A Coupled-Circuit Representation of IGBT Module Geometry for High di/dt Switching Applications

by

Blake Kenton Rablah

B.Eng., The University of Victoria, 2004

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

Master of Applied Science

in

The Faculty of Graduate Studies

(Electrical & Computer Engineering)

The University Of British Columbia

January 5th, 2007

Abstract

A coupled-circuit element representation of an IGBT module from Westcode Semiconductors Inc. has been developed, simulated, and experimentally confirmed at TRIUMF in Vancouver, BC. This model can be simulated in PSpice to predict the distribution of current within the IGBT module under high di/dt switching conditions. The goal of the simulations is to determine if the module is an acceptable candidate for implementation in a thyratron replacement switch for high-power pulse-power applications.

The model developed has been shown to agree well with frequency domain measurements, and also with finite element method (FEM) simulations and boundary element method (BEM) simulations of the device.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank TRIUMF, and in particular Ewart Blackmore, for their continued support, funding and lab-space, dating back to my original co-op placement at the facility. I would also like to thank Mike Barnes and Gary Wait for their mentorship and guidance, without which this work would not have possible.

Chapter 1

Background

In the field of high-power pulsed-power, high speed switching is the foundation upon which all else is built. Switches in this field are often required to withstand blocking voltages in the tens of kilovolts and switch several kiloamps in less than a microsecond. One application in which these switches are in frequent deployment is in high energy physics labs around the world. These switches drive kicker magnets that deflect charged beams of particles, or drive klystrons that accelerate these particles.

New developments in semiconductor research are leading to faster solid state switches that are encroaching on other established pulsed-power technologies. However, many of these switches have been designed for larger commercial markets such as traction applications, and have not been designed explicitly with pulsed-power in mind[2]. Pulsed power places an extra burden on the designers of semiconductor switches. Specifically, the high-frequency content of the pulses and the large di/dt and dv/dt that result make the geometry and resulting parasitic components of the devices a particular concern. This puts designers of pulsed semiconductor devices and modules in territory that has historically been the concern of VLSI packaging designers.

1.1 Intended Application

The work described in this document was begun in response to the need for a semiconductor based switch for high pulsed currents with very short risetimes. One potential application of such a switch was for the yet-to-be-funded International Linear Collider to drive klystron modulators. It is envisioned that this semiconductor-switch would be able to turn-on approx-

imately 3kA in approximately 300ns, for a peak di/dt >10kA/μs, and would be able to be connected in series configurations to permit blocking voltages of at least 75kV[3].

Other thyratron replacement designs cite similar requirements[4], so these figures can be considered general design goals. However, in order to be considered sufficient for application to kicker magnet applications, the requirements become somewhat more demanding. Table 1.1 details the design objectives for which the Westcode module is being considered.

Table 1.1: Anticipated requirements for a thyratron replacement

Criteria	**Requirement**
Blocking voltage	Stackable to >65kV
Peak current (I_{pk})	>8kA
Risetime (to I_{pk})	<1μs (min. 10kA/μs)
Pulsewidth (τ_{pulse})	1μs $\rightarrow$ DC

1.2 Available Technology

1.2.1 Thyratrons

At the present time, the main devices used for high-power high-speed switching are hydrogen gas filled thyratrons. Similar in design to the more common vacuum tube, a thyratron is a gas-filled device, with an anode and a cathode separated by one or more grids. Switching is achieved when the gas inside the thyratron ionizes and goes from being strongly insulating, to conducting in a very short time. Some thyratrons are capable of blocking over 100kV, and are able to switch tens of kiloamps in just tens of nanoseconds[5]. Typical kicker magnet applications require current rise-times in the neighbourhood of 30ns to peak currents above 3kA. This results in a di/dt in excess of 100kA/μs. As an example of one such application: a TRIUMF-designed kicker-magnet for the LHC at CERN featured a thyratron driver that switched a 54kV PFN into a 5Ω load: the current rises to a maximum

of 5.4kA with a 10% to 90% rise-time of 30ns, corresponding to a di/dt of 180kA/μs[6].

As impressive as these numbers are, thyratrons exhibit some very significant drawbacks when compared to semiconductor devices:

Spontaneous Turn-on Thyratrons are known to have a finite probability of spontaneous turn-on. Furthermore, this probability goes up considerably as the voltage across a thyratron rises, and also with the length of time this voltage is held. Designers of thyratron circuits sometimes employ fast resonant power supplies to quickly ramp up the voltage on the thyratron shortly before switching, minimizing the thyratron's exposure to high voltage and reducing the chance of spontaneous turn on.[7]

Switching Control The second disadvantage of thyratrons is the lack of control while switching. Thyratrons latch into the on-state, and will remain latched until they reach zero-current, and the plasma inside the tube is exhausted[5].

Lifetime Thyratrons also have a limited life span that is dependent both on the number of pulses and the total charge transferred. For any remotely high-frequency high-current application, this can make them expensive to deploy. One such application: as klystron modulators for the International Linear Collider (ILC), would require a team of individuals dedicated to thyratron replacement and adjustment (see below)[8].

Maintenance In addition to lifetime concerns, thyratrons require frequent adjustment over their lifespan to maintain their switching characteristics. This adds both to the complexity and cost of a thyratron deployment.

1.2.2 Solid State Thyratron Replacements

Even the most advanced solid state devices still have limited voltage and current ratings in comparison to high power thyratrons. Nonetheless, since

approximately 1990, considerable effort has been made to produce a solid state replacement for hydrogen gas filled thyratrons. Companies like ABB and Dynex Semiconductor now manufacture thyratron replacement modules constructed from series connected SCR or IGBT devices[4] [9]. However, these devices fall far behind gas thyratrons in terms of switching speed, and as such are not yet suitable "replacements".

SCR Based Solutions

Both ABB and Dynex, as well as some other manufacturers have developed thyratron replacement stacks from SCR modules. These stacks are capable of di/dt values in excess of 20kA/μs, and are capable of very high peak currents, with 90kA advertised for one Dynex device (Dynex part #: PT85QWx45).

However, while SCRs fix the lifetime and spontaneous turn-on issues of Thyratrons, they still exhibit some significant drawbacks. SCRs also latch in the on-state and exhibit a similar lack of control during switching as thyratrons. Also, thyristor SCRs require considerable gate currents to switch on. This is especially true when fast switching is desired, where gate currents can exceed 200A[3]. The necessity of a large gate current places strong demands on the gate drive circuit; this can be problematic to implement when each gate drive circuit is floating several kilovolts apart from the next.

IGBT Based Solutions

IGBTs are bipolar power switching devices that feature a MOSFET-like insulated gate. IGBTs have begun to take over in industry sectors previously dominated by SCRs, especially for medium power applications. The ability to turn-off, even under load, and the simplicity of the gate drive circuit are the two main selling points for IGBT based switches. Additionally, IGBTs can self-limit their current when short circuited[11], which is of interest to a designer looking for some degree of fault tolerance in their design.

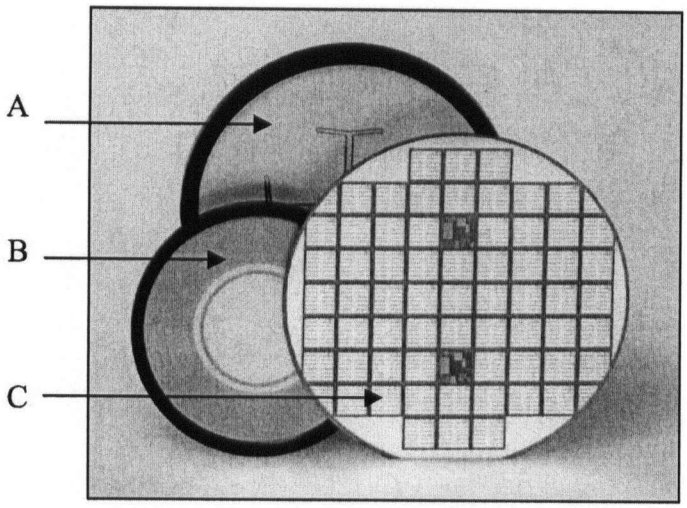

Figure 1.1: Silicon wafer comparisons - A = SCR Thyristor, B = IGCT, C = IGBT chip wafer

Controlled switching is also possible with IGBT modules, and active control of series connected modules has been achieved [12].

There are some significant shortcomings associated with current IGBT technologies that limit their applications in the high-power pulsed-power field. While breakdown voltages of individual IGBT die are starting to encroach on territory typically occupied by SCR thyristors, they have relatively lagged in current handling capacity. Part of the reason for this stems from a limitation in the IGBT silicon manufacture that sets the maximum die size at $\approx 1\text{cm}^2$; many modern SCRs and IGCTs can be constructed on single 91mm diameter wafers (Figure 1.1 Figure 1.1 [10]).

Due to the limitation in die size, high-power IGBT devices are constructed in modules of several IGBT die in parallel. This introduces the problem of ensuring uniform current sharing. Under DC or slow transient conditions, current sharing is dominated by the relative on-state resistances of the various IGBT die within a module. However, under high speed switching conditions, it has been shown that the internal layout of IGBT modules has a marked effect on the current distribution within the module, and cor-

respondingly also the performance and lifetime of the module[2]. Further analysis of IGBT module geometry has identified that parasitic inductances and coupling to the gate circuit within the packaging are largely to blame for current imbalances in some types of module[13][14].

IGBTs are of continued interest for designers of a thyratron replacement switch, but it is apparent that more attention needs to be paid to IGBT module design if such an application is to be realized. ABB manufactures a thyratron replacement that uses stacked IGBT modules, but it's relatively slow switching speed (maximum of $1.5\text{kA}/\mu\text{s}$) is a testament to the challenge that faces designers of IGBT based solutions[15].

It is possible to boost di/dt performance of a slower switch by adding a saturating magnetic core in series with the switch. This "saturating reactor" will initially appear very inductive and limit the rise of current, until saturation occurs, which results in an effective short circuit. The disadvantage of this approach, to pulsed power applications, is that it adds a considerable amount of parasitic inductance to the circuit and can contribute to impedance mismatches in the system. Additionally, as the characteristics of the non-linear magnetic material are fixed, lower operating voltages will result in slower saturation of the magnetic material, and therefore a significantly reduced di/dt: limiting the useful dynamic operating voltage of the switch.

1.3 Switch Selection

1.3.1 Electrical Criteria

Electrically, the switch must be capable of meeting the requirements already outlined in Section 1.1, specifically the requirements for kicker-magnets outlined in Table 1.1.

1.3.2 Form Factor

IGBT modules come in several different types of packages. Two of the main approaches to package design are shown in Figure 1.2; on the left is a "brick

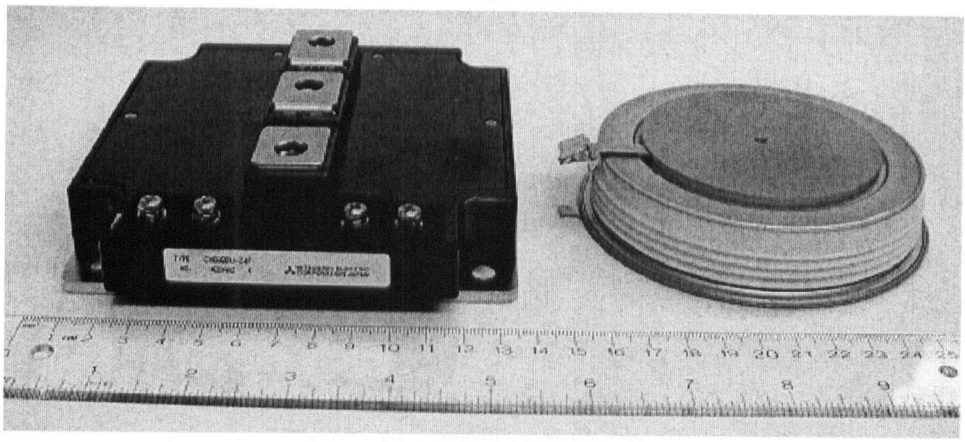

Figure 1.2: Comparison of IGBT module packaging options

type" package manufactured by Mitsubishi, on the right is a "hockey puck" style of press-type package from Westcode.

Brick Type Modules In the brick-type package the IGBT die are flow-soldered onto a PCB that contains traces for the collector, emitter and gate. Additional interconnections are made with individual wire bonds. This type of arrangement, the most common for power IGBT modules, offers easy connection to external circuits via bolt-on terminals, and is relatively cheap to manufacture. However, for pulsed power applications, this ease of connection comes at a price. [2][13][14] illustrate the difficulties associated with achieving uniform current distribution using this design. Additionally, low inductance series configuration is difficult, as these modules take up a considerable amount of space, and both emitter and collector terminals are on the same side of the device.

One significant advantage of this type of module is that they can be configured such that the return path has a minimal effect on the current distribution within the module. As much of the work detailed in this document deals with this particular subject, it is apparent that this is not of small importance.

Press Pack Modules Devices of this design, are constructed such that the emitter and collector terminals make up the top and bottom of the module. The term 'press pack' comes from the way in which individual die are sandwiched inside the module between the collector and emitter terminals. Electrical contact is achieved by compressing the module under a considerable amount of force (between 20kN and 70kN, depending on the design and application).

Series connection is relatively simple as modules can be stacked one upon the other, and the entire structure compressed to make electrical contact both between and within each individual module. This series configuration also lends itself well to low inductance designs, for the return current path can be constructed to surround the stack: creating a coaxial structure.

One aspect where press pack designs have a distinct advantaged is cooling. Heat is conducted away from both sides of the silicon by the emitter and collector contacts. Cooling solutions can then be applied at these terminals.

Another advantage, perhaps more important to pulsed power application, of the press pack design concerns the mode of failure for the IGBT die. Press pack IGBTs fail to short circuit. This means that, with proper redundancy in a series stack of modules, that a failure in one module (or device within a module) is rarely catastrophic, and the stack usually remains functional[16]. In contrast, brick-type IGBTs often fail violently, severing the wire-bonds within the module and resulting in an open circuit[2]. For even a few individual IGBT die to fail in this fashion would render the entire stack useless.

1.3.3 Prior Work

SLAC Experiments

TRIUMF's initial investigations into IGBT based thyratron replacements began in conjunction with researchers at the Stanford Linear Accelerator Center (SLAC)[14]. In their attempts to develop a switch to replace thyratrons for klystron modulators they experienced catastrophic failure in some IGBT modules when they were exposed to soft short-circuit conditions[2].

In their attempt to determine the cause of the failures, they performed FEM analyses of the geometry of the device and attributed the failures to an asymmetry of path inductances: contributing to an imbalance of device currents.

TRIUMF Current Distribution Analysis

TRIUMF continued the analysis from where SLAC left-off. The TRIUMF contribution was to perform a coupled-circuit element extraction of the IGBT module, and to simulate the conditions which lead to the failure of the device in PSpice. The results of this analysis showed conclusively that, while asymmetric path inductances were a factor, it was inductive coupling to the module's internal gate trace that was largely to blame for the current imbalance and resulting failure of the module. [14] and [17] detail this process.

1.3.4 Westcode Module Specifics

One of the other modules available to the SLAC researchers, and to TRIUMF was a prototype Westcode module designed specifically for pulsed-power applications. This device features a redesigned gate trace, and other, undisclosed improvements for high di/dt switching. The device is rated at 5.2kV, and 900A DC. However, for the intended application of a thyratron replacement, it is likely to be operated at 60% of the rated blocking voltage, and pulsed at almost 10x the DC current (3.1kV and >8kA, respectively).

Figure 1.3[18] shows a rendering of the internals of this specific module. The image is an exploded view of the device which illustrates the 14 IGBT "cassettes" that house the IGBT silicon die (detailed in the inset, bottom-right), the copper stand-off "posts" that locate the die, and the aforementioned redesigned gate trace.

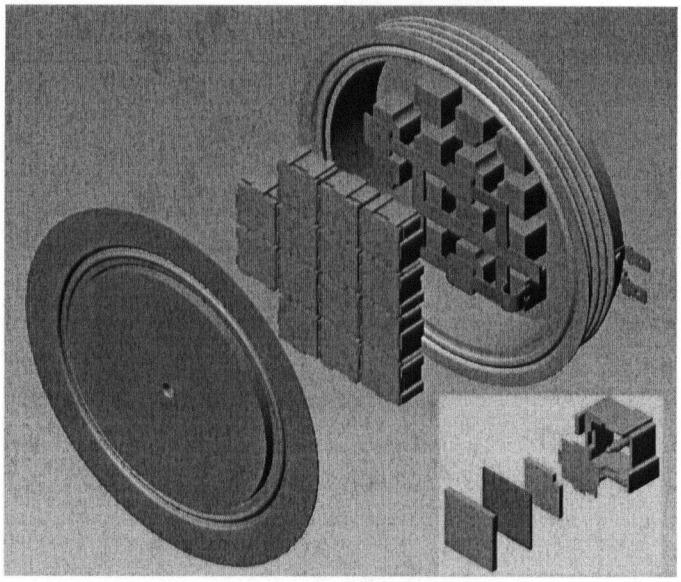

Figure 1.3: Rendering of a Westcode press pack IGBT module intended for pulsed power applications

In the course of the research detailed in this thesis, it was necessary to dismantle one of these IGBT modules. Figure 1.4 is a series of photos documenting this process.

There are 14 IGBT die within the Westcode module. For the rest of this document they will be referred to by number, as labelled on the illustration in Figure 1.5.

