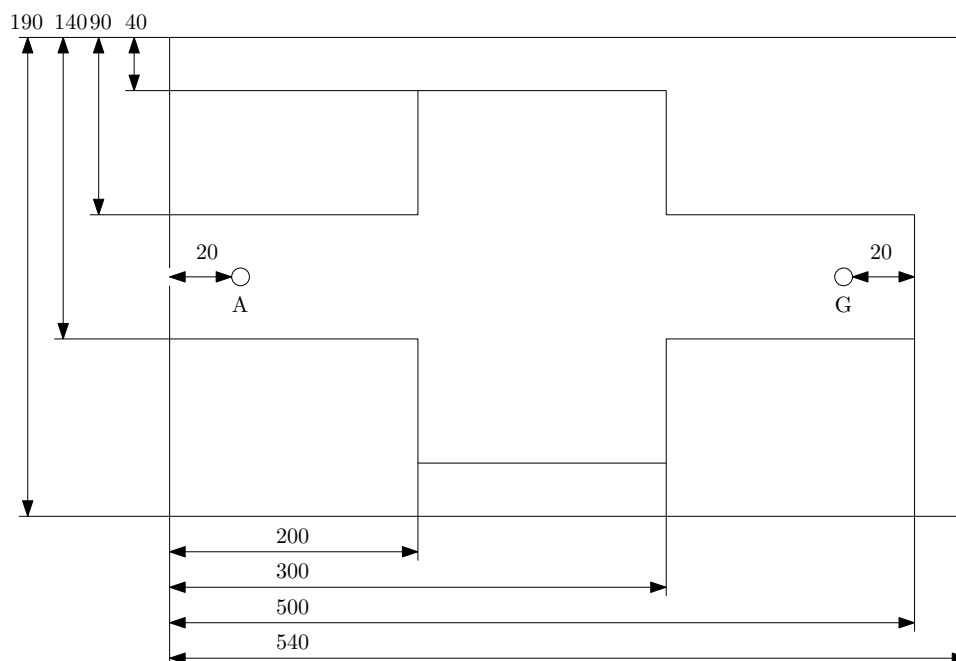
Figure 4.18: Expansion Duct Set-Up

The used 2D configurations are shown in Figure 4.17 and the setup itself in Figure 4.18. The spatial discretization was 1 mm. These are comparable cases, but with different widths for the duct border. This was done to show the change in efficiency of both the interpolated and the fractional latency technique with different ratios of fast and slow materials in the system. The case is presented in full detail in Chapter 5, where a physical realization is compared to simulation. As described there, the utmost left yellow dot is the "A" microphone position and the utmost right yellow dot is the "G" microphone position.

The adjustments were made in the dense boundaries (red material, aluminum). This leaves the main medium (blue material, air), where the measurements are made, at the same parameters across all simulations. The error incurred is thus limited, as the dense boundary has five orders of magnitude higher impedance compared to air, and is thus nearly an infinitely dense boundary (open circuit). The reflection coefficient between air and the boundary thus doesn't change much and little energy can actually enter the boundary material and propagate in it. Thus, the change in its parameters has negligible effect on the simulation results, which are obtained in the air medium only. This is confirmed in the plots, where the interpolated and non-interpolated (normal) cases overlap, as well as in Chapter 2, where this is discussed in more detail. The use of significant oversampling for these reference cases aids in this agreement. It must be noted that the used oversampling is a result of previously established error criteria, and a valid simulation bandwidth of 25 kHz was maintained for each case.

The differences between the precise and adjusted materials are shown in Table 4.3, where

127

	Air	Al (correct)	Al (adjusted)
ρ	$1.1198 \times 10^3 \frac{g}{m^3}$	$2.6989 \times 10^6 \frac{g}{m^3}$	$2.6989 \times 10^6 \frac{g}{m^3}$
k	$6.99746 \times 10^{-9} \frac{ms^2}{g}$	$1.45138 \times 10^{-14} \frac{ms^2}{g}$	$1.584727 \times 10^{-14} \frac{ms^2}{g}$
a	$3.45383 \times 10^2 \frac{m}{s}$	$5.05261 \times 10^3 \frac{m}{s}$	$4.83537 \times 10^3 \frac{m}{s}$
z	$4.13769 \times 10^5 \frac{g}{sm^2}$	$1.36365 \times 10^{10} \frac{g}{sm^2}$	$1.30502 \times 10^{10} \frac{g}{sm^2}$

Table 4.3: Materials Used in the TINA Simulations

ρ is the compressibility, k the compressibility, a the wave speed, and z the impedance. These parameters are acoustic ones, but relate directly to their electrical equivalents of permeability and permittivity.

Comparison of Simulation Parameters and Results

The three simulation cases, repeated for each used computation method, are summarized below. The normal case, which is the non-interpolated simulation using the adjusted material parameters from Table 4.3, served as the base line for the comparisons. In all cases, the normal case case was used as the comparison base line.

The frequency-domain averaged error was computed as a basic error criterion to help assess the simulation approaches. It is a percentile number, computed by obtaining the magnitude error at each frequency point, at microphone location A, and compared to the normal case, summing these together, and normalizing to the number of spectral components.

Memory use was determined using the Linux "top" program, and the indication for data size was recorded. This is the total memory that the program uses for data, thus non-executable code, including the stack and any portion swapped to disk. This served as an indication of how much memory the simulator requires to solve a particular case.

All cases were executed on an Intel Core 2 Duo E8400 3 GHz system, running a stock Kbuntu 9.04 "Jaunty Jackalope" 64 bit Linux distribution on kernel 2.6.28-15. The solver operated in a single thread and was confined to one core. The average core usage was over 99 %, indicating nearly unique usage of that core for the TINA solver process.

Full Border

In this case, the border comprised roughly one half of the total simulation volume. This meant that the latency, which mainly applies on the air portion of the simulation, could only act on a comparatively limited portion of the simulation area. A speed gain of roughly two times was obtained, and a reduction in memory use of about 15 %, compared to the normal case. Compared to the interpolated cases, the performance was better due to the absence of per-cell

Full Border	Normal	Interpolated	Latency ratio 10	Latency ratio 100
# Mesh cells	88400	88400	88400	88400
Air/Aluminum cells ratio	2/1	2/1	2/1	2/1
Latency/Interpolation ratio	N/A	10	10	100
Spatial discretization air	14 Cells/λ	14 Cells/λ	14 Cells/λ	14 Cells/λ
Spatial discretization alu.	193 Cells/λ	202 Cells/λ	202 Cells/λ	202 Cells/λ
Δt air	7.31182e-8 s	6.99744e-8 s	1.02366e-6 s	1.70609e-7 s
Δt aluminum	7.31182e-8 s	6.99744e-8 s	6.99744e-8 s	6.99744e-8 s
Δt Simulation/Base	7.31182e-8 s	6.99744e-8 s	2e-5 s	2e-5 s
# Time steps base	1312942	1371931	4800	4800
Computation time	31732 s	36395 s	18714 s	24236 s
Effective cells/s	3.7e6	3.3e6	6.5e6 eqv.	5.0e6 eqv.
Memory use	59 MB	60 MB	50 MB	55 MB
% FD avg. mag. err. vs nor.	N/A	0.06 %	13 %	8.5 %
Speed-up vs normal	N/A	0.9 x	1.7 x	1.3 x
% Memory use vs normal	N/A	+ 2 %	- 15 %	- 7 %

Table 4.4: TINA Simulation Results and Parameters Full Border
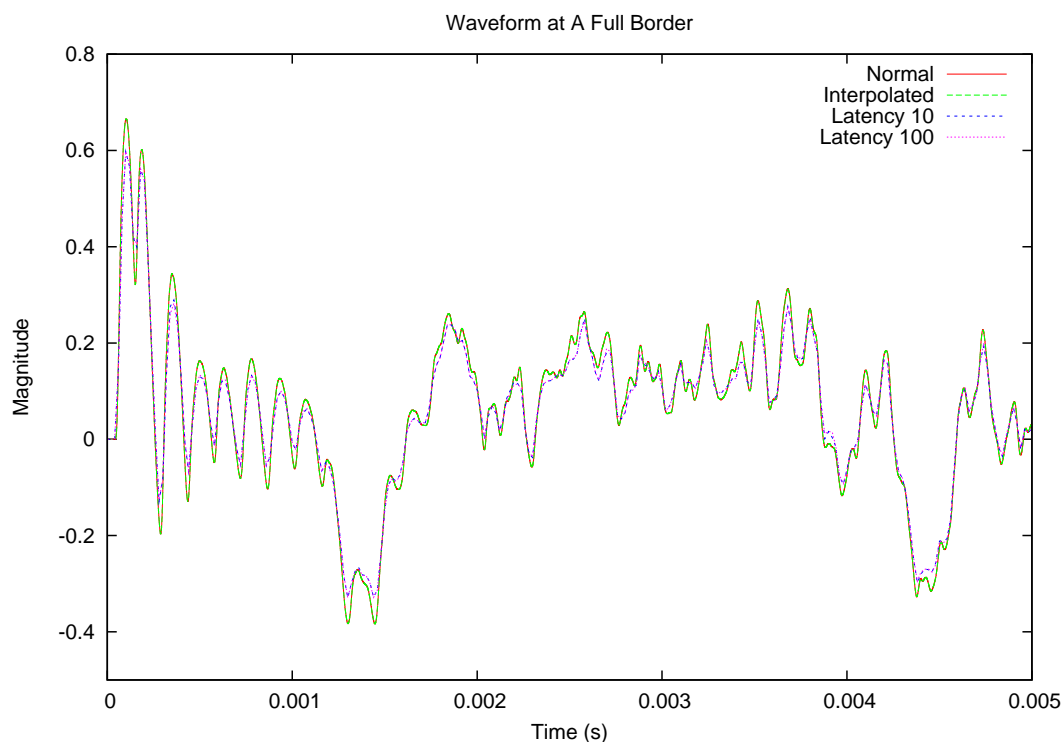
Figure 4.19: Full Border Time Domain Results at A

Transfer Function A->G Magnitude Full Border

(a) Magnitude

Transfer Function A->G Phase Full Border
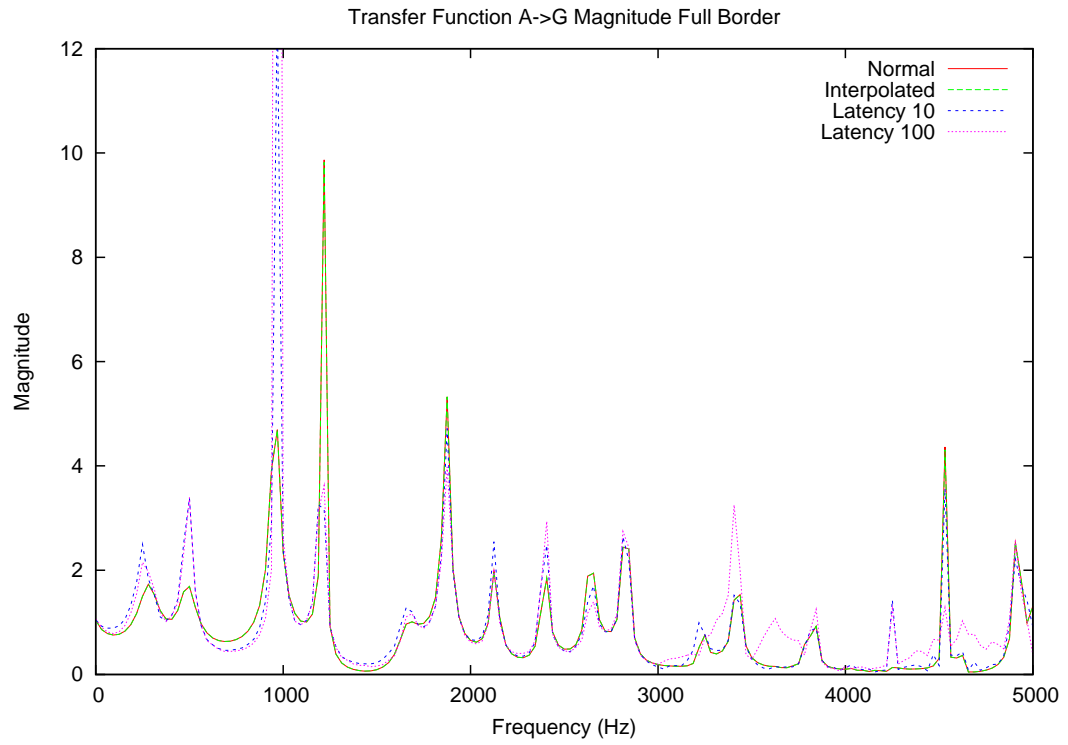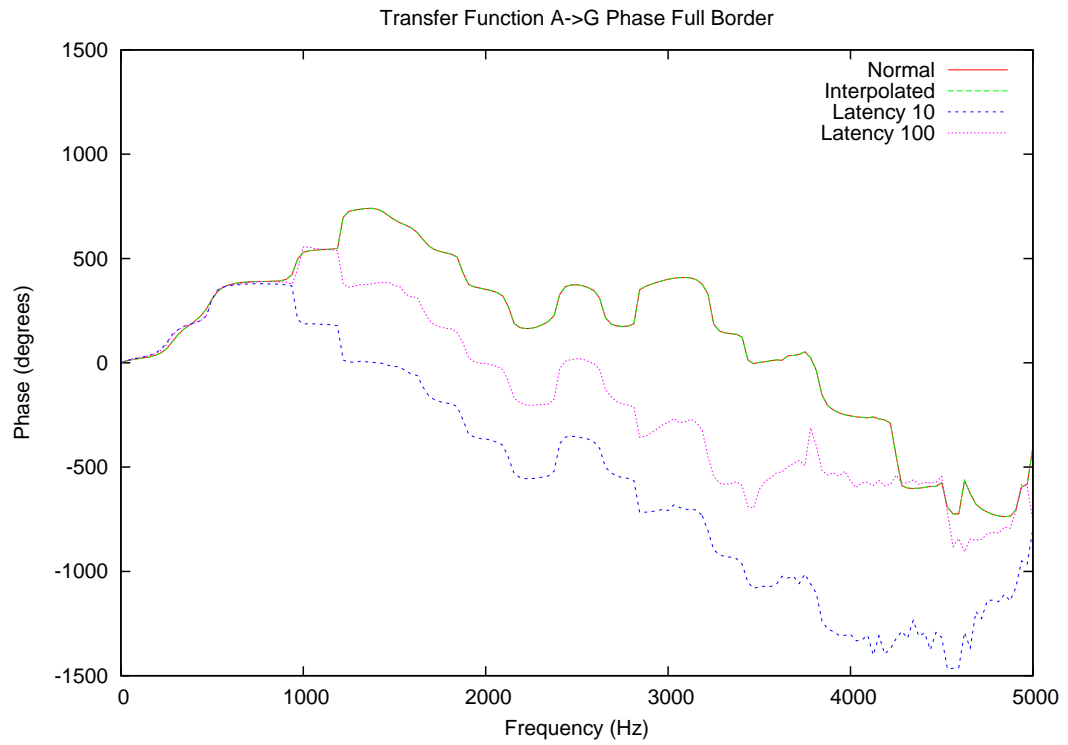
(b) Phase

Figure 4.20: Full Border Frequency Domain Transfer Function A → G

Medium Border	Normal	Interpolated	Latency ratio 10	Latency ratio 100
# Mesh cells	88400	88400	88400	88400
Air/Aluminum cells ratio	6/1	6/1	6/1	6/1
Latency/Interpolation ratio	N/A	10	10	100
Spatial discretization air	14 Cells/λ	14 Cells/λ	14 Cells/λ	14 Cells/λ
Spatial discretization alu.	193 Cells/λ	202 Cells/λ	202 Cells/λ	202 Cells/λ
Δt air	7.31182e-8 s	6.99744e-8 s	1.02366e-6 s	1.70609e-7 s
Δt aluminum	7.31182e-8 s	6.99744e-8 s	6.99744e-8 s	6.99744e-8 s
Δt Simulation/Base	7.31182e-8 s	6.99744e-8 s	2e-5 s	2e-5 s
# Time steps base	1312942	1371931	4800	4800
Computation time	35774 s	46628 s	6090 s	16947 s
Effective cells/s	3.2e6	2.6e6	19.9e6 eqv.	7.2e6 eqv.
Memory use	75 MB	77 MB	50 MB	62 MB
% FD avg. mag. err. vs nor.	N/A	0.06 %	7.0 %	1.5 %
Speed-up vs normal	N/A	0.8 x	5.9 x	2.1 x
Memory use vs normal	N/A	+ 3 %	- 33 %	- 17 %

Table 4.5: TINA Simulation Results and Parameters Medium Border
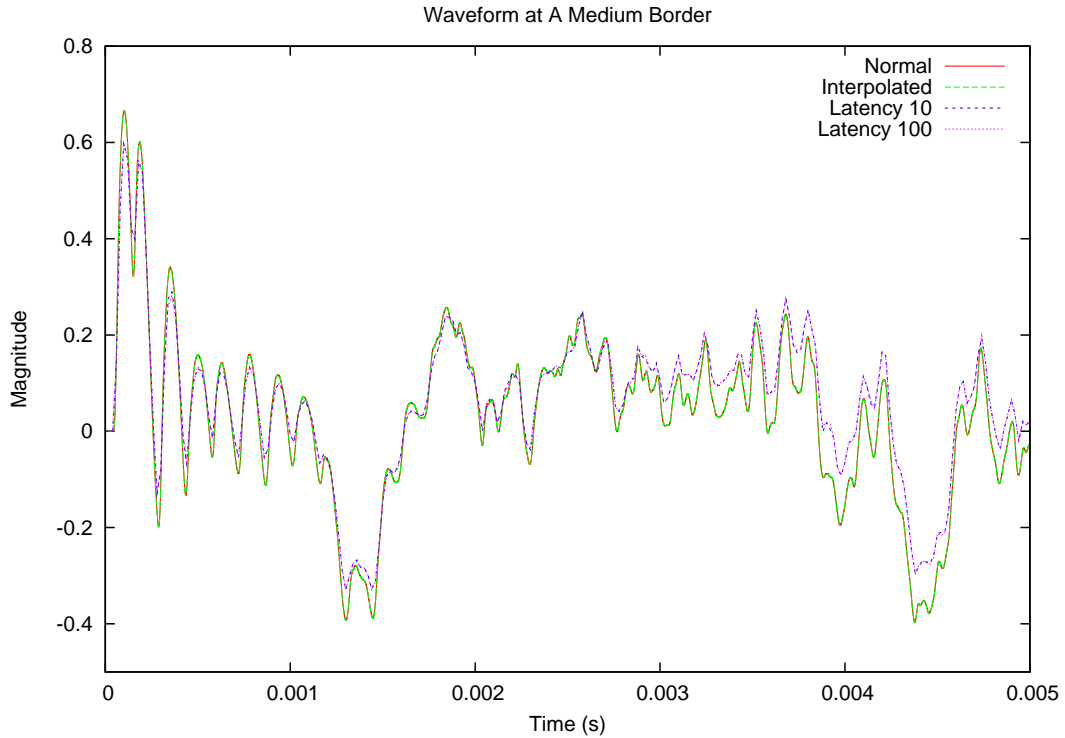
Figure 4.21: Medium Border Time Domain Results at A
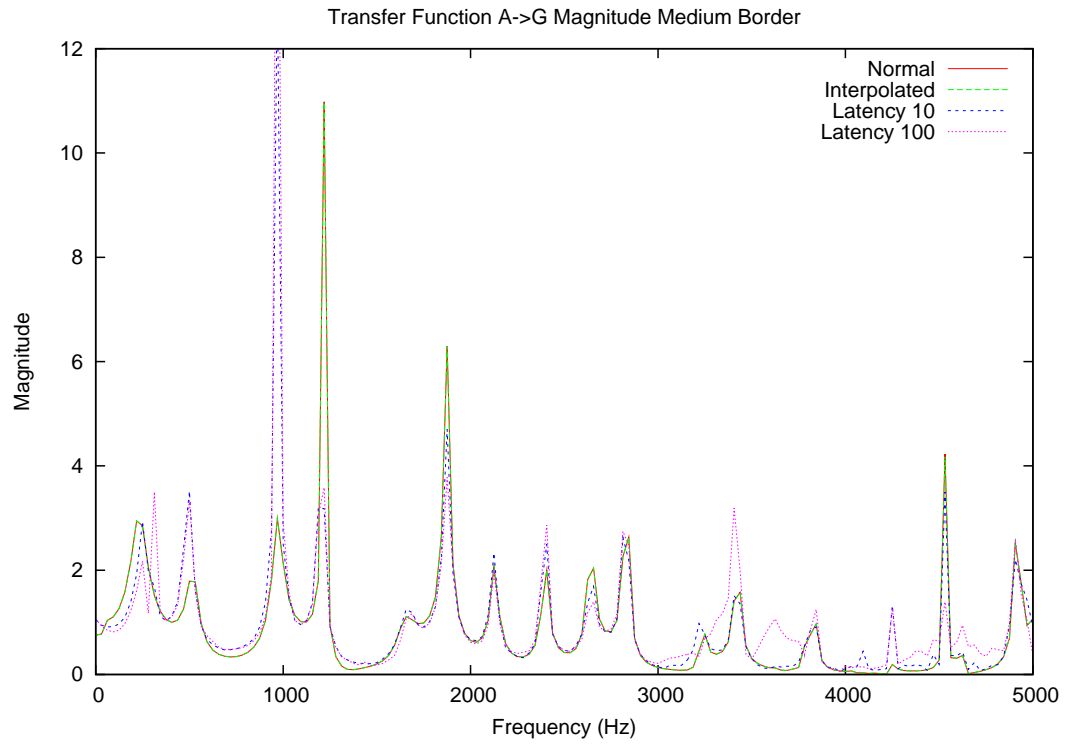
(a) Magnitude
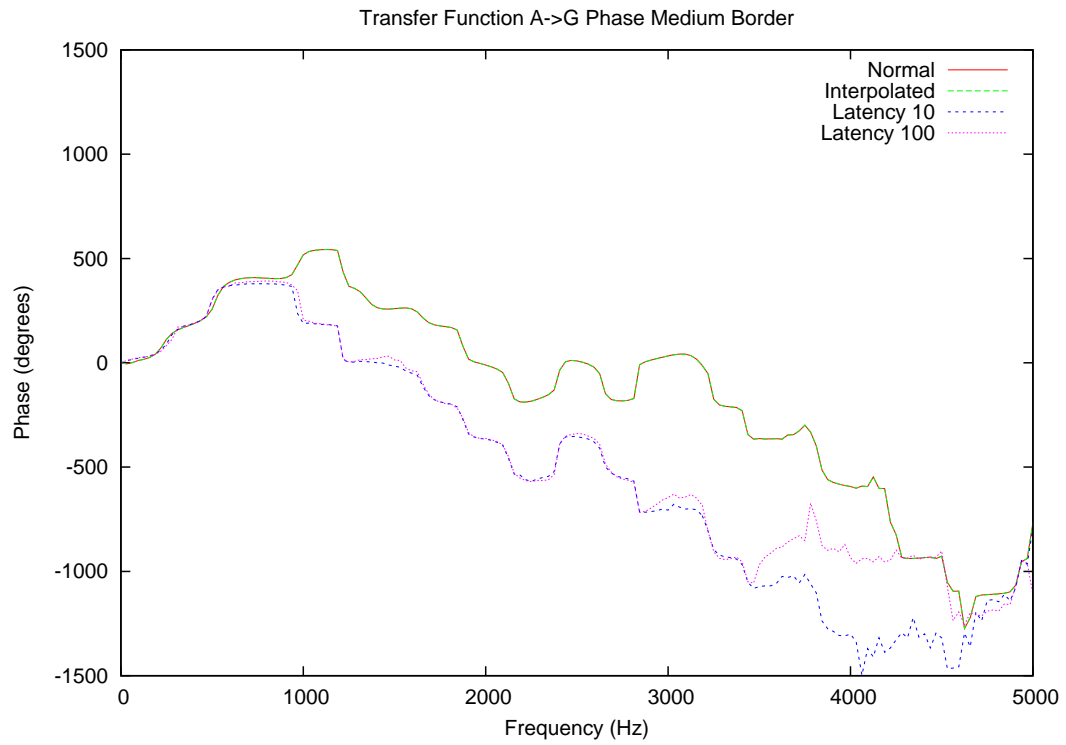
(b) Phase

Figure 4.22: Medium Border Frequency Domain Transfer Function A → G

Thin Border	Normal	Interpolated	Latency ratio 10	Latency ratio 100
# Mesh cells	88400	88400	88400	88400
Air/Aluminum cells ratio	60/1	60/1	60/1	60/1
Latency/Interpolation ratio	N/A	10	10	100
Spatial discretization air	14 Cells/λ	14 Cells/λ	14 Cells/λ	14 Cells/λ
Spatial discretization alu.	193 Cells/λ	202 Cells/λ	202 Cells/λ	202 Cells/λ
Δt air	1.02366e-6 s	1.02366e-6 s	1.02366e-6 s	1.70609e-7 s
Δt aluminum	7.31182e-8 s	6.99744e-8 s	6.99744e-8 s	6.99744e-8 s
Δt Simulation/Base	7.31182e-8 s	6.99744e-8 s	2e-5 s	2e-5 s
# Time steps base	1312942	1371931	4800	4800
Computation time	37872 s	49410 s	2379 s	14603 s
Effective cells/s	3.1e6	2.5e6	51e6 eqv.	8.3e6 eqv.
Memory use	79 MB	82 MB	50 MB	63 MB
% FD avg. mag. err. vs nor.	N/A	0.06	-435e6	- ∞
Speed-up vs normal	N/A	0.8 x	16 x	2.6 x
Memory use vs normal	N/A	+ 4 %	- 37 %	- 20 %

