


Current Sharing and Redistribution in High Power IGBT Modules



Dissertation submitted to the University of Cambridge
for the Degree of Doctor of Philosophy

by
John C. Joyce

Clare College

May 2001

Declaration

The work presented in this dissertation was carried out in Cambridge University Engineering Department, during the period from October 1996 to January 2001, under the supervision of Dr. Patrick Palmer.

The author wishes to declare that, except for commonly understood and accepted ideas or where reference is made to the work of others, the work in this dissertation is his own, and includes nothing which is the outcome of work done in collaboration. It has not been submitted in part, or in whole, to any other university for a degree, diploma or other qualification.

J. C. Joyce
Clare College
May 2001

Acknowledgements

My thanks go to my supervisor, Dr. Patrick Palmer, for his direction, advice and encouragement during the research which led to this thesis. Funding from EPSRC, GEC Plessey Semiconductors (latterly Mitel Semiconductor and then Dynex Semiconductor) and Cambridge University Engineering Department is gratefully acknowledged. Funding via the Office of Naval Research's "Virtual Test Bed" project during the writing up of this thesis is also much appreciated.

A number of colleagues have provided help and assistance in a number of ways. In particular I am grateful to Dr. Bernard Stark for his patience in winding magnetic field coils, reading my writings, and instruction in the use of Silvaco's software suite; to Dr Haile Rajamani for his help, support and encouragement; to Dr Ben Gordon for his seemingly endless knowledge of all matters electrical and electronic; to Mick Furber and John Grundy for their help with design and construction of apparatus, and generally making the whole undertaking more bearable; and to those other staff at the Engineering Department who helped in various ways.

Finally, my thanks to my family and friends for their support and understanding.

Version history

1.00 (13/05/2001) Submitted version

1.01 (27/08/2001) Corrected version

Abstract

The Insulated Gate Bipolar Transistor (IGBT) is now the favoured power semiconductor device for most medium power applications. One major difficulty with current IGBT manufacturing technologies is the inability to produce chips rated at more than 100 amps both reliably and cost-effectively. To achieve higher current ratings, a number of chips are connected in parallel to form a module; up to 32 chips operating in parallel have been reported.

The aim of this investigation is to examine what measures can be taken to ensure that current is evenly distributed amongst the chips in such a module, during both conduction and switching. An even inter-chip current distribution will make optimum use of the chips in a module, enhancing reliability. The current distribution within existing modules is investigated, and tests indicate that the turn-off transient can produce uneven current distributions and switching losses. Attention then turns to investigation of the matching of chips for selective assembly into power modules. It is found that matching on the characteristics currently measured can improve transient performance. It is also seen that gate resistors, and their tolerances and materials, are important aspects of module construction.

The focus then moves to susceptibility to differential oscillations, a problem encountered when paralleling MOSFETs. Analysis of a simple circuit model of the IGBT in its saturated region leads to a method for predicting the minimum gate resistor required to avoid sustained differential oscillations during turn-off. Finally, the device simulator ATLAS enables the investigation of the parallel operation of IGBTs under transient operation, and in particular the effects of structural differences between chips caused by processing variability. Certain aspects of device processing are found to be more critical than others.

Indexing Terms

The following may be useful for indexing this thesis.

Circuit stability

Circuit simulation

Current distribution

Electric current measurement

Insulated Gate Bipolar Transistors

Multichip modules

Power semiconductor devices

Semiconductor device models

Abbreviations and Symbols

C_{CE}	Collector-Emitter capacitance
C_{GC}	Gate-Collector capacitance; also known as Miller capacitance
C_{GE}	Gate-Emitter capacitance
C_{ISS}	Input capacitance
C_{OSS}	Output capacitance
C_{RSS}	Reverse transfer capacitance
DBC	Direct Bond Copper (a method of mounting IGBT chips)
DG-EST	Dual Gated Emitter Switched Thyristor
DMOS	Dual-diffusion Metal Oxide Semiconductor
di/dt	Rate of change of current
dV/dt	Rate of change of voltage
EMI	Electromagnetic Interference
EST	Emitter Switched Thyristor
GCT	Gate-Commutated Thyristor
GTO	Gate Turn-Off thyristor
I_C	Collector current
I_L	Load current
IGBT	Insulated Gate Bipolar Transistor
IGCT	Integrated Gate-Commutated Thyristor
IGT	Insulated Gate Transistor (original name for the IGBT)
MAGT	MOS Assisted Gate-Triggered Thyristor
MathCAD	A mathematical package
MBGT	Mos and Bipolar Gated Transistor
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
NPT	Non Punch Through (also known as 'non buffer layer')
PCB	Printed Circuit Board
PSPICE	A circuit simulation package

PT	Punch Through (also known as 'buffer layer')
SCR	Silicon Controlled Rectifier (original name for the thyristor)
V_{CE}	Collector-Emitter voltage
V_{GE}	Gate-Emitter voltage
$V_{GE(th)}$	Threshold voltage

Contents

Declaration	ii
Acknowledgements	iii
Abstract	iv
Indexing Terms	v
Abbreviations and Symbols	vi
List of Figures	xii
Chapter 1. Introduction	1
The Development of Power Electronics	
The Thyristor	
The Gate Turn Off Thyristor (GTO)	
The Insulated Gate Bipolar Transistor (IGBT)	
Other Recent Developments	
Why Develop Better Power Devices?	
Chapter 2. The IGBT	10
Structure	
Equivalent Circuit	
Fabrication	
Steady State Operation	
Transient Operation	
High Power Modules	
Plastic Modules	
Controlling the IGBT	
Development of Gate Drive Techniques	
Simplest Gate Drive	
Active Snubber	
Closed-Loop Collector Voltage Control	
Other IGBT Developments	

Summary	
Chapter 3. Current Measurement	35
Introduction	
Switching IGBTs	
Initial Experiments	
Measurement Techniques	
Existing Techniques	
Current Transformer	
Tektronix Current Measurement System	
LEM Converter	
Resistive Shunt	
Rogowski Coil	
Other Approaches	
Magnetic Field Coils	
Compensated Field Probes	
Integrator Design and Testing	
Calibration	
Measurement System Characterization	
Initial Measurements with the Field Probe	
Summary	
Chapter 4. IGBTs in Parallel	51
Introduction	
Discrete IGBTs in Parallel	
IGBT Chips in Parallel	
Gate Topology	
Parameter Measurements	
Matched Module	
Mismatched Module	

Gate Drive Techniques	
Baseline Modules	
Gate Resistors	
Summary	
Chapter 5. Stability Analysis	68
Characteristic Equation	
Determining Model Parameters	
Results	
Variations in Device Parameters	
Variations in Stray Impedances	
Alternative Module Topologies	
Stability Analysis with Circuit Simulation	
Oscillators	
Discussion	
Chapter 6. Device Simulation	90
Introduction	
Device Simulation Software	
Simulation Circuit Design	
Model Parameter Measurements	
Gate Resistors	
Gate Oxide Thickness	
P-Base Doping	
Gate Oxide Thickness and P-Base Doping Combined	
Gate Drive Techniques	
Intra-Chip Issues	
Summary	

Chapter 7. Discussion	114
Current Measurements	
Stability	
Transient Current Distribution	
Module Design Issues	
Future IGBTs	
Chapter 8. Concluding Remarks and Further Work	124
Conclusions	
Recommendations for Further Work	
Appendix A. References	127
Appendix B. Characteristic Equation Derivation	136
Appendix C. Publications	143

List of Figures

1. Diagrammatic representation of two cells from a typical vertical IGBT structure and single 1600V half-cell approximately to scale
2. Basic equivalent circuit diagram for the IGBT
3. Fabrication of a DMOS IGBT (part 1)
4. Fabrication of a DMOS IGBT (part 2)
5. The top of the finished IGBT half-cell
6. Simulated basic IGBT steady-state characteristics
7. Simplified schematic for a typical IGBT application
8. Typical measured IGBT turn-on waveforms
9. Typical measured IGBT turn-off waveforms
10. Typical IGBT and diode arrangement on DBC tile, and four tiles made up into a 'half bridge' IGBT module
11. Circuit diagram of the half bridge module
12. Simple IGBT gate drive circuit diagram and typical waveforms
13. Active snubber IGBT gate drive circuit diagram and waveforms
14. General principle of active collector voltage control and typical waveforms
15. Step-up (boost) converter arrangement and turn-off waveforms
16. Module for initial experiments, and typical waveforms
17. Pearson current monitor, LEM converter and IGBT module
18. Tektronix current probe and IGBT module for size comparison
19. Measurement of current via the B-field
20. Theory of the compensated field probe: two coils around a conductor
21. Practical realization of the compensated field probe
22. Field compensation
23. Integrator circuit diagrams
24. Bode plot of measurements of the integrator's response
25. Calibrating probes

-
26. Early measurement result for a two-chip IGBT: overall view and magnified
 27. Two chips in parallel in a module, forming a single switching element
 28. Individual chip gate voltage experiments, with the module as supplied, and with chip gates tied
 29. Measurements with minimum gate resistance, showing differential oscillations just beginning between the two chips.
 30. V_{CE} measurements
 31. Collector current (I_C) vs collector voltage (V_{CE}) with different gate voltages (V_{GE})
 32. Matched module. Good current sharing is evident.
 33. Matched module, heated to approx 100°C. Sharing is still good.
 34. Matched module, one chip only heated to approx 100°C, the other remaining at approx 30°C.
 35. Mismatched module, showing difference in gate behaviour and switching losses.
 36. Driven with a simple 'hard switched' gate drive, the IGBT pair exhibits some redistribution.
 37. The same module operated with an 'active snubber' type gate drive and a closed-loop collector voltage control gate drive. In both cases, the current imbalance lasts for much longer, and the currents do not converge during the current fall. Consequently the difference in switching losses is much greater.
 38. Initial measurements with a 'baseline' module with a normal hard switched gate drive
 39. Chip gates tied; the redistribution has gone
 40. The baseline module in figure 38, but with the gate resistors replaced with two 6.2Ω metal film resistors. Current sharing during turn-off is much improved.