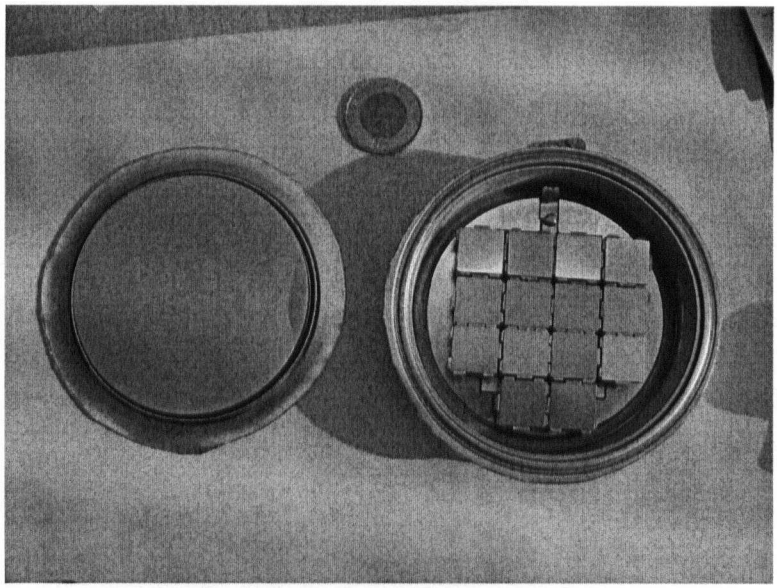

(a) Opened Westcode IGBT module showing 14 parallel IGBT "cassettes".

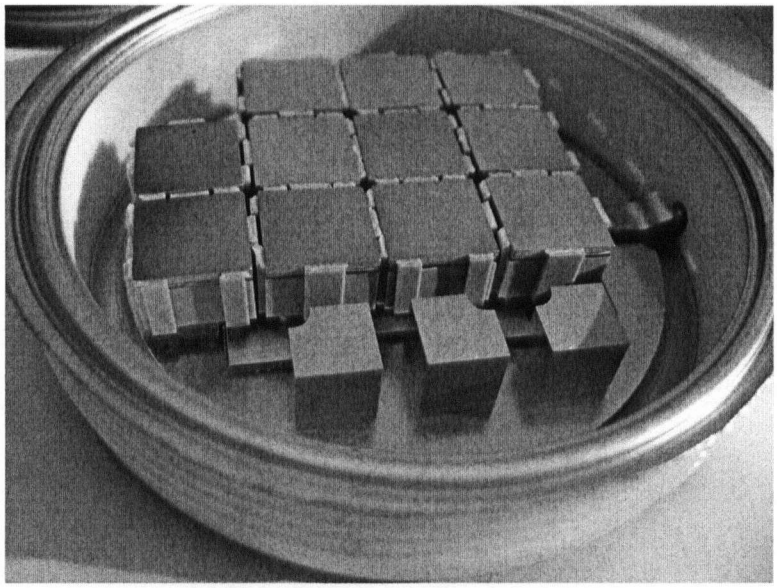

(b) Detail of Westcode IGBT internals. Visible are the copper posts, IGBT die (in cassettes), and the gate trace.

Figure 1.4: Photos of a dismantled Westcode module

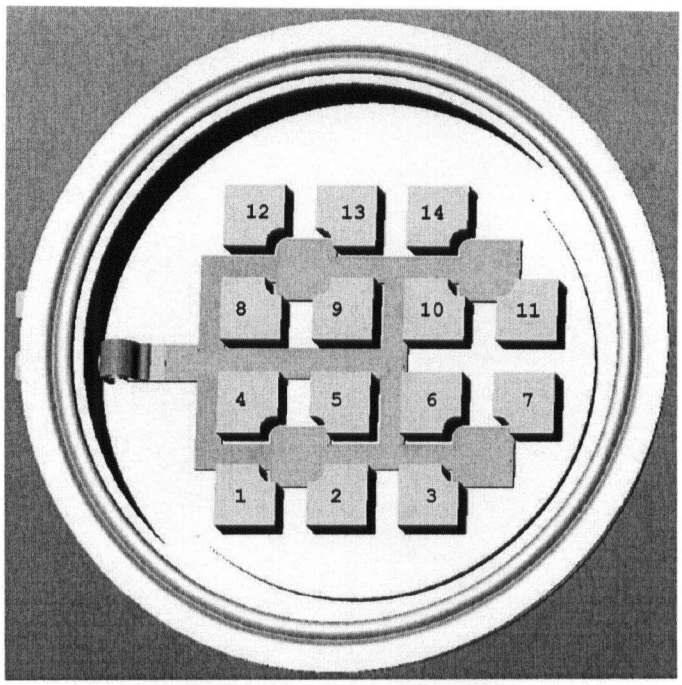

Figure 1.5: Rendering of the internals of a Westcode press pack IGBT module with numbered current paths

1.4 Focus of this Thesis

Westcode publishes in their documentation the assertion that, as all of the paths through the IGBT module are of equal-length, that they exhibit equal inductances. Figure 1.6 is an equivalent circuit of the IGBT module, including parasitic inductances, as supplied in their documentation[19]. In this figure these parasitic inductances are lumped into just two parameters, one at the emitter and one at the collector of the device.

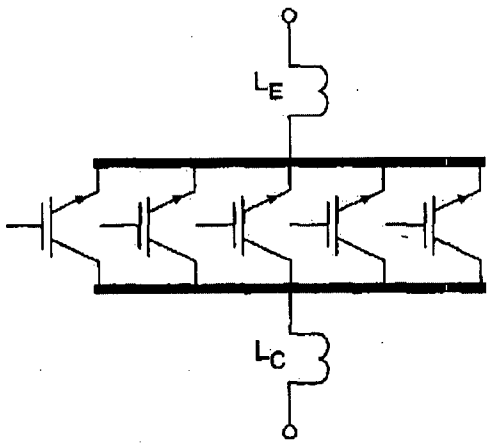

Figure 1.6: Westcode provided equivalent circuit of press pack module

When assessing the suitability of this module for a thyratron replacement at TRIUMF, this assumption was called into question. It was suspected that there might be significant issues with uniform current sharing under pulsed conditions, and that these issues might also be strongly influenced by the geometry of the return current path.

The research documented in this thesis covers the creation, simulation and testing of a coupled-circuit model of the Westcode module that explicitly takes into account the inductive coupling that exists, both between the separate current paths through the device and to the return path.

It is hoped that this circuit model will be of great help in the future for the process of prototyping a thyratron replacement switch. One of the

biggest advantages of such a circuit is that it allows for new geometries to be quickly and easily tested in PSpice, along with the external circuit and any non-linear switching elements. As PSpice is already used to simulate existing pulsed switches[20], as well as their loads[6], this is the preferred method of implementing the characteristics of the IGBT module

Chapter 2

Simulation of Current Distribution

In order to ascertain how the current is distributed within the IGBT module, I undertook to model and simulate the internal structure of the Westcode press pack IGBT module.

Simulation Objective The goal of these efforts was to produce a versatile circuit model of the Westcode IGBT geometry that could be employed in PSpice to predict the current distribution under a wide variety of conditions in a minimal amount of time. Additionally, as the requirements of a thyratron replacement would likely involve pushing any switch to it's physical limits, an accurate model of the IGBT geometry would be of great benefit later in the thyratron-replacement design process to avoid costly failures.

2.1 Simulation Methods

Each electrical path within the hockey-puck IGBT module consists of an IGBT die and a copper post. The ideal software tool to analyse and produce a current distribution for this geometry would be a 3D electromagnetic code, capable of also simulating the non-linearity of the semiconductor die. However, to the author's knowledge, no such code is available. Hence the circuit analysis code PSpice, the de-facto standard for non-linear circuit simulation, is used to simulate the electrical current paths within the IGBT module.

For this, inductance and resistance values are required for each of the electrical paths in the IGBT module: these values are derived from an elec-

tromagnetic code.

Initially a 3D FEM code was used to simulate each IGBT die (as a linear element with the appropriate on-state resistance, see Appendix A) and the copper posts of the module. However, meshing of the model was overly problematic: due both to the excessive number of elements required to perform a high-frequency simulation, and the difficulties associated with meshing mating regions with very different material properties. This approach was abandoned.

Hence, instead, a 2D code is used to model the IGBT. Firstly, the copper post was modelled, with the same current distribution throughout its length, and the inductance per unit length and resistance per unit length, of each post, was determined. The same procedure was used to calculate the inductance per unit length of each on-state IGBT die; and the results are used to assign component values to a PSpice model of the IGBT module to provide a complete model. Such parasitic extractions are performed in the frequency domain.

In order to produce the circuit model, and to be assured of its accuracy, I undertook to model and simulate the Westcode module via several different approaches. As described in the subsequent sections, different software tools were used, and employed several different modelling techniques. These tools were organized into discrete "workflow" which produce a current distribution, given a geometric model of the IGBT module, under certain conditions.

2.1.1 Software Tools

Different software packages were used throughout this process, these are detailed in Table 2.1. Using different software packages with dissimilar solving techniques to solve the same geometry provided us with confidence in the final outcome. This is especially important in our case, where physical measurements to determine the current distribution within the IGBT module, especially under switching conditions, while not impossible, are themselves difficult and prone to error.

Software Package	**Usage**
Vector Fields Opera 2D	• 2D-FEM Modelling. • AC steady state and transient simulation.
Integrated Engineering Software (IES) OERSTED	• 2D-BEM/FEM modeling. • AC steady state simulation. • Circuit parameter extraction.
Fasthenry	• 3D-Magnetoquasistatic Multipole modeling. • Circuit parameter extraction.
OrCAD PSpice	Flexible time domain and frequency domain simulation of extracted circuit parameters.
Matlab	• Scripts to convert impedance matrixes into resistances and coupled inductors. • Scripts to extract current densities from simulation tool log files.

Table 2.1: Modelling and simulation software.

2.1.2 Simulation Geometries

Return Path Geometry

As has been pointed out previously in this document, the geometry and effect of the return path on the current distribution should not be neglected. To attempt to produce the least biased current distribution within the simulated Westcode module, and to lower the inductance of the circuit, a coaxial structure was modelled.

For practical reasons, a fully coaxial return path is unfeasible: there is no way to easily access the device, and the connection of gate drives would become problematic. As such, an array of six rods in a circular shape was simulated, approximating a coaxial structure. The number six was chosen because it has been shown that this number of return rods minimizes the inductance for a geometry[21]. Adding additional paths to further approximate a fully coaxial structure yields little to no benefit to the overall inductance.

The rods are spaced some distance away from the IGBT module (10cm-12cm), as this is indicative of the spacing required for the potential application of a stack of modules, where the voltage difference between the top of the stack and the return path may exceed 75kV.

2D Model

In the following simulations, a 2D model of the Westcode module appears as a cross-sectional slice, viewed from above, of the IGBT module's 14 forward paths, surrounded by the six return paths. Figures 2.2 and 2.5 illustrate how this was implemented. With the first of these taking advantage of the symmetry of the arrangement, and the second modelling a complete 2D 'slice'.

3D Model

The 3D model of the Westcode IGBT module described in Section2.2.4 is an extruded version of the 2D model described above. The device is modelled

as having 100cm long posts, rather than the actual length of ≈1cm. The reason this was done was to minimize the effect of fringing of magnetic field that might be simulated at the ends of the rods. In the actual device these fields would be excluded by the highly-conductive copper emitter and collector contacts.

2.2 Simulation Workflow

Using various permutations of the tools from the section above, I developed several different workflows to determine the current distribution within the Westcode module. Each of these workflows was capable of producing a solution to one or more scenarios from of a set of test cases. Exploiting the features and capabilities of the different software packages, we were able to cross-check the solutions between workflow techniques, and software packages.

The current distribution within the IGBT was examined for the following test scenarios:

a  Frequency domain solution with single excitation frequency.

b 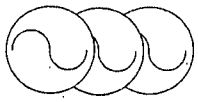 Frequency domain solution with multiple excitation frequencies (AC-sweep analysis).

c 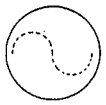 Time domain solution with a single-frequency sine wave drive.

d  Time domain solution with an arbitrary current pulse.

e 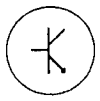 Time domain solution with non-linear circuit elements.

2.2.1 Workflow 1: Opera 2D Field Solution

This technique involves using the Opera2D software package for both modeling, and for solving for the individual die currents at discrete AC drive frequencies.

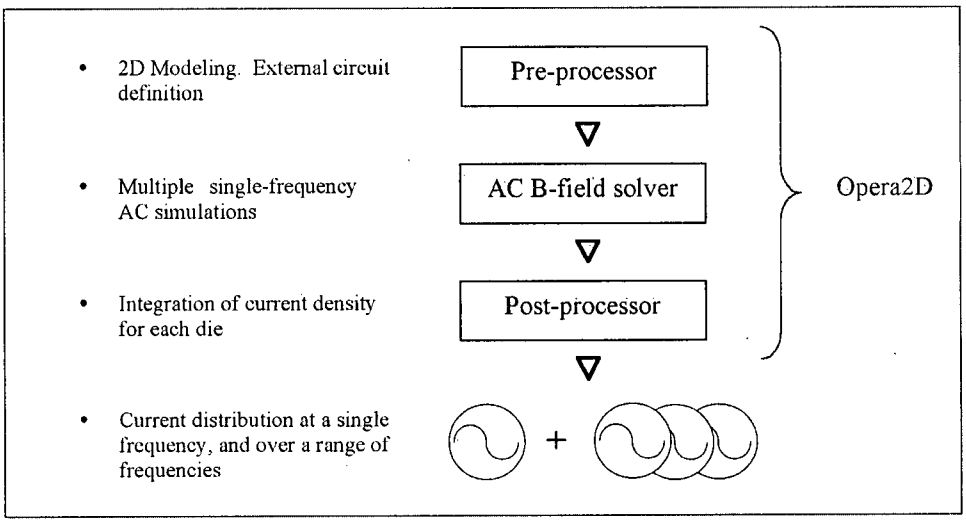

Figure 2.1: Opera2D Field Solver workflow

The geometry simulated is a cross sectional slice of one half of the symmetrical IGBT module (top, centre); complete with half of the corresponding return path (surrounding), as shown in Figure 2.2. Normal **B** field boundary conditions were applied to the top edge of the model, and the solver set to regard this as the axis of symmetry. The solver considers the device to be infinitely 'long', however a length in the Z-direction is specified for the device, so that the results may be scaled accordingly. Being a 2D model, the material simulated is assumed to be homogeneous over the entire length.

Simulating just half of the IGBT module allowed for a denser mesh to be used while still retaining the same total number of elements. This is important when doing a high-frequency analysis on a highly conductive material, as the skin effect crowds current to a very thin depth along the edge of each conductor, and correspondingly small elements must be used.

Equation 2.1 illustrates how the skin depth (δ) depends inversely on both the frequency of simulation (f) and the conductivity (σ). For the materials under consideration (silicon, copper), the relative permeability (μ_r) is near unity, so the total permeability ($\mu = \mu_r \mu_0$) will be fixed.

$$\delta = 1/\sqrt{2\pi f \mu \sigma} \tag{2.1}$$

As an extreme example: for a copper bar ($\sigma = 5.96 * 10^7 Sm^-1$) with a 10MHz drive frequency, δ is a mere $15\mu m$. With an accepted minimum of three elements per skin-depth[22], this amounts to an extremely high element count, even when when a graded mesh is employed (as it is in this model).

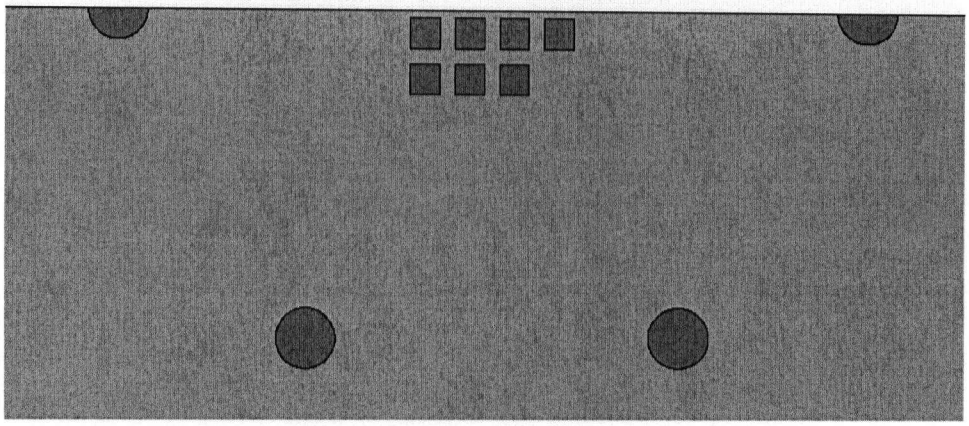

Figure 2.2: Modelled geometry - Opera 2D

This workflow performs the FEM analysis of the module at a specific frequency while assuming that the geometry being modelled is part of a larger external circuit (detailed in Figure 2.3). The resistance in the external circuit is arbitrarily chosen to be large compared to the impedance of the module. The voltage source can then be scaled to give the appropriate current. In this case, a peak current of 10kA for a complete module was modelled: although, as the model is linear, and we are only concerned with the *relative* distribution of current, the actual value of the peak current, as

modelled in the electromagnetic simulation, is not important

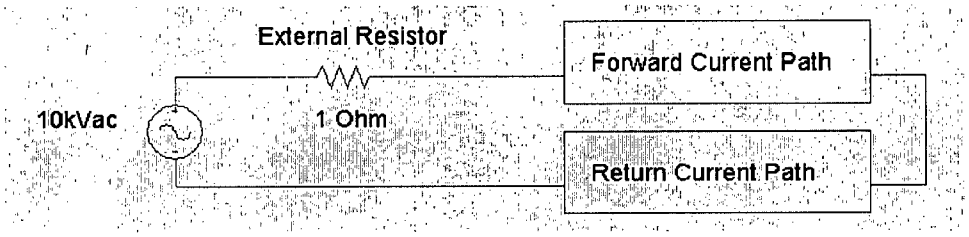

Figure 2.3: External circuit for Opera2D AC analysis

I ran the simulation with a series of frequencies starting at 1Hz, and extending to 10MHz. Using Opera2D's built-in scripting ability, the current density for each of the die was integrated and extracted at each simulated drive frequency. Matlab was then used to parse the resulting Opera2D log-file, retrieve the die currents, and then normalize these to the total simulated module current.

Simulation Time To perform a single-frequency analysis on this model takes just over five minutes to simulate on a Pentium IV 3.0GHz machine with 2GB of RAM. An additional few minutes is required to perform the post-processing and obtain the current distributions from the model.