Table 4.6: TINA Simulation Results and Parameters Thin Border
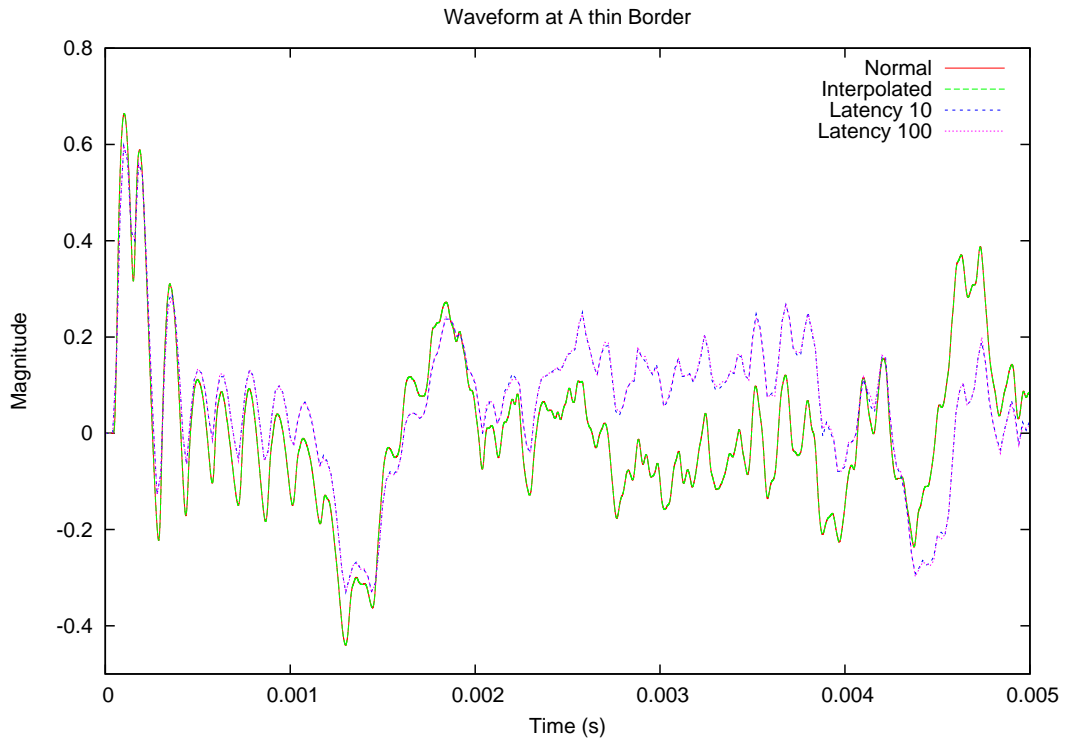
Figure 4.23: Thin Border Time Domain Results at A

Transfer Function A->G Magnitude Thin Border

(a) Magnitude

Transfer Function A->G Phase Thin Border
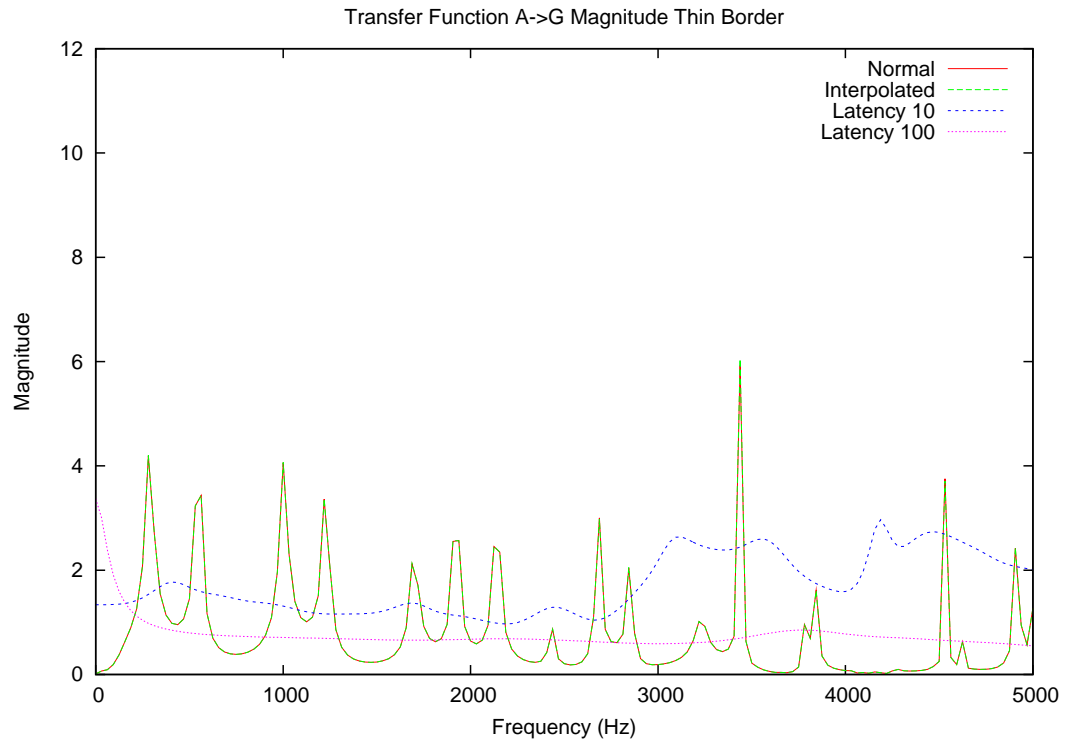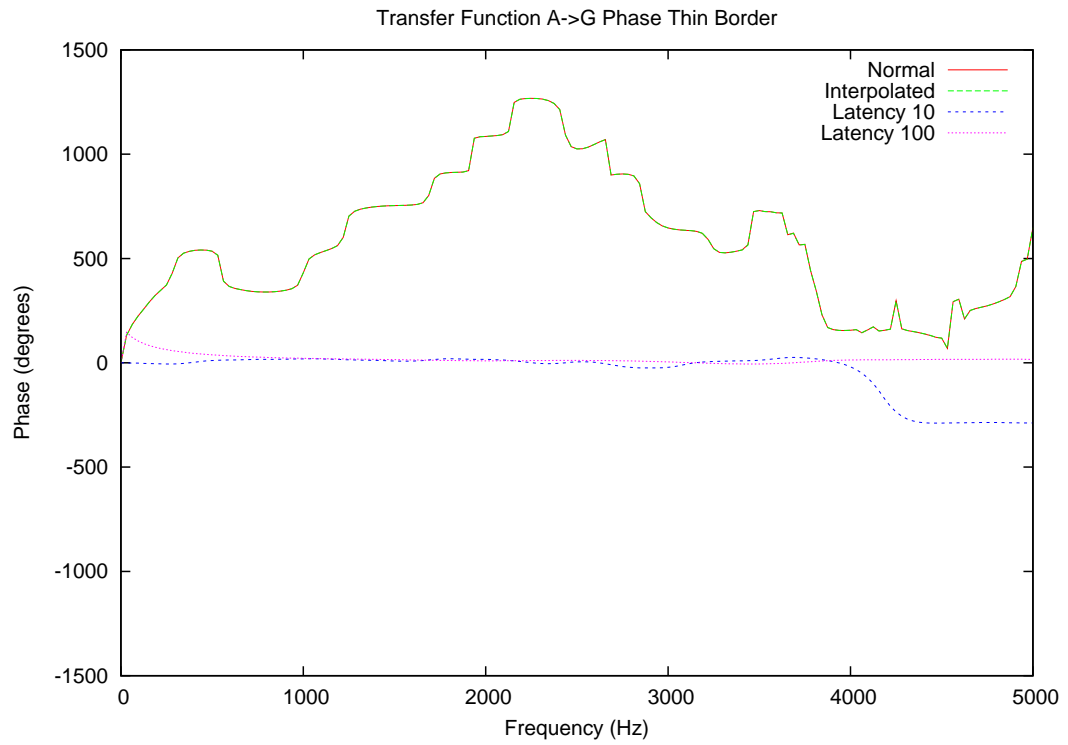
(b) Phase

Figure 4.24: Thin Border Frequency Domain Transfer Function A → G

interpolation and reduced data storage operations. Table 4.4 summarizes the results. With less border cells, larger speed-gains are expected, as will be illustrated in the next two cases.

Accuracy wise, the fractional latency consistently underestimated the amplitude of the signal, in time domain, but closely matched the signal phase. This may be due to the more limited master bandwidth of this case, which was that of the base area, and thus 25 kHz, while the test signal had components up to 5 kHz. In EMTP, it is usually recommended to use one tenth to one fifth of the simulation bandwidth for accuracy. Thus, in this case, one fifth was used.

In frequency domain, the results are relatively close, except for a strong, exaggerated resonance around 1 kHz, where the latency case significantly over-estimates the amplitude. This resonance causes a shift in the unwrapped phase, which otherwise fairly closely parallels the reference results. The origin of this resonance may be explained by the somewhat different simulation environment that exists due to the interpolation boundaries that result from fractional latency, as opposed to distributed interpolation used before. These may result in a slight discontinuity which, for certain frequencies and angles of incidence, may appear as a different impedance. Thus, complex border-structure-material interactions may result, which could give rise to spurious resonances in the complex structure. In fact, the resonance at 1.2 kHz is under-estimated. This lends some credibility to the above explanation, where the slight change in the system due to the border caused a shift in the predicted main resonance frequency. Still, the averaged error was about 13 % for the ratio 10 latency case.

Of interest is that the ratio 100 latency case, which is supposed to be more accurate than the ratio 10 case, consistently performed worse visually in the plots, but better for averaged error. Later in this chapter, it is found that the ration 100 case is also less stable. This is likely due to numerical or implementation issues, due to the many repeated computations in the presence of these errors, as is explained later in this chapter.

Medium Border

With a thinner border, there are now three times more air cells in the simulation, which can be solved using fractional latency, and thus operate at a slower rate. This is very noticeable in the computation time and memory use, with a nearly 6 times speed-up and 1/3 less memory required. Visually, the error in the computation increased, although phase performance remains good. The averaged frequency-domain errors are smaller than before, although one may wonder how usable this simple metric is in evaluating these results.

It is of interest that the thinner boundary to the volume did not result in a much different system response. Any changes are likely due to the presence of a second interface, from aluminum to air, where in the previous case the aluminum was terminated at infinity, thus a infinite aluminum boundary around the air cavity. As such, the propagation of energy in this boundary

is different. However, due to the large impedance mis-match between the air and aluminum, only a small amount of the acoustic energy in the air ends up interacting with the boundary, thus is transmitted into from the air volume, and vice-versa.

This observation allows simulations to be modified in a way that reduces the computational load, but still maintains reasonable accuracy. When the material impedances are sufficiently different, one could substitute large sections of computationally expensive materials by other types, and make better advantage of the speed-ups and memory use reductions possible with fractional latency. This is, in fact, the reasoning behind a common simplification in TD-FD and TLM methods, where sufficiently high or low impedance boundaries are replaced by open circuits or shorts. These require virtually no computation, and may thus improve computational performance.

There is, however, an advantage of leaving at least some of the original material in place when the behavior of the system right at the boundary itself is of interest. Doing so avoids the the extreme conditions of open or short circuit conditions at the location where the observations must be made.

Thin Border

Although thin borders have proven to work in many different simulations done in preparation for this thesis, this particular case causes the simulation to quickly become unstable when fractional latency is used. In Figure 4.23, the simulation can be seen to diverge much more significantly, compared to the other cases. Soon after the time window shown the computation became unstable and the results oscillated to increasingly higher values.

Through further experiments, by allowing the previous computations to continue much beyond the time window of interest, similar instabilities were found when more and more cells were of the medium requiring faster computation, in this case air. The instability appeared sooner, the more air cells there were in the simulation. From this, it seems that the stability issues are related to solution iterations, or more precisely, to the number of iterations summed over all the latency areas in the computation. A further hint to this is the visually worse performance of the ration 100 latency cases, which include ten times more solution cycles for those areas, resulting in earlier onset of instability. The issue is discussed further in the following section, together with a possible solution which will also be demonstrated.

4.4.8   Issues with Fractional Latency in TINA

The current implementation of fractional latency in TINA still has a number of issues, which will be discussed here. In essence, the main problem appears to be a long-term stability issue,
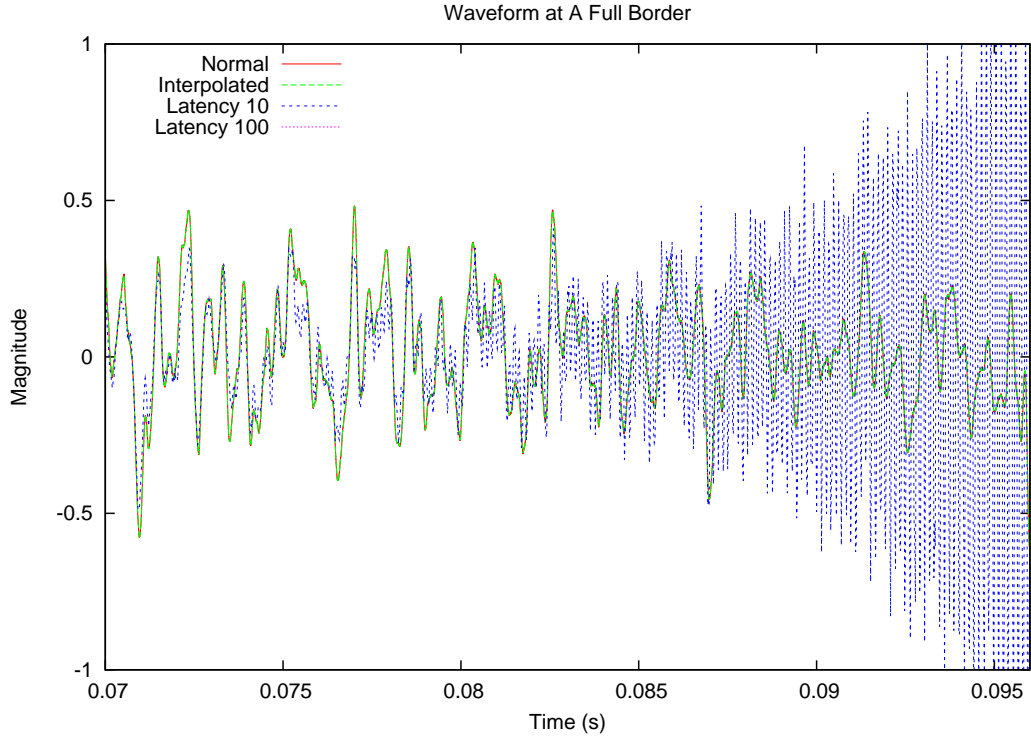
Figure 4.25: Fractional Latency Long-Term Stability

probably related to numerical rounding. A possible solution will be demonstrated. A secondary issue is not a problem of fractional latency or TINA, but a potential user error that needs to be highlighted.

Long-Term Stability Issues

When the prior simulations are continued for a long time, in most cases the fractional latency computations became unstable and unbounded. This happened repeatably with the main case under study in this thesis, although many other test configurations used did not show the problem. Those cases were smaller, and contained far fewer cells in the mesh. The total number cells and latency areas solved appear to have a significant influence on the long-term stability of the computation, as illustrated in the cases above.

From Chapter 3, we find that the more a line is over-sampled, thus the more points per wave length, the closer the poles move to the unit circle. Small errors can thus push them into instability more easily than less sampled cases, where a poles are further away from the unit circle, and there are thus more losses in the system. This is consistent with the observations, where smaller, coarser cases were never found unstable and larger, more detailed cases became

much more readily unstable.

A typical example is shown in Figure 4.25, which is the same case as used before, with the full border. All studied latency cases eventually became unstable, but the ratio 100 case became unstable much sooner than the ratio 10 case, at about 0.04 s versus 0.08 s, respectively, in the full border case.

At this time, the origins of this behavior can only be speculated upon until a full stability analysis of the fractional latency method is completed. However, a number of educated guesses can be made.

- The TINA implementation used float precision for the variables and computations in order to conserve memory. Thus, as the energy in the system decays, numerical round-off errors between very small numbers can easily result in sizable errors when divisions and multiplications are made.

- Cumulative errors in the computations, caused by insufficient accuracy in the interpolation, may result in run-away. This is possible, as the ratio 100 case performs many more operations than the ratio 10 case. The ratio 10 case, however, has a lower bandwidth, and thus introduces more dampening for high-frequency oscillations.

- The loss-less simulation does not allow the energy in the errors to be dissipated. Adding some losses could be beneficial. This approach is persued later in the next section.

- There is an implementation error in the algorithm. It is quite easy to, somewhere in the many interpolations and history manipulations, be one time step off in the history sources and thus compute with the wrong values, or have an interpolation factor computed between incorrect time stamps. This will, over time, cause significant error.

- The border structure in the thin border case was only one cell wide. Although each cell in the simulation was chosen to be of a size that insures the desired spatial discretization, e.g. ten cells per wavelength, if the structure has only one cell in it, there effectively a spatial under sampling. In the medium border case, the border is ten cells wide, and was more stable. However, the normal and interpolated cases were stable in all cases, even in the presence of thin borders.

- The fractional latency algorithm itself is unstable. This is rather unlikely, as interpolation was shown to be dissipative. However, the look-ahead, although using known and unchanging values, is essentially a violation of causality.

It is, however, the author's opinion that the instability is a result from a combination of cumulative interpolation error, due to numerical round-off caused by insufficient precision of

the "float" variable type, combined with the use of a loss-less, thus non-dissipative, medium within the latency areas. This may push the effective latency area's transfer function poles beyond the unity circle.

To offer a qualitative basis for this idea, from Figures 3.2 and 3.3, we find that the more interpolated line models are oversampled, the closer their poles approach the unit circle. Ideally, for the loss-less line model, they would have to precisely be on the unit circle. We also found that interpolation is dissipative, thus has a stabilizing effect.

However, if round-off errors are present and accumulate, it is possible, when the case is sufficiently over-sampled, and thus has its poles very close to the unit circle, that those errors cause some of the poles to exceed the stability limit. Thus, the more time steps and cells involved in a computation, the more errors could accumulate, and the higher the likelihood of instability.

When using interpolation in the line models for synchronicity, all cells have internal interpolation, but also some dissipation. When using latency, the cells within a latency area are all implemented using loss-less models, and interpolation only happens on the area borders. The latency interpolation in the borders naturally causes errors to and from the neighboring areas. These erroneous signals propagate through the loss-less medium within the latency areas and, at each interaction with the border, accumulate more error. Since there is only some dissipation in the borders, and the propagating energy only interacts with them as it reflects, further errors are added with each reflection. There is no distributed dissipation, as in the interpolated case, which could aid with stability.

In the simulations shown above, the wider the border, the less frequent the propagating energy in the area needs to interact with a border, and the less the errors accumulating event occurs, for each base-area time step. The thinner the border, the more frequent the interactions occur within the latency area for each base-area time step, and the quicker the errors accumulate. The simulations show the stability and precision degrading with thinner borders. This also happens when more time steps within an area are used, which results in more events in the latency area per base-area time step, and thus more errors. The latter is illustrated by comparing the ratio 10 and ratio 100 cases.

A number of experiments could be performed to determine and confirm the cause of the instabilities. First, increasing the numerical precision, by transitioning to "double" as a variable type, may give insight regarding the first two potential causes. If the computations are more stable using this type, numerical round-off is likely to be a factor. Second, the algorithm can be run an a small case and a couple of time steps computed by hand. The results could be verified to locate an implementation issue, and validate the third potential cause.

The last cause is hard to validate, as it requires a complete, formal study of the algorithm

and its stability. An abbreviated case could be composed, based on a single 1D transmission line fitted with latency cells on both ends, and the study of Chapter 3 could be repeated. The case is, however, inherently more complex due to the variable interpolation factor and it may be hard to find a closed-form solution.

This reduced-complexity study should be sufficient to validate the fractional latency algorithm, as the TINA solution itself was not changed to implement fractional latency, and the same circuit solutions are still used. The output interpolation does not affect the computation itself, so it does not need to be considered regard the stability of the method.

It would also be interesting to implement a lossy line model in the simulation and use it in conjunction with fractional latency to introduce some dissipation in the latency areas and observe stability in the presence of such losses.

Finally, regarding the determination of locating any implementation errors is that the TINA algorithm itself, it was found to be extremely robust. During the research, many serious errors were discovered and fixed, however, they only had minor effects on the simulation results. Given the long computations required to trigger the instability in many cases (over seven hours for the full border case), and the literally millions of time steps over more than 88000 cells that this involves, it is hard to locate such issues, especially if they are caused by small round-off problems or off-by-one issues in the history manipulations and interpolation functions.

4.4.9 Losses to Aid Stability

In order to help dissipate the energy associated with the errors in the fractional latency simulation, a simple loss model was added to the TINA simulator, and used in both the normal and latency 10 thin border cases discussed earlier. Since the addition of losses makes the resonances less steep, due to the decreased Q factor of the resonant structures, one must evaluate this case on its own. Hence, a comparison with the normal case is made under the same loss condition.

To implement the losses, the distortion-less line model of Appendix A.4 was used. The attenuation function α was set to 10^{-3}, which resulted in a loss multiplication factor of about 0.9990005. Applying these losses, the simulation was found to be in agreement with the normal case, and the instability disappeared completely. Under these conditions, the fractional latency, using the distortion-less line model, was eleven times faster than the normal case. In Figure 4.26, we see that the simulation is now stable and in agreement with the normal case over the long term. The negative going peak near the end is the beginning of the next impulse of the BLImp sequence. Figure 4.27 shows a detailed time-domain comparison between the normal and fractional latency case, while Figure 4.28 shows the magnitude and phase responses.
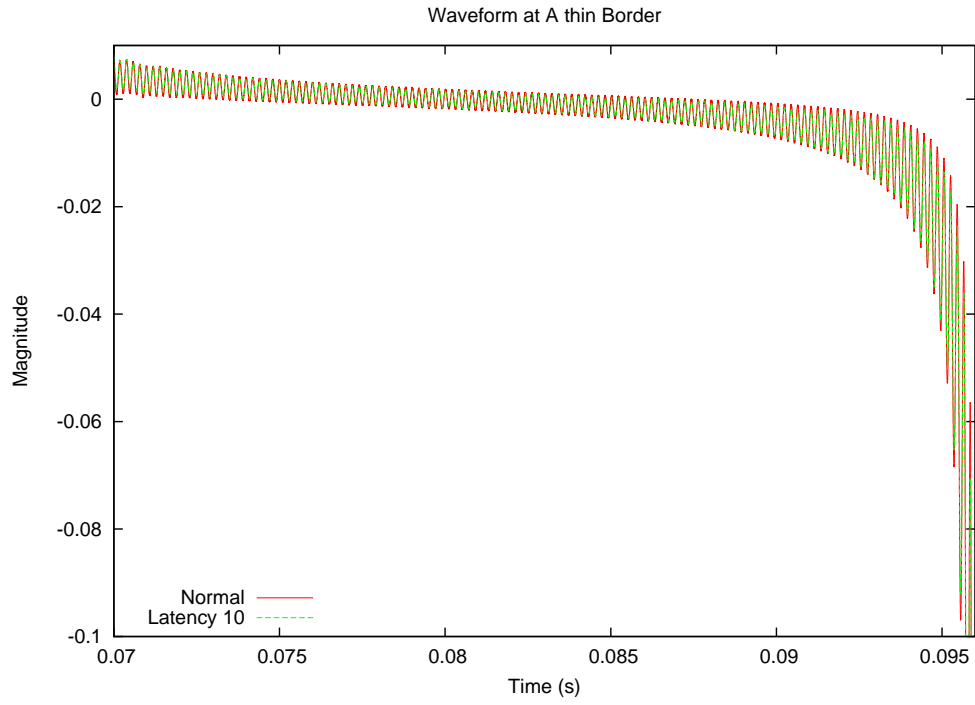
Figure 4.26: Thin Border Time Domain Results at A, Long-Term
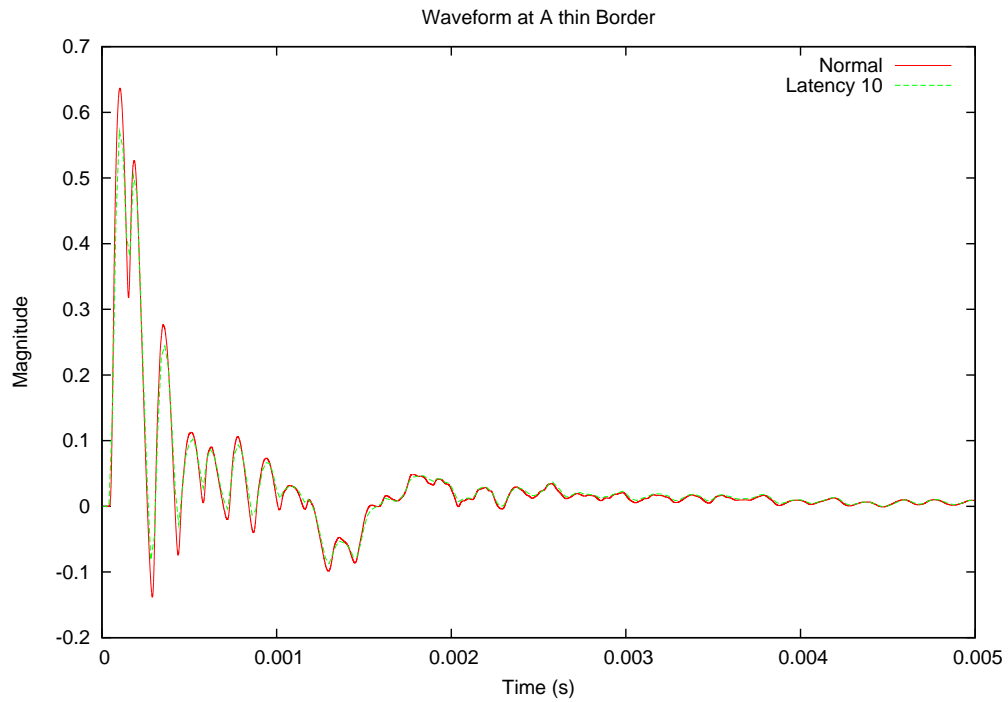
Figure 4.27: Thin Border Time Domain Results at A
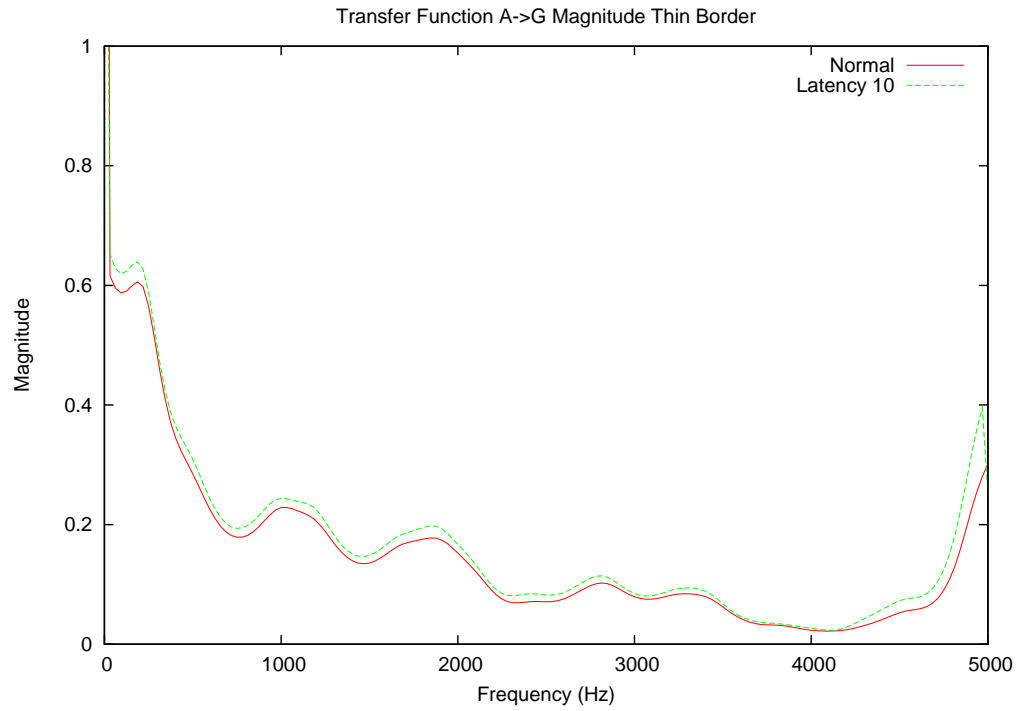
(a) Magnitude
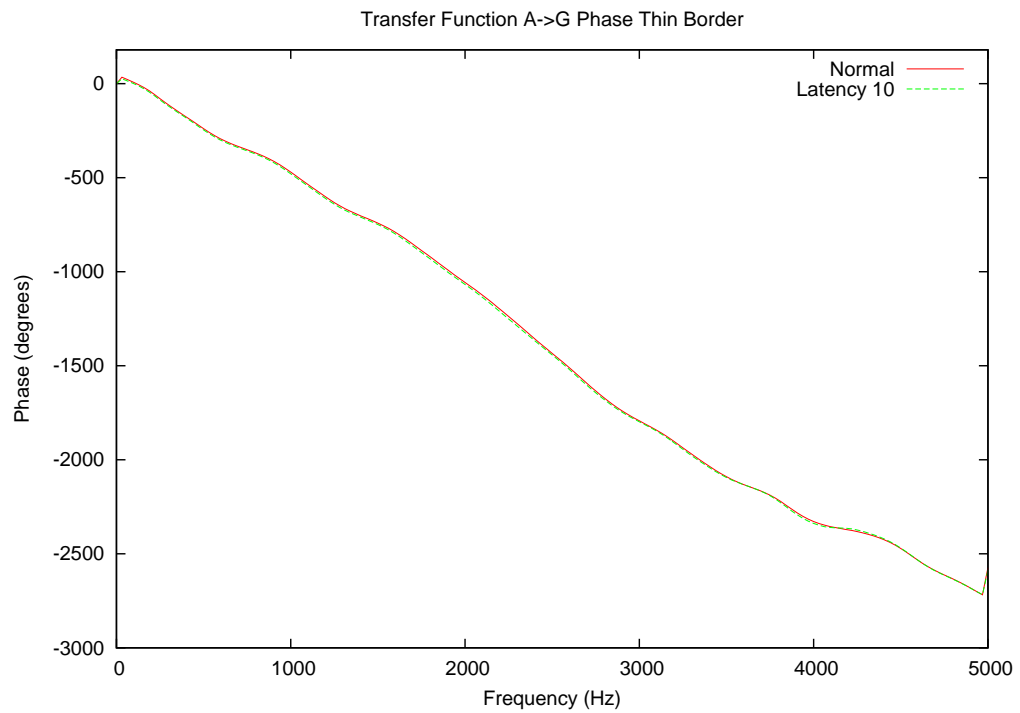
(b) Phase

Figure 4.28: Thin Border Frequency Domain Transfer Function A → G

From this, we may conclude that adding losses does suppress the oscillations and this lends credibility to the notion that numerical rounding and its resulting error may be at the basis of the instabilities observed in the loss-less, and thus non-physical, cases.