-
41. Resistance of a gate resistor varying with temperature over a typical working range.
 42. A different baseline module, with gates as supplied and tied
 43. Differential oscillations seen whilst switching a module at a much lower current than its rating.
 44. The small-signal equivalent circuit of a single IGBT
 45. Two IGBTs in parallel, with parasitics and resistive load.
 46. Differential mode circuit of a single IGBT for analysis
 47. IGBT circuit with bias and AC source and small-signal equivalent circuit to determine C_{ISS}
 48. IGBT circuit with bias and AC source and small-signal equivalent circuit to determine C_{OSS}
 49. IGBT circuit with bias and AC source and small-signal equivalent circuit to determine C_{RSS}
 50. Device capacitances as a function of collector voltage, with $V_{GE}=0V$ and $V_{GE}=7V$
 51. Minimum value of R_G required at different collector voltages.
 52. Minimum value of R_G required for different values of C_{CE} with varying C_{GC} , calculated with Kassakian & Lau's simplified models, and compared with full model
 53. Minimum value of R_G required for different values of C_{GE} with varying g_m , calculated with Kassakian & Lau's simplified models, and compared with full model
 54. Minimum value of R_G required for different values of L_E with varying L_G , and different values of R_E with varying L_C .
 55. (a) Closed-loop oscillator; (b) Loop opened and AC signal source added
 56. Single IGBT equivalent circuit, as used in the MathCAD analysis, but split for analysis of transfer function
 57. Bode plot of the response of the circuit in figure 56 for different values of R_G

-
58. Nyquist diagram corresponding to figure 57, and close-up showing +1 point
 59. Nyquist diagram for $R_G=2.0, 2.2, 2.4$ and 2.6Ω showing +1 point
 60. Bode plot of the response of the circuit in figure 56, with $L_G=200\text{nH}$
 61. Minimum value of R_G required for different values of L_E , with varying L_G
 62. Small-signal model for two parallel IGBTs with uneven emitter inductance
 63. Frequency response of the circuit of figure 62 with the left-hand IGBT excited
 64. The emitter end of a typical IGBT model
 65. Schematic diagram of one IGBT chip and its associated parasitic elements
 66. The circuit used for simulations with the Silvaco software suite
 67. The simulated circuit, with measured and estimated parasitic components
 68. Simulation of two identical chips in parallel, but different gate resistors (6.2Ω and 6.5Ω)
 69. Measured turn-off of parallel chips with different gate resistors
 70. Switching waveforms for two chips in parallel with a large difference in gate resistors
 71. Electron currents in the two IGBTs during normal conduction
 72. Electron currents at $t=2.5\mu\text{s}$, during phase 2
 73. Electron currents at $t=2.7\mu\text{s}$, the end of phase 2
 74. Electron current densities at $t=2.9\mu\text{s}$, the end of phase 3
 75. Electron current densities at the end of phase 4: both MOS channels are virtually off
 76. Hole and electron currents plotted separately for the same turn-off as figures 63 through 67
 77. Effect of gate oxide thickness on threshold voltage $V_{GE(th)}$ and on-state voltage $V_{CE(sat)}$
 78. Switching waveforms for two chips in parallel with different oxide thickness

-
79. Switching waveforms with the oxide thickness increased to 4%
 80. Effect of p-base doping on threshold voltage ($V_{GE(th)}$) and on-state voltage ($V_{CE(sat)}$)
 81. Switching waveforms with different p-base doping
 82. Collector current and voltage waveforms and gate waveforms
 83. Collector current and voltage waveforms and gate waveforms with chip gates tied with a 1Ω resistor
 84. Simulated transfer characteristics (V_{GE} vs I_C at fixed $V_{CE}=100V$) of the two chips
 85. Switching waveforms for chips with different oxide thickness and p-base doping, and active snubber
 86. Switching waveforms for chips with different oxide thickness and p-base doping, gates linked with a 1Ω resistor, and an active snubber
 87. Single IGBT chip and bond wires, showing asymmetry in gate current path arrangements
 88. Equivalent circuit for single IGBT chip shown in figure 87
 89. Proposed IGBT turn-on rig
 90. Diagrammatic representation of a typical three-chip tile



Introduction

The Development of Power Electronics

In the earliest days of electricity, electrical appliances were limited to simple items such as lighting, heating and motors. Resistance banks, rheostats and contactors provided the required control. Motor-generator sets enabled power conversion after a fashion, although this came at a price in terms of cost, efficiency, noise and space.

The arrival of the thyristor in 1957, then called the Silicon Controlled Rectifier (SCR), made attractive a number of previously difficult options for power control, conversion and distribution. High voltage DC links employing thyristors to convert between DC and AC power could interconnect AC power grids, allowing power to be moved around as required between generators and consumers to match changes in demand. Thyristor choppers (a simple circuit for producing a variable DC voltage from a fixed one) replaced contactors and resistance banks in applications such as traction, yielding reliability improvements, at the expense of more complicated electronics.

More recently, the diversity of electrical power requirements has broadened enormously, with larger variable-speed motor drives employing induction machines rather than the DC motors they historically used. This trend is beginning to extend both downwards in the power range towards domestic applications, and upwards, with drives of 20MW being produced. Induction motors are inherently much more robust than DC machines; they require little maintenance, the number of mechanical parts is small, and there are no components which can be expected to wear out and require replacement in normal service.

Whilst DC machines have long been controllable over a wide speed and power range with the crude apparatus referred to earlier, and later with thyristors, variable speed induction motor drives require much more sophisticated electrical power control. For these, a variable voltage, variable frequency power supply is required, preferably sinusoidal, and it was the lack of an economical method for producing this which had prevented the wider application of variable speed induction machines and drives. The arrivals of the bipolar transistor and GTO were crucial in opening up the mass market for such machines, and the IGBT is further expanding the possibilities. To some extent, the development of power semiconductors and of motor drives have driven each other.

Power electronics has other important applications in energy generation and conversion: novel energy generation technologies, such as windmills and photovoltaics, require power electronics to control their energy conversion for these generators to operate at optimal efficiency. The generated power must also be turned into a form which can be distributed and used. Another application is the recovery of waste energy, for example by enabling electric trains supplied from AC supplies to recover kinetic energy during braking, returning it to the electricity supply. Both of these operations were difficult or impossible in the absence of power semiconductors.

A brief review of the development of the main power semiconductors is instructive.

The Thyristor

Thyristors were the first commercially available semiconductor power devices, appearing in 1957. A thyristor enables a large current to flow through its anode and cathode terminals by applying a short pulse of a much lower current to its gate terminal. The device remains switched on, even after the gate current is removed, until the external circuit causes the anode current to fall to zero. They are slow to turn off ("commutate"): the external circuit must remove the anode current and apply a reverse voltage to the anode for the turn-off time (typically 10 μ s) to force commutation. Consequently, forced turn-off arrangements can be complex, and so thyristors are now generally found in very high power circuits (such as AC to DC converters) where commutation occurs naturally.

Early uses in traction in the UK included an AC electric locomotive, which was fitted with a fixed ratio transformer, diode rectifier and a thyristor chopper, instead of the then standard tap-changing transformer. Subsequently these thyristor choppers were used in lower-power traction drives too, but the thyristor's period as the device of choice for such applications was curtailed by the emergence of variable speed induction motor drives, along with the GTOs to drive them.

Other thyristor applications include load resonant drives, where an inductive load has an associated capacitor bank which acts to commutate (turn off) the thyristors in the controlling circuit without any external control circuitry, and controlled rectifiers, which can produce a variable voltage DC supply from an AC one. In recent times, thyristors have become confined to the highest power applications, such as high voltage DC links between AC power grids, due to the arrival of the GTO and the IGBT. They remain the cheapest and biggest power device currently available, being available with ratings approaching 10,000A. They can also be run in series for high voltage ratings.