Summary With proper implementation, this workflow should be able to produce some of the most accurate simulations of the current density under steady state conditions. However, the length of time required to do a complete FEM simulation at each simulation frequency, combined with the inability to simulate a device that is heterogeneous along its length limits the utility of this workflow to verifying the results produced by other workflow.

2.2.2 Workflow 2: IES OERSTED Field Solver

IES OERSTED features an integrated AC steady state field solver which utilizes the boundary element method, finite element method, or a hybrid combination of the two. IES has integrated their software such that modeling, simulation, and analysis are all performed in the same interface.

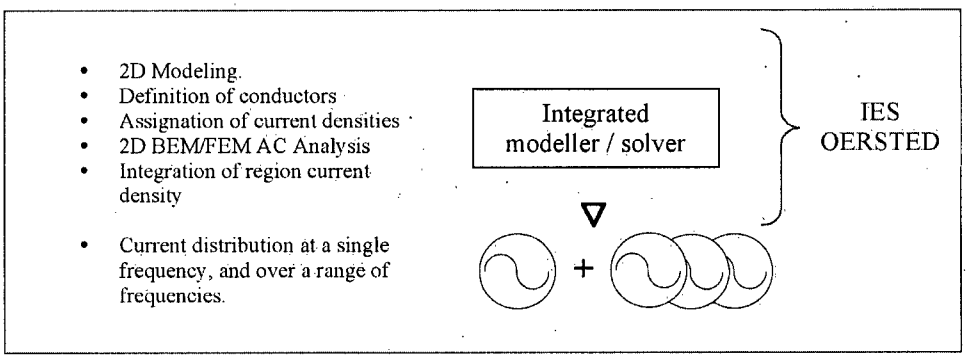

Figure 2.4: Workflow 2 - IES OERSTED Field Solver

To simulate the IGBT in OERSTED, a two-dimensional model representing a complete cross-section of the IGBT module was created that did not take account of symmetry (Figure 2.5). Otherwise, the geometry represented by this model was identical to the Opera2D rendering. This was possible as OERSTED's use of 1D boundary elements at the periphery of each conductor negated the necessity for a large number of FEM elements resulting from skin effect (see Section 2.2.1 for an example) . As with the Opera2D simulation, the module simulated must be assumed to be made of homogeneous material along its length.

Once the geometry was established, it was a relatively simple matter to assign the total current density in the forward and return conductors, define the drive frequency and set the self-adaptive solver running. Once the solver had finished, the current density was produced for the entire module and integrated over individual regions.

To obtain the current distribution over a range of frequencies, it is necessary to set up a parametric analysis within IES OERSTED. OERSTED

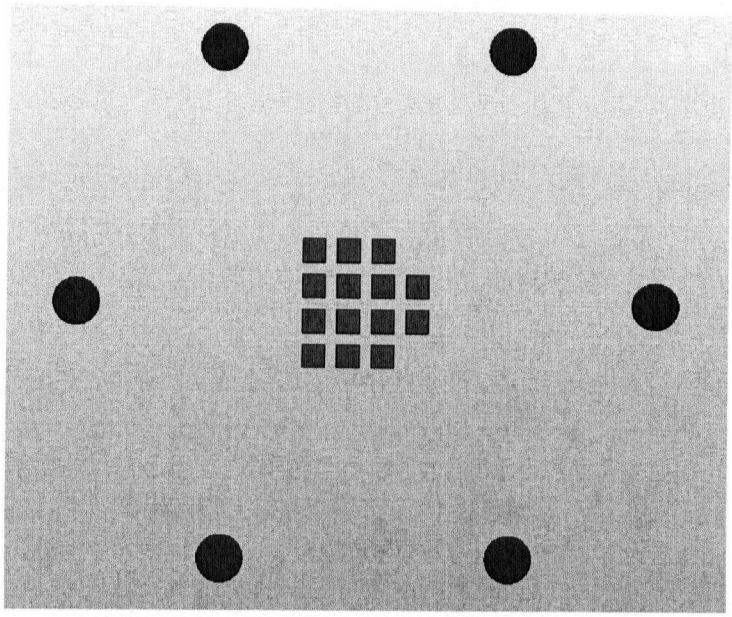

Figure 2.5: Modelled geometry - IES OERSTED, IGBT die are shown in red, return paths in blue

has an advantage over Opera2D in that a number of post-processing options can be specified to be run for each parametric case without any additional scripting required. For our model, the post-processing options calculated the total current present in each path at each specified drive frequency.

Summary The OERSTED field solver is very similar to the Opera2D field solver, and has the same set of advantages and disadvantages. The ease of use of the IES unified environment is a benefit, but, as with Opera, this method is useful mainly as a comparison tool for the other workflow.

2.2.3 Workflow 3: OERSTED Extraction with PSpice Simulation

This method employs IES OERSTED as an extraction tool that produces an impedance matrix at a single extraction frequency. This matrix is converted into mutually coupled circuit parameters which can then be analysed in any PSpice circuit.

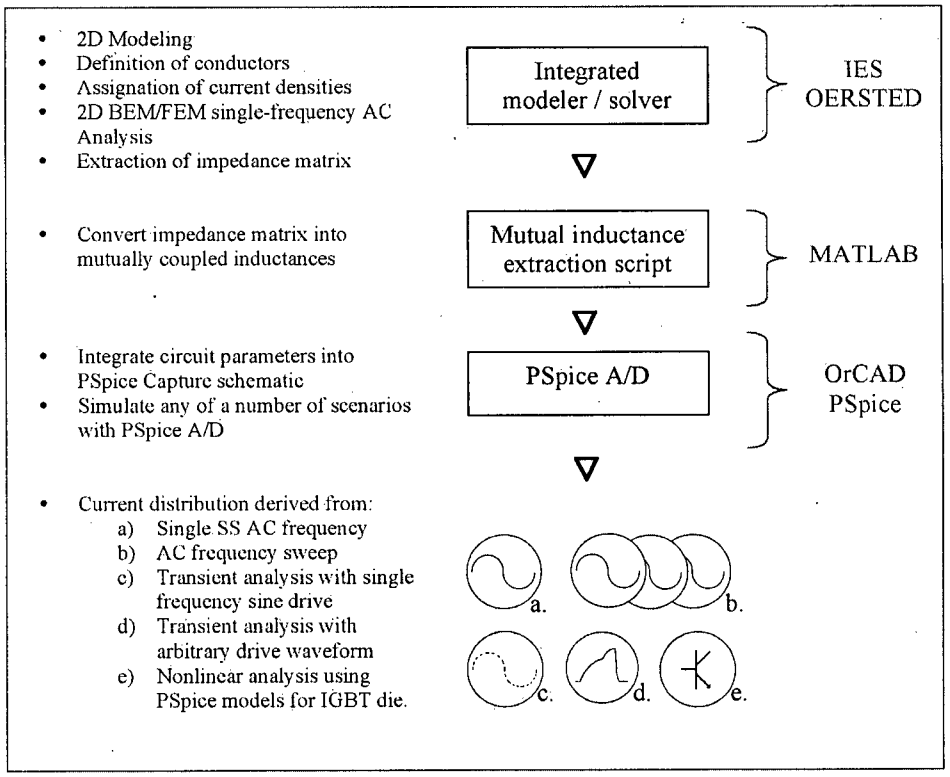

Figure 2.6: Workflow 3 - OERSTED extraction and PSpice simulation

The circuit parameter extraction solver uses nearly the same geometric model in OERSTED as the field solver. The only major difference to this model is a result of the method that OERSTED uses to produce circuit parameters: For x forward paths, and y return paths, OERSTED produces an $x * x$ impedance matrix. This matrix represents the complex impedances

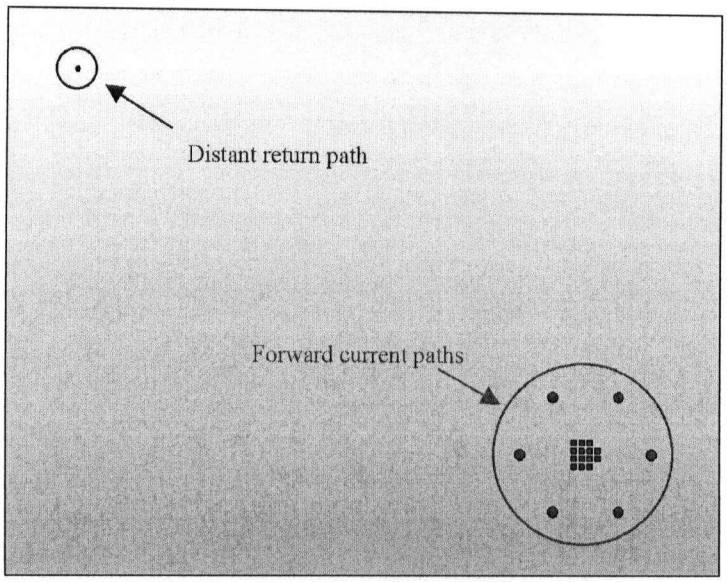

Figure 2.7: OERSTED geometry for circuit parameter extraction

of the forward conductors *in the presence of the return conductors.* As a result, the return paths are not explicitly represented and expressed as circuit parameters (this is in contrast to Fasthenry; see Section 2.2.4 - below). To render the PSpice model more flexible, it was decided to model all of the conductors in the existing geometric model as forward current paths, and include a separate return path which is relatively distant from the model (Figure 2.7). This permits the actual return path to be modelled from within the circuit simulator, and allows some alterations to be made without requiring re-extraction of parameters.

The impedence matrix, as well as circuit parameters, are produced using BEM/FEM/Hybrid analysis (as with the OERSTED field solver, above) at a single extraction frequency which is determined at the time of simulation.

Extracting circuit parameters at a single frequency, and then using them as linear elements over a wide range of frequencies in PSpice can be problematic, as these values are, in fact, frequency dependent (this effect is covered in Appendix B). Inductance ladders can be constructed to simulate this

effect and have been employed successfully by others[23]. However, construction of these inductance ladders adds considerably to the complexity of the PSpice model and may not be considered necessary. As shown in the aforementioned appendix, it is possible to achieve greater accuracy from a single set of extracted circuit elements by combining resistive and inductive components generated from the same model at different extraction frequencies. The resistive portion is taken from a low frequency extraction, and the inductive portion from a high frequency extraction.

The process of rendering the extracted impedance matrix into coupled circuit parameters is detailed in Appendix C. This conversion process generates a PSpice `.include` file that contains `.PARAM` statements for each of the self inductance and resistance values. Also included in this file are the coupling coefficients between each inductor. Attaching this file to a PSpice schematic with appropriately labelled resistors and inductors produces a circuit that electrically mimics the geometric effects of the different current paths.

For a composite structure such as the Westcode module, where different materials (silicon, copper, molybdenum) are layered to make up the complete module, modelling can be accomplished by performing a separate 2D extraction for each distinct layer in the device. The resulting `.include` files can be included in the PSpice simulation, with the circuit elements for each path are arranged in series within the schematic, thus emulating the previously attempted 3D simulation.

The circuit parameters extracted are per unit length. To adapt these values to represent the real module, the resistances and inductances are linearly scaled accordingly (terms RatioCuR & RatioCuL in Figure 2.8). The coupling terms are expressed as a ratio of the self- and mutual-inductances, and therefore remain independent of the length of the extruded model and need no adjustment.

Simulation Time Each single frequency parasitic extraction in OERSTED 2D takes just over 11 hours on the same machine as used for Section 2.2.1.

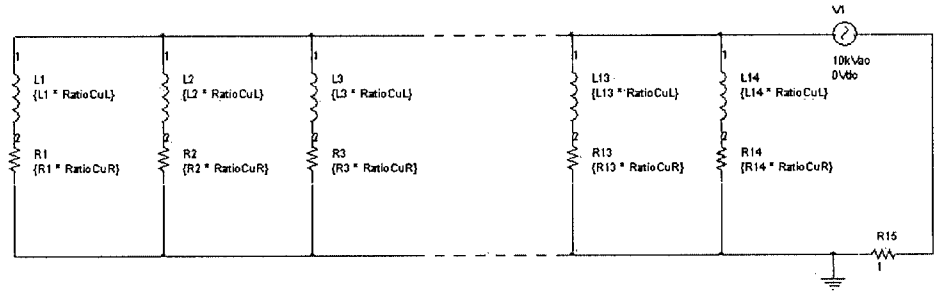

Figure 2.8: PSpice AC Analysis with OERSTED Circuit Parameters

Summary The real advantages of this particular workflow are in the flexibility that PSpice simulation offers:

- It is possible to represent a device that is made of composite materials along its length (as is the case with the Westcode module) by modelling 2D slices of the various materials separately. The resulting circuit parameters are then arranged in series within PSpice and scaled to the appropriate length to simulate the entire device.
- Both AC and transient analyses can be performed much faster than with FEM/BEM analysis alone.
- Non-linear circuit elements can be introduced to represent the IGBT die portions of the module, and switching behaviour can then be explicitly assessed.
- In non-linear analyses the effect of inductive coupling to the gate traces from main current paths can also be simulated. This has been shown to have the potential to significantly alter the current distribution during switching[14].

However, there are drawbacks inherent with this workflow:

- The technique described above to model a composite device by layering the results from independent simulations, works only for specific

geometries. In the case of the Westcode module, it is effective because virtually all the current flow is orthogonal to the 2D plane(s) modelled. This means that there will be very little (magnetic) coupling between layers, and that each set of extracted inductances can be considered coupled only within themselves. For other geometries, this may not be the case.

- Extracting circuit parameters at a single frequency, and then using them as linear elements over a wide range of frequencies may lead to inaccuracy in the simulation. This can be mitigated with additional processing of the extracted circuit elements.

2.2.4 Workflow 4: Fasthenry Extraction with PSpice Simulation

This workflow uses the "free of charge" 3D parasitic inductance extraction tool Fasthenry to produce impedance matrixes. These matrixes are then converted into PSpice resistors and coupled inductors which can then be inserted into a PSpice schematic and simulated[24].

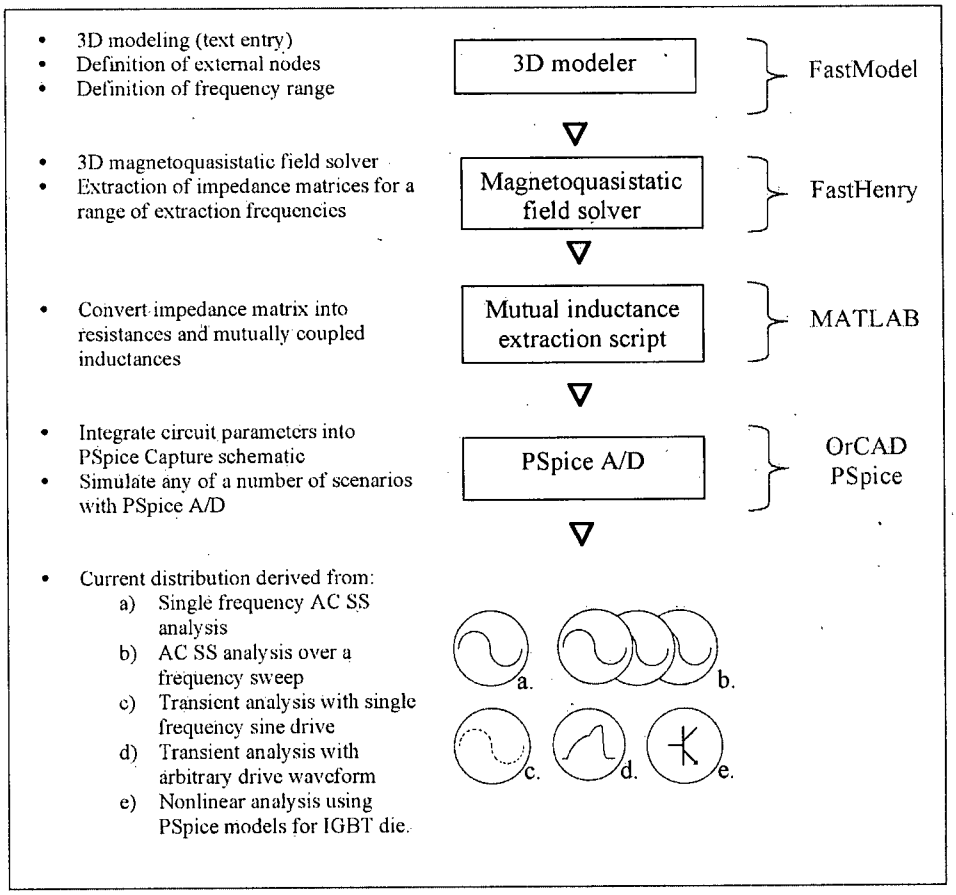

Figure 2.9: Workflow 4 - Fasthenry extraction and PSpice simulation

Fasthenry is a magnetoquasistatic multipole parastitic inductance extraction program developed in the mid-90's at MIT[24]. Models are created

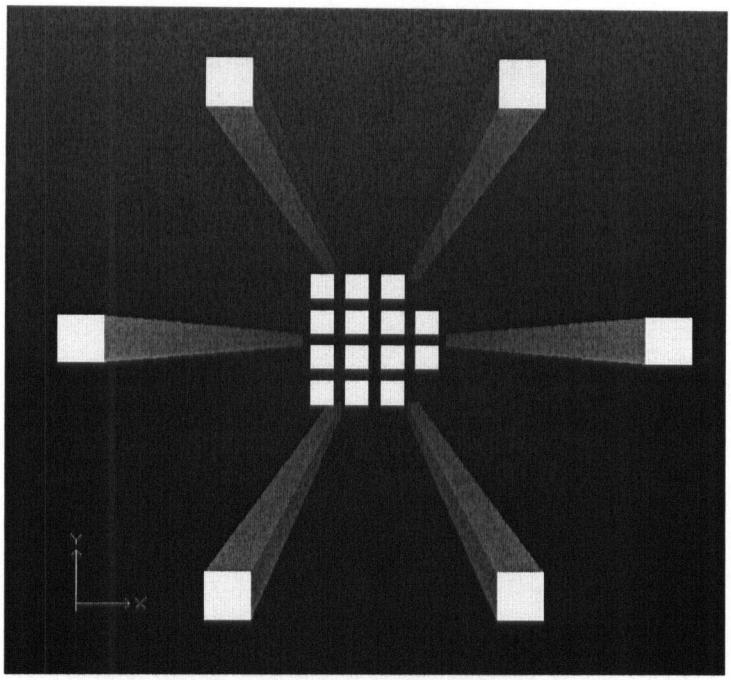

Figure 2.10: Fasthenry 3D geometry for circuit parameter extraction

from a text input file by defining nodes and then connecting these nodes with rectangular conductors. The software also has the capability to create conducting planes by meshing together conductors.