Under-Sampling the Solution

When setting the simulation band width, and thus the simulation time step for the base area and output interpolation, what is really configured, with respect to the output generation, is the rate at which the faster sub-systems are observed. It is thus very possible to under-sample this output generation and encounter temporal and spatial aliasing issues while the simulation itself is internally correct. Visually, the time-domain output will not be able to follow faster variations. The resulting samples themselves, however, are correct. This is illustrated in Figure 4.29, which is an almost under sampled case.

In this figure, the latency data points were shown, joined by a line. The normal and interpolated cases were strongly oversampled to give precise results. It can be seen that the sample points obtained from fractional latency are quite close to the correct values. However, the bandwidth of the simulation was too low to catch the sharp peaks (high frequency components). The signal in the plot had components up to 25 kHz, which is equal to the Nyquist bandwidth for the simulation.

Thus, when using fractional latency with a local fast implementation, one must take care to use a sufficiently fast sampling of the simulation to prevent under sampling the system. The Nyquist bandwidth is insufficient, as the fractional latency acts as a low-pass filter. In practice, oversampling by a factor of about five to ten gives good results, as was done in the rest of this chapter, where the signal bandwidth was limited to 5 kHz, compared to the simulation bandwidth of 25 kHz.

4.5   Conclusions

In this chapter, we introduced the concept of fractional latency and showed two implementations thereof, with and without extrapolation. It was demonstrated how to use the method to achieve synchronicity and achieve a computational performance gain over the previously shown interpolated line model approach by moving the interpolation out of the individual model cells and move it to the boundary of a latency area.

The method operated though achieving synchronism in an asynchronous, event-driven computation and, in time-decoupled methods, allows operation without the need for an explicit extrapolation by means of transmission-line look-ahead. It is the novel combination of absolute
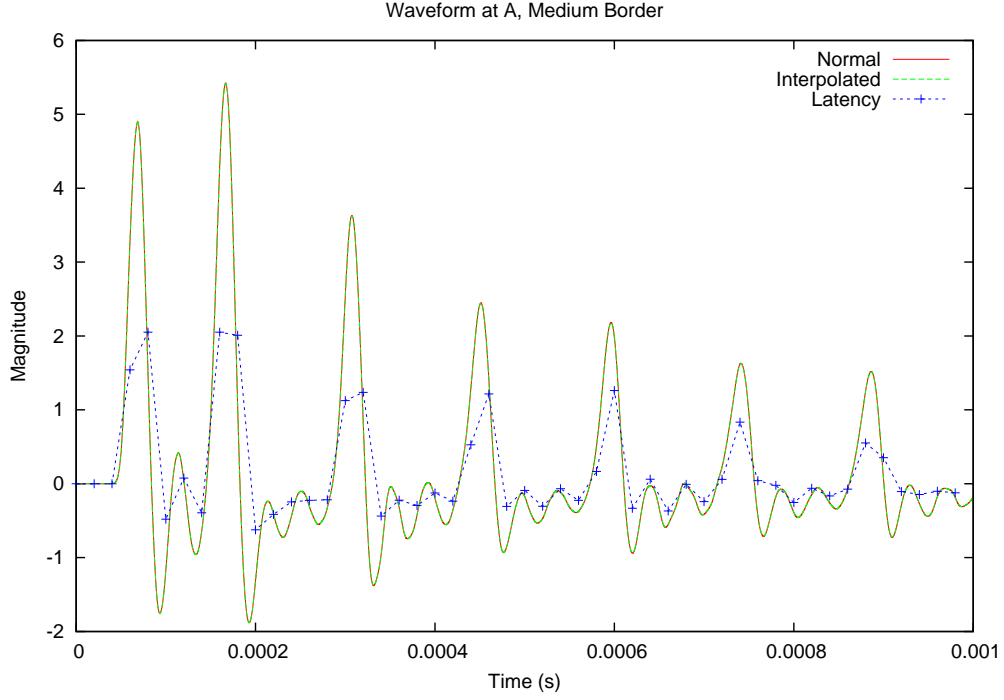
Figure 4.29: Nearly Under-Sampled Fractional Latency Case

time-stamp base synchronization and transmission-line look-ahead concepts make fractional latency possible. These concepts allow the extension of traditional latency approaches, which only allow integer ratios between the time steps.

Future work would focus on a detailed study of the stability and frequency-domain properties of the algorithm, as well as improving the implementation to insure no algorithm errors exist and investigate the long-term stability behavior by using higher precision variables in the solver code. Also, the use of dissipative media and integration rule switching, should be investigated to further aid long-term stability by introducing a level of dissipation in the simulation.

Chapter 5

Experimental Results

> In desperation I asked Fermi whether he was not impressed by the agreement be-
> tween our calculated numbers and his measured numbers. He replied, "How many
> arbitrary parameters did you use for your calculations?" I thought for a moment
> about our cut-off procedures and said, "Four." He said, "I remember my friend
> Johnny von Neumann used to say, with four parameters I can fit an elephant, and
> with five I can make him wiggle his trunk." With that, the conversation was over.
> *Freeman Dyson "A meeting with Enrico Fermi", Nature #427, page 297)*

In this chapter, various two-dimensional systems were constructed and measured to validate the
TINA method. The systems varied in size from roughly fifteen centimeters to over two meters.
This allowed a reasonable range of different simulation parameter ranges to test the simula-
tor with, and used cell sizes ranging from about one square centimeter down to one square
millimeter. These cell sizes are consistent with the acoustic wavelengths in the media for the
audible range, as well as the spatial discretization accuracy required to faithfully represent the
systems in the simulation. Also, a novel method to compute impulse responses in simulation
will be shown, based on periodic, band-limited impulses computed in the time domain.

5.1 Two-Dimensional Systems in a Three-Dimensional World

The TINA simulator, as presented in this thesis, can only compute one and two dimensional
systems. This is not a fundamental limitation. However, we limited ourselves to two dimen-
sions, as a result of the focus of the research. In order to perform experiments, we thus en-
counter an interesting issue: how may one perform a two-dimensional experiment in a three-
dimensional world?

This is, of course, not possible, but a good approximation may be made if we limit one
spatial dimension so that it becomes very small compared to the wave length. In practice,
when using acoustics, this means that we can design a planar experiment if we constrain the
geometry with a sufficiently dense barrier in one of the spatial dimensions.

When this condition applies, two-dimensional plane-wave propagation will occur. Even more, the system becomes infinite in extent in the constrained dimension. A point source thus becomes an infinitely long line source due to the mirrored point sources introduced by the boundaries. The original geometry of the system has become a slice projection. This means that the usual $1/r$ acoustic pressure (amplitude) degradation, for a spherical wave front emanating from a point source, is not applicable. It is replaced by $1/\sqrt{r}$ instead, which applies for cylindrical wave fronts from linear sources.

The validity of this, that the 2D simulations exhibit cylindrical waves, was implicitly verified with the vocal tract model lateron. This model was radially symmetric, and thus not a flat 2D model. The simulation results for this model show poor agreement with the measurements, while the other models, being better 2D approximations, conformed reasonably well.

5.2 The Band-Limited Impulse Response Method

To reduce the computation time of the simulation to the strict minimum, the impulse response of the systems was obtained, as this requires the simulation to take only as long as the inverse of the lowest frequency in the pulse. Since it is impossible, in an EMTP simulator, to exceed the Nyquist frequency of the simulation when latency techniques are used, special care must be taken to ensure no excess spectrum exists, as temporal and spatial aliasing may occur, ruining the results. Also, if the impulse could be made truly periodic, a leakage-free spectral analysis of the time-domain signal would be possible by matching the FFT time window with the period of the impulse.

Traditionally, band-limited impulse responses are obtained by filtering a regular impulse, or by the use of Gaussian pulses. Both, however, result in signals with some spectral energy in the tails. Although this amount may be low, it is non-zero.

In this thesis, a novel method was used that produces perfect, band-limited impulses with any desired spectrum that can be matched precisely to the FFT parameters used for output processing and obtaining the spectrum from the time-domain simulation.

The signal required for this is little known, but is discussed by Moore in his book on computer music [58]. It is, in essence, the time-domain sum of all desired frequency components – stated differently: the conversion of a box-car spectrum into the time domain. The method shown by Moore uses a closed-form, time-domain solution for this sum of sines that is very efficient to compute in the time domain. From the expressions below, we may see it requires only three sine or cosine series development computations for any number of consecutive, summed
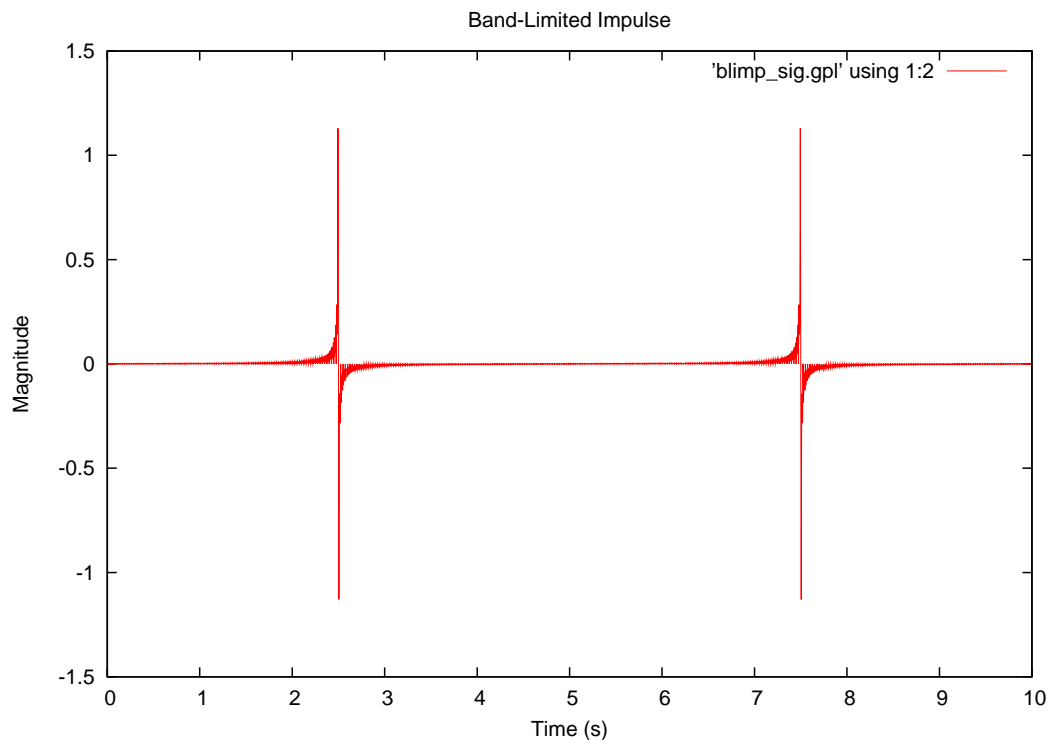
Figure 5.1: Time Sequence of a Band-Limited Impulse Signal

frequency components:

$$\sum_{k=1}^{n} sin(k\theta) = sin\left[\frac{(n+1)\theta}{2}\right]\frac{sin(n\theta/2)}{sin(\theta/2)} \tag{5.1a}$$

$$\sum_{k=1}^{n} cos(k\theta) = cos\left[\frac{(n+1)\theta}{2}\right]\frac{sin(n\theta/2)}{sin(\theta/2)} \tag{5.1b}$$

Although the use of band-limited pulses for obtaining frequency responses is not new, we have not seen the use of this type of perfect spectrum impulse, computed in the time domain, for impulse response measurements in simulation. Figure 5.1 show an example of a sequence of these band-limited pulses.

It must be noted that this is not an effective method to obtain impulse responses in practical measurement, compared to, for example, a frequency sweep, as the full energy of the injection is condensed in a narrow pulse. This results in relatively little energy that can be practically used in the experiments, and thus poor signal to noise ratio, unless many pulses are averaged. Also, practical amplifiers and loudspeakers frequently have problems faithfully reproducing such narrow impulses without significant harmonic distortion, which would again deteriorate

the experiment.

However, when used in simulation, we have found this method to give very good results. Since each frequency component in the signal may be perfectly matched to the FFT frequency bins used in further processing, as the proposed band-limited impulse method is fully periodic, leakage-free signal processing is possible, which improves the observed signal-to-noise ratio in the computation.

Although a time-domain computation was used to obtain the impulse signal, it can be derived from the frequency domain as well through analytical means using transform pairs, or through a suitable FFT transform.

5.3 Experiments

A number of experiments were performed to evaluate the TINA method. The first three sets were mainly exploratory, in order to learn how to design a suitable test configuration. It is the last set-up, the expansion duct of Section 5.3.6 that was not only the best experiment, but also showed the closest agreement with the simulation. Still, the earlier tests are included here to show some of the pit falls and limitations of 2D systems and simulations and the care that must be taken when considering the termination of the 2D experiment into free-space.

5.3.1 Equipment

The various experiments were performed using the same set of basic equipment. Loudspeakers, microphones, and their associated pre-amplifiers varied, but due to the use of transfer functions in all computations and comparisons, the influence of the signal chain was canceled out.

FFT Analyzer

The same FFT analyzer was used for all experiments, using the same settings. The device was a SR770 Stanford FFT Network Analyzer, by Stanford Research Systems.

The following FFT parameters were used:

- Span: 12.5 kHz

- Line width: 31.25 Hz

- Acquisition time: 32 ms

- Start frequency: 0 Hz

- Center frequency: 6.25 kHz

- Averaging: 16

To prevent time aliasing due to the finite acquisition time, we must verify the maximum size of the environment that can be safely analyzed with the given time window. Such a criterion is given in [17], page 32. We quote:

> "For non-periodic signals time aliasing arises when the duration of each record is similar to or less than the impulse response of the system under investigation, causing cross talk corruption in the signal processing.
>
> Time aliasing may be avoided by selecting the duration of each record to be much larger than the acoustic propagation times within the ... system, that is

$$t \gg \frac{2x_l}{c_0}$$

> t is the sample record length in seconds x_l is the distance from the sample to the furthest microphone c_0 is the sound velocity in meters per second

For our system, with a spatial extent of two meters for some cases, we find that our time window *t* must be much larger then 11 ms. We meet this criterion, just barely, by a factor of three. Since all other systems are much smaller, we may consider the requirement met."

Microphones and amplifiers

Two different Brüel & Kjaer condenser microphones and amplifiers were used. For all experiments we used 1/2" types and the Nexus Conditioning Amplifier, except for the measurements in Section 5.3.5 and 5.3.6, where we used a 1/4" model (type 4135) and the type 2609 microphone amplifier. For all measurements, the amplifiers were set to no weighting, and care was taken that they were not driven in saturation.

Power amplifier and loudspeaker

The amplifier used was an Alesis RA100. The loudspeaker was a Realistic midrange tweeter, cat. # 40-1288A fitted with a custom funnel to create a point source. For Section 5.3.6, a RadioSchack 4" midrange speaker cat. # 40-1197 was used on the source plate with a 1mm wide hole. Overdrive distortion was evaluated by FFT analysis of the output signal and avoided.

Air humidity

In order to determine the speed of sound, we had to measure the air temperature and humidity. This was done with a wet and dry bulb thermometer, HB Psycro-Dyne, cat. # 23015. The fan was allowed to run for one minute before the readings were taken. The results for the different days of measurement were found to be very similar, specifically: 22 degree Centigrade, 60 % relative air humidity. This yielded in an estimated speed of sound of 345 m/s. It was found, however, that the simulations in the most accurate system, the expansion duct of Section 5.3.6, showed much better agreement with a speed of sound of 330 m/s, which represents a 4 % deviation from the value found through the standard equation. This value was consequently used. The error could be in part due to inaccuracies in the experimental configurations, as well as the measured environmental parameters.

5.3.2 3D Vocal-tract model

Initially, we thought that the TINA method possessed circular symmetry in 2D. This would have been a perfect match for the vocal-tract model, which is also circularly symmetric (this is why it was studied). However, closer investigation revealed that TINA does not have such symmetry, but rather extends in an infinite plane. As such, the vocal-tract cross-section would appear as an infinitely wide "corrugated roof" as opposed to the desired profiled tube.

We can deduce this from the assumptions made in the 2D model. In order for the acoustic energy not to spread above or below the plane, we either need perfect reflection, or equal acoustic pressure above and below the plane. As such, a point source becomes a line source. In the other experiments, this situation was emulated by using wooden boards as reflectors, so that an infinite-plane approximation could be achieved and the sources in effect became line sources perpendicular to the plane.

Experimental set-up

The physical measurement setup is described in Anderson [2]. In order to render the vocal-tract model in the computer, a special plug-in for the TINA simulator was written. It allows any binary TIFF image to be drawn as a structure. Using this model, we used an optical scan from the physical model, processed it into a black and white binary image, and then rendered it in the simulator as aluminum. The process is illustrated in Figure 5.2. The top image is the scan. The second image is the binary outline, where black areas are rendered as a specified material in TINA. The third image is the actual simulation structure as used by TINA. The green border is the environment termination, the blue air, red and yellow aluminum, light-blue an ideal absorber, and off-white the acoustic source.
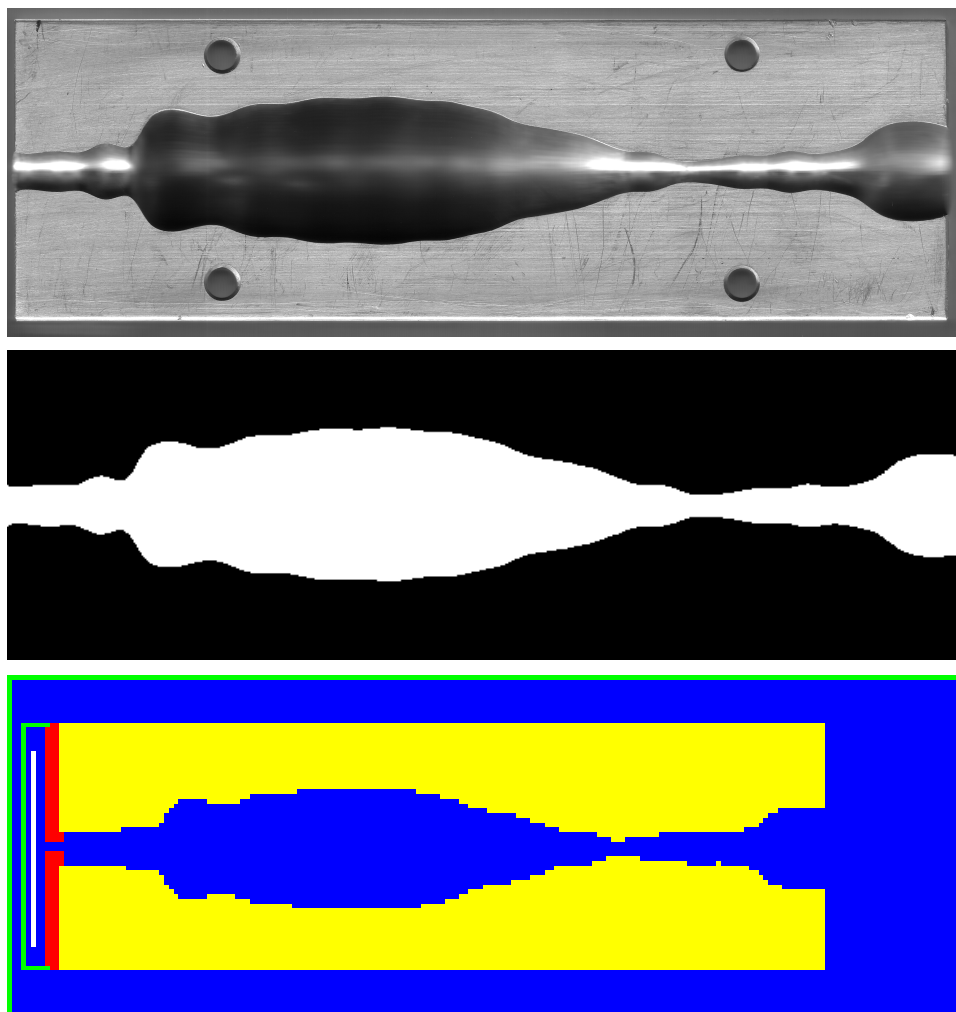
Figure 5.2: TINA Vocal-tract-model processing

Comparison with measurements

First, it must be noted that the measurements were done with a 31.25 Hz frequency step size. For the simulations, however, a 100 Hz interval was used for the vocal tract model. As such, dense clusters of resonances were blurred together in the simulation results.We do seem to capture the general trend of the resonances (Figure 5.3). However, the magnitudes and frequencies of the resonances are very different. These errors are mostly due to the incorrect symmetry assumptions, as well as reflections from the simulation boundary. The measurement point at the end of the vocal tract is very close to the boundary, thus any reflections from it will have a significant effect on the simulation output in both magnitude and phase due to interference effects. The phase plot is completely unusable, as the errors in resonant frequency also cause large changes in the phase. With correct behavior, the simulated phase response would be akin to the envelope of the measured phase.
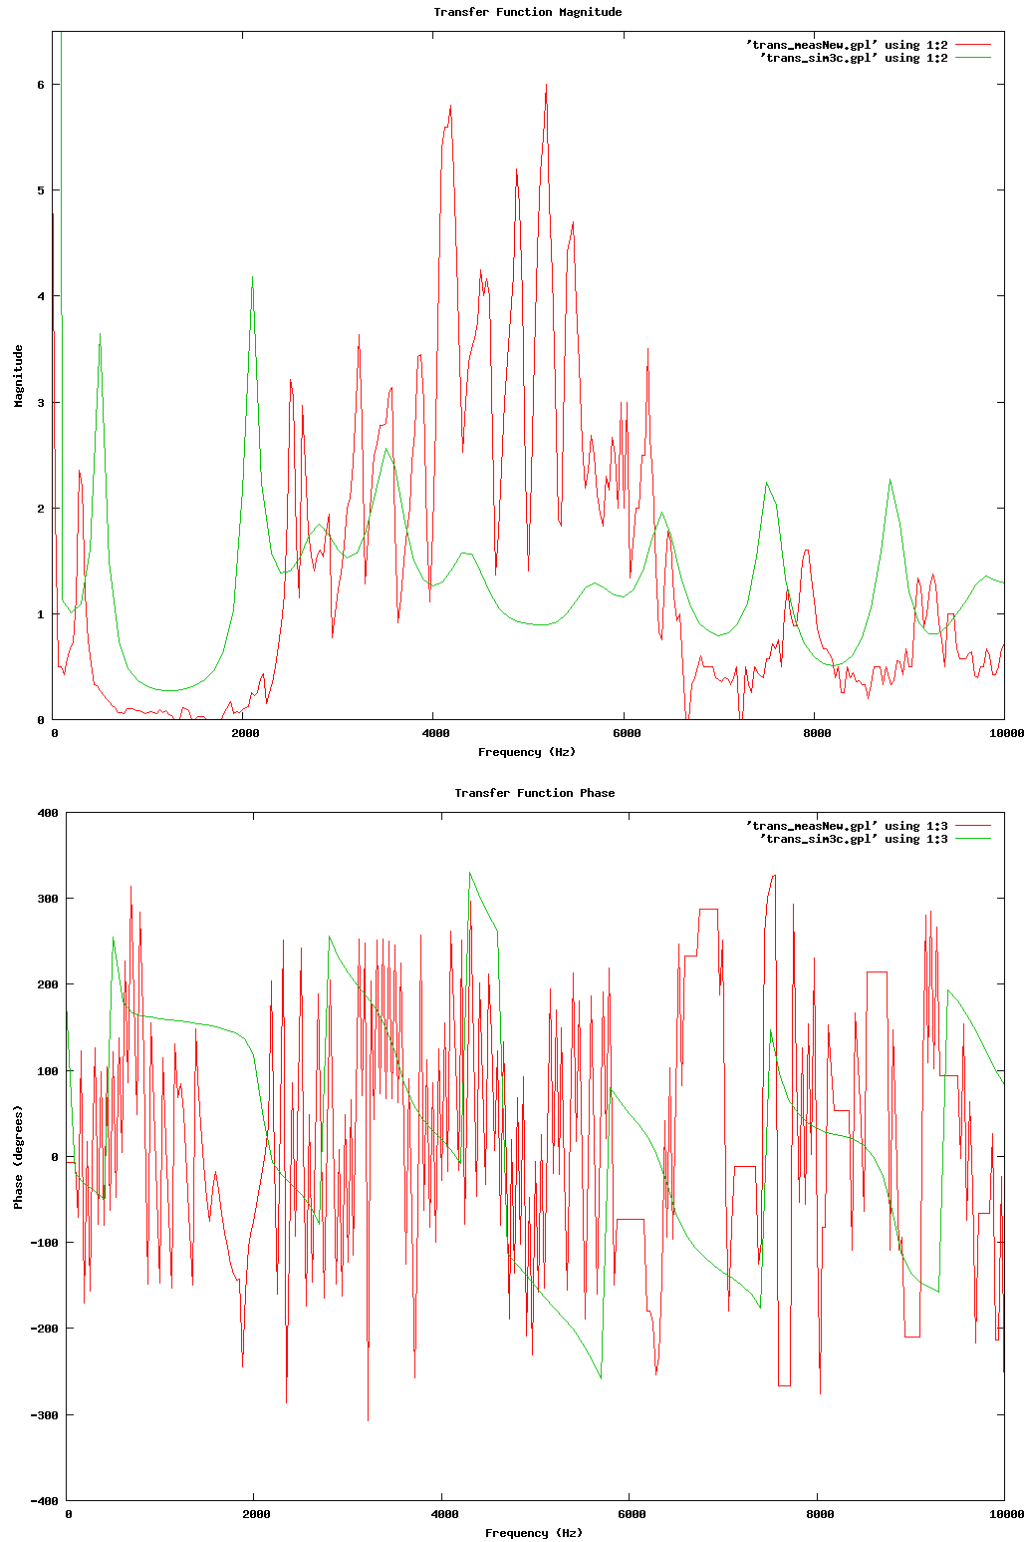
Figure 5.3: Comparison of Measured and Simulated Acoustic Pressure for the Vocal Tract

Due to the erroneous symmetry, the various resonances seen in the model do not agree well, as the simulated case results in wholly different structure than the physical model. In addition, the cells in the simulation were 1 mm by 1mm. There is an error of 1 cell in the drawing functions due to roundoff. Since some features in the model are only 3 mm in size, this was a significant factor. However, compared to the wavelength, this error is still relatively small, and the response of the system should still be quite similar. The overall trend is still present, however.