The major disadvantage of thyristors, apart from their inconvenient turn-off, is the need for snubber circuits to limit di/dt at turn on and applied dV/dt . Snubbers are bulky, lossy, and add expense to a converter circuit. Furthermore, the added complication inevitably means that making the whole system reliable is more difficult. Meanwhile, controlled-thyristor rectifiers cause very noisy voltage waveforms, which is unpopular for power distribution; the current waveform they produce is also fairly poor, requiring large filters.

The Gate Turn-Off Thyristor (GTO)

The main development from the thyristor was the GTO [1]. It is similar in many ways, but with the added feature that it can be turned off via its gate. As an extension of the thyristor, it was the first fully controllable high voltage solid state switch, i.e. a switch which could be turned both on and off at will. However, turn off is demanding on the gate control circuit, with up to a third of the device current needing to be taken out through the gate for around 10 μ s to turn it off. To make large devices, processing technologies had to be significantly improved to avoid local 'hot spots' during switching. Despite these

problems, it enabled variable speed induction motor drives to completely displace DC machines as the choice in traction applications, the lower maintenance and greater reliability of an induction machine being major advantages. The GTO appeared initially in lower-power drives on light vehicles, but soon spread to larger drives, notably in high-speed trains such as the TGV.

It has had less impact on industrial drives in general, as it is significantly more expensive than the thyristor. Consequently it offers little extra for applications such as synchronous motor drives. Like the thyristor before it, snubbers are required to limit di/dt and dV/dt . They are not normally run in parallel, as severe current redistribution between devices occurs at turn off. Careful design and manufacturing control are required to ensure that even single devices do not self-destruct: local 'hot spots' can be generated due to redistribution within a single device during switching. It took many years of development to achieve a reliable GTO scaled up to the same physical size as the thyristor, and even then, lower current densities than those achieved by the thyristor are seen. Series operation is possible, although the thyristor is generally favoured for such very high power applications due to its higher current ratings.

The Insulated Gate Bipolar Transistor (IGBT)

The IGBT is a device which combines features of the MOSFET and the bipolar transistor, and was invented in the early 1980s by two independent groups [2] [3]. Its gate drive requirements are inherited from the MOSFET, and are much more modest than those of the GTO: the gate appears to its driver as a capacitor, which simply has to be charged to around +15V to turn the device on, and discharged to turn it off. Also inherited from the MOSFET is an active region where the gate can control the device current; this has important implications for its behaviour as a switch, and this is utilised in some applications, including the running of multiple devices in series. From its bipolar features it gains lower conduction losses than those of a similarly rated MOSFET. External di/dt limiting is unnecessary, at least as far as the IGBT itself is concerned, and dV/dt and di/dt limiting can be achieved (if required) through the gate's controlling ability. Consequently, the snubbers mentioned

before can be dispensed with, improving reliability and reducing both physical size and cost of a system.

The switching delay for the IGBT is much shorter than for the GTO, with a switching transition typically taking a few hundred nanoseconds to complete, as compared to around 10 μ s in the GTO case. This allows higher switching frequencies to be used, with the possibility of moving this frequency out of the audible range. Paralleling is possible, and indeed is implicitly assumed to be viable, as large IGBT switches are constructed from a number of smaller chips, typically rated at 75 to 100A, operating in parallel. Inherent robustness to short circuit conditions, which can be experienced due to faults of one kind or another, is an added feature: the IGBT limits its own current to around ten times its normal working current, and can turn off safely from such a condition, provided that the condition is detected sufficiently quickly.

Disadvantages include higher losses in the device itself than in the thyristor or GTO, reducing the efficiency of the whole circuit and making heatsinking arrangements more demanding. They are also currently only commercially available up to ratings of around 4.5kV and 1500A. Long-term reliability is yet to be fully established, although the earlier style plastic packaged devices had a number of failure modes, including some which could result in debris being fired off in unpredictable directions [4]. Another problem with the plastic type package is that some failure modes involve the power device failing as an open circuit, rather than the usual short circuit failure experienced by GTOs and thyristors. This requires some careful design if unsafe failures are to be avoided, particularly in current-fed applications. The recent adaptation of the press-pack style modules used for GTOs and thyristors to suit the IGBT offers a solution to these problems.

The IGBT seems ideal for use in variable-speed induction motor drives, such as are required for traction applications, with its fast switching, simple gate drive, and lack of any requirement for separate snubbers.

Other Recent Developments

The bipolar transistor and the MOSFET both warrant a mention as power semiconductor switches, although neither could usefully achieve the device

ratings required for the applications discussed here. Scaling up the bipolar transistor for high current applications is possible, by operation in parallel [5] and by using the well-known Darlington configuration. Both of these approaches are inefficient and inconvenient to interface with logic level circuitry. By contrast, the MOSFET has always offered the convenience of simple drive circuit requirements and fair current capability, particularly when operated in parallel [6]. However, high voltage MOSFETs have remained elusive, with prohibitive conduction losses, although at the time of writing viable 600V devices are just becoming available [7].

Attempts have been made in more recent years to combine as many of the various desirable aspects of power devices as possible, with different devices yielding different compromises. Many attempts have been made to improve the IGBT's performance, and in particular to reduce its conduction losses. Some improvements have come from geometrical variations in the device, such as the trench gate IGBT [8] and the Injection Enhanced Gate Transistor (IEGT) [9]. An alternative approach has been to give a thyristor, with its inherently lower conduction losses, the ease of MOS gate control. One such device which became a commercial product is the MOS Controlled Thyristor (MCT) [10]. However, its success has been limited, with few devices being produced, due to severe current redistribution via a feedback mechanism [11]. Being a thyristor, it also lacks the active gate control possibilities of the IGBT.

More recently, much consideration has been given to dual-gate devices such as the MBGT [12]. This has similar switching properties to the IGBT, but an extra gate connection allows the on-state voltage (and hence power loss) to be reduced to thyristor-like levels, at the expense of a second gate drive. It is switched on in an IGBT fashion with its MOS gate, and a short time later its second (bipolar) gate is switched, causing it to behave like a thyristor, reducing its conduction losses. Turn off is the reverse of turn on, with the bipolar gate being switched off first, followed by the MOS gate. The bipolar gate characteristics are similar to those of a GTO gate, in that a substantial current must be drawn out during turn-off; however, its requirements are much less demanding than in the GTO case, as switching is much quicker. The MBGT has

seen a little use in specialised applications (although called the MAGT - MOS Assisted Gate-Triggered Thyristor) [13], but more general development and widespread application were hindered by the differing natures of the two gates and the development of more attractive alternatives.

A development of the GTO has been the implementation of unity gain turn-off, in which the entire device current is removed through the gate terminal. Termed the Gate Commutated Thyristor (GCT) [14], it is a very similar device to the GTO. The rapid removal of the entire device current eliminates the 'hot spots' which have proved troublesome in GTOs, and also the redistribution effects which have prevented the operation of devices in parallel. Such a turn-off is very demanding on the gate drive circuit, and the gate drive was soon integrated with the power device to produce the Integrated Gate-Commutated Thyristor or IGCT [15]. It provides an 'easy to use' solution, with the application designer no longer required to design a suitable gate drive.

The Emitter Switched Thyristor (EST) [16] family of devices was another attempt to give the thyristor the speed and convenience of MOS gate control: a low-voltage MOSFET in the 'cascode' connection with the thyristor, a topology well known from earlier semiconductor devices, could facilitate turn-off. Combining an EST and an IGBT in parallel in a single semiconductor device produced the Dual-Gate Emitter Switched Thyristor (DG-EST) [17]. The IGBT section is used to control switching, and the EST section provides good conduction performance. At turn-on, the IGBT is switched on first, and the EST section is turned on a short time later. Turn-off is achieved by turning off the EST, moving the device current to the IGBT section, which can then control turn-off. Conduction losses are reduced to those of the EST, and switching is controlled in the familiar IGBT fashion; the two functions can be performed by what are effectively separate devices optimised separately for their respective roles. Similar solutions have been realised in the past with discrete devices. However, the presence of two gates and associated control requirements is clearly unattractive, and methods have been proposed to overcome this [18].

Countless other devices have been proposed with a range of advantages; a summary of many of them is presented by Stark [18]. Inevitably each one is a

compromise, few see the light of day as viable power devices, and even fewer reach production. Meanwhile, development of the IGBT and its variants continues, with lower losses and higher power handling being the perennial aims. At the time of writing, only the IGCT and the variants of the IGBT are appearing in production quantities for use in the applications under discussion here.

Why Develop Better Power Devices?

The question could be asked as to why there is a need to continue developing power devices when most requirements can already be met. With the IGBT covering most medium power requirements, the GTO and the thyristor covering higher powers, and the MOSFET for low powers, it seems as though the problems are solved. However, GTOs and thyristors still demand bulky passive protection components, and IGBTs remain relatively inefficient due to their high power dissipation when conducting. The voltage rating of IGBTs remains a problem, a practical upper limit presently being around 3.3kV; whilst higher voltage parts are available, switching frequencies fall dramatically, to similar frequencies to the long-established GTO, and the trade-off moves in favour of the GTO with its much lower conduction losses. The quest for higher switching frequencies is driven by the desire to improve power quality and eliminate the audible noise so often associated with high power operation in current frequency ranges.