Augmenting the text entry system is another, currently free of charge, software product called FastModel [1]. FastModel provides a real-time display of the geometry as it is entered, and can also be used to configure and launch FastHenry. The text entry system, while ill-suited to geometries not constructed entirely of straight-segments, is surprisingly efficient in creating simple models.

External nodes are defined representing ports for the network solver. The solver commences to analyze the effect of an AC drive applied to one port, on all the other ports. What contributes to make FastHenry unique among our solvers is its use of "partial inductances". Partial inductances

[1] Available from http://www.fastfieldsolvers.com/

are a mathematical abstraction that permits inductances to be calculated without an explicitly defined return path[25]. Instead, the 'return path' is considered to be infinitely far from the conductor, and the conductor's self- and mutual-inductances are calculated accordingly. Defining conductors in this fashion requires that the mutual terms be considered whenever PSpice simulations are run, as it is through mutual coupling that the effect of the actual return paths are considered.

The result of a simulation in FastHenry is an $n * n$ complex impedance matrix, where n represents the number of external ports present in the model. Following from the example given in OERSTED circuit parameters section (section 2.2.3, above): for a FastHenry model with x forward paths, and y return paths, $n = x + y$. Each impedance matrix is valid for a single extraction frequency (see Appendix B).

To convert the impedance matrix to a PSpice include file, a MATLAB script was written. This script is virtually identical to the script used to convert the OERSTED derived impedances and also produces a similar PSpice `.include` file.

Simulation of the Fasthenry extracted circuit elements in PSpice also follows the same procedure, and they can be used in the same circuits as those generated by OERSTED. While Fasthenry is a 3D simulation program, and therefore able to simulate heterogeneous structures, it is also possible to combine the extracted circuit elements from independent simulations in the same way as with the OERSTED workflow. For the Westcode geometry, we chose to implement the model using the latter of these techniques. This permitted us to examine the silicon and copper regions of the module separately within PSpice and yet avoid forcing Fasthenry to solve the much larger matrix that would result from attempting to include these (largely uncoupled) regions together in one composite model.

Simulation Time The time taken to produce an impedance matrix in Fasthenry varies with the number of ports, conductors and filaments in the model. For our model (20 ports, 20 filamented conductors) of the Westcode module this came to $\approx$**15m** per extraction. For each additional material

layer, or desired frequency, a separate extraction must be performed.

PSpice is able to produce either transient or AC solutions to simple circuits that make use of these extracted circuit elements in <1 minute on average.

Summary As this workflow employs PSpice as a simulation tool and uses identically formatted circuit parameters to the previous workflow it follows that it shares all the corresponding advantages. Additionally, being a 3D solver, Fasthenry is capable of modelling geometries that are difficult or impossible for 2D solvers to render[14].

Circuit elements derived by Fasthenry are, like those from OERSTED, produced for a single frequency and have similar limitations and mitigation techniques.

2.3 Simulation Results

This section compares the results of simulations run using the techniques and models described above. As the simulation techniques are themselves experimental, much of the focus of this section is on comparing the results of these techniques. We are especially interested in determining the dependability of the two workflows which use PSpice as a simulation engine (Workflows 3 and 4), for it is these techniques that are the most flexible.

2.3.1 AC Steady State Results

The following simulation results compare the AC magnitude of the current distribution, as derived by the above techniques, over a swept frequency range from 1Hz to 10MHz.

The first section below considers a hypothetical scenario of a "pure copper IGBT module; this simulation allows us to directly compare the results from all of the workflows. The second section details results from a more complex model, which includes a small slice of a material with the resistivity of silicon to represent the sandwiched IGBT die.

"Copper Only" Frequency Sweep

Field Solvers Initially the module was modelled using the FEM and BEM field solvers as a hypothetically pure copper device. It was anticipated that these tools would provide the most accurate distribution of the current for this limited model as they perform a complete simulation for each frequency step specified. Figure 2.11 presents the results from the field solvers.

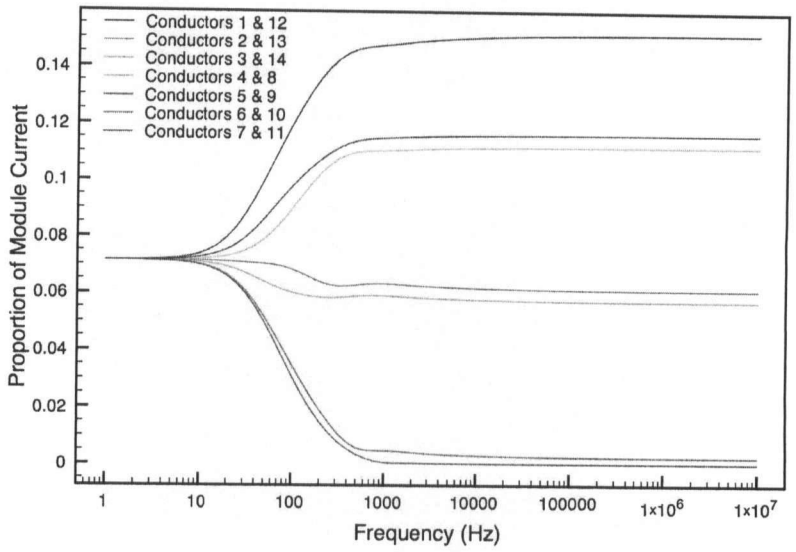

(a) Opera 2D (Workflow 1) prediction for an "all-copper" Westcode IGBT module.

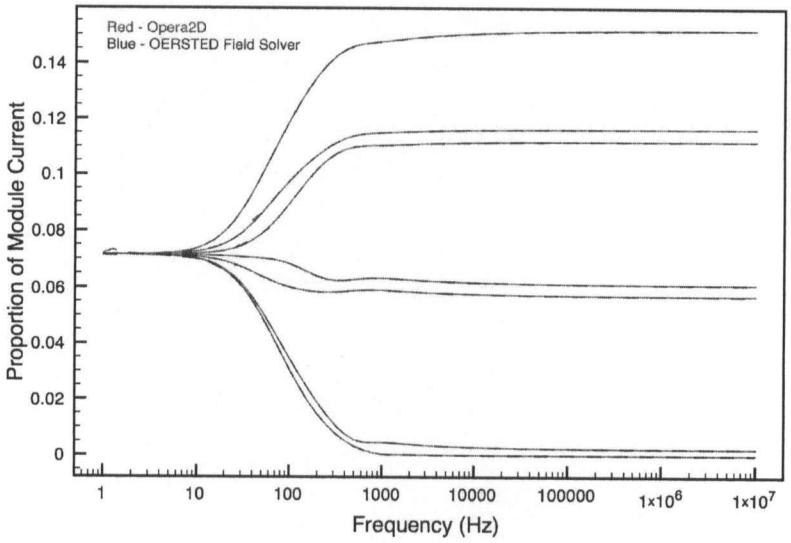

(b) Comparison: Opera 2D (Workflow 1) and OERSTED Field Solver (Workflow 2).

Figure 2.11: Field solver predictions of steady state current distribution for an "all copper" module

Figure 2.11(a) is a plot of the proportion of the module current predicted to flow through each path by Opera 2d as a function of drive frequency (Note: due to symmetry within the module, each plotted line represents two current paths, refer to Figure 1.5 for numbering). At low frequencies (<10Hz, in this case), the distribution is near-uniform and each conductor assumes $1/14^{th}$ of the total module current (≈7.14%). As the frequency increases, the inductive portion of the impedance begins to dominate, and the distribution becomes increasingly less uniform; conductors near the edge of the device take more current as the drive frequency increases while innermost conductors take less. Above a certain frequency (≈1 kHz in this example) the current distribution becomes almost entirely determined by the relative values of the coupled inductances of the module. At the higher frequencies the innermost conductors (Conductors 5 & 9, and 6 & 10) conduct little to no current, while the conductors at the corners of the device (Conductors 1 & 12) conduct almost twice as much as they do at low frequencies.

The second portion of this figure, 2.11(b), illustrates the very good agreement between the two types of field-solver. Additionally, as the OERSTED model explicitly defines all 14 forward paths through the module without taking advantage of symmetry, and yet still presents only seven discernible lines on the plot, it can be assumed that the assumption of symmetry in the Opera 2D model was correct. As such, *for the remainder of this document, reference will only be made to the lower half of the IGBT module (Conductors 1 → 7)*, with the symmetric top portion assumed to exhibit the same characteristics.

The only readily apparent difference between the distribution as predicted by Opera 2D and the OERSTED field solver, is at very low frequencies (≈2Hz). Here, some anomalies are present in the OERSTED simulation. This is most likely attributable to the solver's use of 1D boundary elements at the surface of the conductors to solve for what is likely to be, at these low frequencies, a bulk distribution of current. For such distributions, FEM modelling is much more appropriate. However, as the discrepancy is small, localized and explainable, it is not considered to be problematic.

Circuit Solvers The plots in Figure 2.12 compare the distribution obtained through PSpice simulation of extracted circuit parameters to the field solver solution. Figure 2.12(a) presents the distribution obtained from simulating Fasthenry extracted circuit elements, and Figure 2.12(b) the results using OERSTED derived parameters[2].

The following observations can be drawn from Figures 2.11 & 2.12:

- The four plots all share the same general shape when transitioning from one extreme to the other.
- Additionally, this transition takes place over a similar frequency range for all four techniques: approximately two decades.
- All four workflow produce similar distributions at simulation drive frequencies below approximately 10 Hz; the predicted distributions also plateau to similar values above approximately 10 kHz.
- If the distribution predicted by the field solvers can be assumed to be the most correct, then the two PSpice solved techniques underestimate the magnitude of the imbalance in this hypothetical "copper only" module for approximately one and a half decades spanning upwards in frequency from approximately 10 Hz.
- In both PSpice simulations, Conductors 2 and 4 are predicted to take a considerably higher proportion of the module current between 100Hz and 1kHz (see circled regions). This is notable, as it is the only location where the shape itself of the distribution, as it transitions from resistively dominated to inductively dominated, differs significantly between workflow.

[2]It is of note that the extracted circuit elements simulated to produce the following plots (and, unless otherwise specified, all subsequent PSpice-derived results) are amalgamations of two separate extractions. One of these extractions was performed at a low frequency (1Hz), and the other at a high-frequency(1MHz). The circuit elements used combine the resistive values from the low frequency extraction with the inductances and coupling terms of the high frequency extraction, as it was found that this particular combination gave the best 'fit' over the simulated frequency range without resorting to ladder networks. Applendix B details the rationale and methodology behind this decision.

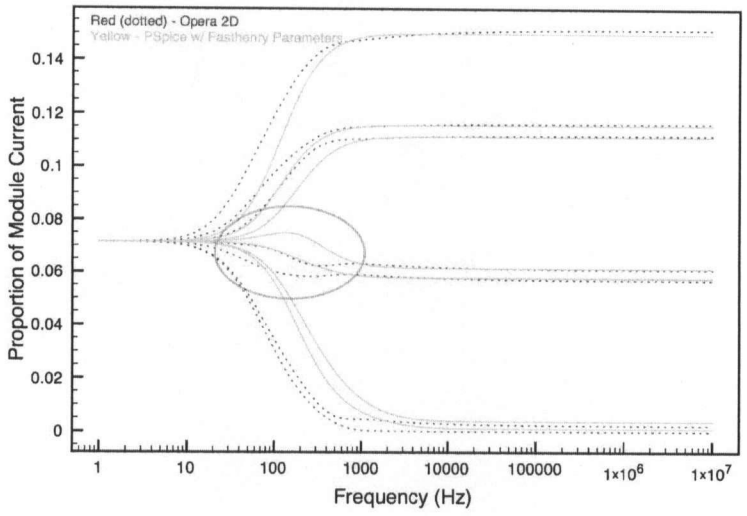

(a) Comparison: Opera 2D field solver (Workflow 1) with PSpice model using Fasthenry extracted parameters (Workflow 4).

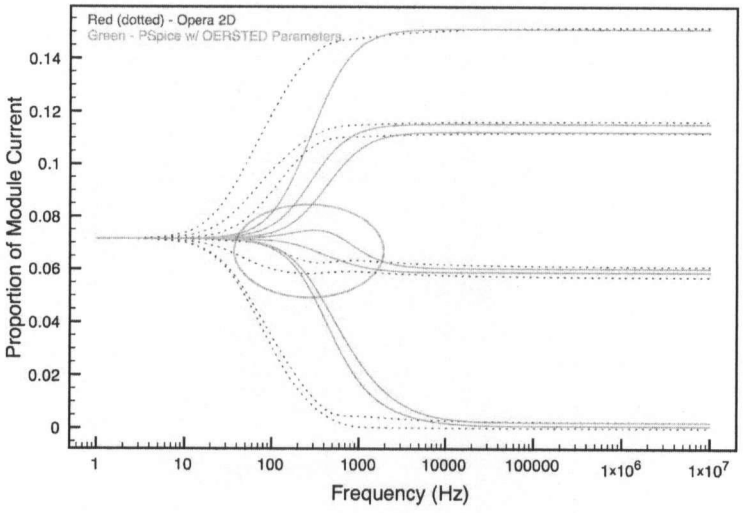

(b) Comparison: Opera 2D field solver (Workflow 1) with PSpice model using OERSTED extracted parameters (Workflow 4).

Figure 2.12: PSpice simulated predictions of steady state current distribution for an "all copper" module

At the extreme ends of the spectrum, the different workflow appear to agree well. In Figure 2.13 a comparison between the simulation techniques is made at 1Hz, and then again at 1MHz. As these particular distributions represent the lower and upper asymptotes, respectively, they are applicable over the majority of the spectra, excepting the two decades of transition mentioned above.

The low frequency distribution in Figure 2.13(a) is identical to well within $1/100^{\text{th}}$ of 1% of the total module current. It is interesting to note that, at low frequencies, the PSpice predicted currents are more uniform between die than the field solver values. This is possibly a result of accumulated error from integration in the field solvers, but is too small to be of any concern.

At 1MHz, Figure 2.13(b), the predicted current distribution is also similar across the various workflow. The largest spread is for the predicted current through Conductor 4, and is less than 0.2% of the total module current.

Choice of Workflow If the field solver solutions produced by Workflow 1 and Workflow 2 are to be considered accurate representations of the current distribution within this hypothetical pure-copper IGBT module, then they can be used to determine which of the PSpice simulation workflow is best suited for continued investigation of the Westcode module.

There is relative insignificance in the variation of predicted current distribution in the low and high frequency plateau regions of the AC sweep analysis. Therefore it follows that the choice of extracted parameters should be those that most closely follow the transition from resistively dominated to inductively dominated. *It was thus decided to use the Fasthenry parameters, adjusted as described in Appendix B, to perform further analysis of the Westcode IGBT module.*

Fasthenry is also attractive as an extraction tool as it is much faster than OERSTED in performing extractions, making it easier to prototype new designs. Additional attractions of FastHenry are its ability to handle 3D geometries, and its low price (free).

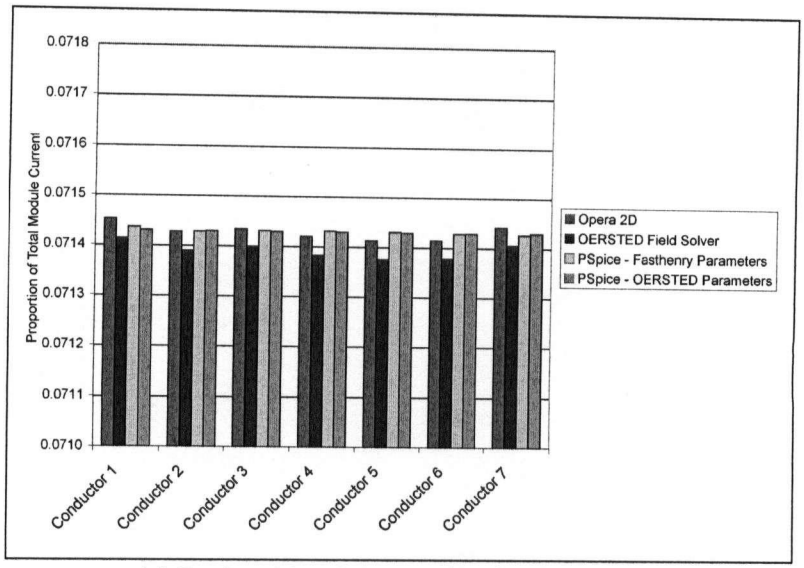

(a) Predicted low frequency (1Hz) distribution.

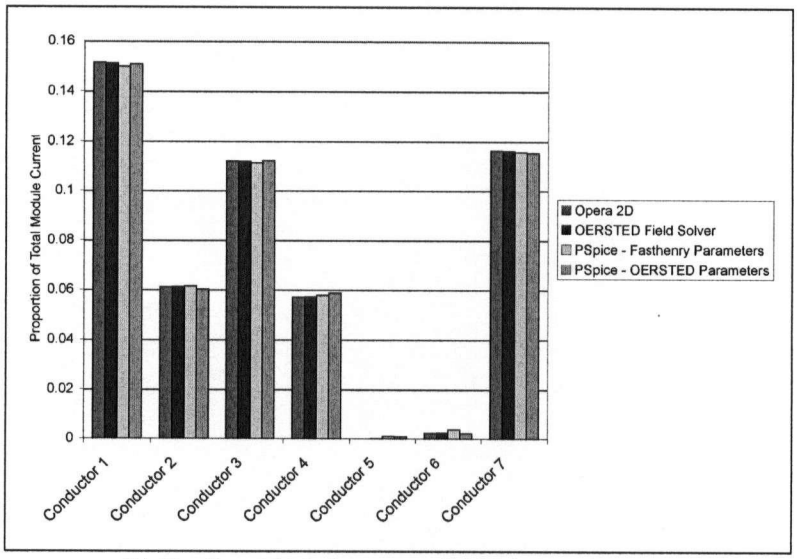

(b) Predicted high frequency (1MHz) distribution.