5.3.3 2D Room With Source at an Edge

In order to emulate 2D plane waves as they appear in TINA, we need to construct a suitable 2D experimental approximation. A practical method to such an approximation is to spatially confine the acoustic waves in a narrow duct. If the constricting boundaries chosen are highly reflective with respect to the acoustic medium, a 2D approximation is obtained, as long the frequencies under investigation are below cutoff for all propagation modes of the duct.

The reflective boundaries have the effect of giving the system infinite extent in the Z-axis. That is, a point source is converted into a line source in the Z axis, assuming the 2D plane is in the X-Y directions, due to the reflections.

In this first experiment, the source was placed at the edge of the structure. This was, in retrospect, not a good choice, as the discontinuity of the boundary to free space introduces partial reflections. The experiment was subsequently refined, the results of which are described in section 5.3.5.

For each of the two experiments, two configurations were measured:

- An open room: the waves were only restricted in the Z-direction. The X and Y dimensions were left unbounded. (The boundaries were the end of the experimental structure, where the waves terminate into free space.

- A tapered duct: the waves were now restricted in the plane by two more boundaries. A tapered shape was chosen to provoke resonances at different frequencies. The nonrestricted dimensions of the duct terminated into free space and/or the 2D room approximation.

Experimental set-up

The system was built out of plywood panels, spaced 3.8 cm apart using small wooden blocks, as illustrated in Figures 5.4 and 5.5 (cover board removed to show duct). This physical distance allowed for the placement of the microphone and point source in-between the panels. However,

the distance also puts an effective upper limit on the frequency for which the two-dimensional approximation is valid. Furthermore, the lateral extent of the plywood was only 100 cm from the source on the Y-axis, and 221 cm on the X-axis. As such, the limiting dimensions are the Z-axis spacing of 3.8 cm and the Y-axis extent of 1 m. The waveguide cut-off frequency can be found from equation 5.2. *n* and *m* are the propagation modes, while *a* and *b* are the spatial dimensions.

$$f_c = \frac{c}{2}\sqrt{\left(\frac{n}{a}\right)^2 + \left(\frac{m}{b}\right)^2} \tag{5.2}$$

Investigating the fundamental modes (1,0) and (0,1), with an acoustic wave speed of 345 m/s, a 3.8 cm by 100 cm wave guide has cut-off frequencies of about 4500 Hz (for 3.8 cm) and 170 Hz (for 1 m). Using these results, we choose 4000 Hz as an upper limit and 200 Hz a lower limit for the two-dimensional approximation.
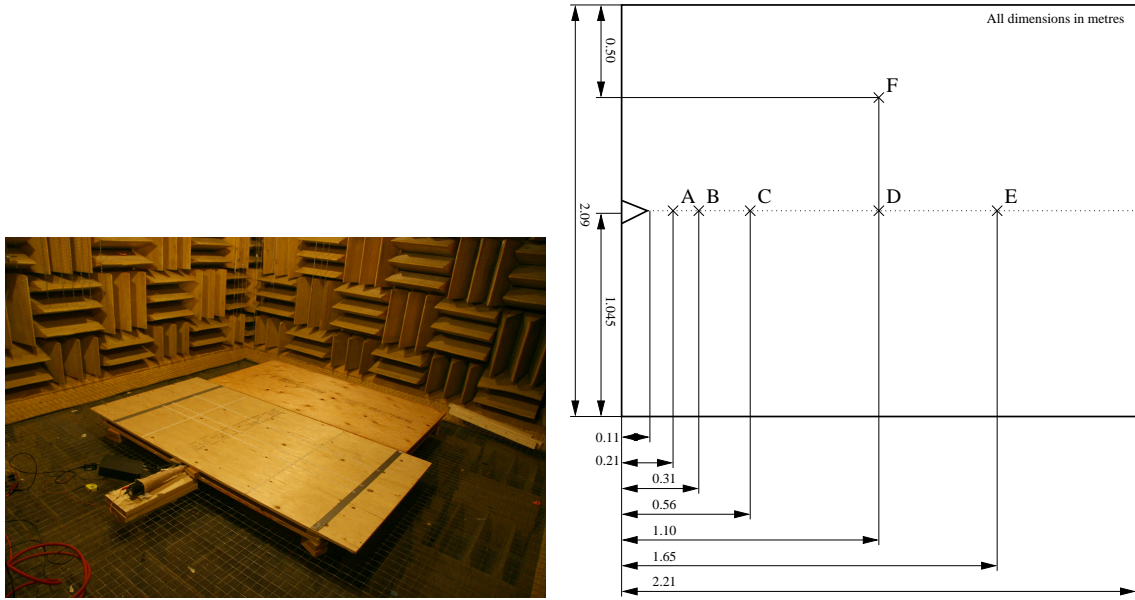
Figure 5.4: Experimental Set-Up With Edge Source - Open Plane

5.3.4   Comparison with measurements

The transfer functions between point A and various other points in the system were computed from the raw measurement and simulation data. To reduce noise, the measurements were aver-
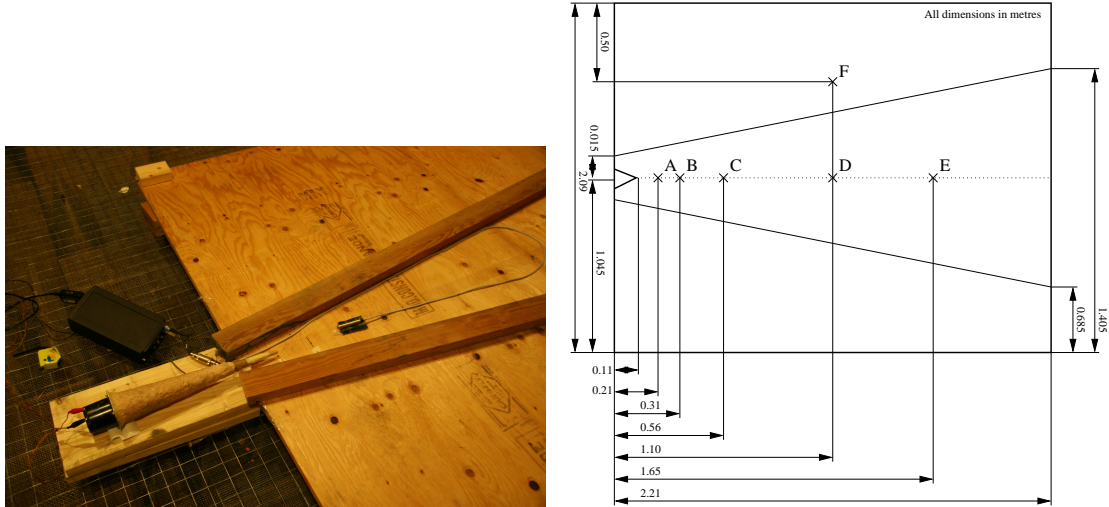
Figure 5.5: Experimental Set-Up With Edge Source (Cover Removed) - Duct

aged over 16 acquisitions. Also, the simulation was done with an infinite-plane approximation, while the measurements have impedance discontinuities at the boundaries that could not yet be modeled, as this requires three-dimensional geometry to describe the end of the wooden panels. However, the simulation boundary itself still exhibits some reflection, due to a matching problem. These reflections cause significant distortion in the simulation output, which is clearly visible in the phase plot. All these factors significantly distort both the measurements and the simulation.

For the C-location, the magnitude and phase responses were over-layed for direct comparison (Figures 5.6 and 5.7). The results were poor, given that an open plane should yield a perfectly flat response. However, the plane was of finite extent, and thus exhibited a plane to free-space discontinuity, yielding reflections.

For the duct, however, the situation improved somewhat, as the duct enclosed the source (Figure 5.5), effectively removing the 2D plane boundary discontinuity on that side from the physical system, as well as the sides. However the other end of the duct still terminated in free space. The situation is similar in the simulation. In the simulations, the duct spatial dimensions were also truncated wrt the physical system.

For higher frequencies, the duct resonances became so close together that they, for the measured result, blended in a continuous bump. The simulation, due to the absence of losses, shows a strongly exaggerated magnitude response for these frequencies. Still, from the open plane simulations, we can see that the reflections in the simulation are in the order of 0.2 - 0.4 relative magnitude. As such, the noise floor is so high that the majority of the duct data is essentially meaningless, except for the strong resonances at the higher frequencies. Thus, we must conclude that the higher frequency bump, the only valid part in the simulation, exists in

the measurements, but the experiment can only be considered indicative at best.

Figure 5.6: Comparison of Measured and Simulated Acoustic Pressure (Point A→C) - Open Plane
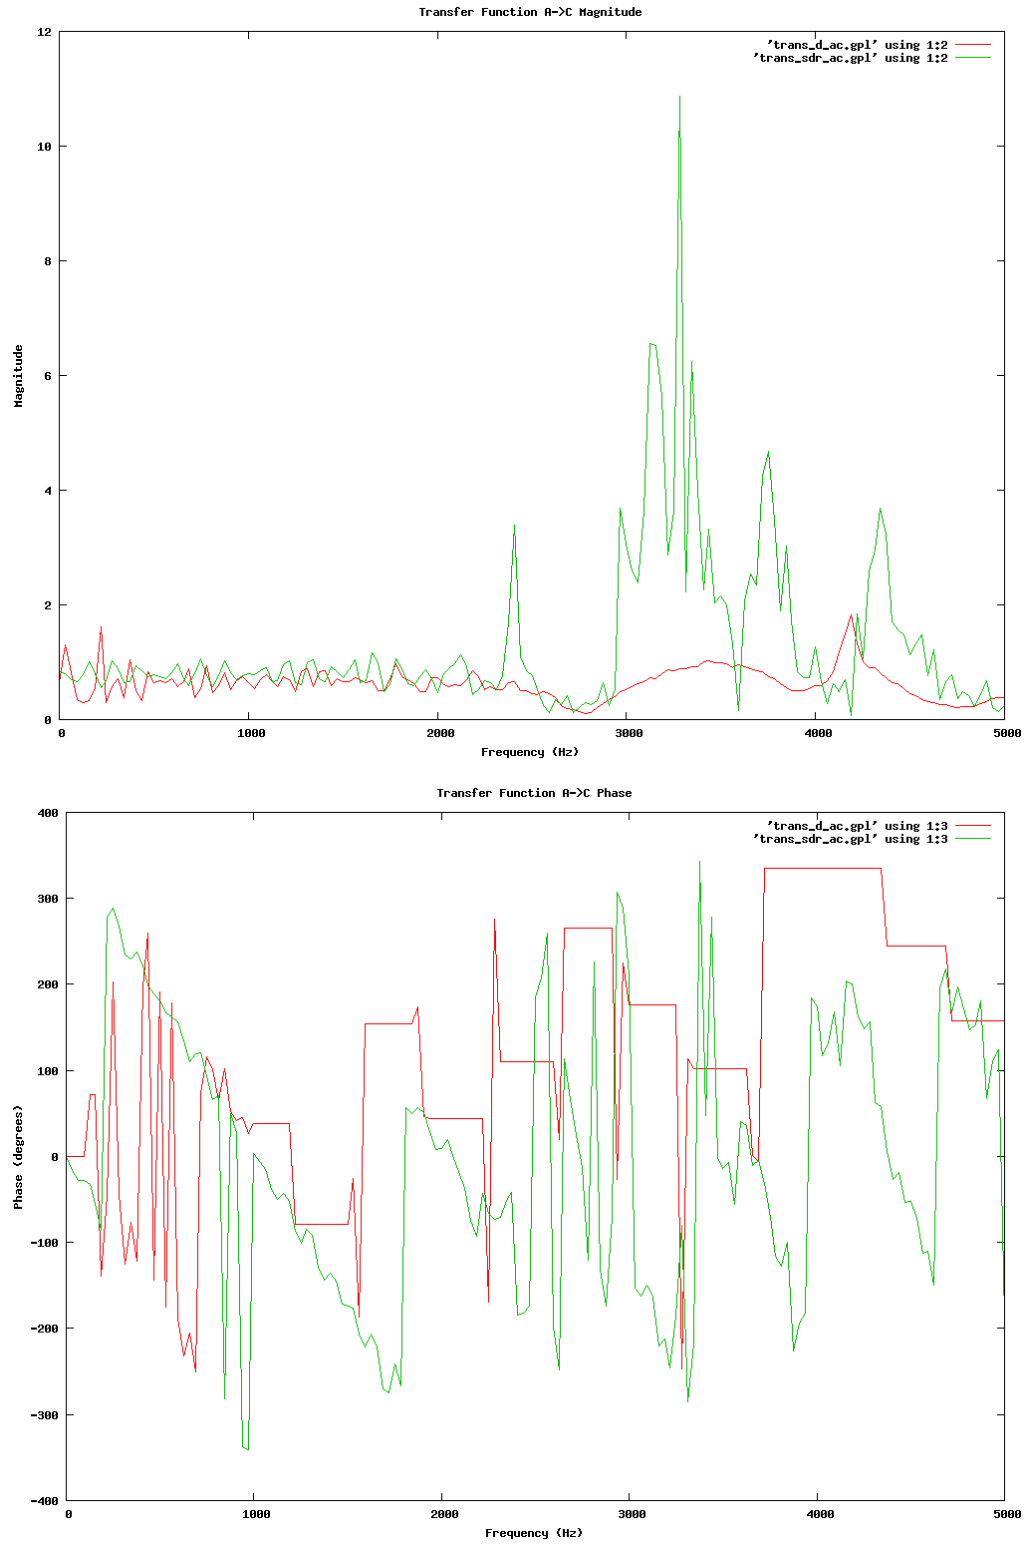
Figure 5.7: Comparison of Measured and Simulated Acoustic Pressure (Point A→C) - Duct

5.3.5   2D Room With Central Source

In order to prevent the problems associated with having the source close to the boundaries in the previous experiments, the source was next placed in the center to improve on the 2D infinite-plane approximation. Also, the source was now placed at the wide end of the duct to reduce the influence of the impedance discontinuity at the end. In the new configuration, the discontinuity into free space is now reduced by the area exposed by the narrow end of the duct (Figure 5.9). Future improvements in the simulations should then yield results closer to the measurements.

Experimental Set-Up

The system was constructed in the same way as in Section 5.3.3. Figures 5.8 and 5.9 (cover board removed to show duct) illustrate the differences in the configuration. The valid frequency range for the experiment is now somewhat different for the Y-dimension. It is now 120 cm, which puts the cut-off frequency slightly lower at 140 Hz. However, for consistency in the data, we kept the previous 200Hz to 4000 Hz valid frequency range. A further addition was a custom frame to position the loud speaker in the setup. It allowed the source to be positioned vertically, and the depth of the funnel tip into the setup adjusted to the desired depth, which was midway in the clearance between the wooden panels.
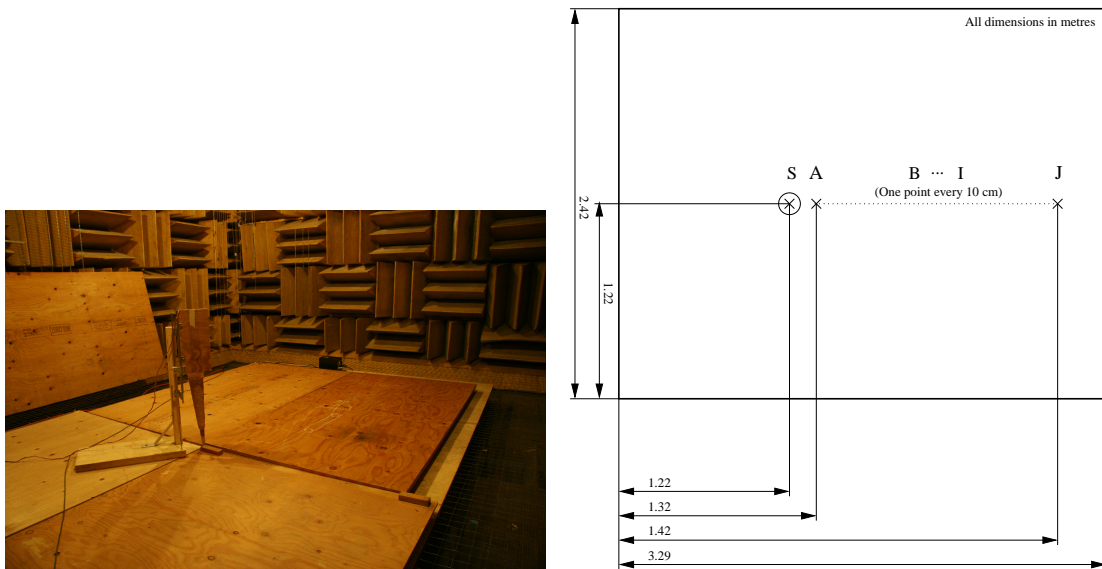
Figure 5.8: Experimental Set-Up With Central Source - Open Plane
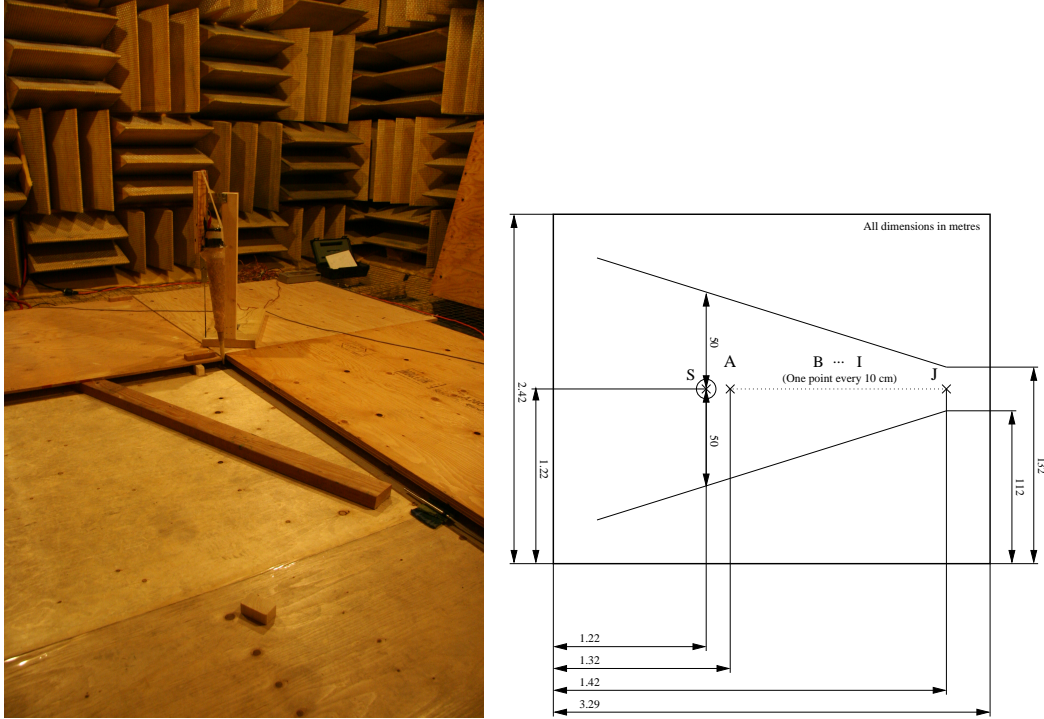
Figure 5.9: Experimental Set-Up With Central Source (Cover Removed) - Duct

Comparison With Measurements

The transfer functions were computed in the same way as in Section 5.3.3. Once more, boundary discontinuities in both the measurements and simulations existed. The measurements averages agree reasonably well with the simulations. We attribute remaining errors to the boundary reflections, both in the physical set-up and the simulation. Also, the resonances in the simulated duct were not dampened, due to the absence of losses, which resulted in significant magnitude deviations.

The duct simulations do show most of the expected resonances, but the magnitudes are off due to the undamped system. The phase response is in reasonable agreement with the measurements. One has to, however, keep in mind that phase jumps occur in the plots. These wild flips are generally due to a small increase or decrease that causes the phase calculations to wrap around. If one mentally unwraps the phase, the results are close.