Meanwhile, adjustable speed drives remain desirable in many areas of industry where they are not yet viable, for assorted economic reasons; productivity gains may be possible, or output may be more easily matched to a varying demand. For example, a pump required to produce a varying flow of some liquid could be operated at a varying speed to suit demand, rather than throttling the pump or recirculating some of the pumped liquid. Substantial energy savings have been achieved by using a variable speed drive in such applications.

The spectrum of applications continues to broaden into both higher and lower powers. Development of power semiconductors appears set to continue in the quest for ever greater efficiency, with better solutions to the perennial trade-offs

between conduction losses, switching losses, switching frequency, voltage rating and current rating. This study, however, concentrates on the well-established IGBT, and in particular the operation of IGBT chips in parallel within a module to achieve a single high current switch. Obviously such an approach must rely on each of the IGBT chips carrying an equal share of the total current for optimum performance (i.e. maximum current rating); clearly there will be implications for reliability if sharing is not uniform and the module is not derated suitably.

2

The IGBT

A basic outline of the IGBT's structure, fabrication and electrical properties are given here to set the scene for the investigation which follows.

Structure




Figure 1. Diagrammatic representation of two cells from a typical vertical IGBT structure (left) and single 1600V half-cell approximately to scale (right)

Figure 1 shows a vertical section through one and a half cells from a representative conventional vertical DMOS (Double diffusion MOS) NPT (Non Punch Through) power IGBT structure. A typical power IGBT consists either of many thousands of these cells laid out in a regular matrix, or of stripes. Cells are formed by rotating a unit half-cell around its emitter edge; stripes by extruding the cross-section shown. The emitter metal on the top of the chip covers the whole device, whilst the gate connection runs underneath this to each cell, and is formed from polysilicon. There are three external contacts, namely collector, emitter, and gate. The right-hand diagram gives a guide to the scale of the device: a 1600V device of this type would be around $300\mu\text{m}$ thick, the top $5\mu\text{m}$ of which would be occupied by the top structure, and the unit half-cell being around $15\mu\text{m}$ wide.

The device described here is an NPT (Non Punch Through) type: it is designed so that the n^- base supports the collector voltage by accommodating the entire depletion region when the device is switched off. A separate family of IGBTs, the 'Punch Through' (PT) design, has a much thinner n^- base, designed so that the depletion region 'punches through' (extends right through it) when blocking high voltages. An additional thin n^+ layer on the collector side of the p^- base, known as the 'buffer layer', prevents the depletion region from reaching the collector p doping. The thinner n^- base gives a much lower on-state voltage and hence lower conduction losses, but other characteristics of the PT device are less desirable, notably slower switching and higher tail current.

Equivalent Circuit

Figure 2 (below) shows a simple equivalent circuit representation of the IGBT. It can be thought of as a $p-n-p$ bipolar transistor and an n -channel MOSFET in a Darlington configuration, the two devices sharing two of the four semiconductor regions. The relevant regions are labelled on figure 1. The MOSFET also has a parasitic $n-p-n$ bipolar transistor, although in modern designs this is inactive, due to its base-emitter junction being shorted by the IGBT's metal emitter contact. Early IGBT designs suffered 'latch-up' of the thyristor created from the two bipolar transistors, but design improvements have eliminated this.

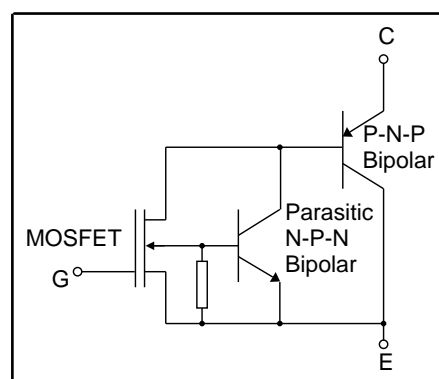


Figure 2. Basic equivalent circuit diagram for the IGBT

Some confusion can be caused by the naming of the IGBT terminals and the internal bipolar transistor's regions. In some documents the external terminals are referred to as 'drain' and 'source' due to the similarity of the vertical MOSFET, and occasionally even 'anode' and 'cathode' as traditionally applied

to the other major power devices, but the manufacturers' convention of 'collector' and 'emitter' is applied in this thesis.

Fabrication

The process of fabricating IGBT chips is a complicated one with many subtleties and possible variations on a basic theme. Here the fabrication of an IGBT of the type pictured in figure 1, with a trench emitter, is described; the 1800V device used for much of the simulation and practical work in this investigation is such a device. In essence the fabrication is identical to that of a vertical high-voltage DMOS power MOSFET, apart from the additional collector p doping.

The entire process is termed a 'self aligning' one, as only one mask is required to define the pattern of cells in the active IGBT areas of the wafer. Once the mask has been applied, all the subsequent processing stages should be automatically aligned with each other, as will become clear. The importance of this is that the channel length and doping - the main defining characteristics of any MOS device - are independent of the masking, and hence slight variations in cell sizes or shapes due to limitations of the lithography process can be tolerated. The situation is somewhat complicated by the requirement to produce edge termination around each chip to enable it to block high voltages, but this part of the device and its associated processing are not relevant to device conduction or transient operation, the two areas of interest here. A brief description of the key stages in the fabrication process is given below, along with sketch diagrams of the top few μm of what will become a single half-cell after processing. More specific details of the processes involved and IGBT processing in general can be found in [19], and the materials themselves and their physics in [20].

The bare silicon wafer has a faint n type doping (figure 3a) which will form the wide drift region (n-base) which occupies most of a high voltage device. A blanket (i.e. covering the whole wafer) n type doping is applied to what will become the emitter side of the wafer. A thin film of oxide is then grown, also covering the entire wafer, later to form the MOSFET section's thin insulating gate oxide. This oxide layer is covered with polysilicon (figure 3b), which is

then highly p doped, and this will form the gate electrode and connection. The high level of doping reduces the gate resistance, which is desirable to minimize delays to the gate signal during switching; there is usually only a single gate connection to a chip, and with the capacitive nature of the IGBT gate, any resistance in the gate connection between cells will slow down signal propagation and lead to non-uniform switching.

A photo-resistant coating is applied, followed by exposure to light using a suitable mask, defining the cell pattern. The photo-resist is then developed and the surface of the wafer etched to remove the polysilicon and oxide in the required pattern of cells, half of one being seen in figure 3c.




Figure 3. Fabrication of a DMOS IGBT (part 1)

The next stage is the formation of the MOSFET channel. A thick layer of oxide is deposited on the whole wafer, and then etched away anisotropically (i.e. in a particular direction) to leave an 'oxide spacer' defining a 'doping window' (figure 4a); this will control the application of the channel's p doping (the p-base). After application of the p-base doping (figure 4b), the spacer is removed by etching, and the n+ doping of the MOSFET source can now be introduced. A further etching step removes an amount of the silicon, including the unwanted part of the n+ well, making way for the emitter metal contact (figure 4c).

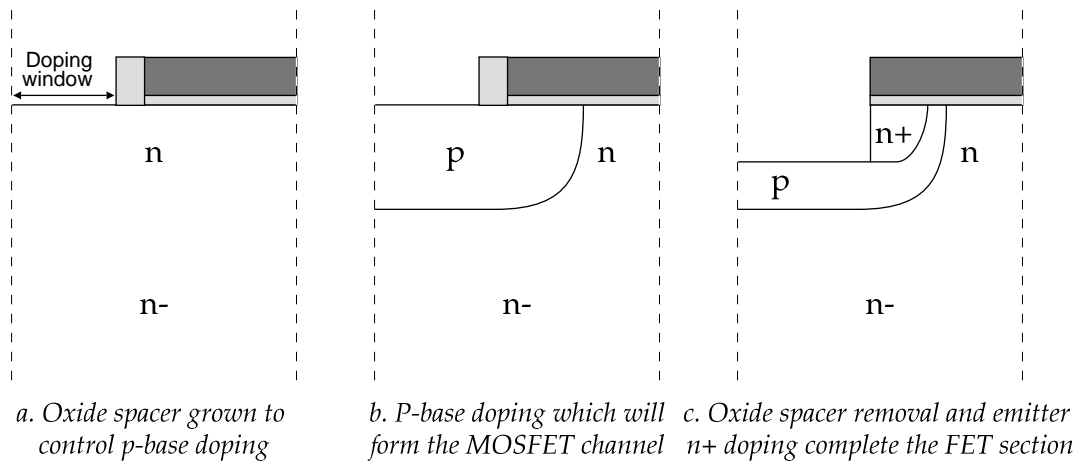


Figure 4. Fabrication of a DMOS IGBT (part 2)

The final processing step on the top of the device is to apply the emitter metallisation, but first an additional oxidation is required to prevent this from simply shorting to the gate connection. Another oxidation and selective etch process forms this insulation, and then the emitter metallisation can be applied to the whole device, as shown in figure 5. The collector doping and metallisation, on the opposite side of the wafer, are the final stages in the process to form the IGBT's active area.