Figure 2.13: Comparison of workflow predicted distributions at low frequency and high frequency AC drives

2.3.2 Resistive Module Results

With a simulation methodology chosen, it was decided to perform simulations more representative of the actual IGBT module characteristics, while retaining a simple, linear model.

This involved performing a second set of parasitic extractions of the same geometry used for the above copper module, this time with the resistivity of the simulated material adjusted to match that of an IGBT die in the linear on-state. Initially, as neither detailed specifications, nor a dismantled device was yet available, an estimation was made of this value from manufacturer datasheet information then combined with a estimate as to the IGBT die dimensions. Later, the on-state resistance was determined experimentally, the dimensions of the die measured and the models revised. This process is detailed in Appendix A.

As described in Sections 2.2.3 and 2.2.4, the coupled parameters obtained from this second set of extractions are combined with the previously extracted copper elements. The composite structure is then simulated in both the frequency domain, with an AC sweep, and in the time domain, with current pulses of varying rise-times.

Since the measured and calculated values obtained for the IGBT linear on-state resistance exhibit some differences, and are only applicable to IGBT die already conducting a significant amount of current ($>\approx$40A/die), a range of die resistances was analysed. These were derived by scaling the resistive portions of the originally extracted parameters in PSpice to produce an appropriate value for the die resistance. This provides a convenient method to quickly asses the effect on the current distribution of changes in on-state resistance (for example, resulting from an altered gate voltage), or to examine the distribution earlier in the switching process when the IGBT die has not fully turned-on. This approach is also considered valid because the length of the copper post is significantly longer than the length of the IGBT die: hence the effect of the self-inductance of the post is much greater than that of the die. Therefore, the slight change in inductance of the die (at a given frequency), resulting from the altered resistivity, is a comparatively

small effect.

2.3.3 AC Steady State Results

In Figure 2.14(a), the effect of changing the value of the simulated "die" resistance on the steady state current distribution is examined. The resistance values simulated were chosen to represent the possible range of IGBT on-state resistances (the 12.5mΩ, and 25mΩ traces), and to characterize the transition from the copper-only distribution (the 0.25mΩ, and 2.5mΩ traces).

From this plot, it is apparent that each 10-fold increase in die resistance results in the shifting of the transition point, from a resistively determined distribution to an inductively determined one, by one decade. This is confirmed by examining the phase angle of the module. The frequency at which the phase angle becomes 45°, indicating the inductive portion of the impedance is equal to the resistive portion, also shifts linearly as the die resistance is increased. This is borne out by theory:

$$\theta = \tan^{-1}(2\pi f L/R) \tag{2.2}$$

Therefore a increase in R produces the same effect on the phase θ as a proportional increase in f. This frequency also corresponds to the point at which the first sign of divergence in the module current is visible in Figure 2.14(a).

An immediate corollary from this is: increasing the on-state resistance of each path through the IGBT module improves the distribution of current throughout the module. An additional conclusion that can be reached from looking at these plots is that the value of the simulated die resistance has no effect on the final high-frequency current distribution; this distribution appears to be inductively determined by the geometry.

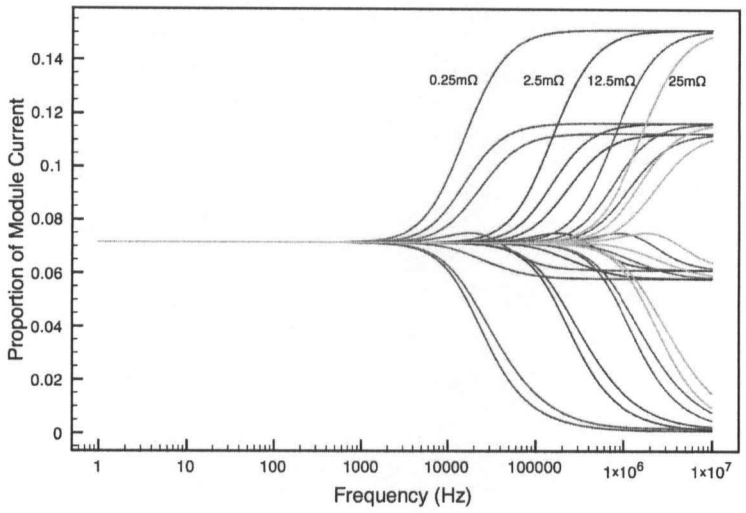

(a) Effect on current distribution of varying resistance of Si region. Resistances shown on plot are per-path.

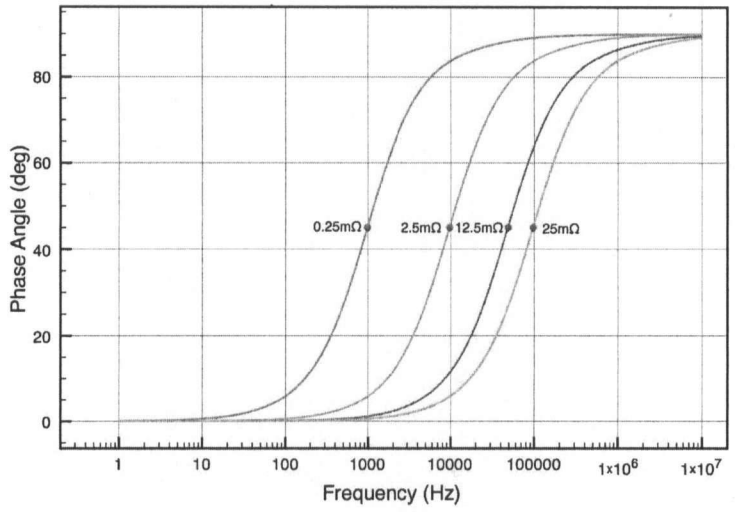

(b) Phase angle of module impedence as a function of Si region impedance. Red dots indicate 45° phase angle

Figure 2.14: Effect of introducing resistive Si region to model current distribution

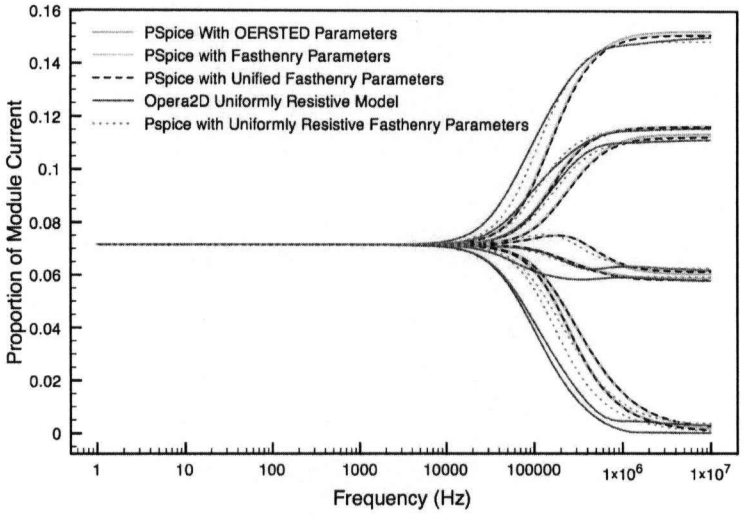

Figure 2.15: Simulation tool comparison with 2.5mΩ Si region resistance

Figure 2.3.3 revisits the choice of workflow. On this plot the current distribution is analysed for the following:

- Workflow 3 & Workflow 4, the two PSpice workflows, using OERSTED and Fasthenry as extraction tools respectively, are compared. Both are using so-called "mixed" parameters with inductive and resistive values combined from 1Hz and 1MHz extractions, respectively.

- A PSpice simulation of a single set of Fasthenry-extracted parameters that represents the combination of the copper and silicon regions. This simulation models the silicon region as a separate slice within the Fasthenry 3D model of the IGBT module.

- An Opera 2D predicted current distribution with a uniform conductivity calculated to reflect a total resistance of 2.5mΩ (the average resistance of the post and IGBT die) for each current path. The value of the conductivity used to create this was 434Sm^{-1}.

- A PSpice simulation using circuit elements extracted from a Fasthenry

model with this same value of conductivity. In this model, the IGBT die and post combination is modelled as being one homogeneous volume.

Workflows 3 & 4 and the Unified Fasthenry model, all produce nearly identical current distributions. This hints that the shift in the frequency of divergence, observed between workflows illustrated in Figure 2.12, is likely a result of differences in the extracted resistance values for each copper conductor. Given that any resistance value chosen to represent the IGBT die, will swamp the extracted value for the copper post, this source of error can be considered to be insignificant for this particular application of these circuit elements.

The Opera 2D and PSpice simulations that use homogeneous uniform resistivities, both predict divergence at a lower frequency than the other simulations. These can be viewed as less-correct. The rationale for this is as follows: The current in regions simulated as highly conductive copper is able to redistribute itself within each conductor; it will fringe to the outermost edge. When a uniform resistivity is assumed, the skin depth for this same region is much larger, and the current will fringe much less. This redistribution of current has the effect of reducing the coupling between the IGBT posts and to the return path (by approximately 5%, as predicted by Fasthenry), shifting the divergence point downward in frequency.

2.3.4 Pulsed Time Domain Results

This section investigates the simulated current distribution within the IGBT as it is driven by a current pulse of varying rise-times. The pulse is generated in PSpice by discharging a 5Ω coaxial cable through the PSpice model of an IGBT module, and into a matched load. Figure 2.16 shows a circuit diagram of this idealized pulse forming network as implemented in PSpice.

The length of the pulse is determined by the delay time of the transmission line. In this case, the line is comprised of two 2.5μs segments connected together to form one 5μs long cable which will deliver a pulse 10μs wide. The rise and fall times of the pulse are determined by the transit time (`TTRAN`)

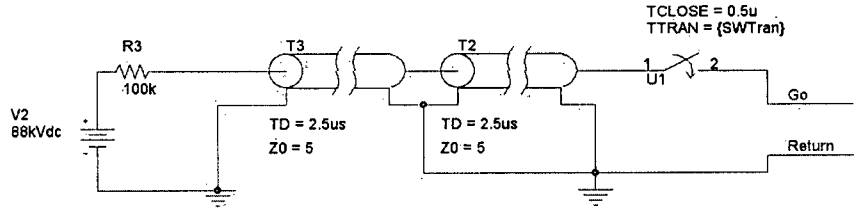

Figure 2.16: Pulse forming circuit used for time domain pulsed distribution analyses

of the switch U1. This parameter was adjusted for each simulation case to produce the requisite 10%→90% rise-time. The amplitude of a current pulse from a PFN into a matched load is 1/2 of the PFN voltage divided by the characteristic impedance. Here, an 8kA pulse was desired, and was obtained after a small (10%) increase in the PFN voltage from the anticipated 80kV to accommodate for a voltage drop established across the 100kΩ resistor by the initial DC bias-point calculation.

Figure 2.17 shows the simulated current in each conductor during turn-on when a module with 12.5mΩ on-state resistance is pulsed by an 8kA pulse with a 10%→90% rise-time of 300ns. Again, as in the simulations described in the previous section, a linear resistive element was used in place of the IGBT die.

As predicted by the frequency domain measurements, Conductor 1 takes a significant amount more current than the other die during the switching period. In fact, the current in this conductor is predicted to overshoot what would (ideally) have been it's maximum current by ≈180A, (30%).

The 8kA pulse requires the switch to conduct a peak di/dt of 30kA/μs. This is an ambitious figured that, as pointed out in the background portion of this document, has not yet been realized with IGBT switches.

The plots in Figure 2.18 illustrate effect that altering rise-time has on the current distribution within the simulated Westcode module. Also investigated in these plots is the effect of adjusting the simulated on-state resistance of the IGBT die.

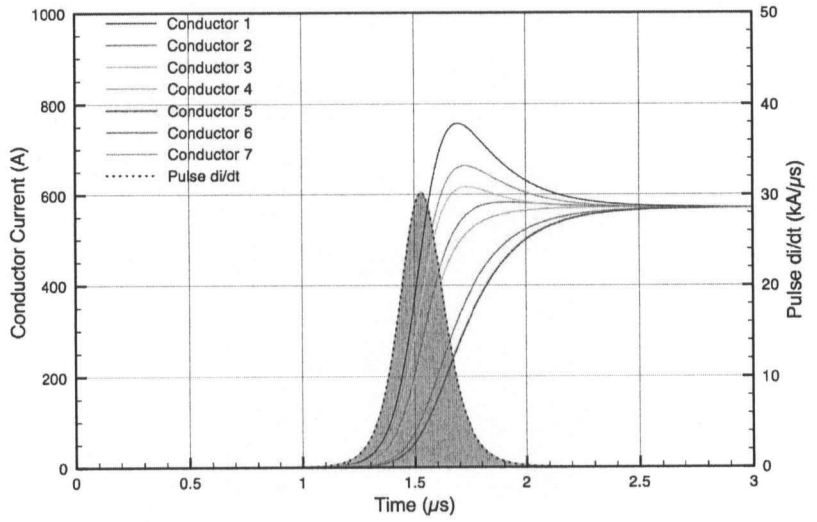

Figure 2.17: Current distribution for simulated module with 12.5mΩ on-state resistance and a 10% to 90% rise-time of 300ns

Figure 2.18(a) confirms the results of the frequency domain analyses in Figure 2.14. The linear relation to the on-state resistance observed in the frequency domain plots is again seen here: doubling the on-state resistance permits the rise-time to become halved, while still maintaining the same distribution.

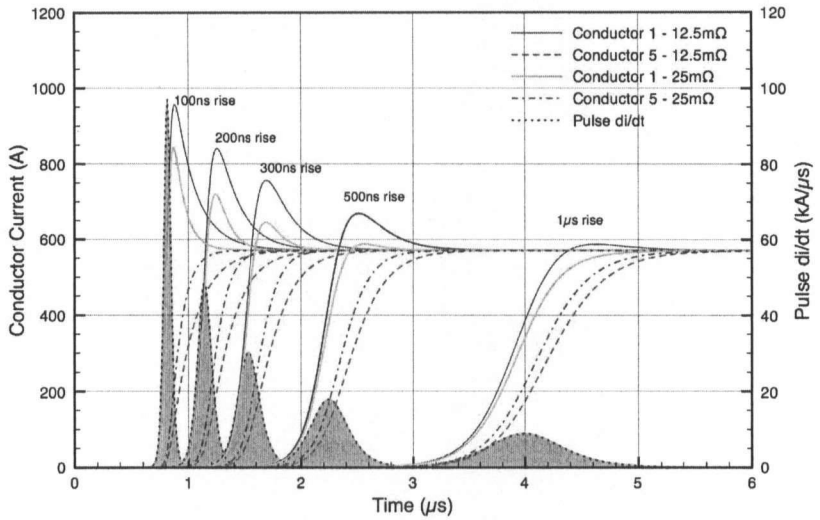

(a) Comparison: Module current spread for different pulse rise-times. Data for 12.5mΩ and 25mΩ on-state resistances shown. Solid traces indicate most-conducting current path, dashed traces indicate the least-conducting path.

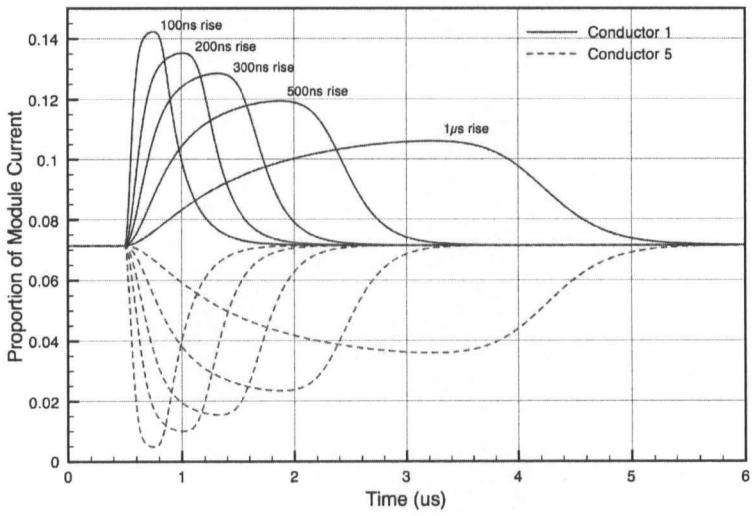

(b) Relative proportion of die current within module for various pulse rise-times. 12.5mΩ on-state resistance data shown.

Figure 2.18: Effect of varying rise-time and on-state resistance on current distribution for a 8kA pulse

This follows from the relationship between the rise-time and edge bandwidth of a signal. For a single pole circuit such as this, the equivalent edge bandwidth for a given 10%→90% rise-time can be estimated from the following approximation[26]:

$$BW \approx \frac{0.35}{T_r} \tag{2.3}$$

Here, BW represents the equivalent edge bandwidth, and T_r the rise-time of the pulse. This approximation is a rule-of-thumb used for determining the choice of measurement equipment for measuring pulses. What is important to note is again the linear relationship between the rise-time and frequency content of the pulse.

In Figure 2.18, we can see that, as the rise-time of the pulse decreases, the relative distribution of current within the device tends towards the inductively dominated distribution: for a 100ns pulse, the distribution is nearly completely inductively determined.

Up until this point, we have assumed that the geometry of the model was fixed. However, as it is a simple matter to investigate in PSpice, illustrating the effect of reducing the height of the copper region in the IGBT module is presented in Figure 2.19.

In this plot the numbers are indicative of the copper stand-off post height in the simulated device relative to the actual post height of the Westcode module. Shortening the height of each post has the effect of reducing L in Equation 2.2 by an equivalent amount. This shifts the frequency of divergence higher, and has the corresponding effect on the rising-edge distribution of the pulse.