The boundary reflections are still a significant source of error, in the range of 0.2-0.3. However, due to the strong resonances, the signal to noise level in these simulations is much greater than before, and valid conclusions on the resonance frequencies of the system, as well as phase response, could be obtained.
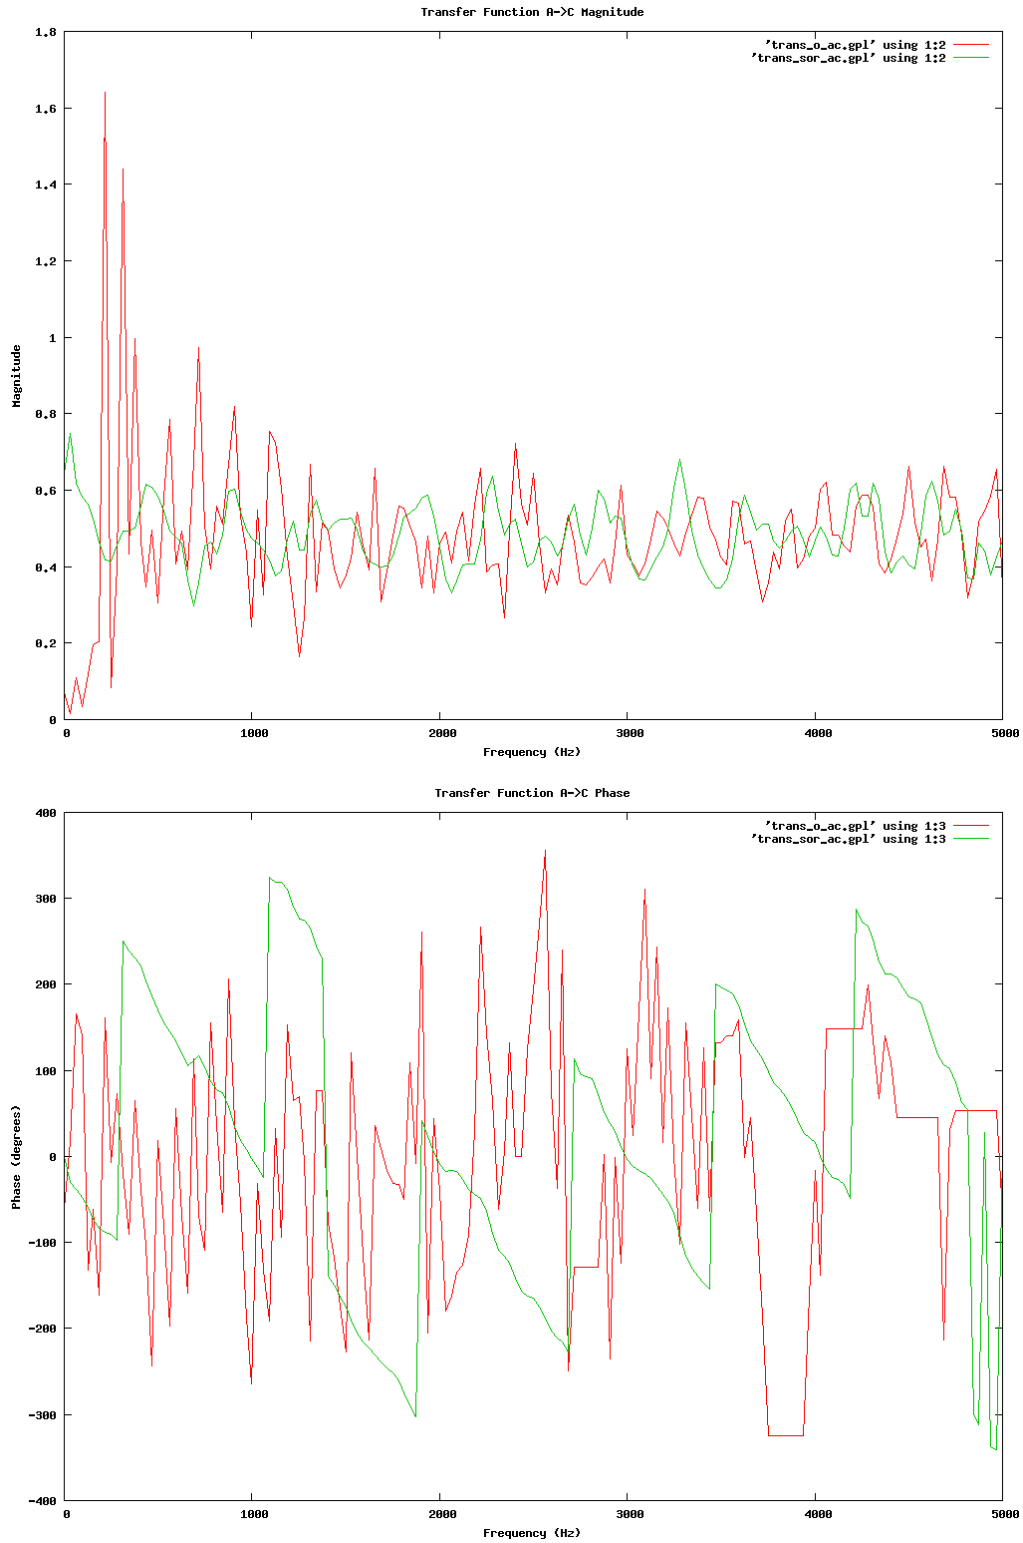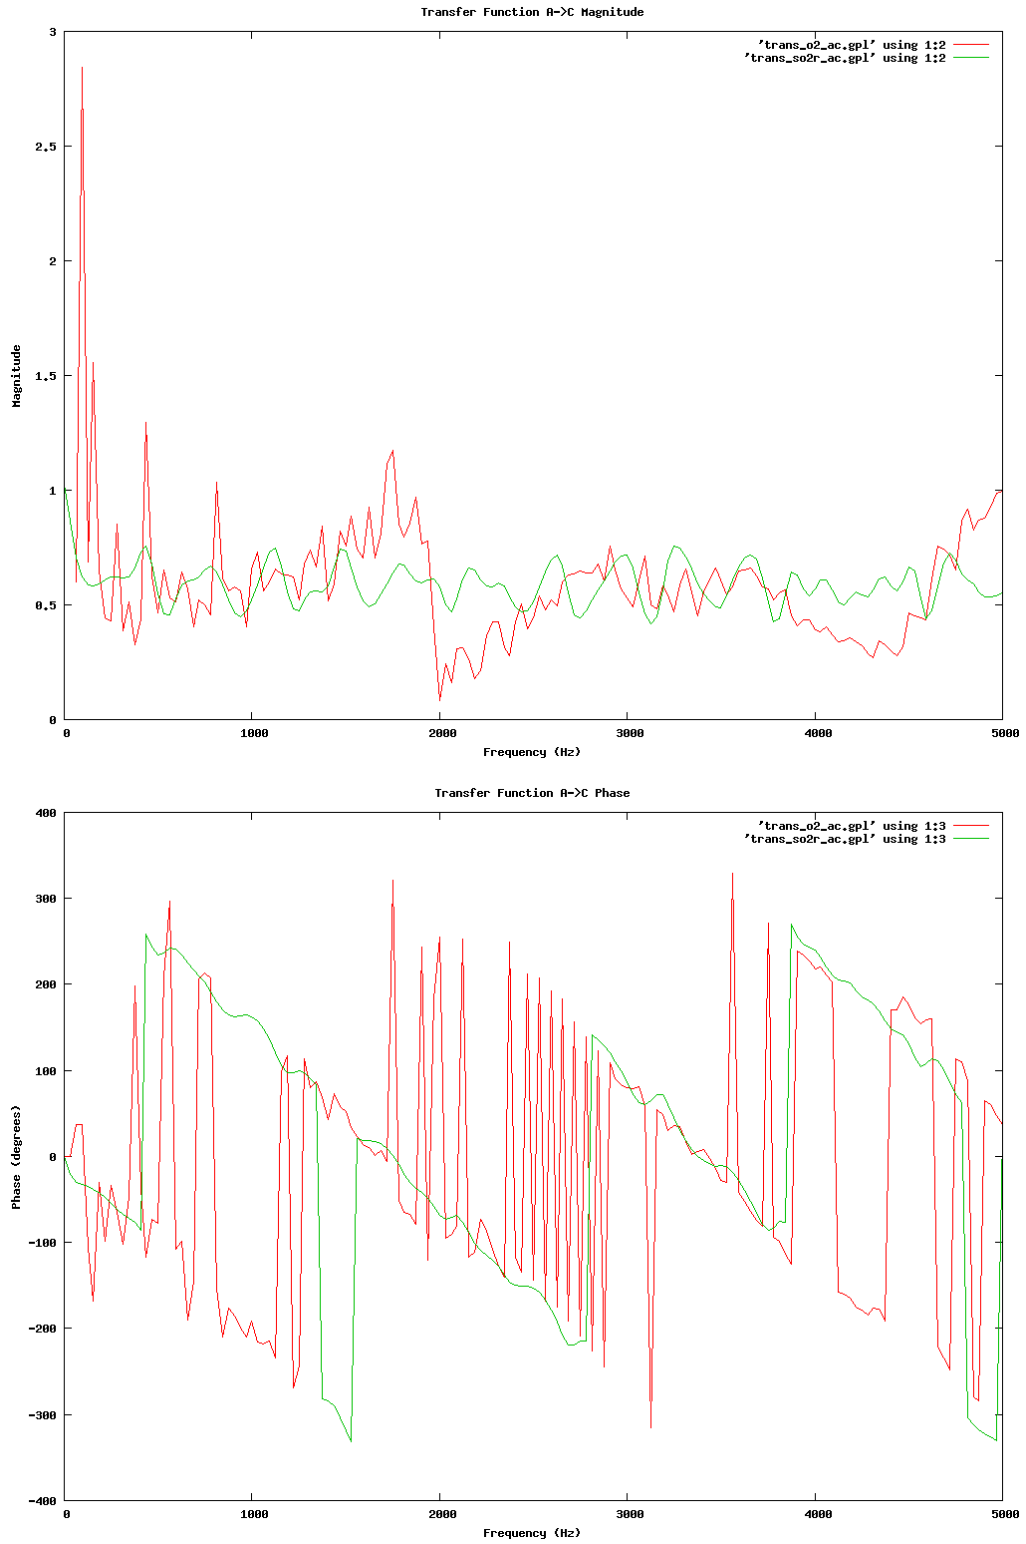
Figure 5.10: Comparison of Measured and Simulated Acoustic Pressure (Point C→A) - Open Plane
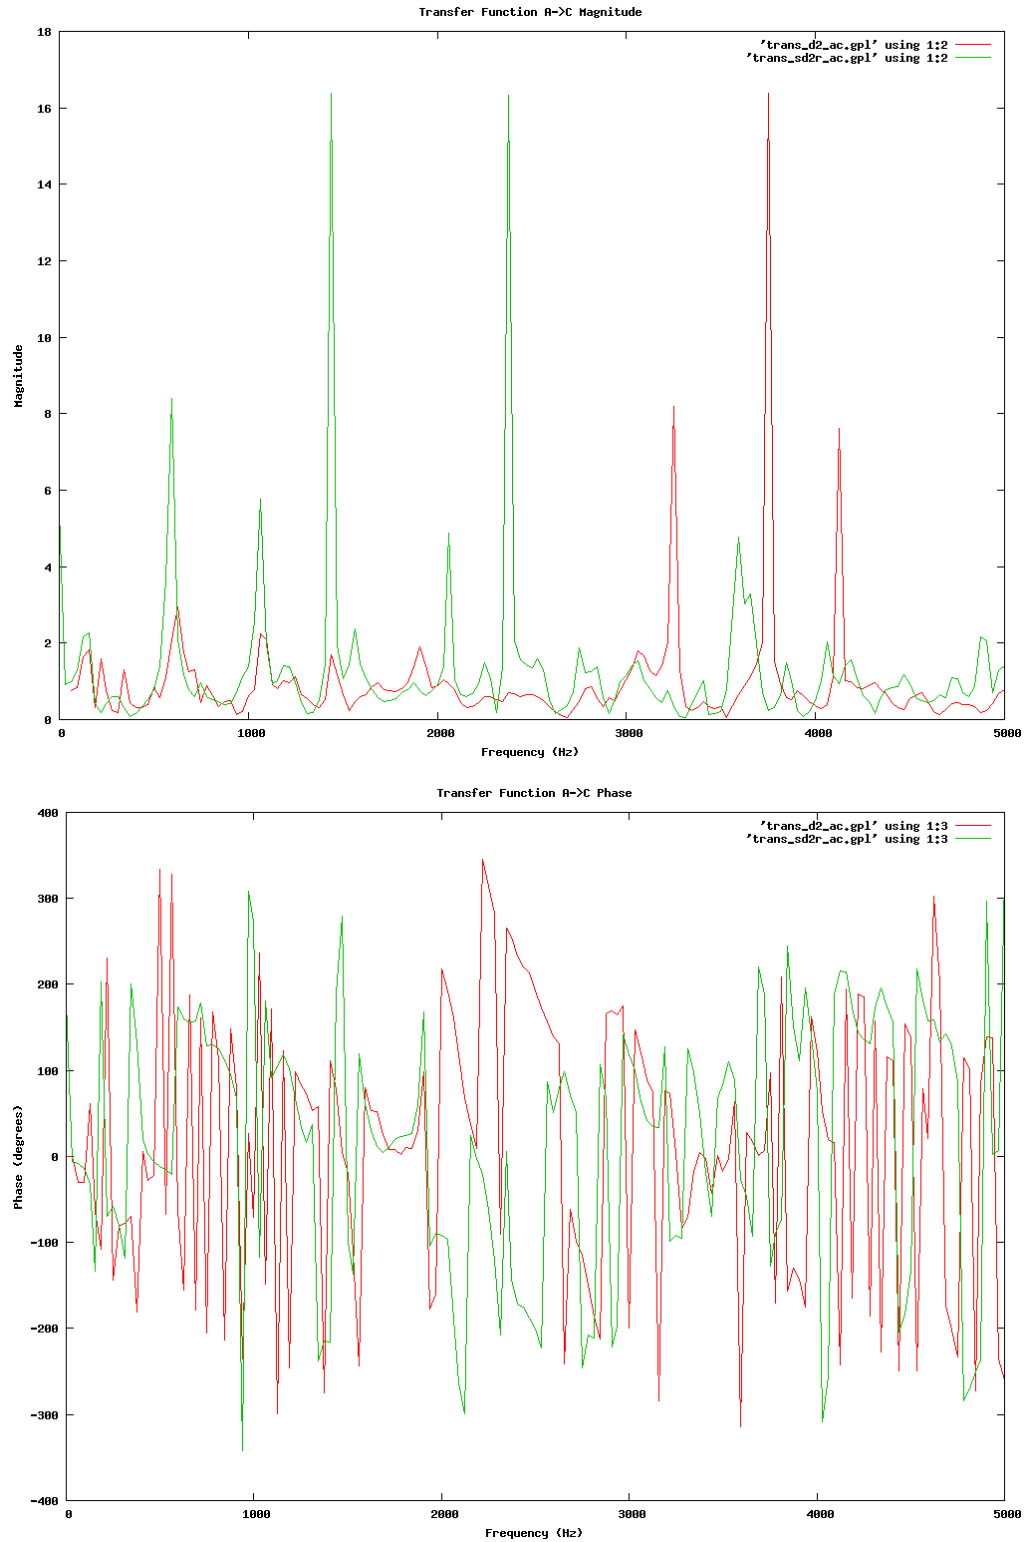
Figure 5.11: Comparison of Measured and Simulated Acoustic Pressure (Point C→A) - Duct

161

5.3.6   Expansion Duct

Using the experience gained from the preceding experiments, a final, high-precision experiment was devised to validate the TINA method. Stiff boundaries from vocal tract combined with well-defined and simple, planar geometry of duct experiments, into a closed system, so boundary issues were avoided. The configuration gave good agreement with the simulated data.

Experimental Set-Up

The system, shown in Figure 5.12, consisted out of a duct, cut in a 2 mm thick plate of aluminum clamped between a sandwich of two 2 mm thick aluminum plates, backed by plywood sheets. The thin duct allows a much higher usable frequency range for the 2D approximation, while the stiffness of the metal, clamped by a multitude of bolts, and loaded with the mass of the plywood, prevented the strong spurious resonances that plagued the large setups described earlier in this chapter.
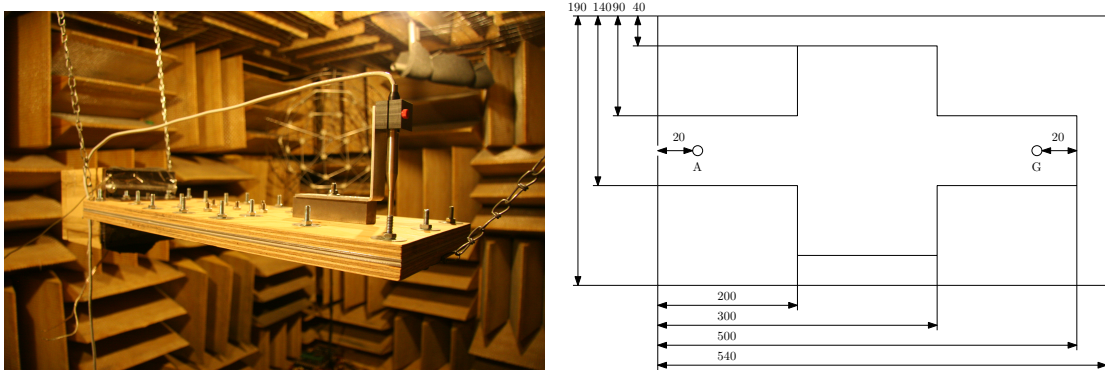
Figure 5.12: Expansion Duct Experimental Set-Up

The loudspeaker injected its signal through a 1 mm wide, 2 mm tall slit in a thick steel plate, from the edge of the configuration. The system is entirely closed, so no discontinuities to free-space, nor absorbing boundary condition issues arise in measurement and simulation. The microphones were inserted through holes in the top, while each unused hole was plugged with an aluminum insert to maintain continuity in the duct. Due to excess radiation from the back of the loudspeaker interfering with the measurements, it was fitted with a box made of 1" thick particle board, sealed with foam against the speaker mounting plate. This reduced the speaker interference to roughly 2 % of the average measured signal.

Lastly, the physical size of the microphone itself was now taking into account, which is a sizable portion of the wave length for the higher frequencies evaluated. Thus, in the simulation,

Transfer Function A->G Magnitude
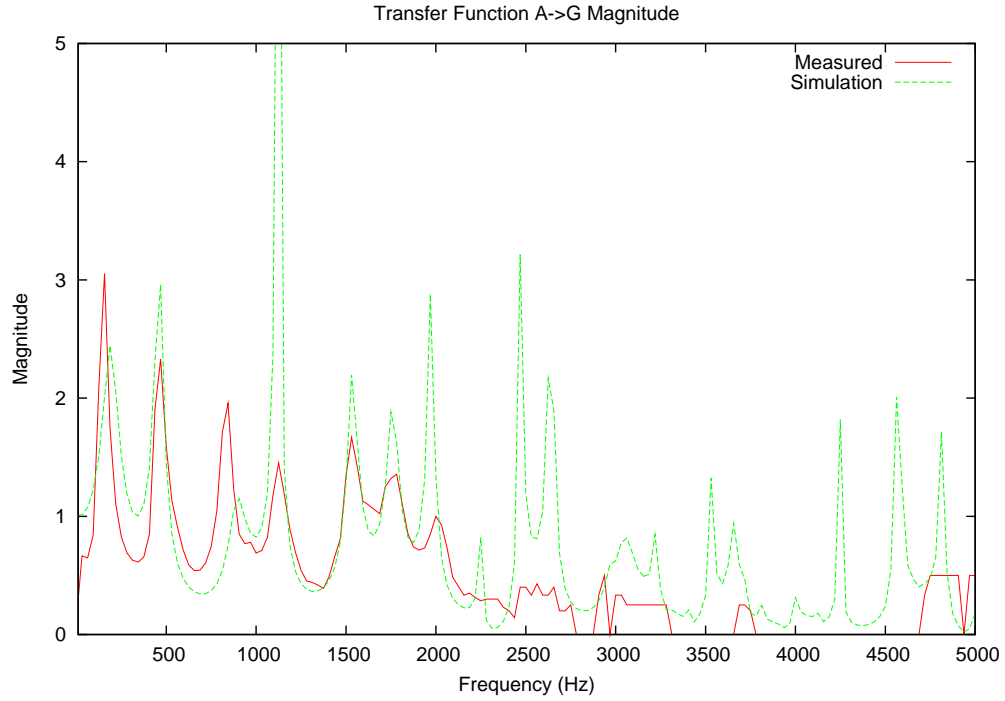
Figure 5.13: Expansion Duct Magnitude Plot

the acoustic pressure over an area comparable to that of the physical microphone was averaged to obtain the result for that location.

Comparison With Measurements

Except for the magnitude overestimation, due to the absence of losses in the simulation, the results from Figure 5.13 show excellent agreement. There were some issues with the accuracy of the duct itself, due to a problem with the water jet cutter, with the duct dimensions in error by roughly 1 mm. The simulation, however, was performed with a 2 mm cell size, thus relatively coarse. As such, the errors in the physical dimensions still below the spatial discretization resolution of the computation.

The measurements themselves presented some issues due to the strong filter effect of the tiny speaker aperture, resulting in insufficient acoustic energy in parts of the spectrum. This may be observed from the FFT analyzer screen shots in Figure 5.14. Still, the simulation follows the trends visible in this reduced signal.
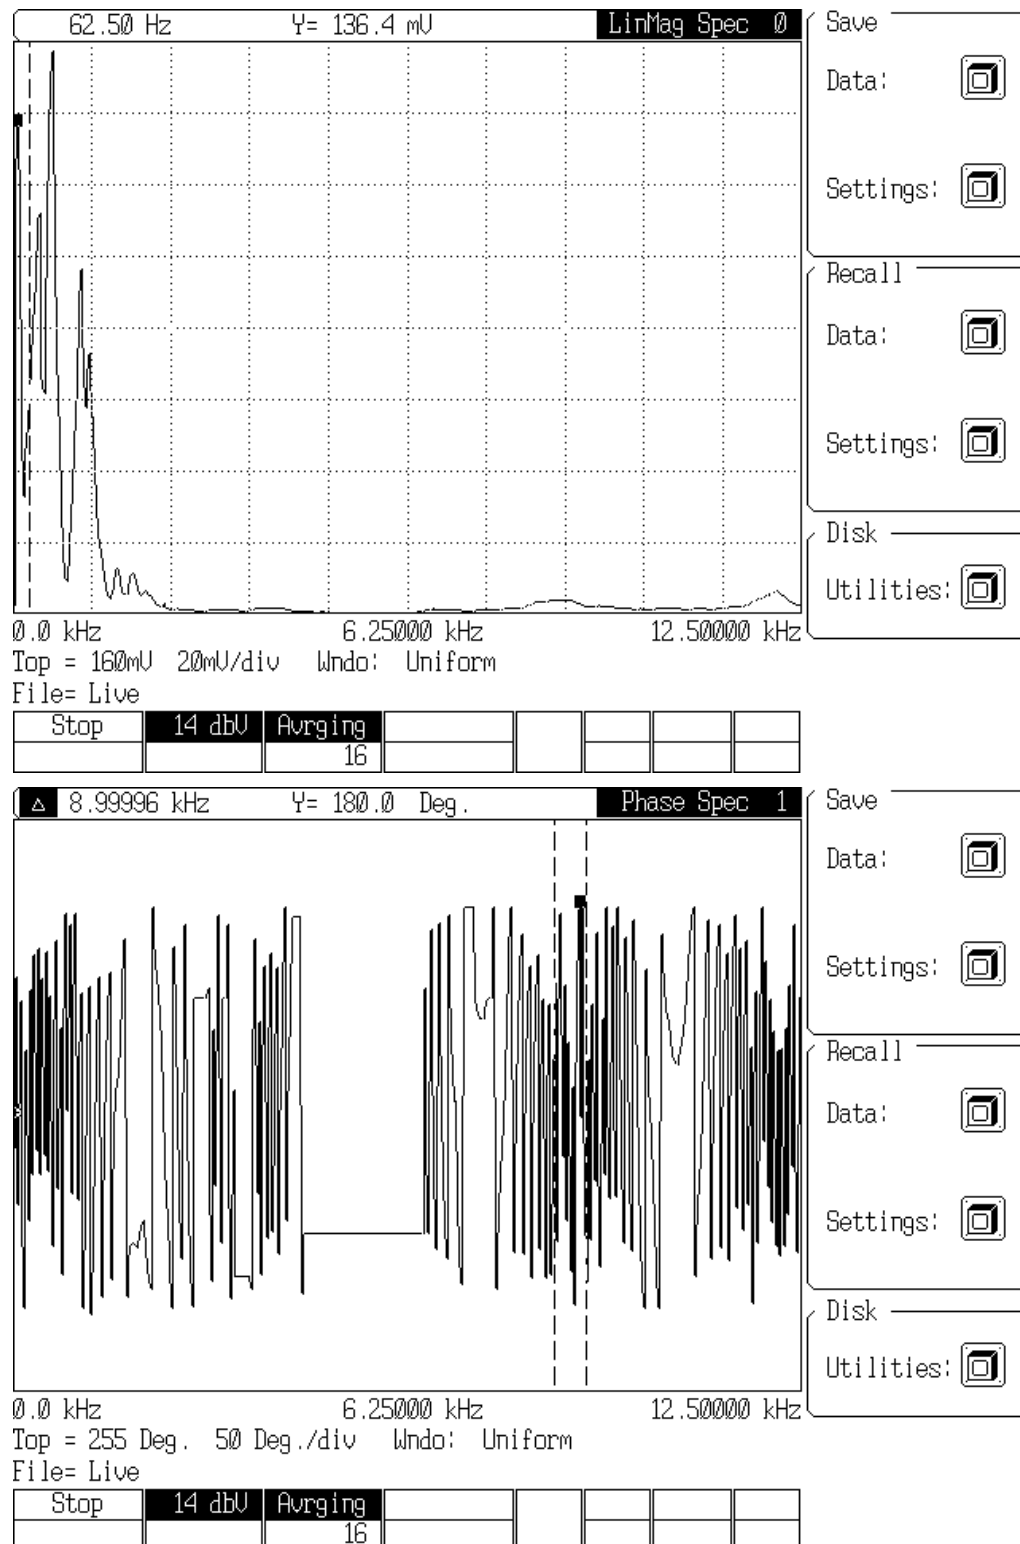
Figure 5.14: FFT Analyzer Screen Shots

5.4 Magnitude Over-Estimation

In all simulations shown, the magnitude of the simulated spectra is significantly higher than those in the measurements. Also, for higher frequencies, an increasing shift in the system resonances can be observed.
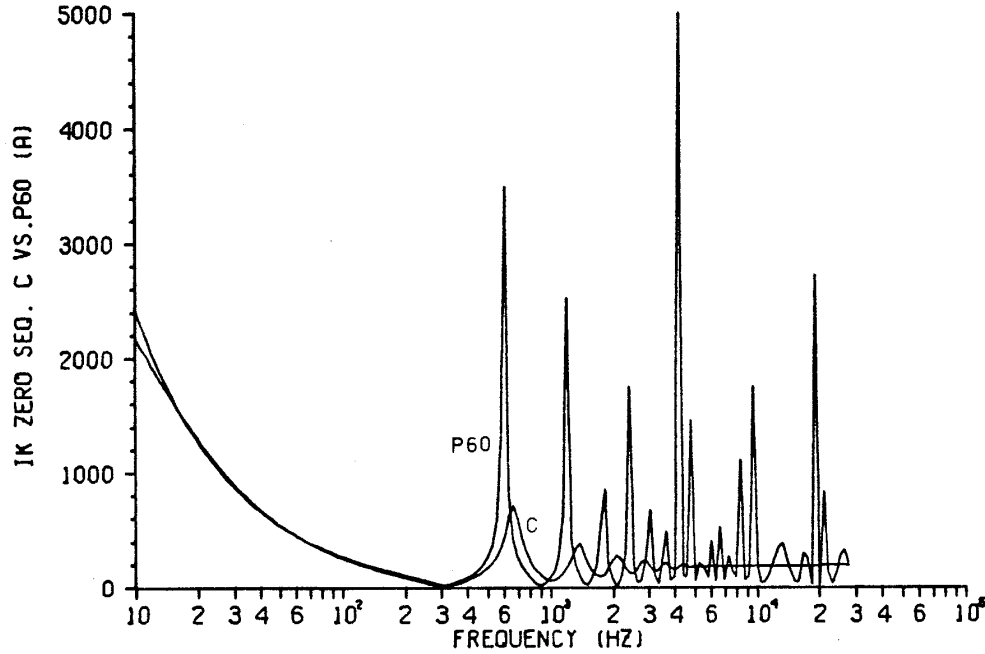
P.5.2(a): S/C Response, zero seq. Exact vs. 60-Hz.
Mid to high frequencies.

Figure 5.15: Illustration of Loss-Less Model Errors

This behavior is typical of the use of loss-less models in the computation. This phenomena was studied by Martí [50, 51] during the development of the frequency-dependent line model in the EMTP. A relevant plot from [50] illustrating this issue is reproduced, with permission from the author, in Figure 5.15. In this figure, a solution for a 60 Hz power system with frequency-independent parameters for 60 Hz (P60) is compared with a computation using exact parameters (C). It was found that the addition of correct, frequency-dependent losses resolved the mismatch.

In the computations shown in this thesis, the use of the distortion-less line model did result in a more accurate computation, but due to the frequency-independent nature of those losses, the improvement was limited. Frequency-dependent losses are important for acoustic phenomena in air, as the attenuation for 1 kHz is approximately 5dB/km, for 10 kHz approximately 95 dB/km, and for 100 kHz approximately 4000 dB/km. Thus, these losses are both significant

and far from constant over frequency. Future work would definitely require the adaptation of the frequency-dependent transmission line model to the TINA simulator.

5.5 Conclusions

The TINA method works quite well, in that the simulation results give a reasonably close agreement with the measured data, when those measured systems present a good 2D approximation. The resonances are over-estimated, which is due to the absence of losses in the current TINA media models. Without losses, resonances exhibit a too high Q-factor, and thus excessive magnitude. However, the predicted frequencies show good agreement.

Future work would include the development of lossy media models and better techniques of determining the acoustic parameters of the media in use. Also, the band-limited impulse method for transfer function computation requires further development, give the highly encouraging experiences with this method throughout the thesis work in obtaining clean transfer functions from the simulations.

Chapter 6

Conclusions and Future Work

> Before I came here, I was confused about this subject. Having listened to your
> lecture, I am still confused. But on a higher level. *Enrico Fermi*

In this chapter, we will summarize the main contributions of the thesis, as well as out-line the
most pertinent future work.

6.1 Conclusion

The TINA method proved to work reasonably well. Although the method requires more com-
putation than the basic TLM method it is directly related to, it offers increased flexibility and,
when fractional latency is used, may even out-perform the basic TLM method in computational
speed. Synchronicity through interpolation performs well, accuracy wise. The constituent, 1D
loss-less line models were characterized analytically and shown to be numerically stable. Since
the TINA method is an EMTP-derived circuit solution, it is also conserves energy. From those
studies, accuracy requirements were formulated.

Through a number of experiments, the TINA approach was validated. Comparisons with
EMTP showed the algorithm to be implemented correctly, while comparisons with various
measurements showed the accuracy, but also the limitations of a 2D simulator. Also, to perform
the comparisons, the band-limited impulse method for transfer function calculation proved an
excellent approach, yielding clean spectra with a high signal to noise ratio.