Figure 5. The top of the finished IGBT half-cell

A number of details are omitted in this description; for example, the guard rings around the edges of the chips must be produced, and the gate polysilicon must be brought to the surface in order for a contact to be formed. The wafer must be cut up into individual chips ready for testing and packaging. These processes however are of little consequence in understanding the relevant steps in the formation of the IGBT itself, and are thus omitted.

Steady State Operation

With the collector voltage negative with respect to the emitter, the bottom p-n junction J1 (the bipolar section's base-emitter junction; see figure 1) is reverse

biased, and no current can flow. A depletion region forms, primarily in the low-doped n^- base. This mode of operation (i.e. reverse blocking) is not normally encountered, as the voltage which can be supported by the thin p-emitter used by most modern devices is quite low. In virtually all power IGBTs, the presence of an antiparallel freewheel diode means that a reverse voltage will only appear across the IGBT whilst that diode turns on. A reverse-blocking IGBT is possible, and indeed necessary for certain applications such as the matrix converter. However, to achieve this, the simple blanket doping and metallisation of the collector side would need to be replaced with processing to produce edge termination; such double-sided processing is complicated and expensive.

With a positive collector voltage, but no gate voltage, the bipolar section's base-collector junction J2 is reverse biased instead, and again no current can flow; once again a depletion region forms, mainly in the n^- base. Forward conduction begins when the collector is positive with respect to the emitter, and the gate voltage is raised high enough to invert a layer of the p-base below it, forming a conducting channel. The threshold voltage $V_{GE(th)}$ indicates the minimum gate voltage for this to occur. Electrons can now flow through the MOSFET section and into the bulk of the device, eventually leaving via the collector; meanwhile, the bottom junction (J1) is forward biased and the collector current injects holes into the bulk. The hole and electron currents settle to a ratio dictated mainly by the ratio of their mobilities, given the requirement for charge neutrality and the dominance of the drift current.


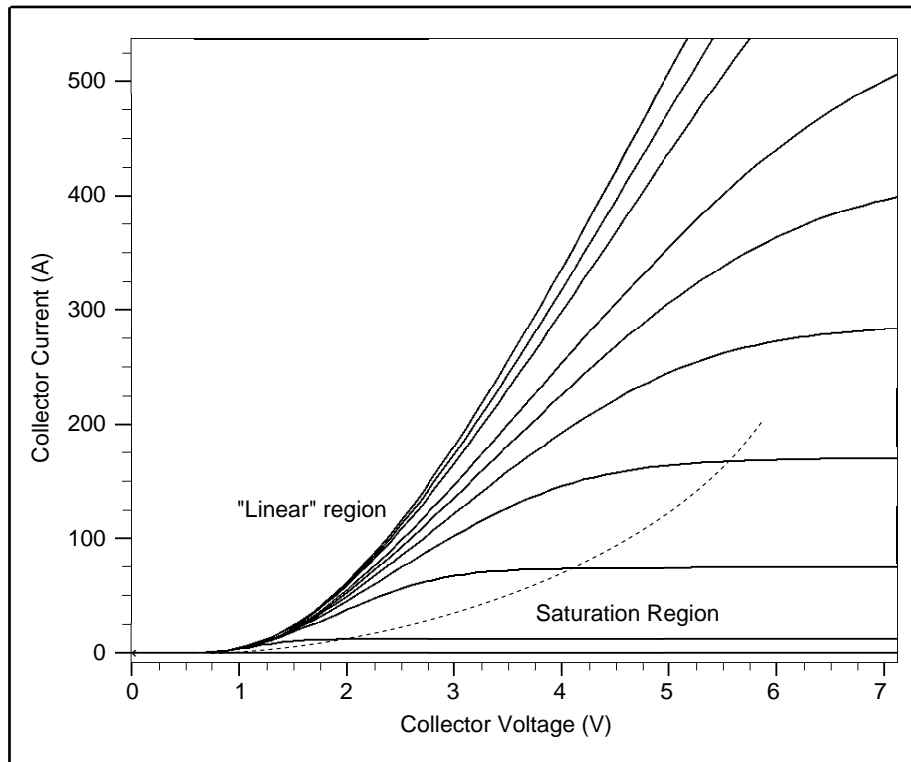


Figure 6. Simulated basic IGBT steady-state characteristics: output characteristic (V_{CE} vs I_C for different V_{GE}) (top) and transfer characteristic (V_{GE} vs I_C at high V_{CE}) (below)

The steady-state characteristics are shown above; they bear a fair resemblance to those of a power MOSFET, with the two different regions of operation shown. The 'linear' region is the normal region for operating the IGBT, where

the gate bias is high and the collector voltage is low. In this region, the conducting channel in the MOSFET section is sufficiently deep that it does not constrict the flow of electron current, and so changing the gate voltage has little effect. A slight reduction in collector voltage may be achieved by increasing the gate voltage, as the resulting deeper channel has less resistance; however in this region of operation the device voltage is dominated by other effects, such as the wide n-base resistance and the bottom p-n junction's (J1's) forward voltage drop, so the difference is small. A disadvantage of increasing the gate voltage is that the short circuit current is much higher.

The term 'linear region', inherited from the MOSFET, is misleading here, as in the IGBT case the collector current (I_C) is not proportional to the collector voltage (V_{CE}) in this region, as can be seen from the characteristics above. Whilst the bulk of the device appears resistive, there is the extra forward biased p-n junction (J1) in the current path, with its non-linear current-voltage characteristic. This forward-biased junction will also inject more carriers as the collector current increases, reducing the resistance of the n- base.

If the magnitude of the gate bias (and hence the depth of the channel) is small relative to the device current, the IGBT will be in its saturation region, with the gate bias controlling the device current by restricting the flow of electron current into the n-base. Whilst the IGBT is rarely operated in its saturation region for any length of time, as it causes high instantaneous power dissipation in the device, the device passes through the saturation region during switching. Consequently the characteristics of the device in this region will be of particular interest when transient operation of parallel devices is considered.

The foregoing explanation treats the IGBT as having two distinct regions of operation, namely linear and saturation, with a definite transition. The characteristic shown in figure 6 shows that this transition is in fact gradual, and inspection of the general characteristic equations (given in e.g. [21]) shows these two regions are merely two extreme situations of a single general case.

The saturation region provides the IGBT with its ability to automatically limit the current flowing in the event of a fault, a very useful facility not available in the thyristor family of devices. For a given gate voltage, the collector current is

limited to a maximum value largely independent of collector voltage, as can be seen from the transfer characteristics above.

Transient Operation

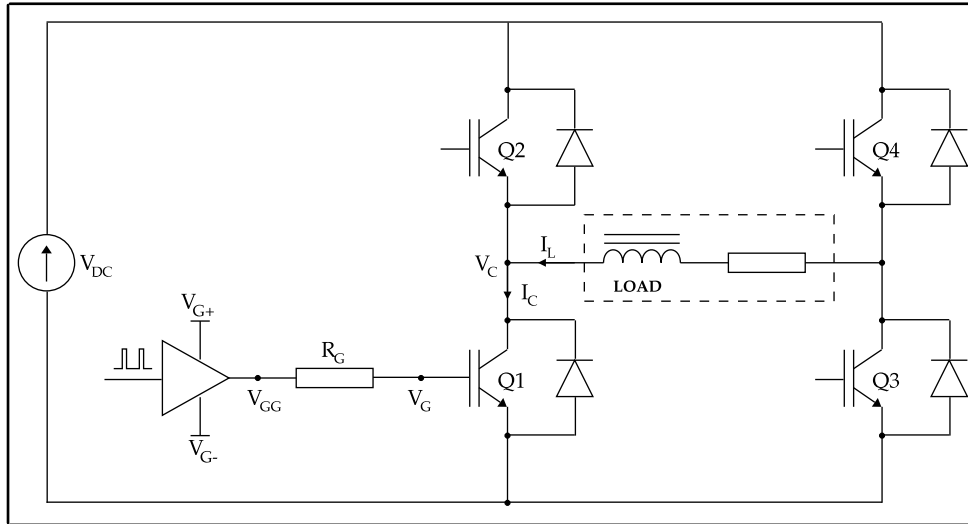


Figure 7. Simplified schematic for a typical IGBT application

The circuit diagram above represents a typical application for the power IGBT: a bridge converter. Here a full-bridge with four power switches is shown; a three-phase bridge is perhaps more common, with six switches, but the principle is the same. The four switches in the bridge converter shown experience similar conditions to each other, so studying the action of one is sufficient to characterise device behaviour. Here the action of Q1 switching on and off, with the load current I_L moving between Q2's antiparallel diode and Q1, is described. A simple gate drive arrangement, with a resistance of R_G , is shown for Q1; obviously Q2 through Q4 will have similar arrangements.

Turn-on is initiated by the gate drive circuit raising its output (V_{GG}) to V_{G+} . A typical set of measured turn-on waveforms is shown in figure 8, and these bear some similarities with those of a MOSFET. The turn-on sequence can conveniently be divided into a number of phases.