Conclusion The results of simulations performed of the Westcode IGBT module present a largely consistent picture of the current distribution within the device. The results of the frequency and time domain simulation corroborate each other well, and the different workflow produce distributions that are largely in agreement.

The sensitivity of the predicted current distribution within the module

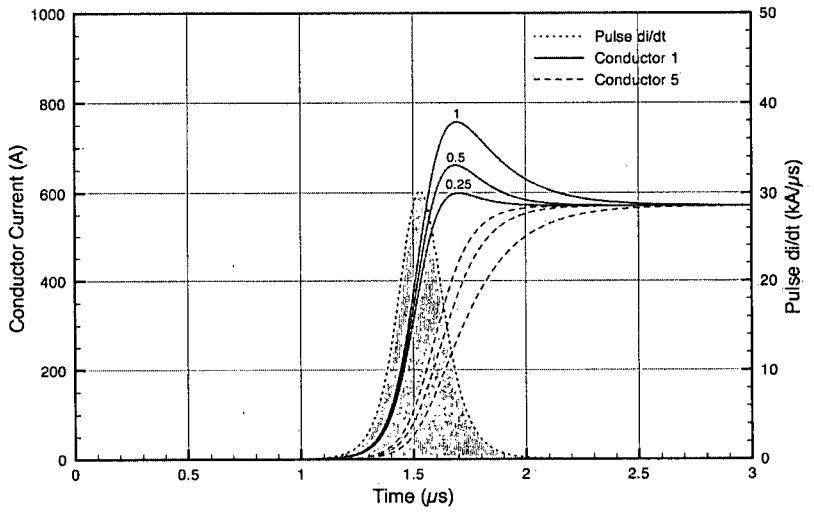

Figure 2.19: Predicted current distribution for simulated 12.5mΩ module with scaled post-heights.

has been established with regards to: pulse rise-time, die on-state resistance, AC drive frequency and copper post height. It has been found that there is a linear relationship between these factors and the uniformity of the distribution. In the case of the pulse rise-time and the die on-state resistance, this relationship is direct. For the remaining two factors, AC drive frequency and copper post height, the relationship is inverse.

Chapter 3

Experimental Verification

The final step in confirming the simulated model was to subject the physical module to experiments to determine the current distribution within the module.

Measurement Goals As the IGBT module is a compact, closed device, acquiring the current distribution is not a simple procedure. This chapter details the equipment and procedures necessary to produce results that can be compared to the simulations.

3.1 Method of Analysis

3.1.1 Frequency Domain

It was decided to initially perform measurements in the frequency domain. While the focus of the research is inherently a time-domain problem, in that it is concerned with transients occurring during fast switching, a frequency domain analysis of the problem is advantageous for two major reasons:

1. It greatly simplifies the process of acquiring measurements.
2. It also aides interpretation of the measurements: allowing the experimental results to be directly compared to simulation results.

One of the reasons frequency domain measurements are easier to perform than time domain measurements is related to our choice of measurement equipment. Rogowski coils produce an output voltage that is proportional to the *derivative* of current flowing through the 'loop', and as such usually require some form of integration.

However, when the current source is a single frequency AC drive, the output of the Rogowski coil is directly proportional to the current, albeit phase shifted and scaled linearly by the frequency (see Section 3.1.3).

Measurements are further simplified with a frequency-domain analysis, in that the capacitive energy storage and fast gate drive circuits required for fast turn-on of the IGBT module are not required.

While this document is only concerned with the frequency domain analysis of the Westcode module, it is important to note that all of the equipment (e.g. the test jig and the Rogowski coils) have been explicitly designed for use in any subsequent time domain analyses that may be performed.

3.1.2 Resistive Elements (IGBT Replacements)

Experimental determination of the IGBT on-state resistance revealed that a minimum of approximately 40A was required, per die, to put the IGBT into the linear region of its V/I relationship (Applendix A). This was unfeasible, as it would require a stable DC current of almost 600A be flowing through the complete Westcode module for the entire time that measurements were being performed. While current supplies of this magnitude are available at TRIUMF, the amount of power that would be dissipated by the module (in the order of 300W+), plus the complexity of measuring what would be a relatively tiny AC signal in the presence of such a large DC current, forced the development of a different solution.

To introduce a series resistance for each post, and to attempt to assure consistency, it was decided to solder four parallel low-tolerance (5%) surface mount resistors between two layers of copper foil. These sit upon the copper posts within the IGBT module. Uniform contact is assisted by bending up the corners of each resistive element so that the force of the top brass spacer (Figure 3.7(c)) comes in contact with each of them as it is placed upon the module and the corners flatten. Figure 3.1 illustrates 14 of these deployed in the Westcode module. It is worth noting that the module is not compressed in any way with the resistive elements in place.

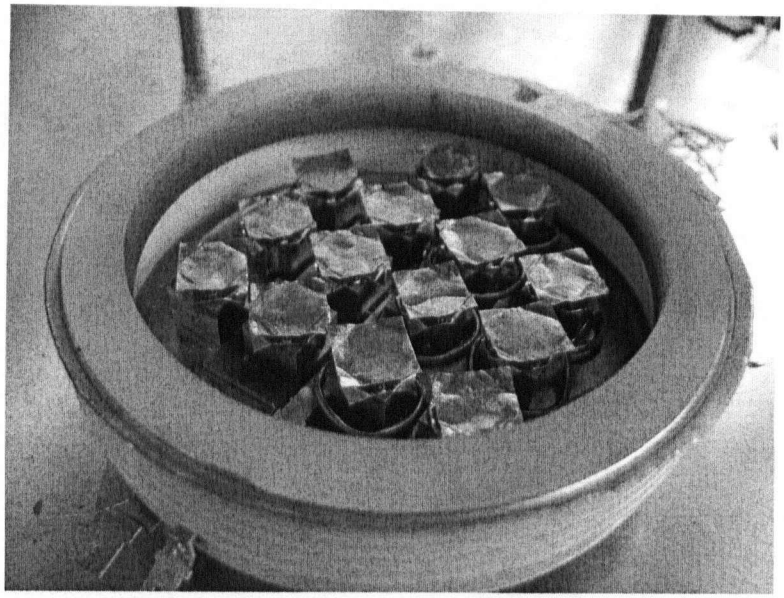

Figure 3.1: Resistive elements to replace IGBT die, in place within the Westcode module

3.1.3 Current Sensing

Measurements required a means of sensing how the current was distributing itself within the IGBT module. The requirements called for a current sensor with an inside diameter that was large enough to be able to fit around each post (approximately 1cm across), but an outside diameter small enough to permit two adjacent sensors to share the space between each post (inter-post spacing $\approx$ 5mm). Additionally, as it is hoped to perform time domain measurements of the IGBT switching characteristics in the future, it was a requirement that the current sensor must be able to fit in the very limited space between each IGBT cassette and the gate trace in the Westcode module (a space of just over 1mm). To the author's knowledge, no such sensor is commercially available.

Rogowski Coils

Rogowski coils are air-cored current transducers. They consist of a series of windings around a non-ferritic former, usually flexible. Rogowski coils are a favourite current sensing tool for those wishing to measure high-power transients as they do not saturate with large currents, and can be constructed in such a fashion so that they may be 'opened' and wrapped around an in-place conductor (i.e. a bus-bar)[27].

Rogowski coils exploit Ampere's law, in that the windings around the former act to perform a piece-wise closed line integral of the magnetic field flowing around a conductor. It further follows from Ampere's law that it doesn't matter what shape this path takes, as long as it completely circles the conductor.

Unlike a conventional CT, rogowski coils are not self-integrating, and produce an output voltage proportional to the derivative of the current flowing through the loop.

$$I_{drive} = I \sin(\omega t)$$

$$V_{coil} = K I \omega \cos(\omega t)$$

Where K is a single frequency constant related to the Rogowski coil response.

Rogowski coils have been employed for very similar applications to this in the past[28], and can be fabricated without highly specialized equipment. The photographs in Figure 3.2 detail the Rogowski coils that I constructed, and those in Figure 3.3 show the coils in-place within the module.

The output of the Rogowski coil can be expressed as the following:

$$V_{coil} = \frac{\mu_0 A N}{l} \frac{di}{dt}$$

Where A is the cross-sectional area of the former, N is the number of winding turns, and l the length of the former. The coils made at TRIUMF had 480 turns, a cross sectional area of just under 1mm^2, and a length of 4.8mm. Seven such coils were constructed.

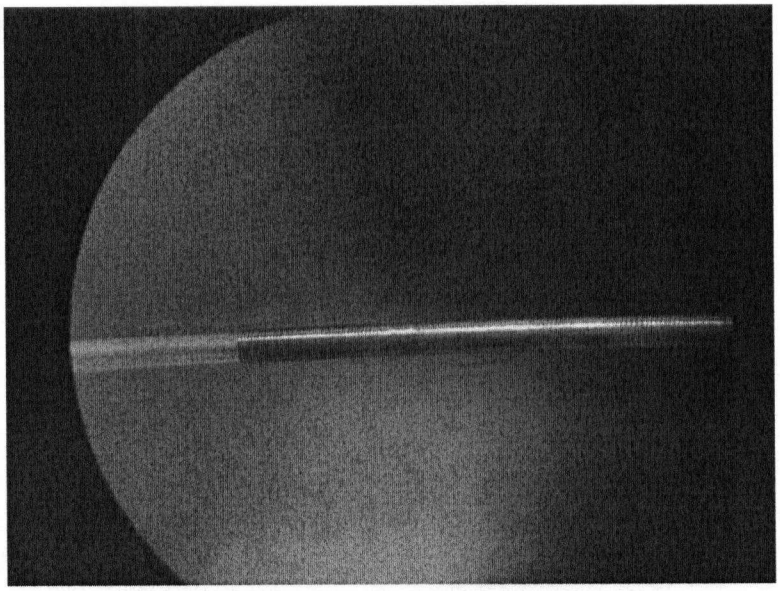

(a) Winding of a Rogowski current-sensing coil.

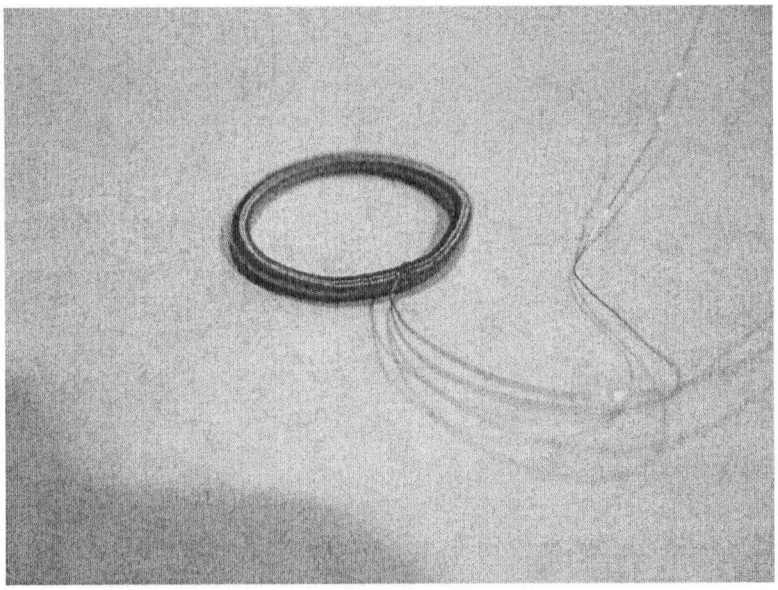

(b) Finished Rogowski coil.

Figure 3.2: Rogowski coil construction

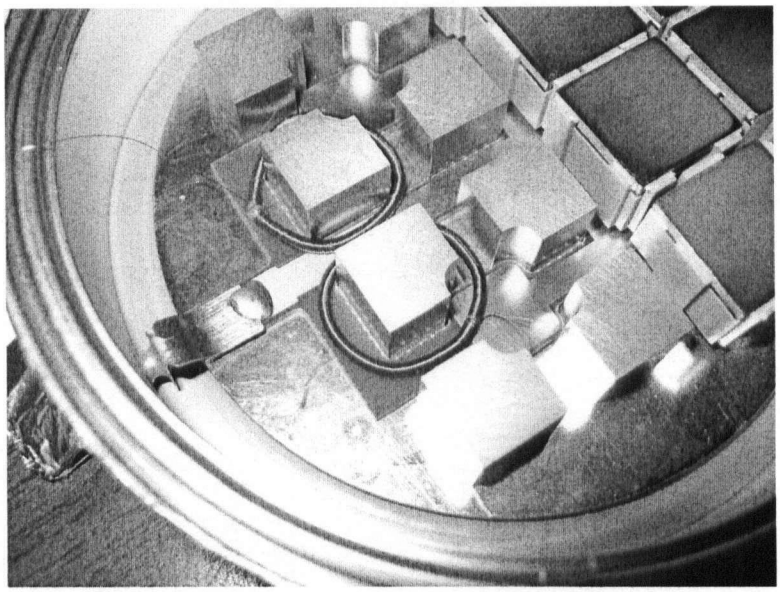

(a) Rogowski coils in place within the Westcode module.

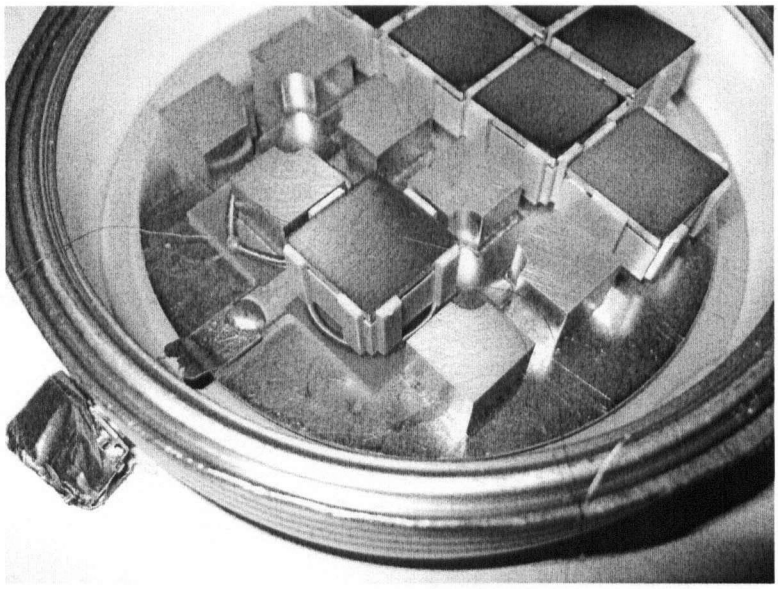

(b) Rogowski coil fitting underneath IGBT cassette.

Figure 3.3: Rogowski coils in-place within Westcode module

Instrumentation Amplifier

As the output from the Rogowski coils is extremely small, especially at low frequencies, an instrumentation amplifier was designed with a gain of ≈10,000. Figure 3.4 shows the schematic for the instrumentation amplifier circuit used. The circuit implemented had R_1 =51.1kΩ, R_G =10.5Ω; and R_2 & R_3 =10kΩ. The expression for the gain of this amplifier is:

$$G = 1 + 2\frac{R_1}{R_G}$$

With the values listed above, the gain of the amplifier is equal to 9963. Four identical channels of this amplifier were built. A photo of the constructed four-channel amplifier is shown in Figure 3.5.

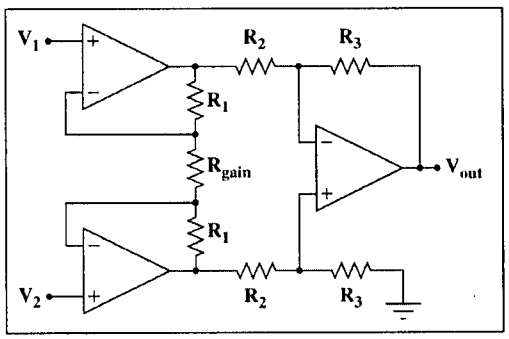

Figure 3.4: Instrumentation amplifier circuit diagram

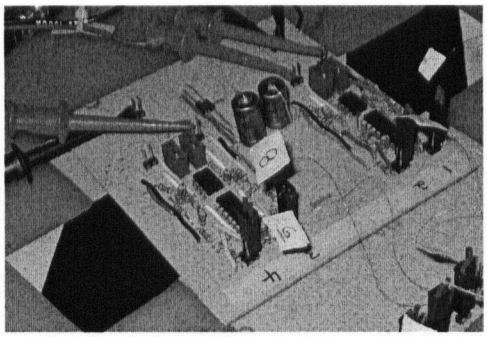

Figure 3.5: Photograph of the instrumentation amplifier used for AC measurements

3.2 Test Jig

3.2.1 Requirements

In all of the simulations described in this document, a geometry for the IGBT and the corresponding return path is presented. This design is intended to reflect a potential electrical layout for a thyratron replacement switch. As such, it is designed to:

- Exhibit low inductance: the geometry of the return path is intended to act as a coaxial structure.
- Produce an unbiased distribution of current within the module.
- Allow for large blocking voltages that would be attained by series connections of modules.

The test jig also needs to exhibit these same electrical characteristics.

It is also necessary that the test jig help to facilitate measurements on the IGBT module. This entails having a relatively short assembly/dissassembly cycle, allowing sufficient space for possible measurements, and, as the process will involve some relatively high-frequency signals (with components >1MHz), a readily accessible common ground for all measurement and supply equipment. Another concern is that the materials used be non-magnetic

to prevent non-linearities. Additionally, to be a general purpose apparatus, the test jig must be mechanically able to compress an IGBT module with sufficient force to make electrical contact, as per manufacturer's specifications[19].

3.2.2 Design

Prevessin Design

After considering the above criteria, and examining the apparatuses already in use to compress presspack IGBT modules and other similarly packaged semiconductors[4, 9], it was decided to use a modified version of a design already in use at CERN Prevessin[1].