All in all, the loss-less TINA method performed as expected. Any resonances found were
over-predicted, due to lack of losses, but the frequencies were correct. The addition of losses
will resolve this issue. Also, the fractional latency technique worked quite well, except it
showed poor long-term stability. This is likely due to numerical issues, and further work will
be required to identify and fix these issues, as the method shows high promise, being able
to obtain good results at significant, order of magnitude, computational speed increases. The
addition of losses did resolve the long-term stability issue

6.2 Contributions

- The development of TINA, a TLM-like method that allows the direct use of EMTP-style transmission line models in the mesh, as well as any other circuit in its cells. Each such cell is interfaced with the simulation though Thévenin equivalents, and performs its own internal solution. Through the pervasive use of transmission-lines, each such cell is time-decoupled from the next, allowing each cell to be solved independent from the other.

- The detailed, analytical study of the EMTP loss-less line model, both in its ideal and interpolated form. Accuracy equations were derived, allowing the estimation of the numerical precision of the model under a given set of operating conditions. This study has not yet been reported in the literature and showed, amongst others, where the EMTP rule of thumb for the interpolated line, that its traveling time must be five to ten times the simulation time step, originates. The models were shown to be stable. The interpolated model was shown to be dissipative, as a low-pass filter, with the error also depending on where in the interpolation interval the result needs to be found. Also, the termination conditions of the line were found to have a significant contribution to the total error.

- Latency techniques were expanded to fractional ratios. Use of the new fractional latency technique results in an event-drive, asynchronous computation that requires special techniques to not only solve, but also extract information from. The concepts of absolute time and transmission-line look-ahead were introduced. The latter is an exact prediction technique based on the time-decoupling offered by the transmission lines.

- A new technique to obtain impulse responses in simulation, based on band-limited, periodic impulses. By matching these pulses, and their spectral content, to the FFT used to obtain the spectrum from the time-domain TINA solution, a leakage-free spectrum may be obtained with any desired spectral resolution. Also, since there is no spectral leakage, the signal to noise ratio is high. A computationally efficient method to compute the pulses in the time domain is also shown.

- The 2D TINA method, and thus by extension also the 2D TLM method, was verified for a number acoustic experiments. It was found that good, 2D approximations require much care to construct, but that the 2D computations could give good agreement with them. Losses in the experimental configurations were relatively high, compared to the loss-less computations, resulting in over-estimation in the simulation transfer functions, compared to the measured ones. However, the predicted resonance frequencies were well in agreement.

6.3 Future Work

- The addition of lossy, frequency-dependent, and non-linear material cells based on EMTP line models.

- Extension of TINA to 3D. This will require parallelization of the method to maintain reasonable computation times.

- Parallelization of the TINA method. A parallel TINA solver core was written, using OpenMP, but was not fully completed.

- Parallelization of the TINA method on a GPU. The self-contained, time-decoupled cell solution appears an ideal fit for computation on a GPU, which has literally hundreds of small arithmetic units that can perform simple operations on a small data set. These are an ideal fit for the cell computations, which are repetitive, simple operations to solve the internal circuit and the line models and do not require much, if any, branching.

- Identification of methods to find the required model parameters for the lossy models, from real-world materials. Maybe impedance tube experiments could be used for this purpose [17].

- Further development of the error expressions for the interpolated line models under fully generalised termination conditions.

- Further study of the stability, dispersion, and dissipation of more advanced line models, such as the CP (constant-parameter) and FD (frequency-dependent) line models.

- Further study of the fractional latency technique to characterize it better and determine the cause of the long-term stability issues.

- Investigate the use of integration rule switching to criticaly damp the long-term instabilities found in detailed fractional latency simulations.

- Further development of the theory of the BLImp (band-limited impulse response) method, and comparison with regular steps and impulses, as well as Gaussian pulses. The method has already been extended to include impulse, steps, and ramps, but not properly characterized yet, and evaluated for its effects on the stability of a computation, or in the presence of non-linear elements in a computation.

- Further improvements in the experimental configurations to provide better 2D approximations and to obtain more precise material parameters on-site, such as the speed of sound and losses.

Bibliography

[1] D.A. Al-Mukhtar and J.E. Sitch. Transmission-line matrix method with irregularly graded space. *IEE Proceedings H on Microwaves, Optics and Antennas*, 128(6):299–305, December 1981.

[2] Peter Anderson. Vocal-tract acoustics: an experimental approach. Technical Report Mech 543 project report, The University of British Columbia, May 2007.

[3] Howard Anton and Chris Rorres. *Elementary Linear Algebra*. John Wiley & Sons, Inc., USA, 9th edition, 2005.

[4] Jean-Pierre Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185–200, October 1994.

[5] Jean-Pierre Bérenger. A multiwire formalism for the FDTD method. *IEEE Transactions on Electromagnetic Compatibility*, 42(3):257–264, August 2000.

[6] L. Bergeron. *Water Hammer in Hydraulics and Wave Surges in Electricity*. John Wiley & Sons, New York, USA, 1961.

[7] L. V. Bewley. *Traveling Waves on Transmission Systems*. General Electric associates. John Wiley & Sons, Inc., New York, USA, 2nd edition, 1951.

[8] William E. Boyce and Richard C. DiPrima. *Elementary Differential Equations and Boundary Value Problems*. John Wiley & Sons, Inc., New York, USA, 4th edition, 1986.

[9] M. Celuch-Marcysiak, W.K. Gwarek, Zhizhang Chen, M.M. Ney, and W.J.R. Hoefer. Comments on 'a new finite-difference time-domain formulation and its equivalence with the tlm symmetrical condensed node' (and reply). *IEEE Transactions on Microwave Theory and Techniques*, 41(1):168–172, January 1993.

[10] Gary Ka-Yue Chan. Prediction of low-frequency sound-pressure fields in fitted rooms for active noise control. MASc thesis, Department of Mechanical Engineering, The University of British Columbia, Vancouver, British Columbia, Canada, December 2007.

[11] Zhizhang Chen, Michel M. Ney, and Wolfgang J. R. Hoefer. A new finite-difference time-domain formulation and its equivalence with the TLM symmetrical condensed node. *IEEE Transactions on Microwave Theory and Techniques*, 39(12):2160–2169, December 1991.

[12] Weng Cho Chew and William H. Weedon. A 3D perfectly matched medium from modified maxwell's equations with stretched coordinates. *Microwave and Optical Technology Letters*, 7(13):599–604, September 1994.

[13] D. Christensen. *Ultrasonic Bioinstrumentation*. John Wiley and Sons, USA, 1988.

[14] Christos Christopoulos. *The Transmission-Line Modeling Method: TLM*. IEEE Press, Piscataway, New Jersey, USA, 1995.

[15] Christos Christopoulos. Multi-scale modelling in time-domain electromagnetics. *International Journal of Electronics and Communications*, 57(2):100–110, 2003.

[16] Christos Christopoulos. *The Transmission-Line Modeling (TLM) Method in Electromagnetics*. Number 7 in Synthesis Lectures on Computational Electromagnetics. Morgan & Claypool Publishers, USA, 2006.

[17] Convenor. Acoustics - determination of sound absorption coefficient and impedance in impednace tubes - part 2: Transfer-function method. Technical Report ISO/CD 10534-2, R.W. Guy Working Group 14 TC43/SC2, 1998.

[18] Donard de Cogan. *Transmission Line Matrix (TLM) Techniques for Diffusion Applications*. Gordon and Breach Science Publishers, Amsterdam, The Netherlands, 1998.

[19] Tom De Rybel, José R. Martí, and Murray Hodgson. Analytical validation of time-step interpolation in transient insular nodal analysis. In *Acoustics'08 Paris*, number ACOUSTICS2008/177 in NS30-4aNSc, page 5, Paris, France, 29 June - 4 July 2008. Acoustics'08.

[20] Hermann W. Dommel. *EMTP Theory Book*. Microtran Power Systems Analysis Corporation, Vancouver, British Columbia, Canada, 2nd edition, 1996.

[21] Fredrik Edelvik. A new technique for accurate and stable modeling of arbitrarily oriented thin wires in the FDTD method. Technical Report 2002-016, Uppsala University, Uppsala, Sweden, April 2002. `http://www.it.uu.se/research/publications/reports/2002-016/`.

[22] Albert Einstein. *Relativity, the Special and the General Theory*. Three Rivers Press, New York, USA, 1961.

[23] Richard P. Feynman. *QED: The Strange Theory of Light and Matter*. Penguin Books, 80 Strand, London, WC2R 0RL, England, 1990.

[24] GotoBLAS. `http://www.cs.utexas.edu/users/flame/goto/`.

[25] Jacques Hadamard. *Lectures on Cauchy's Problem in Linear Partial Differential Equations*. Yale University Press, New Haven, Connecticut, USA, 1921.

[26] Jonathan L. Herring and Christos Christopoulos. Multigrid transmission-line models for solving electromagnetic field problems. *Electronics Letters*, 27:1794–1795, September 1991.

[27] Jonathan L. Herring and Christos Christopoulos. Solving electromagnetic field problems using a multiple grid transmission-line modeling method. *IEEE Transactions on Antennas and Propagation*, 42(12):1654–1658, December 1994.

[28] Richard Holland and Larry Simpson. Finite-difference analysis of EMP coupling to thin struts and wires. *IEEE Transactions on Electromagnetic Compatibility*, EMC-23(2):88–97, May 1981.

[29] J. A. Hollman and J. R. Martí. Real time network simulation with PC-cluster. *IEEE Transactions on Power Systems*, 18(2):563–569, May 2003.

[30] Mark N. Horenstein. *Microelectronic Circuits and Devices*. Prentice-Hall Inc., Englewood Cliffs, New Jersey, USA, 2nd edition, 1995.

[31] Peter Johns and Raymond Beurle. Numerical solution of 2-dimensional scattering problems using a transmission line matrix. *Proceedings of the IEE*, 188(9):1203–1208, 1971.

[32] Nicolai M. Josuttis. *The C++ Standard Library: a Tutorial and Reference*. Addison-Wesley, USA, 1999.

[33] L. Kinsler, A. Frey, A. Coppens, and J. Sanders. *Fundamentals of Acoustics*. John Wiley & Sons, USA, 3rd edition, 1982.

[34] Heinz Otto Kreiss. Über sachgemässe cauchyprobleme. *Mathematica Scandinavica*, 13:109–128, 1963.

[35] Erwin Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, USA, 4th edition, 1979.

[36] Gabriel Kron. Equivalent circuits to represent the electro-magnetic field equations. *Phys. Rev.*, 64(3-4):126–128, August 1943.

[37] Gabriel Kron. Equivalent circuit of the field equations of Maxwell - I. *Institute of Radio Engineers Proceedings*, 32:289–299, May 1944.

[38] Prabha Shankar Kundur. *Power System Stability and Control*. The EPRI Power System Engineering Series. McGraw-Hill, New York, USA, 1994.

[39] Karl S. Kunz and Larry Simpson. A technique for increasing the resolution of finite-difference solutions of the Maxwell equation. *IEEE Transactions on Electromagnetic Compatibility*, EMC-23(4):419–422, November 1981.

[40] Peter D. Lax and Robert D. Richtmyer. Survey of the stability of linear finite difference equations. *Communications on Pure and Applied Mathematics*, 9(2):267–293, February 1956.

[41] G. Ledfelt. A stable subcell model for arbitrarily oriented thin wires for the FDTD method. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 15(5-6):503–515, September 2002.

[42] E. Lelarasmee, A.E. Ruehli, and A.L. Sangiovanni-Vincentelli. The waveform relaxation method for time-domain analysis of large integrated circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 1(3):131–145, July 1982.

[43] L. R. Linares and J. R. Martí. Sub-area latency in a real-time power network simulator. pages 541–545, Lisbon, Portugal, 1995. International Conference on Power System Transients.

[44] Stanley B. Lippman, Josée Lajoie, and Barbara E. Moo. *C++ Primer*. Addison-Wesley, Upper Saddle River, New Jersey, USA, 4th edition, 2005.

[45] Hendrik Lorentz, A., Albert Einstein, Hermann Minkowski, and Hermann Klaus Hugo Weyl. *The Principle of Relativity, a Collection of Original Memoirs on the Special and General Theory of Relativity*. Dover Publications, Inc., 31 East 2nd Street, Mineola, New York 11501, USA, 1952.

[46] Philip Cooper Magnusson. *Transmission Lines and Wave Propagation*. Allyn and Bacon, Boston, USA, 1965.

[47] N. Marcovitz and J. Schwinger. On the reproduction of the electric and magnetic fields produced by currents and discontinuities in wave guides - I. *Journal of Applied Physics*, 22(6):806–819, June 1951.

[48] J.R. Marti and J. Lin. Suppression of numerical oscillations in the EMTP. *IEEE Transactions on Power Systems*, 4(2):739–747, May 1989.

[49] J. R. Martí, L. R. Linares, J. A. Hollman, and F. A. Moreira. OVNI: Integrated software/hardware solution for real-time simulation of large power systems. Seville, Spain, 24-28 June 2002. PSCC'02, Power Systems Computation Conference.

[50] José R. Martí. *The Problem of Frequency Dependence in Transmission Line Modelling*. Phd dissertation, Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, British Columbia, Canada, April 1981.

[51] José R. Martí. Accurate modelling of frequency-dependent transmission lines in electromagnetic transient simulations. *IEEE Transactions on Power Apparatus and Systems*, PAS-101(1):147–157, January 1982.

[52] José R. Martí. *EECE560 Class Notes*. The University of British Columbia, Vancouver, British Columbia, Canada, 1994.

[53] Maxima, a computer algebra system. `http://maxima.sourceforge.net/`.

[54] James Clerk Maxwell. *A Treatise on Electricity and Magnetism*, volume 1 of *Oxford Classic Texts in the Physical Sciences*. Oxford University Press, Great Clarendon Street, Oxford OX2 6DB, UK, 3rd edition, 2002.

[55] James Clerk Maxwell. *A Treatise on Electricity and Magnetism*, volume 2 of *Oxford Classic Texts in the Physical Sciences*. Oxford University Press, Great Clarendon Street, Oxford OX2 6DB, UK, 3rd edition, 2002.

[56] Microtran. `http://www.microtran.com/`.

[57] Gaspar Monge. *Essays on Descriptive Geometry*. 1799.

[58] F. Richard Moore. *Elements of Computer Music*. Prentice Hall, Englewood Cliffs, New Jersey, 07632, USA, 1990.

[59] F.A. Moreira and J.R. Marti. Latency techniques for time-domain power system transients simulation. *Power Systems, IEEE Transactions on*, 20(1):246–253, February 2005.

[60] Fernando A. Moreira. *Latency Techniques in Power Systems Transients Simulation*. Phd dissertation, Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, British Columbia, Canada, 2002.

[61] Robert Oerter. *The Theory of Almost Everything: The Standard Model, the Unsung Triumph of Modern Physics*. Pi Press Books, 1185 Avenue of the Americas, New York, 10036, USA, 2006.

[62] Katsuhiko Ogata. *State Space Analysis of Control Systems*. Prentice Hall, Englewood Cliffs, New Jersey, USA, 1967.

[63] Alan V. Oppenheim, Alan S. Willsky, and S. Hamid Nawab. *Signals & Systems*. Prentice Hall, Upper Saddle River, New Jersey, USA, 2nd edition, 1996.

[64] Alan V. Oppenheim, Ronald W. Schafer, and John R. Buck. *Discrete-Time Signal Processing*. Prentice Hall, Upper Saddle River, New Jersey, USA, 2nd edition, 1999.

[65] Max Planck. Über irreversible strahlungsvorgänge. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, 1:440–480, June 1899.

[66] J. A. Portí and J. A. Morente. A three-dimensional symmetrical condensed TLM node for acoustics. *Journal of Sound and Vibration*, 241(2):207–222, March 2001.

[67] J. A. Portí and J. A. Morente. TLM method and acoustics. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 14(2):171–183, March 2001.

[68] J.A. Porti, J.A. Morente, M. Khalladi, and A. Gallego. Comparison of thin wire models for TLM method. *Electronics Letters*, 28(20):1910–1911, September 1992.

[69] S. H. Pulko, I. A. Halleron, and C. P. Phizacklea. Substructuring of space and time in TLM diffusion applications. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 3(3):207–214, September 1990.

[70] QUCS: Quite universal circuit simulator. `http://qucs.sourceforge.net/`.

[71] Simon Ramo, John R. Whinnery, and Theodore Van Duzer. *Fields and Waves in Communication Electronics*. John Wiley & Sons, Inc., New York, USA, 3rd edition, 1993.

[72] Robert D. Richtmyer and K. W. Morton. *Difference Methods for Initial-Value Problems*. Wiley Interscience, New York, USA, 2nd edition, 1967.

[73] Matthew N. O. Sadiku and Lawrence C. Agba. A simple introduction to the transmission-line modeling. *IEEE Transactions on Circuits and System*, 37(8):991–999, August 1990.

[74] P. Saguet and E. Pic. Le maillage rectangulaire et le changement de maille dans la méthode TLM en deux dimensions. *Electronics Letters*, 17(7):277–279, February 1981.

[75] R.A. Saleh and A.R. Newton. The exploitation of latency and multirate behavior using nonlinear relaxation for circuit simulation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 8(12):1286–1298, December 1989.

[76] Bahaa E. A. Saleh and Marvin Carl Teich. *Fundamentals of Photonics*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 2nd edition, 2007.

[77] Ashely Saulsbury, Fong Pong, and Andreas Nowatzyk. Missing the memory wall: The case for processor/memory integration. In *ISCA'96*, pages 90–101, Pennsylvania, USA, 22-24 May 1996.

[78] Adam Semlyen and Francisco de Léon. Computation of electromagnetic transients using dual or multiple time steps. *IEEE Transactions on Power System*, 8(3):1274–1281, August 1993.

[79] Sundaram Seshu and Norman Balabanian. *Linear Network Analysis*. John Wiley & Sons, New York, USA, 1959.

[80] Albert Shadowitz. *The Electromagnetic Field*. Dover Publications, Mineola, New York, USA, 1975.

[81] Lee Smolin. *The Trouble With Physics: The Rise of String Theory and the Fall of a Science, and What Comes Next*. First Mariner Books, 215 Park Avenue South, New York, 10003, USA, 2006.

[82] Darrol Stinton. *The Design of the Aeroplane*. Blackwell Science, 9600 Garsington Road, Oxford, OX4 2DQ, UK, 2nd edition, 2001.

[83] John C. Strikwerda. *Finite Difference Schemes and Partial Differential Equations*. Society for Industrial and Applied Mathematics, University City Science Center, Philadelphia, Passedena, USA, 2nd edition, 2004.

[84] Steven H. Strogatz. *SYNC: How Order Emerges From Chaos in the Universe, Nature, and Daly Life*. Theia, an imprint of Hyperion Books, New York, USA, 2003.

[85] F.L. Teixeira and W.C. Chew. General closed-form PML constitutive tensors to match arbitrary bianisotropic and dispersive linear media. *IEEE Microwave and Guided Wave Letters*, 8(6):223–225, June 1998.

[86] Paul A. Tipler. *Physics for Scientists and Engineers*, volume 1. Worth Publishers, 33 Irving Place, New York, 10003, USA, 3rd edition, 1991.

[87] M. A. Tomim, J. R. Martí, and Wang L. Parallel solution of large power system networks using the multi-area thévenin equivalents (MATE) algorithm. *International Journal of Electrical Power and Energy Systems*, 31(9):497–503, October 2009.

[88] Marcelo A. Tomim. *Parallel computation of large power system networks using the multi-area Thévenin equivalents*. Phd dissertation, Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, British Columbia, Canada, 2009.

[89] Robert H. Voelker and Ronald J. Lomax. A finite-difference transmission line matrix method incorporating a nonlinear device model. *IEEE Transactions on Microwave Theory and Techniques*, 38(3):302–312, March 1990.

[90] Hermann W. Dommel. Digital computer solution of electromagnetic transients in single- and multiphase networks. *IEEE Transactions on Power Apparatus and Systems*, PAS-88(4):388–399, April 1969.

[91] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Oxford University Press, Amen House, London, England, 1965.

[92] A.J. Wlodarczyk, V. Trenkic, R.A. Scaramuzza, and C. Christopoulos. A fully integrated multiconductor model for TLM. *IEEE Transactions on Microwave Theory and Techniques*, 46(12):2431–2437, December 1998.

[93] D. P. Wlodarczyk, A. J.; Johns. New wire interface for graded 3D TLM. *Electronics Letters*, 28(8):728–729, April 1992.

[94] J. Wlodarczyk. New multigrid interface for the tlm method. *Electronics Letters*, 32(12):1111–1112, June 1996.

[95] Chi Chung Wong and Wing Sun Wong. Multigrid TLM for diffusion problems. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 2(2):103–111, June 1989.

[96] K. Yee. Numerical solution of initial value problems involving Maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, May 1966.

[97] A.I. Zecevic and N. Gacic. A partitioning algorithm for the parallel solution of differential-algebraic equations by waveform relaxation. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 46(4):421–434, April 1999.

[98] Svetlana S. Zivanovic, Kane Yee, and Kenneth K. Mei. A subgridding method for the time-domain finite-difference method to solve Maxwell's equations. *IEEE Transactions on Microwave Theory and Techniques*, 39(3):471–479, March 1991.

[99] V. I. Zubov. *Methods of A.M. Lyapunov and their Applications*. P. Noordhoff Ltd., Groningen, The Netherlands, 1964. Translation prepared under the auspices of the United States Atomic Energy Commission, edited by Leo F. Boron.

Appendix A

Derivation of the 1D Transmission Line Models in Time and Frequency Domain

> Life can only be understood backwards, but it must be lived forwards.
> *Soren Kierkegaard*

EMTP-style transmission line models were used and manipulated throughout this thesis. Given the pivotal role of these models in this research, their origins and derivation will now be discussed.

The EMTP models have a long history, starting with the work by Bergeron in the early thirties on water hammer in hydraulics. He demonstrated that, to a certain degree of approximation, the one-dimensional hydrodynamics model he used has a similar form as the one dimensional wave equation used in electro-magnetics, with similar solutions. Bergeron noted this fact and explored it in [6], where he demonstrated the similarities between water hammer in hydraulic systems and lightning strikes in electrical systems. Bergeron's approach was consequently used by Dommel in his development of the EMTP [90] during the late sixties, due to the elegance and simplicity of the model, when translated in numerical form. This model became the standard model in EMTP, known as the constant parameter, or CP line model [20].

The CP line model is a one-dimensional transmission line approximation that incorporates the distributed line impedance, lumped losses, and a time delay, all independent of frequency. When the losses are set to zero, only the impedance and time delay remain. This is the core of the loss-less, energy propagation model used in this thesis. Since the EMTP line models, including those that have constant parameter and frequency-dependent behaviors [20, 51, 52], have this impedance and time delay model at the core, and the more complicated behavior can be "bolted on" the basic model, only the basic model will be discussed in this appendix for clarity.

A.1   The 1D Wave and Telegrapher's Equations

There are at least two ways by which the propagation electro-magnetic energy in one dimension can be considered. The more general approach derives from Maxwell's equations [54, 55] and

is a field approach, while a simpler and more intuitive approach derives the wave equations in terms of a distributed-parameter circuit of a transmission line, where the behavior of voltage and current waves traveling along uniform conductors is studied. The latter is an equivalent circuit approach. Both methods are well known and extensively treated in many references, e.g.: [7, 46, 52, 71, 80], and only the final results will be given here.
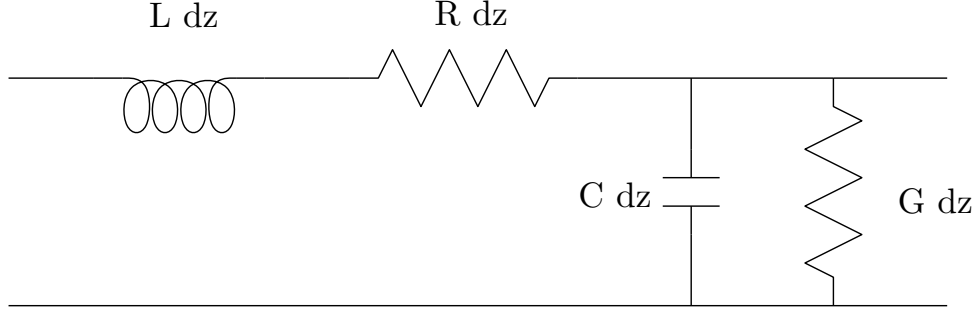
Figure A.1: Differential Length of Lossy Transmission Line

When solving for the one-dimensional, differential circuit equivalent from Figure A.1, we obtain the Telegrapher's equation, which has the same form as the corresponding field solution from Maxwell's equations, but different parameters are involved. This brings up the issue of field to circuit mapping, and is at the core of some of the complexities that occur when computing a field solution using networks. Kron discussed this in his work on electrical networks to approximate the Maxwell field equations [36, 37], and later Johns [18, 31] with TLM.

For the one-dimensional case, there is a direct one-to-one mapping, however, for higher dimensional cases, this is no longer the case and correction factors have to be introduced for some of the parameters to assure correct results. In this thesis, the topic of field to circuit mapping is discussed in Chapter 2 for the proposed method in 1D and 2D, and the correction parameters given where required.

Of importance for our discussion here are the equivalences between circuit and field parameters for the one-dimensional case. This is because the circuit parameters, expressed in terms of distributed resistance, inductance, and capacitance are measurable for a structure, but are hard to define for a bulk material. Conversely, the field parameters of conductivity, permeability, and permittivity are measurable for a bulk material, but not easily found for an arbitrary structure.

Since the TINA simulator is conceived around constructing shapes of interest from bulk materials in order to find the emergent behavior of the structure, it uses field parameters as input, but uses circuit equivalents internally for the component transmission lines that make-up the simulation grid. It thus uses these conversions internally.

Thus, both forms of the wave equation will be given here, as well as the conversion table between the various parameters in the 1D case. From here on, the discussion will, however, concentrate on the Telegrapher's equation form, as it relates more intuitively to the component transmission lines that are at the core of the TINA simulator, and are more familiar within the field of power systems than the field equations are.

From [46], the Telegrapher's (wave) equations for the lossy, distributed circuit in the z-direction, with R is the resistivity, G the conductance, L the inductance, and C the capacitance, as illustrated in Figure A.1 are:

$$
\frac{\partial^2 v(z,t)}{\partial z^2} = RGv(z,t) + (RC + LG)\frac{\partial v(z,t)}{\partial t} + LC\frac{\partial^2 v(z,t)}{\partial t^2} \tag{A.1a}
$$

$$
\frac{\partial^2 i(z,t)}{\partial z^2} = RGi(z,t) + (RC + LG)\frac{\partial i(z,t)}{\partial t} + LC\frac{\partial^2 i(z,t)}{\partial t^2} \tag{A.1b}
$$

while the wave equations from Maxwell's theory, with σ is the electric conductivity, μ is the magnetic susceptibility, and ϵ is the dielectric permittivity are:

$$
\frac{\partial^2 E_x(z,t)}{\partial z^2} = \sigma\mu\frac{\partial E_x(z,t)}{\partial t} + \epsilon\mu\frac{\partial^2 E_x(z,t)}{\partial t^2} \tag{A.2a}
$$

$$
\frac{\partial^2 H_y(z,t)}{\partial z^2} = \sigma\mu\frac{\partial H_y(z,t)}{\partial t} + \epsilon\mu\frac{\partial^2 H_y(z,t)}{\partial t^2} \tag{A.2b}
$$

We notice that there is no counterpart for the resistivity R in the field equations. When neglecting that term, we find that the Telegrapher's equations take the same form as the wave equations and the following correspondences are found:

Circuit		Field
$v(z,t)$	$\sim$	$E_x(z,t)$
$i(z,t)$	$\sim$	$H_y(z,t)$
G	$\sim$	σ
L	$\sim$	μ
C	$\sim$	ϵ

When all losses, thus R and G, or the corresponding σ, are neglected, we find the loss-less wave equation which forms the basis for the loss-less line models derived below. These loss-less, constant-parameter wave equations and resulting transmission-line model is quite universal, and has the same form across many diverse disciplines. Within the this thesis, this is taken advantage of to model acoustical phenomena, as the governing equations in 1D and 2D, for the loss-less and non-dispersive case have exactly the same form as those for electromagnetic fields, e.g.: [13, 33]. Note that this is no longer the case for 3D, as the acoustic waves

do not have polarization, and the model is just a simple extension of the 2D case, while the electro-magnetic model becomes much more complex to accommodate polarization.