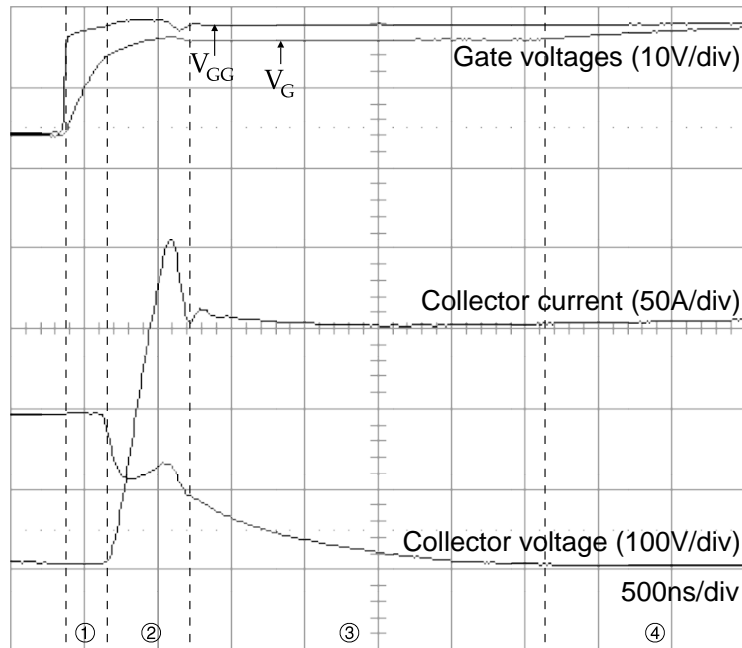


Figure 8. Typical measured IGBT turn-on waveforms

During phase 1, the IGBT's gate voltage (V_{GE}) is below the minimum threshold for conduction ($V_{GE(th)}$), and the gate current is charging the IGBT's (fixed) input capacitance, so the gate voltage rises accordingly. Once the gate voltage reaches the threshold voltage, phase 2 begins and the device starts to turn on, entering its saturation region. The collector current (I_C) begins to rise, the gate voltage's waveform mirroring that of the collector current. The voltage at Q1's collector (V_C) should be held at a constant value of approximately the DC link voltage (V_{DC}) by Q2's antiparallel diode during this period. However, stray inductance in the IGBT's power circuit means that the measured value here is smaller, as a voltage is induced across this inductance by the fast current rise.

The collector current continues to rise, peaking at a value greater than the load current (I_L) due to the diode's reverse recovery charge. The magnitude and shape of the current overshoot are determined by the diode's characteristics and the stray inductances. Although carriers are being injected into the device, the depletion region remains, as the collector voltage is still high; the injected carriers accumulate in the undepleted regions. Any carriers reaching the depletion edge cross the depletion region quickly due to the high electric field associated with the charge in that region.

During phase 3, once the diode has turned off, the collector voltage can fall and the depletion region can shrink. The gate voltage remains approximately constant during this period, and the collector current has dropped back to equal the load current. The (approximately constant) gate current discharges the (increasing) Miller capacitance C_{GC} . This is strongly voltage-dependent, as will be seen later, increasing by orders of magnitude as the collector voltage falls; this also slows down the collector voltage fall.

Finally, in phase 4, the turn-on is complete; the IGBT is in its linear region, the transition being marked by the gate voltage rising above its plateau voltage towards V_{G+} . The collector current is now almost independent of the gate voltage, and all that remains is for its carrier distribution to stabilise. A finite time, of the order of the carrier lifetime (typically tens of μs) is needed for the wide n^- base region to become full of carriers, or 'conductivity modulated'. The collector voltage falls a little more during this period, as the carrier concentration increases and the on-state resistance decreases, although this is not readily discernible from the waveforms here.

Turn-off is accomplished by the gate drive reducing its output V_{GG} to V_{G-} , which is typically either 0V or a few volts below. Once again the switching transition can conveniently be divided into four phases, as seen below in figure 9.




Figure 9. Typical measured IGBT turn-off waveforms

During phase 1, the gate current discharges the gate capacitance, and the gate voltage drops exponentially to the point at which the device enters its saturation region. In phase 2, the gate voltage remains approximately constant at a level dependent on the collector current, namely the 'plateau voltage'. The gate current charges the Miller capacitance C_{GC} , and the collector voltage consequently rises.

Within the device, the MOS channel is shallower due to the lower gate voltage, restricting the supply of electrons to the bulk; meanwhile, in the inductive switching application described here, the external circuit holds the device current constant until its voltage rises. Electrons continue to leave via the collector at the original rate, but they are only supplied at a reduced rate through the constricted MOS channel. With the reduction in the electron supply, the emitter current is now largely holes. As the collector current remains constant, holes are only injected at the original rate. Consequently, the stored charge is removed, starting from the end of the device near the MOS channel. A depletion region forms, consisting of a part of the n-bulk (positively charged) and part of the p-base (negatively charged), and the growing depletion region causes the device voltage to rise. Due to the charge profile in the device during its on-state (seen in later chapters), the initial progress of the depletion region is very slow, but accelerates as it moves through the device. With a large amount of charge stored under the gate during the on-state, the voltage rise is very slow initially, but accelerates through turn-off as the Miller capacitance and stored charge fall.

It will be seen that phase 2 is controlled by two factors; the charge profile within the IGBT during its on-state, and the extent to which the gate drive causes the MOS channel to be restricted, or even completely removed if the gate is driven sufficiently hard. This region also endows the IGBT with its ability to achieve controlled voltage rises during turn-off; the degree of constriction of the MOS channel dictates the rate at which the device is depleted of charge, and hence the rate of growth of the depletion region and the rate at which the device voltage rises.

Once the IGBT's collector voltage has reached the DC link voltage, the freewheel diode can begin to conduct and phase 3 begins. The IGBT current falls at a rate determined by the characteristics of the device itself and the freewheel diode, and any stray inductance. The load current (I_L) moves to the freewheel diode, and the IGBT gate voltage (V_{GE}) decreases. The collector current drops rapidly to a fraction of its on-state value. Finally, in phase 4, the gate voltage decays exponentially to its steady-state value, and the remaining current (the 'tail current'), now all holes, gradually dies away as the remaining carriers left in the device decay.

High Power Modules

The IGBT chip described so far is of little use as it stands; it has to be packaged into a form in which it can be used in a circuit. The package has to provide for both electrical connections capable of carrying the required current, and for adequate heat transfer as well. Initial developments of IGBTs were at ratings of a few tens of amps and hundreds of volts, for which the existing MOSFET and bipolar transistor packages were adequate. As current ratings increased, it was necessary to develop larger packages capable of containing larger and/or multiple devices. Discrete diode chips must be accommodated too for many applications, such as the bridge converter shown in figure 7. Whilst both of these requirements could be met, at least for the smaller early devices, by existing power packages such as the plastic TO247 and the metal can TO3, increasing voltage ratings required greater separation between terminals, and increasing current ratings demanded more, and larger, chips. The inductance of these packages, and in particular the interconnections between them, becomes awkward when they are operated in parallel for higher current ratings.

Two main alternatives were available at the time: either the 'press pack' type of construction which had been used for many years to package the highest power diodes, thyristors and GTOs could be used, or the existing plastic modules could be developed and enlarged to meet the current and voltage rating requirements.

The press pack is attractive in terms of presenting a very low stray inductance and resistance in its internal power circuit, along with heat removal from both

sides of the chips; current ratings of such packages run to thousands of amps and volts. Their application to IGBTs would have been non-trivial for a number of reasons, some of which could have been overcome fairly easily, and some less so. For example, a mounting assembly would be needed to house both the multiple devices required for high current operation, and the discrete gate resistors for each chip, a problem not previously encountered with the existing thyristor devices, where a single large device had been used and no gate resistors were involved.

More problematic was that careful control of device thickness was needed if even pressure on the chips was to be achieved and mechanical stress avoided. Despite these problems, chips designed for pressure mount assembly began to appear in 1995 [22], and press pack style modules using them soon followed. With a suitably designed chip, the technology of incorporating several chips into a press pack was not particularly difficult so much as inconvenient. Aside from the more obvious problems, the traditional round package was not well suited to square chips, and some square packages have been produced.

However, it is the plastic modules which this research focuses on, as these remain the dominant packaging style for the IGBT; the press pack types are just becoming more widely available at the time of writing.

Plastic Modules

One or more IGBT chips, along with antiparallel diode(s), are soldered to a copper layer which is bonded to an insulating tile, usually fabricated from aluminium oxide or aluminium nitride. Copper tracks on the tile provide for the electrical connections to the chip. This method, referred to as 'Direct Bond Copper' (DBC) is illustrated in figure 10. The IGBT collector and diode cathode connections are made through the solder bond, with aluminium wires being used to bond IGBT emitter and diode anode to an emitter track. Another bond wire provides the gate connection, via a small resistor, to a further track on the alumina tile. Separate gate current return ('Kelvin') and collector voltage sensing connections may be provided. Several of these tiles are then connected together with soldered connections to copper bus bars, mounted on a copper base plate, and packaged up into a module with a plastic lid. Thermally conducting silicon gel fills up the space inside the module.