The Prevessin design is shown in Figure 3.6. In this design, which has been developed for pulsed measurements only, energy is stored in ten parallel 50Ω coaxial cables rated at 8kV DC. These cables attach to the bottom aluminium plate of the apparatus, where the outside sheaths of the coaxial cables are also electrically connected. The centre conductors of these ten cables then go on to form the 'return' path that surrounds the DUT. The Prevessin design has several advantages, especially when it comes to time-domain pulsed measurements. These include the direct connection of the energy storage devices, and, most importantly, a single ground comprising the base of the module.

However, as our intention is to perform frequency domain measurements (at least initially), this design is not ideal in that it is difficult to supply from a single source.

TRIUMF Design

Components

Figure 3.7 illustrates the components that make up the test jig assembly used at TRIUMF. As with the Prevessin jig, this test apparatus is designed such that the IGBT module is compressed between two aluminium plates (top and bottom illustrated in Figures 3.7(a) and 3.7(c), respectively) with six brass

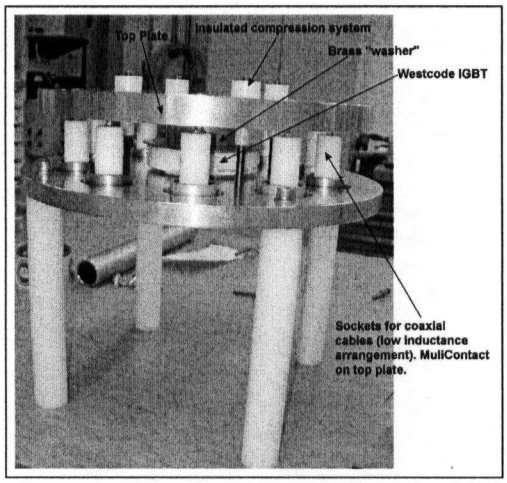

Figure 3.6: CERN Prevessin test jig for presspack IGBTs (Note: "MuliContact"(sic) should read "MultiContact")[1]

conducting rods making up the return path (Figure 3.7(b)). The geometry of this return path is consistent with the simulated models presented in Figures 2.2, 2.5 and 2.10.

Electrical Layout

As mentioned above, the design of the test jig reflects the simulation models, and is also a conceptual design for a thyratron replacement. The frequency domain measurements were carried out at low voltage, and hence did not require HV insulation anywhere in the apparatus to perform.

Figure 3.8 illustrates how this was implemented in the test jig, the arrows show the direction of positive current flow. Further description of the labelled portions of the figure follows:

i. Current enters the jig from the external connection pictured in Figure 3.7(e) (right).

ii. Connection is made to the top brass spacer, which is insulated from the top aluminum plate.

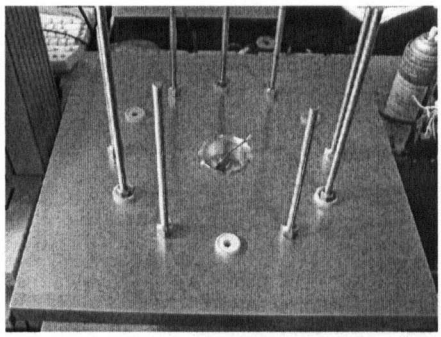

(a) Jig top plate with 6 conducting rods, and 3 clamping rods. Absent: electrical connection plate featured in (f)

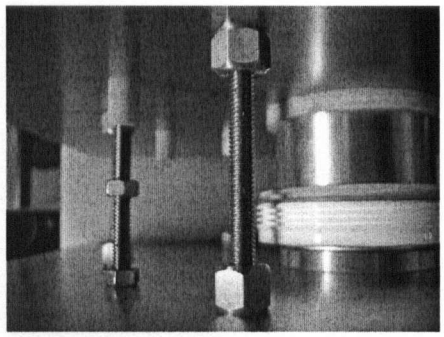

(b) Conducting rod (foreground) and insulated clamping rod (background, left).

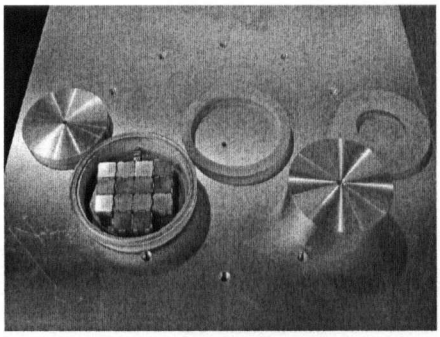

(c) Components comprising module stack: (from left to right) bottom spacer, Westcode IGBT module (with IGBT die), centring ring, top conducting spacer, insulating shim. Background: jig bottom plate.

(d) Insulation for compressing rods.

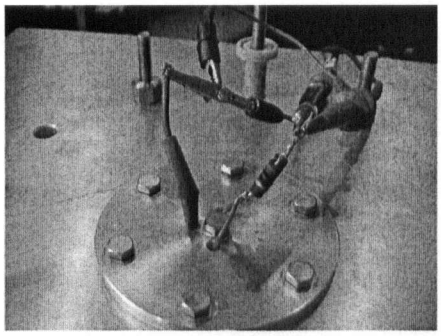

(e) AC source electrical connections to top aluminium plate of jig. Connection point for forward current at right, source return contact at left.

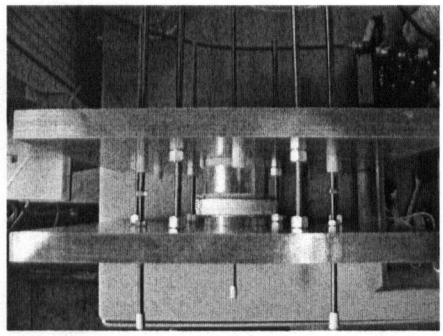

(f) Side view of assembled jig.

Figure 3.7: TRIUMF test jig assembly photos

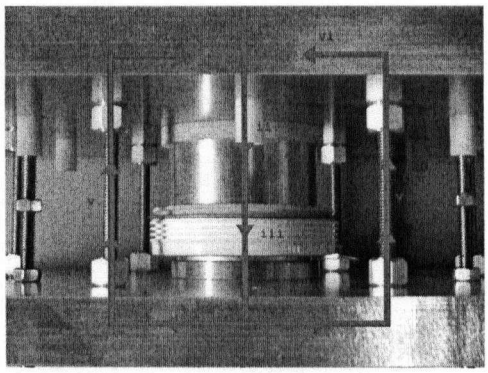

Figure 3.8: Current path through TRIUMF test jig

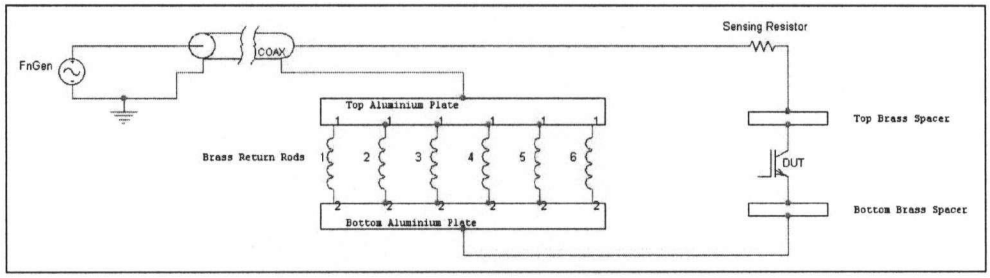

Figure 3.9: Test setup electrical layout

iii. The current gets distributed through the 14 parallel paths within the Westcode IGBT module.

iv. The bottom aluminium plate redistributes the current to the return conductors.

v. The brass rods (two of six shown) return the current to the top plate.

vi. Current goes via the top aluminium plate to the return connector: Figure 3.7(e) (left).

Figure 3.9 illustrates a schematic view of the test layout used for frequency domain analysis.

Mechanical Structure

For measurements that require the IGBT module be compressed, clamping is achieved through the use of either three or, if necessary, six non-magnetic stainless steel rods that are electrically insulated from the top Al plate. These rods are segments of 5/16", 304 grade, stainless steel reddi-rod; the insulation is provided by machined G10 (FR4) epoxy spacers that insulate the rod throughout the entire thickness of the top plate (Figure 3.7(d)). In the TRIUMF design this insulation is not nearly as critical as it is with the Prevessin design because the top and bottom plates are electrically connected at all times via the brass rods. In the Prevessin apparatus, the clamping insulation must withstand the full voltage across the device.

The test apparatus was constructed from non-ferromagnetic materials (i.e. brass, stainless steel, aluminium) such that the return current path would take the same form as the simulated model. For the frequency domain analyses described in this document, the mechanical clamping rods were unused, and the jig was used for the geometry of its return path.

3.3 AC Measurement Results

The measurements were made by exciting the module with a 10V peak AC signal from a 50Ω output impedance function generator, and measuring the output of the Rogowski coils as amplified by the instrumentation amplifier. To remove non-periodic noise, the averaging functionality of the oscilloscope was employed. A 5Ω resistor was placed in series with the test jig and the function generator, and the voltage across this resistor measured for current sensing purposes.

3.3.1 Purely Copper Module

Measurements of the current distribution for the purely copper module proved to be exceedingly difficult. The low-frequency distribution, being dominated by the resistance of each path, is almost entirely determined by the contact resistance to each post. Attempts to get a uniform, or even

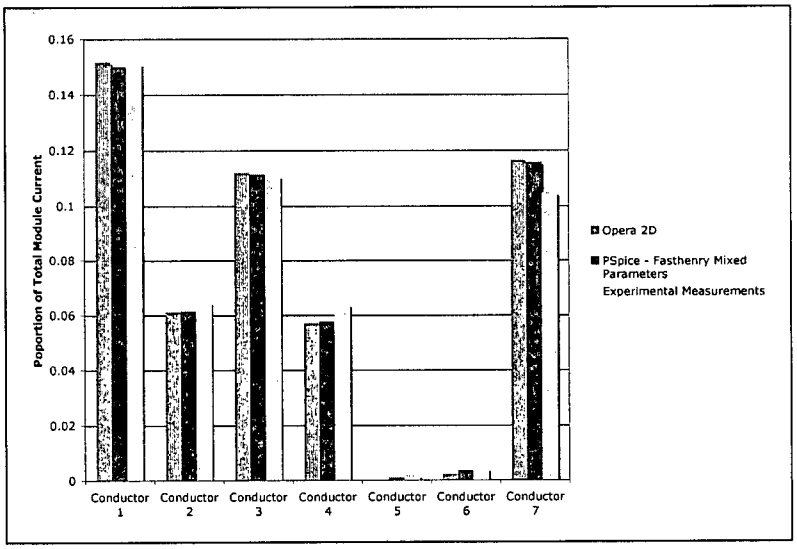

Figure 3.10: Experimental results for high-frequency (2MHz) distribution with copper-only module

consistent, measurement of the low frequency distribution were thwarted by the difficulty of acquiring uniform electrical contact. Attempts were made to improve the distribution by improving the electrical contact to the device (i.e. by the application of conducting paste, and by compressing the module), but these proved fruitless. In some cases these attempts actually made the problem worse as it took even less of an absolute variation in contact resistance to effect the relative distribution of current.

However, even though measurements at low- and intermediate-frequencies were problematic, the copper module did allow us to examine the high-frequency current distribution. This is illustrated in Figure 3.10 where the experimentally determined current distribution at 2MHz is compared with the high frequency asymptote values from simulation.

The agreement for this experimentally determined distribution with the simulation results is good, although not perfect. The largest error is for Conductor 7, where the difference between the simulated and experimentally determined distributions is approximately 1.2% of total module current.

3.3.2 Resistive Module

When attempting measurements on the IGBT module, it soon became obvious that there was an upper-limit of approximately 2MHz to our measurements. Beyond this point, capacitive interference from the module began to dominate the measurements. This was problematic, as simulations showed that, with resistive modules inserted to appropriately reflect the on-state resistance of an IGBT die, the distribution would not fully reach it's inductively determined asymptote by this frequency.

Keeping in mind the potential effect that the contact resistance might have on the distribution, simulations were performed to determine what the maximum value of series resistance could be inserted and while permitting the entire transition from a resistively-dominated to an inductively-dominated distribution to be examined. The results of the simulation pointed to a series resistance of 2.5mΩ per path. The results of measurements with this value of on-state resistance and their comparison to theory are shown in Figure 3.11.

Again, it is likely that the effect of contact resistances played a large role in determining the distribution of current within the module. Figure 3.11(a) shows the comparison to the original PSpice simulation. The low-frequency current distribution is significantly altered by the effect of the contact resistances, which also have the effect of shifting the resistive-inductive transition point upwards in frequency.

A revised PSpice model was made, this time with different values of contact resistance inserted into each path to match the simulation to the experimental low-frequency distribution. The average "contact resistance" added to the PSpice model was 0.6mΩ. The results are shown in Figure 3.11(b).

With the effect of the contact resistances added to the PSpice model, the experimental data correlates much more strongly with the simulation. There are significant discrepancies: most notably the transition point for Conductors 5 & 6, and the high-frequency asymptotes for Conductors 2 & 4. However, it can be said that the agreement between the experimental and simulated results is acceptable.

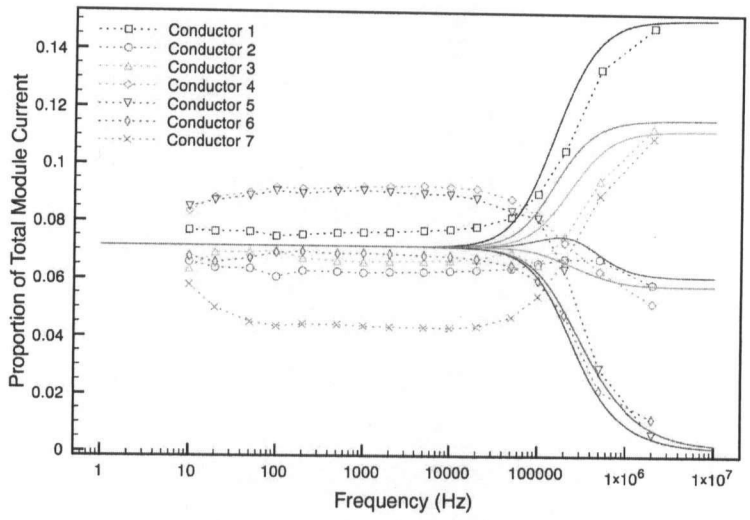

(a) 2.5mΩ comparison with PSpice simulation neglecting contact resistances

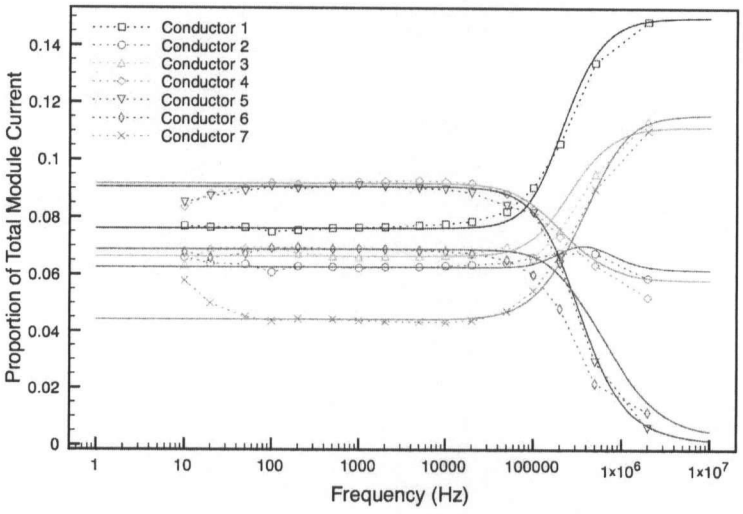

(b) 2.5mΩ comparison after accounting for contact resistances in the PSpice simulation

Figure 3.11: Comparison of simulated and measured current distribution within the IGBT module with 2.5mΩ of series resistance

It is interesting to note that the "bump" predicted by the PSpice simulation for the relative current in Conductor 2 (as highlighted in Figure 2.11) appears to be present in the experimental results. This is interesting as it was not predicted by either of the field solvers, which are considered to be 'more correct'.

Chapter 4

Conclusions and Recommendations for Future Work

This document details the successful design and testing of a methodology to model the geometry of a press-pack semiconductor device as a set of coupled circuit parameters. In the process of creating this methodology, the geometry of a Westcode IGBT module was analysed as to its suitability for application in a thyratron replacement. This analysis concludes that the module design exhibits a high-frequency current imbalance that is inextricably tied to its layout. As has been shown, it is possible to mitigate this imbalance by either increasing the on-state resistance of the IGBT die, or shortening the overall height of the stand-offs upon which these die are placed. However, increasing the on-state resistance also effects losses in the device, and decreasing the post height would increase capacitive coupling to the gate trace from the collector. Additionally, there is a lower limit to the height of the posts imposed by the DC break-down voltage of the the module.

While the model generated appears to be remarkably accurate, especially considering its simplicity (just two coupled inductors and two resistors per current path), there is definitely room to improve its accuracy. This is especially true at intermediate frequencies, where the module distribution is defined by both its resistive and inductive components. In this frequency range, the current model over-estimates the uniformity of the current distribution. I suggest that a method of converting several Fasthenry extractions,

from over a range of frequencies, to a coupled-inductance ladder network be investigated. Given the relatively short amount of time that each Fasthenry extraction consumes (≈15 minutes), and the ability of PSpice to rapidly simulate circuits of far higher complexity than our current model, it is author's opinion that this would be a worth-while endeavour, and would not significantly impact prototyping speed.

Using the existing methodology, there is much future work that can be performed to examine the effects of modifying any combination of internal die layout, or return path configuration. A future model of the IGBT might also introduce non-linear models of IGBT die and consider the effects associated with coupling to the gate trace. Considering the simplicity and speed modifying and simulating the existing model, these should be a relatively simple tasks.

Bibliography

[1] Private communication with M. Barnes June 30, 2006.

[2] R. Cassel and M. Nguyen, "A new type short circuit failures of high power IGBT's," *Pulsed Power Plasma Science, 2001. PPPS-2001. Digest of Technical Papers*, vol. 1, 2001.