A.2 The Ideal, Loss-Less Line Model

By solving the Telegrapher's equation (A.1) for the loss-less, constant-parameter case, the one-dimensional loss-less, constant-parameter transmission-line model can now be obtained [46, 52].

The model is based upon the d'Alembert solution to the wave equation (A.1), and considers the voltages and currents at the terminals of the line. Since the line is lossless, it effectively represents a time delay.

The characteristic impedance of the line, in the loss-less case, is found as: $Z_c = \sqrt{\frac{\mu}{\epsilon}}$ and the wave speed as $a = \frac{1}{\sqrt{\mu\epsilon}}$

A.2.1 Derivation in the Time Domain
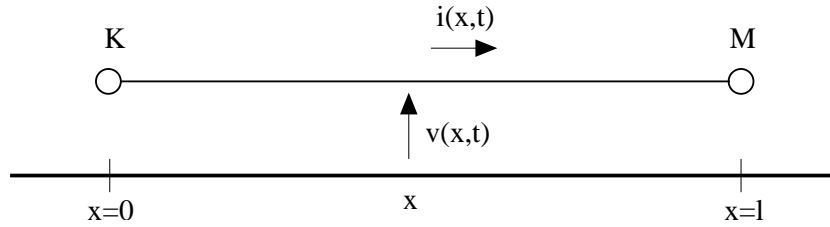
Figure A.2: Transmission Line

The d'Alembert's solution, for a sinusoidal waves, which presumes two waves traveling in the opposite direction.

$$p = p_+ \cos(\omega t - kz) \tag{A.3a}$$
$$p = p_- \cos(\omega t + kz) \tag{A.3b}$$

We can write the voltages and currents at any point at the line as follows [52]:

$$v(x,t) = v_f(x,t) + v_b(x,t) \tag{A.4a}$$
$$i(x,t) = i_f(x,t) + i_b(x,t) \tag{A.4b}$$

With f the forward wave, b is the backward wave. Now, combine both the voltage and the current together in one wave by relating the current to the voltage by means of the line

impedance. The wave speed is given by a.

$$v(x,t) \;=\; v_f(x-at) + v_b(x+at) \tag{A.5a}$$

$$i(x,t) \;=\; \frac{1}{Z_c}v_f(x-at) - \frac{1}{Z_c}v_b(x+at) \tag{A.5b}$$

Defining the traveling time $\tau = \frac{l}{a}$, with l the line length, we can add both equations (A.5), and write their total contribution for the forward perturbation wave for a point on the line (it suffers no reflection regardless of termination):

$$v(x,t) + Z_c i(x,t) = 2v_f(x-at) \tag{A.6}$$

From this, using the traveling time, we can express the terminating end of the line in function of the sending end. With (t) values the present time values and the $(t-\tau)$ values the history terms, we find:

$$v_m(t) + Z_c i_m(t) = v_k(t-\tau) + Z_c i_k(t-\tau) \tag{A.7}$$

A.2.2 Formulation of the EMTP time-domain CP-Line model
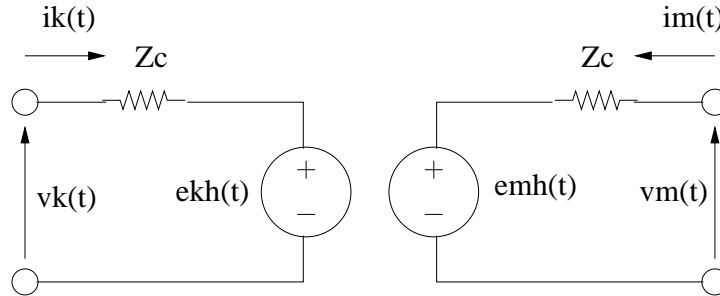
Figure A.3: CP Line Model

Notice that in (A.7), the current flows into one end and out the other, as shown in Figure A.2. This is not convenient in practical use, and the model will now be re-formulated to have the current flow into the model from both k and m sides. The resulting pair of equations, one for each side of the model corresponding to Figure A.3, are:

$$v_m(t) - Z_c i_m(t) = v_k(t-\tau) + Z_c i_k(t-\tau) \tag{A.8a}$$

$$v_k(t) - Z_c i_k(t) = v_m(t-\tau) + Z_c i_m(t-\tau) \tag{A.8b}$$

Re-formulated in a circuit equivalent with history sources we find:

$$e_{kh}(t) = v_m(t - \tau) + Z_c i_m(t - \tau) \qquad e_{mh}(t) = v_k(t - \tau) + Z_c i_k(t - \tau) \tag{A.9}$$

This model has many advantages over a direct solution of Maxwell's equations. It is not only numerically simple, it also time-decouples nodes K and M. It is this time decoupling of the equations that allows the TINA model to function and give it its flexibility. The line model is stable and exact, except for interpolation errors if τ is not an integer multiple of Δt, the simulation time step, as shown in Chapter 3, and no integration rule was explicitly used in the derivation.

A.2.3 Conversion into the Frequency Domain

Since part of the stability analysis was performed in the frequency domain, we will now show how to convert the time-domain form of the loss-less CP-line model into the frequency domain.

Conceptually, this comes down to recognizing that the model is essentially a pure delay. In the frequency domain, the exponential function is the equivalent of the delay in time domain and does not change the signal in any other way.

Thus, starting from (A.8), for valid for the circuit in Figure A.3, and using substitutions of the type $v(t) = \Re\{\overline{V} e^{j\omega t}\}$, we find, for each end of the line:

$$\Re\{\overline{V}_k e^{j\omega t}\} - Z_c \Re\{\overline{I}_k e^{j\omega t}\} = \Re\{\overline{V}_m e^{j\omega(t-\tau)}\} + Z_c \Re\{\overline{I}_m e^{j\omega(t-\tau)}\} \tag{A.10a}$$

$$\Re\{\overline{V}_m e^{j\omega t}\} - Z_c \Re\{\overline{I}_m e^{j\omega t}\} = \Re\{\overline{V}_k e^{j\omega(t-\tau)}\} + Z_c \Re\{\overline{I}_k e^{j\omega(t-\tau)}\} \tag{A.10b}$$

Grouping and re-arranging the terms, we find:

$$\left(\overline{V}_k - Z_c \overline{I}_k\right) e^{j\omega t} = \left(\overline{V}_m + Z_c \overline{I}_m\right) e^{j\omega t} e^{-j\omega\tau} \tag{A.11a}$$

$$\left(\overline{V}_m - Z_c \overline{I}_m\right) e^{j\omega t} = \left(\overline{V}_k + Z_c \overline{I}_k\right) e^{j\omega t} e^{-j\omega\tau} \tag{A.11b}$$

After simplification we find the final frequency-domain form of the loss-less line:

$$\overline{V}_k - Z_c \overline{I}_k = \left(\overline{V}_m + Z_c \overline{I}_m\right) e^{-j\omega\tau} \tag{A.12a}$$

$$\overline{V}_m - Z_c \overline{I}_m = \left(\overline{V}_k + Z_c \overline{I}_k\right) e^{-j\omega\tau} \tag{A.12b}$$

A.3 The Interpolated Loss-Less Line Model

Like most digital, time-marching systems, in EMTP and TINA time is implicit to the solution. The system only knows of the progression of discrete steps and it is the user who assigns meaning to those steps when processing the results. The above time-domain model computes a result at each simulation time step, using past values that were computed at prior simulation times. The amount of time in the past the model refers to is given by the length of the line τ. When that line length is an integer multiple of the simulation time step, a suitable past value exists in the history storage and a solution can be found. However, in a practical system, multiple materials are frequently required for an accurate representation of reality. These materials rarely have wave speeds that work out to integer multiples of each other's propagation time. As such, there is a problem of synchronicity, where different time steps would be required for different material regions within a simulation.

Thus, the models with non-matching transmission time would have to be forced into an integer multiple of the shared simulation time step, which changes their effective transmission time, and thus properties, or the models have to be allowed to operate at their own, local time step.

One way by which the synchronicity between the simulation time step and the model traveling time can be restored is the use of time-step interpolation. This is the method used in EMTP simulations of power systems [20, 52]. This approach can also be used for elements that are not transmission lines, and the simplicity of implementation and computation make it suitable for use in TINA.

A.3.1 The Need for Interpolation

When implemented on a digital computer, the expressions of (A.8) are programmed so that the past values in the equations, referred to by the indices $(t - \tau)$, are available in program memory in order for the results at the current time step (t) to be computed. For a given τ, which is defined by the length of the line l and its wave speed a, we can find the number of history values that need to be stored in memory from the simulation time step Δt:

$$\tau = \frac{a}{l} \tag{A.13}$$

$$history\ depth = \frac{\tau}{\Delta t} \tag{A.14}$$

When the τ of the line is an integer multiple of the simulation time step, this history depth will be an integer number, and the value is available in memory since data is saved in the history storage at each time step.

t(ms)	Vk	Ik	Vm	Im
0	~	~	~	~
0.02	~	~	~	~
0.04	~	~	~	~
0.06	~	~	~	~
0.08	~	~	~	~
0.10	~	~	~	~
0.12	~	~	~	~
0.14	~	~	~	~
0.16	~	~	~	~
0.18	~	~	~	~

(a) History Values in Memory, $\Delta t = 0.02$ ms

t(ms)	Vk	Ik	Vm	Im
0	~	~	~	~
0.03	~	~	~	~
0.06	~	~	~	~
0.09	~	~	~	~
0.12	~	~	~	~
0.15	~	~	~	~
0.18	~	~	~	~

(b) History Values in Memory, $\Delta t = 0.03$ ms

Figure A.4: Non-Interpolated and Interpolated History Tables

However, when the τ of the line is no longer an integer multiple of the time step, we require access to history values that lie in-between exiting data in order to compute the next contribution at the current time step (t). The history memory depth is then rounded-up to the next integer value to accommodate the time-step bound data and interpolation is used to obtain the required history value.

In Figure A.4(a), a simulation time step of $\Delta t = 0.02$ ms and a traveling time of $\tau = 0.10$ ms are used. It can be seen that, for a simulation time (t) of 0.18 ms, the required history value of 0.08 ms is available in the table and the current value can be readily computed.

In Figure A.4(b), a simulation time step of $\Delta t = 0.03$ ms is used. This time, the required history value of 0.08 ms is not available in the table, as it lies in between the available history values of 0.06 ms and 0.09 ms.

When the difference between two such history values is relatively small, it can be seen that the use of interpolation in time between the history values may be of use to find a good approximation for the variable at the required time. In EMTP, and TINA, linear interpolation is used for computational efficiency.

The effects of the rate of change (and thus difference between two points in the history table) on the accuracy of the linear interpolation is studied in detail in Chapter 3. It is found to be stable in all cases and dissipative. As an aside, since the model is now no longer ideal due to the use of interpolation, which works as a low-pass filter, the label "loss-less" no longer strictly applicable. However, with suitable simulation parameters, the errors can be kept small and the model considered loss-less and constant-parameter for practical purposes.

More complex interpolation schemes could be used, at the cost of numerical efficiency and possible stability issues. However, linear interpolation has proven to be sufficient in most cases

and, in those cases where it is not, the use of a smaller time step will alleviate the issues due to increased usable bandwidth and accuracy.

A.3.2 Derivation in the Time Domain

To use the linear interpolation [35], we must first establish the interval over which interpolation must be performed. Since this is at most one time step, the following figure can be established:
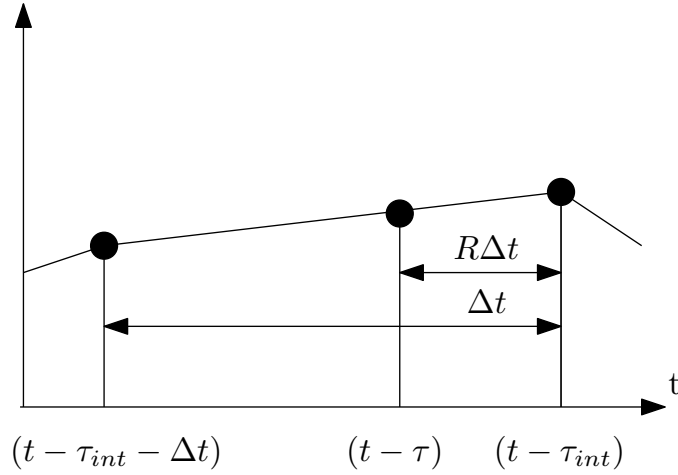
Figure A.5: Interpolation Interval

In this figure, two new variables are established: the interpolation factor R, which varies between $0 \leq R < 1$, and τ_{int}, which is the integer length of the transmission line in simulation time steps, rounded-up. These rounded-up in time, and synchronous with the simulation time step, history values are thus for a line that is longer than the desired one. The history value we actually need to find for the desired line length τ is that at $(t - \tau)$). Their expressions can be readily derived from the figure:

$$\tau_{int} = \left\lfloor \frac{\tau}{\Delta t} \right\rfloor \Delta t \tag{A.15}$$

$$R = \frac{\tau}{\Delta t} - \left\lfloor \frac{\tau}{\Delta t} \right\rfloor \tag{A.16}$$

Applying linear interpolation to the line model from (A.8), the interpolated equation pair

can be found through inspection by linearly interpolating to the voltage and current terms:

$$
\begin{aligned}
v_k(t) - Z_c i_k(t) \;=\; & v_m(t - \tau_{int}) & \text{(A.17a)} \\
& + R\left[v_m(t - \tau_{int} - \Delta t) - v_m(t - \tau_{int})\right] \\
& + Z_c i_m(t - \tau_{int}) \\
& + Z_c R\left[i_m(t - \tau_{int} - \Delta t) - i_m(t - \tau_{int})\right] \\
v_m(t) - Z_c i_m(t) \;=\; & v_k(t - \tau_{int}) & \text{(A.17b)} \\
& + R\left[v_k(t - \tau_{int} - \Delta t) - v_k(t - \tau_{int})\right] \\
& + Z_c i_k(t - \tau_{int}) \\
& + Z_c R\left[i_k(t - \tau_{int} - \Delta t) - i_k(t - \tau_{int})\right]
\end{aligned}
$$

A.3.3 Conversion into the Frequency Domain

The conversion of this model into frequency domain can be done in a similar way as in Section A.2.3 of this appendix. Starting from (A.17) we find:

$$
\begin{aligned}
\left(\overline{V}_k - Z_c \overline{I}_k\right) e^{j\omega t} \;=\; & \left(\overline{V}_m + Z_c \overline{I}_m\right) e^{j\omega(t - \tau_{int})} & \text{(A.18a)} \\
& + R\left(\overline{V}_m + Z_c \overline{I}_m\right) e^{j\omega(t - \tau_{int} - \Delta t)} \\
& - R\left(\overline{V}_m + Z_c \overline{I}_m\right) e^{j\omega(t - \tau_{int})} \\
\left(\overline{V}_m - Z_c \overline{I}_m\right) e^{j\omega t} \;=\; & \left(\overline{V}_k + Z_c \overline{I}_k\right) e^{j\omega(t - \tau_{int})} & \text{(A.18b)} \\
& + R\left(\overline{V}_k + Z_c \overline{I}_k\right) e^{j\omega(t - \tau_{int} - \Delta t)} \\
& - R\left(\overline{V}_k + Z_c \overline{I}_k\right) e^{j\omega(t - \tau_{int})}
\end{aligned}
$$

Re-ordering and simplifying, we find the final version of the frequency-domain form of the interpolated, loss-less line:

$$
\begin{aligned}
\left(\overline{V}_k - Z_c \overline{I}_k\right) &= \left(\overline{V}_m + Z_c \overline{I}_m\right) e^{-j\omega \tau_{int}} \left[1 + R\left(e^{-j\omega \Delta t} - 1\right)\right] & \text{(A.19a)} \\
\left(\overline{V}_m - Z_c \overline{I}_m\right) &= \left(\overline{V}_k + Z_c \overline{I}_k\right) e^{-j\omega \tau_{int}} \left[1 + R\left(e^{-j\omega \Delta t} - 1\right)\right] & \text{(A.19b)}
\end{aligned}
$$

A.4 The Distortion-Less Line Model

Although the focus of this thesis is on loss-less models, there was a need for a simple loss model in the TINA simulator to allow the fractional latency simulations to be stable, as well as find better agreement between the simulations and experiment. Ideally, this model would have

frequency-independent losses and be computationally efficient so its use does not degrade the speed benefits from the fractional latency method.

One such model, called the distortionless line model [20, 46], is of particular interest. The model is based on applying the Heaviside condition to the Telegrapher's equation, and modifies the ideal line model by multiplying it with a single, constant, real attenuation factor. This makes it computationally efficient, as it adds only two multiplications to each constant-parameter (CP) line segment, resulting in a lossy, constant-parameter, thus frequency-independent, model. In other words, these losses attenuate all frequencies equally in magnitude and have no influence on the phase. Hence, the shape of the signal is not distorted, only reduced.

The Heaviside condition is given by the following ratio between the line parameters for the Telegrapher's equation (A.1):

$$\frac{R}{L} = \frac{G}{C} \tag{A.20}$$

The propagation function γ, of which the real part α is the attenuation function and the imaginary part β is the phase function, is given by:

$$\gamma = \sqrt{(R + j\omega l)(G + j\omega C)} \tag{A.21}$$

Also, the phase velocity a and characteristic impedance Z_c are given by:

$$a = \frac{\omega}{\beta} \tag{A.22a}$$

$$Z_c = \sqrt{\frac{R + j\omega L}{G + j\omega C}} \tag{A.22b}$$

Applying the Heaviside condition insures the phase velocity and characteristic impedance of the line are the same as in the loss-less case, and the added losses are purely real and frequency-independent. The solutions are summarized in Table A.1.

	Parameters	Z_c	Phase Velocity a	Attenuation Function α
Loss-Less	$R = G = 0$	$\sqrt{\frac{L}{C}}$	$\frac{1}{\sqrt{LC}}$	0
Distortion-Less	$\frac{R}{L} = \frac{G}{C}$	$\sqrt{\frac{L}{C}}$	$\frac{1}{\sqrt{LC}}$	$R\sqrt{\frac{C}{L}}$

Table A.1: Comparison Between the Ideal and Distortion-Less Line Models

Applying this to the ideal and interpolated lines, the following expressions result, for the

ideal and interpolated models, respectively:

$$v_m(t) - Z_c i_m(t) = \{v_k(t - \tau) + Z_c i_k(t - \tau)\} e^{-\alpha} \tag{A.23a}$$

$$v_k(t) - Z_c i_k(t) = \{v_m(t - \tau) + Z_c i_m(t - \tau)\} e^{-\alpha} \tag{A.23b}$$

$$
\begin{aligned}
v_k(t) - Z_c i_k(t) = \ & \{v_m(t - \tau_{int}) & \text{(A.24a)} \\
& + R\left[v_m(t - \tau_{int} - \Delta t) - v_m(t - \tau_{int})\right] \\
& + Z_c i_m(t - \tau_{int}) \\
& + Z_c R\left[i_m(t - \tau_{int} - \Delta t) - i_m(t - \tau_{int})\right]\} e^{-\alpha} \\
v_m(t) - Z_c i_m(t) = \ & \{v_k(t - \tau_{int}) & \text{(A.24b)} \\
& + R\left[v_k(t - \tau_{int} - \Delta t) - v_k(t - \tau_{int})\right] \\
& + Z_c i_k(t - \tau_{int}) \\
& + Z_c R\left[i_k(t - \tau_{int} - \Delta t) - i_k(t - \tau_{int})\right]\} e^{-\alpha}
\end{aligned}
$$

Thus, the distortion-less models are the same as the ideal models, multiplied by a constant attenuation factor.

A.5 Reducing the Computational Load of the Ideal and Interpolated Line Models

In TINA, where simulations can easily have ten thousands to hundreds of thousands of cells, each with two (1D), four (2D), or twelve (3D) transmission line segments. It is thus important to achieve an efficient implementation of these models to maintain numerical performance as a naive implementation of the models quickly becomes ineffective as the computer's memory is depleted and swapping to hard disk becomes necessary. This greatly reduces the numerical efficiency of the method. In this section, these optimizations are shown that significantly mitigate the issues.

First, a more efficient form of the interpolated and ideal line models is shown, based on a formulation in voltage only, as done in the EMTP simulator. It reduces the memory use, and computational burden, compared to a naive implementation of the models. Second, the use of C-style circular buffers for the line model history term storage, and their performance impact, is discussed.

A.5.1 The Issue with Memory Access

When implementing the ideal and interpolated line equations, (A.8) and (A.17), respectively, there are a multitude of terms to store and manipulate in the model in order to compute the history sources.

Each of these variables must be stored to the history depth of the model. If the model has a τ of, e.g., ten simulation time steps Δt, this means that ten memory spaces per term in the histories are required. More so, these many terms must be accessed from computer main memory, as the on-chip CPU cache is too small to hold the large data sets encountered in TINA simulations.

Given that modern CPUs are so fast that they spend much of their time waiting for data to come from main memory when the required datum is not available in the limited on-chip cache [77], it becomes clear that all unnecessary access to main memory is to be avoided.

The memory and computational load of the models can, however, be reduced by a different formulation and implementation of the line equations. In EMTP, such a formulation is used where the line models are expressed only in voltage terms. This formulation, together with exploiting the time-decoupling in the line model, makes it possible to, instead of storing all the individual history terms, voltages, and currents at both ends of the lines, directly compute the history values themselves and only store these in the history sources.

A.5.2 Ideal Line Model in Voltage Form

The ideal line equations (A.17) are:

$$v_k(t) - Z_c i_k(t) \quad = \quad v_m(t - \tau) + Z_c i_m(t - \tau) \tag{A.25a}$$
$$v_m(t) - Z_c i_m(t) \quad = \quad v_k(t - \tau) + Z_c i_k(t - \tau) \tag{A.25b}$$

The left-hand side of the expressions are the history sources:

$$e_{kh}(t) \quad = \quad v_k(t) - Z_c i_k(t) \tag{A.26a}$$
$$e_{mh}(t) \quad = \quad v_m(t) - Z_c i_m(t) \tag{A.26b}$$

We now substitute the history sources on the left-hand side:

$$e_{kh}(t) \quad = \quad v_m(t - \tau) + Z_c i_m(t - \tau) \tag{A.27a}$$
$$e_{mh}(t) \quad = \quad v_k(t - \tau) + Z_c i_k(t - \tau) \tag{A.27b}$$

Applying KVL on the model (Figure A.3) terminals, we find:

$$v_k(t) \quad = \quad Z_c i_k(t) + e_{hk}(t) \tag{A.28a}$$

$$v_m(t) \quad = \quad Z_c i_m(t) + e_{hm}(t) \tag{A.28b}$$

Re-ordering the above expression:

$$Z_c i_k(t) \quad = \quad v_k(t) - e_{hk}(t) \tag{A.29a}$$

$$Z_c i_m(t) \quad = \quad v_m(t) - e_{hm}(t) \tag{A.29b}$$

Substituting (A.29) in (A.27), we find the required expressions in voltage only:

$$e_{hk}(t) \quad = \quad -e_{hm}(t - \tau) + 2v_m(t - \tau) \tag{A.30a}$$

$$e_{hm}(t) \quad = \quad -e_{hk}(t - \tau) + 2v_k(t - \tau) \tag{A.30b}$$

From this, it can be seen that this formulation requires less terms to be stored on the histories, and also less calculations, compared to the naive form of the models.

By computing the value to place on the sending end of the line at t, which is possible since at any time step, we can use the current values for the history source and node voltage and compute the future history value at the other end of the line $t + \tau$. It is this only a matter of computing this future result at the present time, and placing it on the history buffer. When t has progresses to $t + \tau$, this required value can then be read from the history buffers. The inherent time-decoupling of the model allows for this optimization.

As a result of these manipulations, only two history buffers are used for the entire model, as the node voltages v_m and v_k do not need to be stored. By comparison, the original model required six in a naive implementation where the node voltages and currents needed to be stored, in addition to the histories themselves. The voltage-only expressions allow for a number of numerical optimizations in the implementation of the solutions, reducing the computation load.

A.5.3 Interpolated Line Model in Voltage Form

Using the same procedure as for the ideal line model, we start with the equations for the interpolated model (A.17):

$$
\begin{aligned}
v_k(t) - Z_c i_k(t) \;=\; & v_m(t - \tau_{int}) && \text{(A.31a)} \\
& + R\left[v_m(t - \tau_{int} - \Delta t) - v_m(t - \tau_{int})\right] \\
& + Z_c i_m(t - \tau_{int}) \\
& + Z_c R\left[i_m(t - \tau_{int} - \Delta t) - i_m(t - \tau_{int})\right] \\
v_m(t) - Z_c i_m(t) \;=\; & v_k(t - \tau_{int}) && \text{(A.31b)} \\
& + R\left[v_k(t - \tau_{int} - \Delta t) - v_k(t - \tau_{int})\right] \\
& + Z_c i_k(t - \tau_{int}) \\
& + Z_c R\left[i_k(t - \tau_{int} - \Delta t) - i_k(t - \tau_{int})\right]
\end{aligned}
$$

We now substitute the history sources, (A.26), on the left-hand side:

$$
\begin{aligned}
e_{kh}(t) \;=\; & v_m(t - \tau_{int}) && \text{(A.32a)} \\
& + R\left[v_m(t - \tau_{int} - \Delta t) - v_m(t - \tau_{int})\right] \\
& + Z_c i_m(t - \tau_{int}) \\
& + Z_c R\left[i_m(t - \tau_{int} - \Delta t) - i_m(t - \tau_{int})\right] \\
e_{mh}(t) \;=\; & v_k(t - \tau_{int}) && \text{(A.32b)} \\
& + R\left[v_k(t - \tau_{int} - \Delta t) - v_k(t - \tau_{int})\right] \\
& + Z_c i_k(t - \tau_{int}) \\
& + Z_c R\left[i_k(t - \tau_{int} - \Delta t) - i_k(t - \tau_{int})\right]
\end{aligned}
$$

Applying KVL on the model (Figure A.3) terminals, we find:

$$
\begin{aligned}
v_k(t) &= Z_c i_k(t) + e_{hk}(t) && \text{(A.33a)} \\
v_m(t) &= Z_c i_m(t) + e_{hm}(t) && \text{(A.33b)}
\end{aligned}
$$

Re-ordering the above expression:

$$
\begin{aligned}
Z_c i_k(t) &= v_k(t) - e_{hk}(t) && \text{(A.34a)} \\
Z_c i_m(t) &= v_m(t) - e_{hm}(t) && \text{(A.34b)}
\end{aligned}
$$

Substituting (A.34) in (A.32), we find the required expressions for the interpolated model in voltage only:

$$
\begin{aligned}
e_{hk}(t) \;=\;& 2v_m(t - \tau_{int}) && \text{(A.35a)}\\
& -2R\left[v_m(t - \tau_{int}) - v_m(t - \tau_{int} - \Delta t))\right]\\
& -e_{hm}(t - \tau_{int})\\
& +R\left[e_{hm}(t - \tau_{int}) - e_{hm}(t - \tau_{int} - \Delta t)\right]\\
e_{hm}(t) \;=\;& 2v_k(t - \tau_{int}) && \text{(A.35b)}\\
& -2R\left[v_k(t - \tau_{int}) - v_k(t - \tau_{int} - \Delta t))\right]\\
& -e_{hk}(t - \tau_{int})\\
& +R\left[e_{hk}(t - \tau_{int}) - e_{hk}(t - \tau_{int} - \Delta t)\right]
\end{aligned}
$$

For the interpolated model, similar memory optimizations are achieved as for the interpolated model by calculating the history values at t for the future $t + \tau$. The added complication is that two more storage places are needed to keep track of the $t - \tau_{int} - \Delta t$ values of the node voltage, as well as history buffers that are one deeper. However, with careful implementation of the model, the extra storage for the node voltages v_m and v_k can be avoided.