Figure 10. Typical IGBT and diode arrangement on DBC tile (above) and four tiles made up into a 'half bridge' IGBT module (below)

The chips can be arranged in a number of configurations, typically as a single switching element, a single inverter leg (i.e. a pair of switches), or a complete three-phase inverter power stage (i.e. six switches). The photograph in figure 10 and the circuit diagram in figure 11 show four IGBT assemblies made up into a single inverter leg (half bridge) module, with two IGBTs in parallel in each switching element.




Figure 11. Circuit diagram of the half bridge module

These plastic modules have a number of advantages. They can be assembled into low inductance power circuits, although attention has to be paid to reducing inductance within the module itself. It also lends itself well to the production of 'smart' modules, i.e. devices containing more than just the basic

switching element. Such smart modules offer varying levels of sophistication, from rudimentary fault protection and feedback, to a complete converter power stage in a single package [23] requiring a minimum of external components to produce a working design. More comprehensive solutions yet are promised, for example an entire motor drive system (rectifier, brake chopper, inverter, and protection circuitry) in one package. Meanwhile, protection functions are also being integrated onto the IGBT chips themselves [24], providing comprehensive protection for the chip itself against most situations liable to cause device destruction. Such devices and modules are understandably popular with applications designers.

Early plastic packaged modules suffered a number of reliability issues, in particular due to the number and quality of material interfaces involved. Aluminium wires are bonded to silicon chips, soldered to copper tracks on an alumina tile; these are bonded onto a copper baseplate, either with adhesive or an extra copper track and solder joint, and this is then bolted with suitable thermally conductive paste onto a heatsink. The solder joints were difficult to achieve without voids and consequent local overheating [25], and they suffered from stress cracks. Bond wires were prone to lift off the top of the IGBT chip over a period of time [26], particularly if deep thermal cycling was experienced.

A number of other problems, such as degradation of the aluminium metallisation on the surface of the emitter pads and migration of the thermally conductive paste between module base and heatsink, were seen. The number of different materials involved, with their different coefficients of thermal expansion, inevitably produced these problems with fatigue, although efforts to choose or develop materials with closer temperature coefficients have yielded improvements. Some failure modes could lead to module explosions [4] with substantial quantities of debris ejected. It was several years after the introduction of these plastic modules before the reliability issues were properly resolved, and research continues. The results are now much more satisfactory in terms of long term reliability [27].

Controlling The IGBT

Unlike many of its predecessors, notably the thyristor and the GTO, the IGBT can be closely controlled during switching. As with a MOSFET, its gate terminal controls its behaviour in a relatively well defined fashion; so, for example, the rate of change of collector voltage or current can be monitored and limited via a feedback loop to the gate terminal. Such control can be useful to protect the IGBT itself or other devices in the circuit.

The gate terminal appears capacitive from the point of view of the gate drive circuit: an amount of charge has to be supplied to turn the device on, and removed to turn it off. Device manufacturers' datasheets usually quote input capacitance (C_{ISS}) over a range of collector voltages, and this can be expected to be of the order of 20nF for a 150A IGBT. However, these values are usually quoted at $V_{GE}=0$ (and hence $I_C=0$), and the relevance of these measurements to transient operation is not straightforward, as seen later.

Development of Gate Drive Techniques

In the earliest IGBT applications, gate drive techniques were crude, being developed from those for the power MOSFET. A simple push-pull output stage and a current limiting resistor sufficed to give basic 'hard switched' control. In such circumstances, the rate of change of the collector voltage during switching transitions is governed by the charge and discharge of the IGBT's feedback (Miller) capacitance (C_{GC}) by the gate current. This capacitance is heavily dependent on the collector voltage and current, falling from many nanofarads at low voltage and high current to a few tens of picofarads at high voltage and zero current. Consequently, in the case of hard turn-off, the rate of collector voltage rise is initially low as the large capacitance is charged; as the collector voltage begins to rise, the capacitance falls as the depletion layer grows, and so the rate of rise of device voltage increases.

Snubbers have long been needed in power electronics for a variety of reasons, including protection of the switching device against excessive rates of change of voltage or current. These can cause a power device to be switched spuriously (i.e. without any signal to its gate), possibly causing destruction of the device;

the thyristor and GTO in particular cannot withstand a high rate of anode voltage rise without false triggering. Whilst the IGBT did not demonstrate problems such as these, rapid transients also produce large amounts of electromagnetic interference (EMI) which can affect nearby circuits, particularly low-power ones such as microprocessor systems, where it can produce spurious operation. Given that modern inverters employ microprocessor-based control systems, high levels of radiated interference are clearly undesirable. Regulatory bodies are also increasingly keen to see the levels of electromagnetic noise reduced.

These snubber circuits to control switching transients had traditionally been implemented using large passive components, with inductors being used to control current transients and capacitors to control voltage transients. Such circuits demand bulky, expensive extra components which also incur a penalty in reliability terms; furthermore, in most snubber topologies, a significant amount of wasted power has to be dissipated in either the passive components or the semiconductor device itself. Such approaches were initially adopted and developed for the IGBT [28].

The gate's ability to control switching soon became apparent, and there was a proliferation of schemes making use of this ability. These included crude control of the collector voltage rise being achieved by suitable choice of resistor in the simplest gate drive [29], or even multiple resistors selected by additional logic dependent on the collector voltage or current [30] or gate voltage derivative [31]. The degree of control achieved is not large, and better solutions were sought.

Before long, closed-loop collector voltage feedback controllers became popular, with a number of possible schemes proposed with varying degrees of sophistication. Some offered open-loop [32] or closed-loop [33] control of collector current as well. Such closed-loop collector voltage control schemes also made operation of devices in series feasible [34] without the large, expensive and lossy passive components used in previous series operation schemes to ensure voltage sharing during transient and steady-state operation. A large number of combinations of these basic ideas have been suggested, resulting in rather complex gate drive solutions [35].

The foregoing developments can yield a comprehensive control solution, and yet one which is simple, compact, and inexpensive. Few (if any) extra components are required in the power circuit, and the control circuit can even be integrated into the IGBT module itself [23]. The developments in gate side control of the IGBT can be contrasted with the requirements for the GTO, whose gate drive techniques have also continued to be developed. Achieving comparable levels of control in the GTO is orders of magnitude more complicated as seen in (e.g.) [36]. The IGCT offers an alternative solution: its integrated gate drive and unity gain turn-off relieve the application designer of the need to design and develop the protection schemes and control techniques. However, the application designer is also relieved of most opportunities to tailor operation for specific circumstances should it be required.

Simplest Gate Drive

The simple 'hard switch' method of controlling an IGBT has continued to be developed, and is nowadays conveniently available as a single-chip solution. For the purposes of the investigation the example chosen is the HCPL3120. This combines an optoisolator and push-pull output stage in a single 8-pin IC package, and is capable of delivering a peak current of 2 amps. The manufacturers claim this to be sufficient to drive a 1200V 100A IGBT module, although in practice it has performed satisfactorily with modules rated at 1800V, 400A. The resulting circuit is very simple, with only a few components to be chosen to suit a given application. A suitably designed PCB can be bolted directly to the IGBT module, at least in the case of the plastic package types.

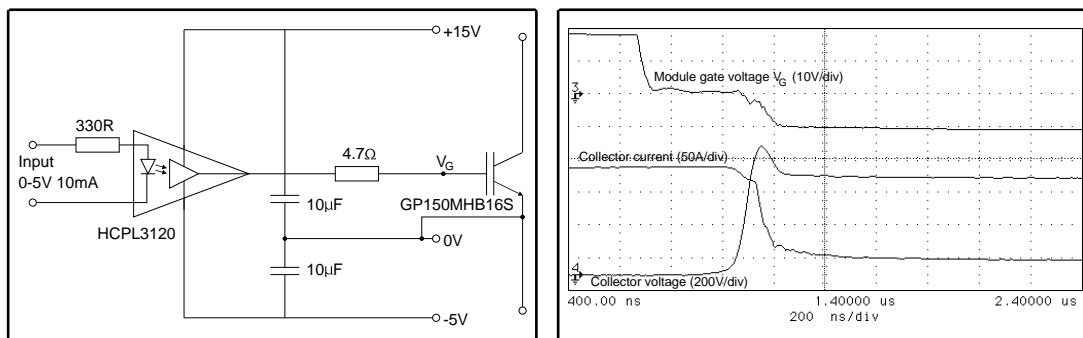


Figure 12. Simple IGBT gate drive circuit diagram (left) and typical waveforms (right)

As it stands, this circuit only provides basic on / off control, but it is very simple, compact, has a modest average power requirement, and should be

reliable due to the low component count. Galvanic isolation between control signal and gate drive is provided by the opto-isolator arrangement within the gate drive chip. The current and voltage rises and falls are governed by the external circuit and the IGBT itself. As noted previously, a modicum of control can be imposed by using a larger gate resistor than the minimum specified by the IGBT's manufacturer; this will slow down the dV/dt during switching in a relatively predictable way. Here the manufacturer's minimum specified value is used to achieve rapid switching.