[3] M. Barnes and G. Wait, "Solid State Switch Alternatives For Klystron Modulators," in *1995 Second modulator klystron workshop: A Modulator klystron workshop for future colliders*, ser. SLAC : 481, A. Donaldson, Ed. The Center, Stanford, Calif, 1996.

[4] A. Dunlop, "Prototype thyratron replacement using semiconductors [pulsed power switch]," *IEE Symposium Pulsed Power 2001 (Ref. No. 2001/156)*, 2001.

[5] E2V Technologies Ltd. Hydrogen thyratrons preamble. [Online]. Available: http://www.e2v.com/download.cfm?type=document&document=613

[6] M. Barnes and G. Wait, "(tri-dn-05-15) pspice analysis of the ags a10 kicker magnets," TRIUMF, Tech. Rep., 2005.

[7] M. Barnes, G. Wait, E. Carlier, L. Ducimetière, U. Jansson, G. Schröder, and E. Vossenberg, "Operation Modes of the Fast 60 kV Resonant Charging Power Supply for the LHC Inflectors," *Proc. of 11 thIEEE Pulse Power Conference, Baltimore, p1309-p1314, June*, 1997.

[8] G. Wait and M. Barnes, "Thyratron Lifetimes, A Brief Review," in *1995 Second modulator klystron workshop: A Modulator klystron workshop for future colliders*, A. Donaldson, Ed., 1996.

[9] A. Welleman, E. Ramezani, and U. Schlapbach, "Semiconductor switches replace thyratrons & ignitrons," *Pulsed Power Plasma Science, 2001. PPPS-2001. Digest of Technical Papers*, vol. 1, 2001.

[10] A. Welleman, J. Waldermeyer, and E. Ramezani, "Solid State Switches For Pulse Power Modulators," *LINAC 2002, Proceedings of*, pp. 709–711, 2002.

[11] M. Nguyen, "Short circuit protection of high speed, high power IGBT modules," *Pulsed Power Conference, 2003. Digest of Technical Papers. PPC-2003. 14th IEEE International*, vol. 2, 2003.

[12] P. Palmer and A. Githiari, "The series connection of IGBTs with active voltage sharing," *Power Electronics, IEEE Transactions on*, vol. 12, no. 4, pp. 637–644, 1997.

[13] G. Leyh, "A critical analysis of IGBT geometries, with the intention of mitigating undesirable destruction caused by fault scenarios of an adverse nature," *Particle Accelerator Conference, 2003. PAC 2003. Proceedings of the*, vol. 2, 2003.

[14] M. Barnes, E. Blackmore, G. Wait, J. Lemire-Elmore, B. Rablah, G. Leyh, M. Nguyen, and C. Pappas, "Analysis of High-Power IGBT Short Circuit Failures," *Plasma Science, IEEE Transactions on*, vol. 33, no. 4, pp. 1252–1261, 2005.

[15] S. Scharnholz, R. Schneider, E. Spahn, A. Welleman, and S. Gekenidis, "Investigation of IGBT-devices for pulsed power applications," *Pulsed Power Conference, 2003. Digest of Technical Papers. PPC-2003. 14th IEEE International*, vol. 1, pp. 349–352, 2003.

[16] S. Gunturi and D. Schneider, "On the operation of a press pack igbt module under short circuit conditions," *Advanced Packaging, IEEE Transactions on*, vol. 29, no. 3, pp. 433–440, August 2006.

[17] B. Rablah, "An analysis of the effects of parasitic inductance in internal IGBT module geometry," 2004, TRIUMF Co-op Student Report.

[18] Westcode Semiconductors Ltd. High power press-pack igbts: Extended range. [Online]. Available: http://www.westcode.com/pigbtfly.pdf

[19] F. Wakeman, K. Billett, R. Irons, and M. Evans, "Electromechanical characteristics of a bondless pressure contact IGBT," *Applied Power Electronics Conference and Exposition, 1999. APEC'99. Fourteenth Annual*, vol. 1, 1999.

[20] L. Ducimetiere, U. Jansson, G. Schroder, E. Vossenberg, M. Barnes, G. Wait, S. Div, and G. CERN, "Design of the injection kicker magnet system for CERN's 14 TeVproton collider LHC," *Pulsed Power Conference, 1995. Digest of Technical Papers. Tenth IEEE International*, vol. 2, 1995.

[21] M. Barnes and G. Wait, "High current 66kv tests on high stability pfn discharge capacitors for cern lhc," *IEEE PULSED POWER CONFERENCE*, vol. 2, pp. 773–776, 1999.

[22] A. Loth and F. Lee, "Two-dimensional skin effect in power foils for high-frequencyapplications," *Magnetics, IEEE Transactions on*, vol. 31, no. 2, pp. 1003–1006, 1995.

[23] S. Kim and D. Neikirk, "Compact equivalent circuit model for the skin effect," *Microwave Symposium Digest, 1996., IEEE MTT-S International*, vol. 3, 1996.

[24] M. Kamon, M. Ttsuk, and J. White, "FASTHENRY: a multipole-accelerated 3-D inductance extraction program," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 42, no. 9, pp. 1750–1758, 1994.

[25] A. Ruehli, C. Paul, and J. Garrett, "Inductance calculations using partial inductances and macromodels," *Electromagnetic Compatibility, 1995. Symposium Record. 1995 IEEE International Symposium on*, pp. 23–28, 1995.

[26] Tektronix. Understanding oscilloscope bandwidth, rise time and signal fidelity. [Online]. Available: www.tek.com/Measurement/App_Notes/Technical_Briefs/bw_rt/55_18024_0.pdf

[27] D. Ward and J. Exon, "Using Rogowski coils for transient current measurements," *Engineering Science and Education Journal*, vol. 2, no. 3, pp. 105–113, 1993.

[28] J. Joyce, "Current Sharing and Redistribution in High Power IGBT Modules;" Ph.D. dissertation, PhD thesis, University of Cambridge, UK, May 2001.

[29] M. Evans, F. Wakeman, R. Irons, G. Lockwood, and K. Billett, "DESIGN CONCEPTS OF A BONDLESS PRESSURE CONTACT IGBT," *EPE'99, Lausanne*, September 1999.

[30] Westcode Semiconductors Ltd. T0900ta52e 5.2kv igbt datasheet. [Online]. Available: http://www.westcode.com/t0900ta.pdf

Appendix A

IGBT On-state Resistance

A survey of Westcode documentation showed that the average on-state resistance, as determined from V-I plots showing the slope resistance of similar IGBT modules in their linear regions, varied from 10mΩ to 15mΩ per die[29][30]. As the module TRIUMF is examining is a prototype device, and, as such, has no datasheet information available, measurements were made to determine the on-state resistance.

Using a single Westcode IGBT die as a switch, the gate was pulsed with a voltage of 15V. The IGBT was set-up to switch a bank of capacitors (10x470μF) into a short-circuit. Measurements of the voltage across the IGBT and the current were made as the charging voltage on the capacitors was increased.

A plot of these measurements is shown in Figure A.1. Also on this plot are low current measurements made using a DC power supply with a maximum output of 1.5A. These measurements are intended to confirm the results of the pulsed measurements.

The measurement of the current through the IGBT was achieved using a Pearson 410 CT, with a Tektronix 7401 oscilloscope used to measure the collector-emitter voltage and read the output from the CT. The current pulses, typically less than 3μs wide, were measured from the mean of the flat-top value. A photograph illustrating the set-up is in Figure A.2 featuring: the capacitor bank (left), clamped IGBT die (right), charging supply (top right) and CT probe (bottom middle).

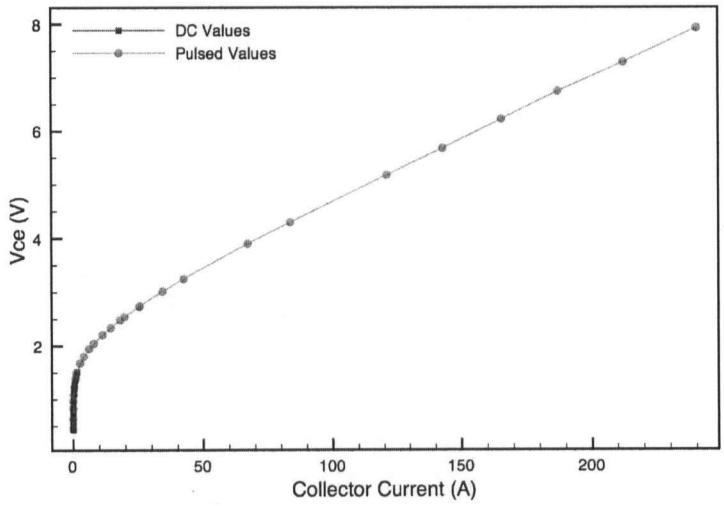

Figure A.1: V-I relationship for Westcode IGBT die to determine on-state resistance (V_{gate} = 15V)

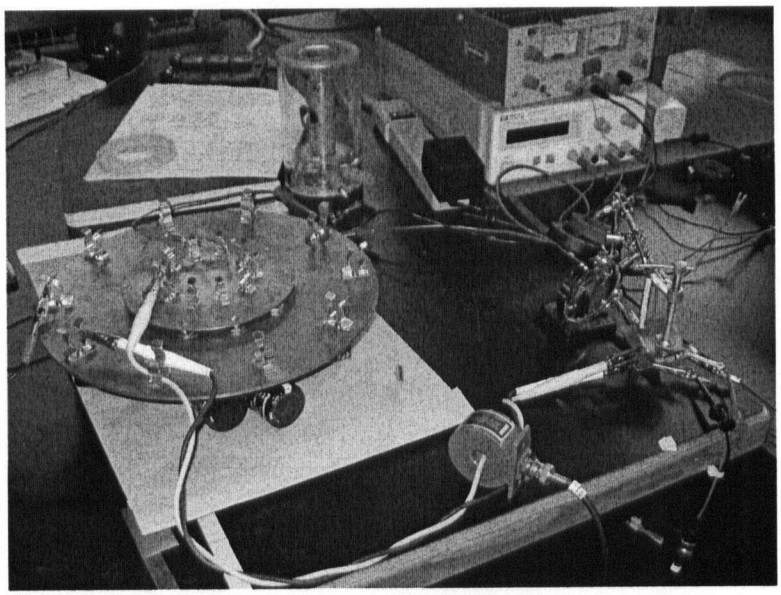

Figure A.2: Measurement set up for IGBT die on-state resistance measurements

The IGBT does not enter the linear region until the collector current is above ≈40A. The slope-resistance of the IGBT was calculated from the slope of the plotted data to be 23.3mΩ. This value is between 1.6x and 2x the values obtained from the manufacturer's curves for commercially available devices.

Given the relationship between the on-state resistance and the internal module current distribution that has been established in Section 2.3.4, it is quite possible that this is an intentional modification on the part of the manufacturer to improve the response of the module to pulsed switching.

Appendix B

Frequency Dependance of Parameters

It is well known that parasitic resistance and inductance values are frequency dependent[22].

As a rule, resistance values increase as the drive frequency is increased. This follows from the skin-effect (Equation 2.1), where, as the current fringes to the outside of the conductor, the cross-sectional area of the conductor available for conduction decreases, increasing the effective resistivity of the conductor material. To a lesser degree, inductance values are also effected by drive frequency. At higher frequencies, the magnetic field is excluded from the inside of the conductor by induced-eddy currents, resulting in a decrease in overall inductance.

To determine the effect that the frequency of extraction might have upon the accuracy of the simulations, Fasthenry was set to run a series of extractions for a purely copper model of the Westcode module from 1Hz to 1Mhz. Plots of the resistance, self-inductance, along with the mutual coupling terms relating to Conductor 1, are shown in Figure B.1.

The resistance of each conductor is highly frequency dependant, and rises about 1.5 orders of magnitude before plateauing near 1MHz. There is variation between the forward conductors ($\approx$2%) with the conductors closest to the inside of the module having a higher extracted resistance.

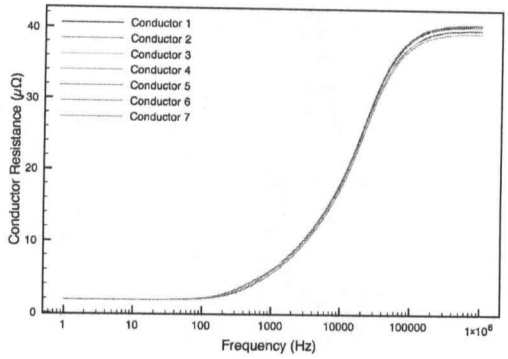

(a) Conductor resistance

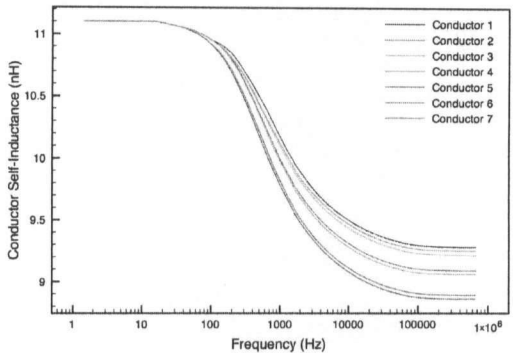

(b) Conductor self-inductance

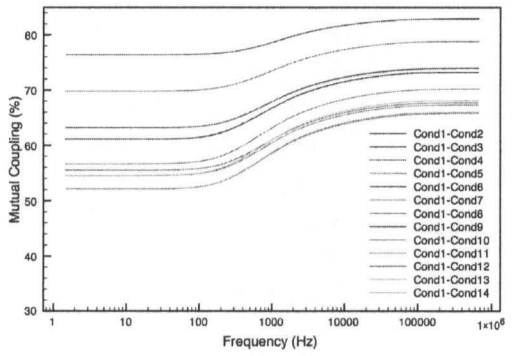

(c) Subset of K-value coupling terms (coupling to Conductor 1 shown)

Figure B.1: Extracted resistance, inductance and mutual coupling terms for a "copper only" model of the Westcode IGBT module as a function of frequency

The self-inductance values decrease, on average, approximately 20% before plateauing. There is also a significant spread between the self-inductances at the high-end of the frequency range (≈5%). It is interesting to note that this high-frequency spread in self-inducatance values acts to enforce a *more* uniform current distribution, as conductors towards the centre of the module have lower self-inductance terms than those at the outside, for higher frequencies. This is due to the magnetic field being excluded from neighbouring conductors, due to induced eddy-currents, and being forced out to the fringes of the module.

From the shape of these plots, it was decided to mix the extracted circuit parameters from high- and low-frequency simulations to produce a single set of circuit elements. Inductance values, as well as the associated coupling terms, would be taken from the high-frequency simulation, resistive values from the low-frequency simulation. The rationale was: at low frequencies, the resistances would dominate and determine the distribution of current, while at high-frequencies, the inductances would be dominant. At both extremes, the distribution would be 'correct'. It was hoped that, for intermediate frequencies, the distribution would be close enough to be accurate.

Figure B.2 shows the extraction frequencies considered, and their effect on the current distribution. For legibility, only Conductor 1 is shown. Opera 2D simulation results are included as a reference for what is to be considered an accurate estimation of the distribution.

From this analysis, it was decided to perform the PSpice simulations with the 1Hz R - 1MHz L parameters, as they provided the best combination of high- and low- frequency fit, as well as an acceptable match over the entire frequency range. If it is deemed necessary to have an improved fit, it is recommended that inductance ladder-networks be constructed.

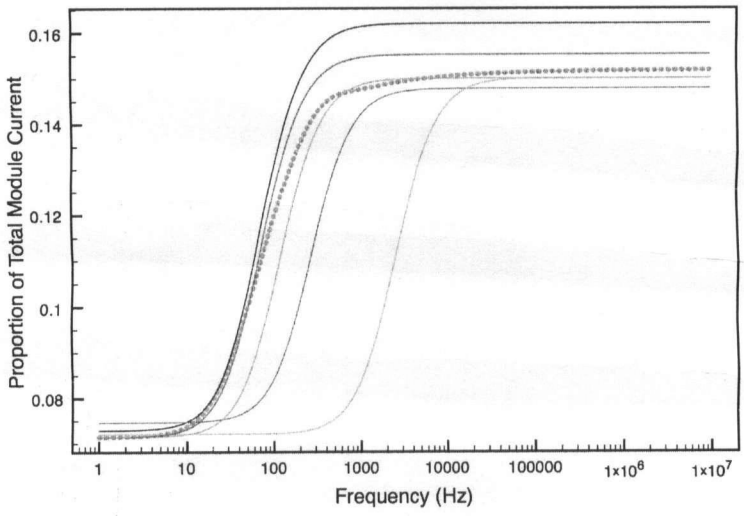

Figure B.2: Different extraction frequencies and their effect on module current distribution

Appendix C

Coupled Inductors From Impedance Matrixes

The following Matlab code performs the conversion of a single-frequency impedance matrix into a set of resistor and inductor values, as well as a set of coupling terms for the inductors.

```
function [M_R,M_L,K_val] = GetMutual(X,freeq)
% Calculates resistance, self inductance, mutual inductances, and coupling values.
% Takes an n X 2*n matrix with separated real and imaginary components in alternat
% columns.

[m,n] = size(X);
M_L = zeros(m);
M_R = zeros(m);
K_val = zeros(m);
%Omega
omega = 2*pi*freeq;

i =1;
j =2;
while i <= m
    while j <= n
        %Set the diagonal of M_L equal to the self-inductances.
        %"  "   "   "   "    M_R    "   "   " resistances.
        if j == 2*i
            M_L(i,i) = X(i,j)/omega;
```

```
            M_R(i,i) = X(i,j-1);
        else
            %set up all the mutual inductances
            M_L(i,j/2) = X(i,j)/omega;
        end
        j=j+2;
    end
    j=2;
    i = i+1;
end

i=1;
j=1;
for i = 1:m-1
    for j = i+1:m
        if sqrt(M_L(i,i)*M_L(j,j)) ~= 0
            K_val(i,j) = M_L(i,j) / sqrt(M_L(i,i)*M_L(j,j));
        end
    end
end

M_R;
M_L;
K_val;
```