As a result of these manipulations, and making use of the time decoupling of the line model, once more only two history buffers are needed for the entire model, as opposed to the six in a naive implementation.

A.5.4 Data Structures for the History Term Storage

The performance of the computation also depends strongly on how the history data is physically stored in the computer memory. Although there are a number of convenient and suitable data structures available in the C++ standard template library (STL) [32], we opted for the use of C-style circular buffers. In spite of required programming and extra computations to manage the data pointers of the circular buffers, in practice they performed much better. This resulted in a far more efficient implementation of these history terms than STL or object-based approaches [44] due to a much lower overhead.

A number of initial relative performance tests of these implementations, using the standard STL libraries in the GCC 4.3.3 and our own implementation of C-style circular buffers, when compared using TINA test cases, showed that the circular buffers were at least one order of magnitude faster and required far less memory, as they did not grow by themselves as the STL containers did. Given these initial test results, only the C-style circular buffers were used from then on in the code development.

The main issue with the STL containers for storage proved to be that, as data was pushed on one end, and data pulled from the other, the memory required kept growing. This happened even thought the data pulled from the container was no longer available, and should thus have been de-allocated. This behavior was further exasperated by the memory allocation methods used in STL, where the memory pool used by a container is doubled whenever the available storage is exceeded. Although this makes the container faster in most cases, as less re-allocation of memory is required, it did cause the machine to rapidly run out of resources in TINA, where histories are pushed on, and pulled off, at every time step. In addition, when the memory was not being depleted, the containers proved to have significant overhead in their operation.

Circular buffers, as implemented in TINA, by comparison only allocate their storage once and move the access pointers around instead. By using modulo operations, an increment or decrement of an access pointer causes it to loop-round when the bounds of the memory storage would be exceeded, causing previously used, but no longer required, storage to be over-written with new data. As such, no memory allocation/deallocation has to occur during computation and memory fragmentation is avoided. The circular buffer does come at the cost of increased computation to obtain the access pointers, but the reduction in memory bus operations and the availability of a fast CPU proved to make this approach significantly more performant, given the easily hundreds of thousands of these buffers in a given simulation.

Appendix B

Eigenvalue Analysis

> We may note that, in these experiments, the sign "=" may stand for the words "is confused with".
>
> *G. Spencer Brown*

Throughout the thesis, eigenvalue analysis is extensively used to obtain the modes, and thus stability information, of a system. They are also used to find system solutions, as eigenvalue analysis is a way to diagonalize a matrix, and thus solve the system of equations governing the system.

The flexibility of eigenvalue analysis has much to do with the nature of physical systems. The system response of linear, time-invariant systems is governed by complex coefficient exponential functions. These encapsulate the sine and cosine behavior of physical systems which, when super-imposed for each mode, as exemplified by the Fourier theorem, can faithfully reproduce the system behavior. These exponential functions represent the relevant modes of the system and, though diagonalization, are encoded in the eigenvalues. Thus, it is no longer required to completely solve the system to evaluate stability, only the eigenvalues need be extracted, as illustrated below in Section B.2.

B.1 Diagonalization on a Matrix

The subject of eigenvalues and their uses, such as matrix diagonalization, is treated by many authors. Here, we will adapt the treatment given by Boyce [8] to introduce the subject, augmented by other sources.

Consider the equation:

$$\mathbf{Q}\mathbf{x} = \mathbf{y} \tag{B.1}$$

In a general sense, (B.1) can be seen as a linear transformation that maps a vector $\mathbf{x}$ into a new vector $\mathbf{y}$, using the mapping encoded in matrix $\mathbf{Q}$. This is a very common type of equation that appears frequently in linear electrical network analysis, as such networks serve to map one set of variables into another by means of a transformation encoded in the circuit. In many cases, such a circuit can be expressed mathematically, e.g.: through nodal analysis, as a matrix, resulting in a system described by the form (B.1).

To find the vectors, we set $\mathbf{y} = \lambda\mathbf{x}$, where λ is a scalar proportionality factor. Now, we seek the solution to:

$$\mathbf{Q}\mathbf{x} = \lambda\mathbf{x} \tag{B.2a}$$

$$\Leftrightarrow (\mathbf{Q} - \lambda\mathbf{I})\,\mathbf{x} = \mathbf{0} \tag{B.2b}$$

This has non-zero solutions if and only if λ is chosen so that the polynomial p of λ:

$$p(\lambda) = det\,(\mathbf{Q} - \lambda\mathbf{I}) = \lambda^n + c_1\lambda^{n-1} + \cdot + c_n = 0 \tag{B.3}$$

Values of λ that satisfy the characteristic polynomial (B.3) are called eigenvalues of the matrix $\mathbf{Q}$. The solutions of (B.2) that are obtained using such a value of λ are called the eigenvectors corresponding to that eigenvalue.

Since (B.3) is a polynomial equation of degree n in λ, there are n eigenvalues $\lambda_1, \ldots, \lambda_n$, some of which may be repeated. If a given eigenvalue appears m times as a root of (B.3), then that eigenvalue is said to have multiplicity m. Each eigenvalue has at least one associated eigenvector.

If all eigenvalues of $\mathbf{Q}$ are simple (synonyms: multiplicity one, distinct), then all n eigenvectors of $\mathbf{Q}$ are linearly independent. If there are repeated eigenvalues, this may not be the case, as the number of linearly independent eigenvectors q for each eigenvalue with multiplicity m is now $1 \leq q \leq m$.

The preceding results can now be used to diagonalize a matrix, which has use in aiding system solution and finding the modes of a system.

This is the diagonalization problem [3]. Given an nxn matrix $\mathbf{Q}$, is there an invertible matrix $\mathbf{V}$ so that $\mathbf{V}^{-1}\mathbf{Q}\mathbf{V}$ is a diagonal matrix? Or, can we find a matrix $\mathbf{V}$ that diagonalizes $\mathbf{Q}$?

Assuming $\mathbf{Q}$ has a full set of n linearly independent eigenvectors $\mathbf{v}^1, \ldots, \mathbf{v}^n$ and corresponding eigenvalues $\lambda_1, \ldots, \lambda_n$, the eigenvectors can be used to form a matrix $\mathbf{V}$ whose columns are the eigenvectors $\mathbf{v}^1, \ldots, \mathbf{v}^n$. Since the columns of $\mathbf{V}$ are linearly independent vectors, $det(\mathbf{V}) \neq 0$. Thus, $\mathbf{V}$ is nonsingular and $\mathbf{V}^{-1}$ exists.

The columns of $\mathbf{Q}\mathbf{V}$ are the vectors $\mathbf{Q}\mathbf{v}^1, \ldots, \mathbf{Q}\mathbf{v}^n$. Since $\mathbf{Q}\mathbf{v}^n = \lambda_k\mathbf{x}^n$, it follows that:

$$\mathbf{Q}\mathbf{V} = \begin{bmatrix} \lambda_1 v_1^1 & \cdots & \lambda_n v_1^n \\ \vdots & & \vdots \\ \lambda_1 v_n^1 & \cdots & \lambda_n v_n^n \end{bmatrix} = \mathbf{V}\mathbf{\Lambda} \tag{B.4}$$

where

$$\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix} \tag{B.5}$$

is a diagonal matrix whose diagonal elements are the eigenvalues of $\mathbf{Q}$. From (B.4) it follows:

$$\mathbf{V}^{-1}\mathbf{Q}\mathbf{V} = \mathbf{\Lambda} \tag{B.6}$$

Thus, if the eigenvalues and eigenvectors of $\mathbf{Q}$ are known, $\mathbf{Q}$ can be transformed into a diagonal matrix. Note that, if $\mathbf{Q}$ has fewer than n linearly independent eigenvectors, there is no matrix $\mathbf{V}$ so that $\mathbf{V}^{-1}\mathbf{Q}\mathbf{V} = \mathbf{\Lambda}$. In this case, $\mathbf{Q}$ is not diagonalizable. Linear independence and diagonalizability are, in fact, equivalent [3].

B.2 Finding the Modes of a System

There is a direct relation between the modes of a linear system, described by an input-output relation or implicitly by a linear PDE, and the eigenvalues.

To show this relation, let us consider the input-output description of a linear system [76]. It is characterized by a linear operator $\mathbf{H}$ that operates on an input vector $\mathbf{x}$ to generate the corresponding output vector $\mathbf{y}$.
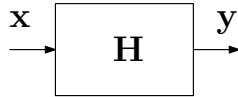
Figure B.1: Linear System

This system is described by the following relation:

$$\mathbf{y} = \mathbf{H}\mathbf{x} \tag{B.7}$$

The modes of such a system are the special inputs that are unaltered (except for a multiplicative constant) when passed through the system. For example:

This system is now described by the following relation:

$$\mathbf{H}\mathbf{x}_n = \lambda_n\mathbf{x}_n \tag{B.8}$$

where n is the number of the mode. The vector $\mathbf{x}_n$ is an eigenvector of the system. The
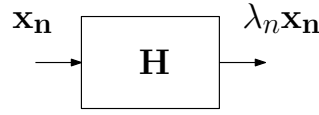
Figure B.2: Linear System with Eigenfunctions Applied

multiplication factor λ_n is an eigenvalue. The above expression is the eigenvalue problem for this system and can be solved using the methods described earlier in this appendix.

Considering a linear dynamic system described by N continuous variables constituting a vector $\mathbf{x}(t)$. The behavior of any of the N variables of this vector is, in general, dependent on all N variables. By describing the system in a new coordinate system, the N variables can be decomposed in N decoupled, one-dimensional systems. These decoupled variables are the modes of the system.

This is where eigenvalue analysis is useful, as it performs this coordinate transformation by means of the transformation matrix $\mathbf{V}$ (and back again, using the inverse) of a diagonal matrix $\boldsymbol{\Lambda}$ containing the eigenvalues, and thus modes, of the system, as follows: $\mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^{-}1 = \mathbf{H}$. Thus, the modes of the system can be readily obtained from the eigenvalues alone.

For a linear system described implicitly by a linear PDE that can be cast in the form (B.8), where $\mathbf{H}$ is a differential operator, the modes are the solutions of the differential equation, and the eigenvalues are called eigenfunctions. These "indestructible functions", which are only scaled by a multiplicative constant, will be of form: $e^{j\omega t}$. The notion of input and output is not meaningful in this context, as the modes are fully given by the solution of only the differential operator itself.

B.3   Computing the Matrix Exponential

Exponentials raised to the power of a matrix appear quite readily in the analysis of linear, time-invariant systems of differential equations. The main reason is found in the importance of the complex exponential as a governing function for the modes of oscillation of real systems. To illustrate the issue, consider a common first order ODE system of equations and its known solution in time domain:

$$\dot{\mathbf{x}} = \mathbf{Q}\mathbf{x} \Rightarrow \mathbf{x}(t) = Ce^{\mathbf{Q}t} \tag{B.9}$$

It is seen that the exponential function is raised to the power of a matrix that contains the system behavior. Such a structure cannot be solved directly. Assuming that the eigenvalues are distinct, thus linearly independent as discussed in Section B.1, there is a standard solution to this problem. The following proof is based on the general procedure as outlined in [62].

Additional steps were added to make the proof self-contained and easier to follow.

First, the exponential function can be expressed as a series:

$$e^{\mathbf{Q}t} = \sum_{k=0}^{\infty} \frac{t^k}{k!} \mathbf{Q}^k \tag{B.10}$$

From the diagonalization process of $\mathbf{Q}$, we can obtain eigenvectors $\mathbf{V}$ and eigenvalues $\mathbf{\Lambda}$.

$$\mathbf{QV} = \mathbf{V\Lambda} \tag{B.11a}$$

$$\Leftrightarrow \mathbf{Q} = \mathbf{V\Lambda V}^{-1} \tag{B.11b}$$

Substituting the above in (B.10):

$$e^{\mathbf{Q}t} = \sum_{k=0}^{\infty} \left(\frac{t^k}{k!} \right) \left(\mathbf{V\Lambda V}^{-1} \right)^k \tag{B.12}$$

Since [3]:

$$\left(\mathbf{V\Lambda V}^{-1} \right)^k = \mathbf{V\Lambda V}^{-1} \, \mathbf{V\Lambda V}^{-1} \, \ldots \, \mathbf{V\Lambda V}^{-1} \tag{B.13a}$$

$$= \mathbf{V\Lambda}^k \mathbf{V}^{-1} \tag{B.13b}$$

Substituting this result in (B.12), and recognizing that the series on the RHS is still an exponential, as shown in (B.10):

$$e^{\mathbf{Q}t} = \sum_{k=0}^{\infty} \left(\frac{t^k}{k!} \right) \mathbf{V\Lambda}^k \mathbf{V}^{-1} \tag{B.14a}$$

$$= \mathbf{V} \left[\sum_{k=0}^{\infty} \left(\frac{t^k}{k!} \right) \mathbf{\Lambda}^k \right] \mathbf{V}^{-1} \tag{B.14b}$$

$$= \mathbf{V} e^{\mathbf{\Lambda}t} \mathbf{V}^{-1} \tag{B.14c}$$

The $\mathbf{\Lambda}$ matrix is a diagonal matrix in Jordan canonical form, containing the eigenvalues of the system. The exponential of this matrix, when the eigenvectors are distinct, is [62]:

$$e^{\mathbf{\Lambda}t} = \begin{bmatrix} e^{\lambda_1 t} & & 0 \\ & \ddots & \\ 0 & & e^{\lambda_n t} \end{bmatrix} \tag{B.15}$$

Using these results, the solution to the original problem of (B.9) can now be expressed as:

$$
\begin{aligned}
\mathbf{x}(t) \;&=\; Ce^{\mathbf{Q}t} &\text{(B.16a)}\\
&=\; C\mathbf{V}e^{\mathbf{\Lambda}t}\mathbf{V}^{-1} &\text{(B.16b)}\\
&=\; C\mathbf{V}
\begin{bmatrix}
e^{\lambda_1 t} & & 0 \\
& \ddots & \\
0 & & e^{\lambda_n t}
\end{bmatrix}
\mathbf{V}^{-1} &\text{(B.16c)}
\end{aligned}
$$

The solution is now described by two constant matrices and one square diagonal matrix containing exponentials to a power of a function. The initial problem of an exponential function to the power of a matrix is thus resolved.

Appendix C

Statements

In keeping with a long-established Dutch and Belgian academic tradition, this thesis is supplemented with a selection of statements. These may be humorous quips, serious thoughts, or general observations regarding the state of the world.

Nullius in verba
("On nobody's authority" or: "take nobody's word for it", "respect the facts")
The Royal Society's motto from the enlightenment era

Much about Vancouver can be explained by the fact that it is a rain forest. Halfway through your first winter, the brain fungus sets in.
Tom De Rybel

There is a theory which states that if anyone ever discovers exactly what the Universe if for and why it is here, it will instantly disappear and be replaced by something even more bizarre and inexplicable. There is another theory which states this has already happened.
Douglas Adams

Most Americans considered Canada to be merely another state that figured out a cute trick to avoid paying taxes to Washington, DC.
Robert Horning

Prediction is very difficult, especially of the future.
Niels Bohr

Mr. Madison, what you've just said is one of the most insanely idiotic things I have ever heard. At no point in your rambling, incoherent response were you even close to anything that could be considered a rational thought. Everyone in this room is now dumber for having listened to it. I award you no points, and may God have mercy on your soul.
excerpt from Billy Madison

Just when you think you've seen it all, someone changes what "It" is.
CmdrTaco

A good scientist is a person with original ideas. A good engineer is a person who makes a design that works with as few original ideas as possible. There are no prima donnas in engineering.
Freeman Dyson

If the aborigine drafted an IQ test, all of Western civilization would presumably flunk it.
Stanley Garn

If something comes with a lifetime warranty and it breaks, does that mean that there'll be assassins on your doorstep?
Tom De Rybel

Inside every cynic there's an idealist desperately yearning to be let out, and when they are let out they're usually a real pain and cause all sorts of trouble.
Chris Boucher

The most exciting phrase to hear in science, the one that heralds new discoveries, is not "Eureka!" (I found it!) but "That's funny ..."
Isaac Asimov

My take on all this is pretty simple: in a country where it is considered a normal, sane, and fun recreational activity to strap two greased sticks to your feet and throw yourself down the side of a friggin' mountain, nobody has the right to call *my* minor peccidillos "unsafe".
Nathan J. Mehl

All wars are civil wars, because all men are brothers ... Each one owes infinitely more to the human race than to the particular country in which he was born.
Francois Fenelon

Distance doens't make you any smaller, but it does make you part of a larger picture.
Yogi Berra

Enjoy yourself while you are still old.
Anonymous

Humanity has advanced, when it has advanced, not because it has been sober, responsible, and cautious, but because it has been playful, rebellious, and immature.
Tom Robbins

I can give you my word, but I know what it's worth and you don't.
Nero Wolfe, "Over my dead body"

A bus station is where a bus stops. A train station is where a train stops. On my desk I have a workstation....
Per Hansson

The aim of an argument or discussion should not be victory, but progress.
Joseph Joubert

I used to think that the brain was the most wonderful organ in my body. Then I realized who was telling me this.
Emo Philips

I will follow the good side right to the fire, but not into it if I can help it.
Michel Eyquem de Montaigne

If I had my life to live over, I'd try to make more mistakes next time. I would relax, I would limber up, I would be sillier than I have been this trip. I know of very few things I would take seriously. I would be crazier. I would climb more mountains, swim more rivers and watch more sunsets. I'd travel and see. I would have more actual troubles and fewer imaginary ones. You see, I am one of those people who lives prophylactically and sensibly and sanely, hour after hour, day after day. Oh, I have had my moments and, if I had it to do over again, I'd have more of them. In fact, I'd try to have nothing else. Just moments, one after another, instead of living so many years ahead each day. I have been one of those people who never go anywhere without a thermometer, a hotwater bottle, a gargle, a raincoat and a parachute. If I had it to do over again, I would go places and do things and travel lighter than I have. If I had my life to live over, I would start bare-footed earlier in the spring and stay that way later in the fall. I would play hooky more. I probably wouldn't make such good grades, but I'd learn more. I would ride on more merry-go-rounds. I'd pick more daisies.
Anonymous

Perhaps a lunatic is simply a minority of one.
George Orwell, "1984"

Facts do not cease to exist because they are ignored.
Aldous Huxley

Most things make sense when you look at them right. It's just sometimes you have to look really, really cockeyed.
Florence Ambrose, Freefall

Homophobia: The irrational fear that gays are going to invade and re-arrange your furniture against your will.
Greywolf Blade

Reality is that which, when you stop believing in it, doesn't go away.
Philip K. Dick

Wisdom is knowing what to do with what you know.
J. Winter Smith

You are never given a wish without also being given the power to make it true. You may have to work for it, however.
R. Bach, "Messiah's Handbook: Reminders for the Advanced Soul"

What we observe is not nature itself, but nature exposed to our method of questioning.
Werner Heisenberg

Clearly the secret of happiness, he reflects quite cheerfully, is a variation on the general principle of banging your head against a wall, and then stopping.
Stef Penny, The Tenderness of Wolves, P.250

The Feynman Problem-Solving Algorithm:

1. write down the problem

2. think very hard

3. write down the answer

Attributed to Murray Gell-Mann

The wind makes dust because it intends to blow, taking away our footprints.
Bushmen folklore (∼1911)

You take the old Goethe much too seriously, my young friend. You should not take old people who are already dead seriously. It does them injustice. We immortals do not like things to be taken seriously. We like joking. Seriousness, young man, is an accident of time. It consists, I don't mind telling you in confidence, in putting too high a value on time. I, too, once put too high a value on time. For that reason I wished to be a hundred years old. In eternity, however, there is no time, you see. Eternity is a mere moment, just long enough for a joke.
Goethe (Hermann Hesse, Steppenwolf)

Ad astra per alia porci (to the stars on the wings of a pig)
John Steinbeck. A professor told him that he would be an author when pigs flew. Every book he wrote is printed with this insignia.

Early to bed and early to rise makes a man stupid and blind in the eyes.
Mazer Rackham (Orson Scott Card, Ender's Game)

For a long time it had seemed to me that life was about to begin – real life. But there was always some obstacle in the way. Something to be got through first, some unfinished business, time still to be served, a debt to be paid. Then life would begin. At last it dawned on me that these obstacles were my life.
Alfred D'Souza

That's the problem with drinking, I thought as I poured myself a drink. If something bad happens you drink in an attempt to forget; if something good happens you drink in order to celebrate; and if nothing happens you drink to make something happen.
Bukowski

The superior man understands what is right; the inferior man understands what will sell.
Confucius

At no time is freedom of speech more precious than when a man hits his thumb with a hammer.
Marshall Lumsden

You can know the name of a bird in all the languages of the world, but when you're finished, you'll know absolutely nothing whatever about the bird... So let's look at the bird and see what it's doing – that's what counts. I learned very early the difference between knowing the name of something and knowing something.
Richard P. Feynman

I don't have to know an answer. I don't feel frightened by not knowing things; by being lost in a mysterious universe without any purpose – which is the way it really is, as far as I can tell, possibly. It doesn't frighten me.
Richard P. Feynman

Dr. Hoenikker used to say that any scientist who couldn't explain to an eight-year-old what he was doing was a charlatan.
Kurt Vonnegut, "Cat's Cradle"

There is one feature I notice that is generally missing in "cargo cult science"... It's a kind of scientific integrity, a principle of scientific thought that corresponds to a kind of utter honesty – a kind of leaning over backwards... For example, if you're doing an experiment, you should report everything that you think might make it invalid – not only what you think is right about it... Details that could throw doubt on your interpretation must be given, if you know them.
Richard P. Feynman

It does not matter how beautiful your theory is, it does not matter how smart you are. If it does not agree with experiment, it is wrong.
Richard P. Feynman

What I cannot create, I do not understand.
Richard P. Feynman, On his blackboard at time of death in 1988; as quoted in The Universe in a Nutshell by Stephen Hawking.

Appendix D

List of Publications

Journal

- Marcelo A. Tomim, Tom De Rybel, and José R. Martí. Multi-area Thévenin equivalents method applied to large power systems parallel computations. *IEEE Transactions on Power Systems*, 2009. Submitted.

- Tom De Rybel, Arvind Singh, Pak Phalmoniroth, and José R. Martí. On-line signal injection through a bus-referenced current transformer. *IEEE Transactions on Power Delivery*, (TPWRD-00171-2009), 2009. Accepted.

- Tom De Rybel, Arvind Singh, John A. Vandermaar, May Wang, José R. Martí, and K.D. Srivastava. Apparatus for on-line power transformer winding monitoring using bushing tap injection. *IEEE Transactions on Power Delivery*, 24(3):996–1003, July 2009.

- Ali Davoudi, Juri Jatskevich, and Tom De Rybel. Numerical state-space average-value modeling of PWM DC-DC converters operating in DCM and CCM. *IEEE Transactions on Power Electronics*, 21(4):1003–1012, July 2006.

Invited Conference

- Tom De Rybel, José R. Martí, and Murray Hodgson. Analytical validation of time-step interpolation in transient insular nodal analysis. In *Acoustics'08 Paris*, number ACOUS-TICS2008/177 in NS30-4aNSc, page 5, Paris, France, 29 June - 4 July 2008. Acoustics'08.

Conference

- Tom De Rybel, Arvind Singh, Pak Phalmoniroth, and José R. Martí. Self-powered on-line signal injection based on a current transformer. In *8th IEEE Electric Power and Energy Conference*, number 1370, page 7, Vancouver, Canada, 6-7 October 2008. IEEE, EPEC'08.

- Tom De Rybel, Marcelo Tomim, Arvind Singh, and José R. Martí. An introduction to open-source linear algebra tools and parallelisation for power system applications. In *8th IEEE Electric Power and Energy Conference*, number 1353, page 8, Vancouver, Canada, 6-7 October 2008. IEEE, EPEC'08.

- Arvind Singh, Tom De Rybel, and José R. Martí. FFT Tutor: A matlab-based instructional tool for FFT parameter exploration. In *Acoustics Week in Canada 6-8 October 2008*, volume 36, pages 82–83, Vancouver, Canada, September 2008. Canadian Acoustical Association, Canadian Acoustics.

- Tom De Rybel, José R. Martí, and Murray Hodgson. Modelling of acoustic wave propagation using transient insular nodal analysis (TINA). In *19th International Congress on Acoustics*, number COM-06-004, page 6, Madrid, Spain, 2-7 September 2007. ICA'07.

- Arvind Singh, Tom De Rybel, José R. Martí, and K.D. Srivastava. Field validation tests of the TLD box for online power transformer winding monitoring systems. In *IEEE Canadian Conference on Electrical and Computer Engineering*, pages 141–144, Vancouver, Canada, 22-26 April 2007. CCECE'07.

- Tom De Rybel, Jorge A. Hollman, and José R. Martí. OVNI-NET: A flexible cluster interconnect for the new OVNI real-time simulator. In *15th Power Systems Computation Conference*, page 6, Liège, Belgium, 22-26 August 2005. PSCC'05.

Theses

- Tom De Rybel. OVNI-NET: A flexible cluster interconnect for the new OVNI real-time simulator. Master's thesis, The University of British Columbia, Vancouver, Canada, March 2005.

- Tom De Rybel. Elektrolysetoestel voor de aanmaak van vaste cyclotrondoelwitten. Master's thesis, Hogeschool Gent, Gent, Belgium, 2002.