Variations on the single gate resistor include use of two resistors and diodes to control turn-on and turn-off independently. Drawbacks include dependence on the particular IGBT in question, which is undesirable for series or parallel operation of devices; additionally, the switching time is increased, which has consequences for switching losses. A variant on this basic circuit uses a zener diode between collector and gate as an over-voltage clamp. Crude series operation of devices is possible in this way.

As mentioned earlier, the HCPL3120 is only recommended for the smaller types of IGBT module. For larger types, a push-pull output stage with a pair of power MOSFETs provides sufficient drive capability.

Active Snubber

A variation on the above circuit uses a simple feedback loop and the controlling capability of the IGBT's gate to regulate voltage rise and fall rates. A separate feedback capacitor (C_F) which swamps the IGBT's Miller capacitance (C_{GC}) at higher collector voltages is added between gate and collector terminals. The initial stage of turn-off is unaffected, as the Miller capacitance is so large, but once the collector voltage begins to rise, the external capacitor dominates. The collector voltage rise is then controlled by the external capacitor during the main voltage rise, relatively independently of individual device characteristics.

The controlling ability of the gate can be clearly seen in the waveforms below; the gate voltage initially falls, but rises once again as the rapid collector voltage begins. The IGBT's channel current is turned back on, restoring some electron current flow, and hence controlling the growth of the depletion layer in the IGBT. The closed-loop nature of the circuit acts in such a way as to give a linear

voltage rise. Unlike the case where the gate resistor is increased to soften the switching transitions, the gate resistor can be kept at the minimum specified by the device manufacturer, so the delay prior to the collector voltage rise is virtually unchanged compared to the simple hard switched case.

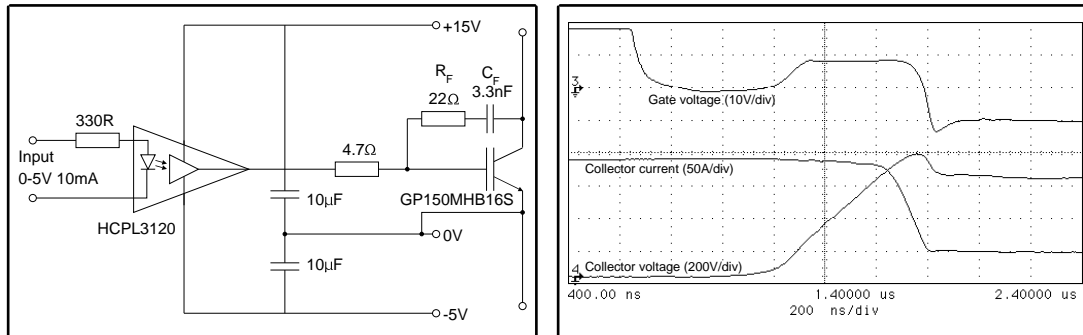


Figure 13. Active snubber IGBT gate drive circuit diagram (left) and waveforms (right)

Some care is needed with the application of this snubber; the external circuit must not apply a large dV/dt to the IGBT's collector, or the active snubber circuit may turn the device on. Such a scenario occurs in a full bridge leg, for example, although use of a sufficiently low impedance gate drive, a resistor in the feedback path (R_F) and negative gate bias during the off state can avoid the problem.

Use of this snubber brings other benefits, such as a reduction in emitted electromagnetic interference, and a smaller collector voltage overshoot at turn-off. Another problem which can be addressed with the active snubber is that of reducing the required ratings of the galvanic isolation between logic circuitry and power device (in this case, the opto-isolator in the HCPL3120). Due to coupling through stray capacitances, this isolation will have a finite dV/dt rating as well as a maximum continuous voltage rating, presenting a problem if high voltages are to be switched with short rise times. With the IGBT's move towards higher voltage operation, such problems are getting more difficult.

Finally, and most significantly for the purposes of the investigation undertaken here, this circuit will provide a useful insight into the consequences of switching parallel IGBTs when use is made of the gate's controlling ability. In common with most of the gate-side control control schemes, the IGBT is kept in its saturation region for an extended period. In other words, it is proposed that investigating the operation of the IGBT module with a simple active snubber

will yield useful information about its operation with the more complicated gate drive schemes, as they all ultimately rely on the device's gate controlled characteristics in its saturation region.

Closed-Loop Collector Voltage Control

One of the more sophisticated gate drive schemes mentioned earlier is to actively control the IGBT's gate in a closed loop feedback system in order to control the collector voltage waveform during transients. The block diagram below shows the salient points of the method: a 'reference' waveform is generated, and the collector voltage is compared with this to generate an error signal. The attenuation B scales the collector voltage to match the reference, and the amplifier gain A is chosen to set the overall control loop gain. An output buffer drives the IGBT gate from this error signal. Provided that errors are kept within reasonable bounds, this enables connection of devices in series for higher voltage applications. The relevance of the technique here is merely to illustrate one of the more sophisticated ways in which use can be made of the gate's controlling ability.

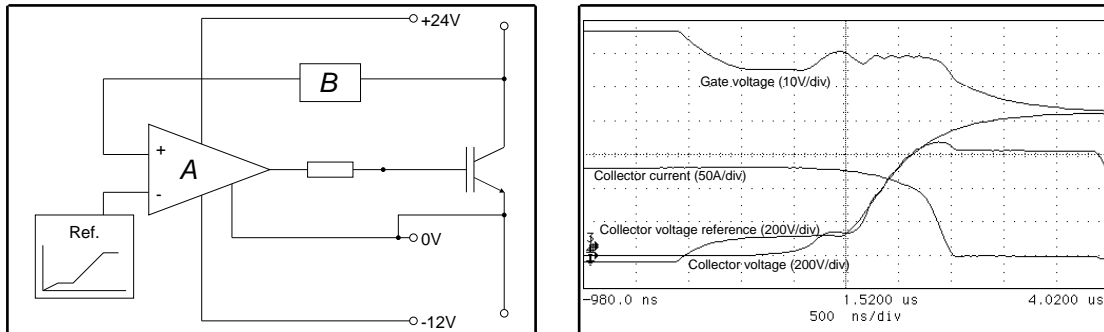


Figure 14. General principle of active collector voltage control (left) and typical waveforms (right)

The particular collector voltage control employed here is a 'two step' method [37]; the reference waveform has two distinct parts to it. The first is a plateau at a value representing a collector voltage of around 100V, and the second is the main voltage rise, in this case a simple R-C ramp. During the first part of turn-off, the gate is driven low until the collector voltage rises to the plateau level mentioned. The active controller then takes over for the remainder of turn-off. The benefit of this method is that the device is taken out of the on-state into its active region, where C_{GC} is much smaller (by an order of magnitude or more), and so a high bandwidth control loop is much easier to implement. Obviously this is particularly important with larger modules.

Once again, the controlling ability of the gate is evident, as is a little oscillation. The particular gate drive used here was not properly 'tuned' to the IGBT module: a number of components have to be changed according to the IGBT's rated current. This complexity is a drawback of the 'two step' method, although for a given application the required values need only be determined once. Nevertheless, the principle of the method is clearly demonstrated here.

Other IGBT Developments

A number of other aspects of IGBT design are worth noting, although they are not considered in detail here. The 'press pack' type modules described earlier are beginning to become commercially available, eliminating numerous problems associated with the plastic package modules at a stroke. In particular the troublesome material interfaces and bond wires are replaced with interfaces whose behaviours are much better understood from thyristor and GTO experience. On the other hand, the integrated solutions with multiple power switches, gate drive circuitry and perhaps even control logic are unlikely to appear in this style of package.

Higher voltage devices remain an aspiration, with the IGBT challenging the GTO and IGCT in ever higher voltage applications. 6.5kV devices have been developed [38], but the trade-off between a low on-state and high switching frequencies is all too apparent. The area of each chip needed for sustaining the high voltage in the off state becomes larger. The device design is thus necessarily a compromise: to keep the on-state manageable at around 4 volts, the typical switching frequency drops to 1kHz or less. The GTO and IGCT remain serious competitors here, with their lower conduction losses.

Trench gated devices have also become available, where the top structure is redesigned so that the lateral MOS channel is replaced by a vertical one. It is then possible to fit more channels per unit area, and hence reduce the on-state voltage of the device. The so-called 'JFET effect', where placing the channels of a conventional vertical device too close together increases the on-state voltage, is also eliminated. The JFET effect comes from the IGBT's middle p-n junction (J2), which is reverse-biased during normal conduction, and consequently has a depletion region either side of it. This depletion region grows with device

current, as the resistive nature of the IGBT's channel and inter-cell volume causes the reverse voltage across J2 to grow. The path for current between adjacent cells is narrowed, increasing the device on-state voltage. Whilst reducing the cell spacing and thus increasing the number of channels in parallel is desirable to reduce the channel resistance, reducing the spacing too far causes the JFET effect to raise the device's on-state voltage. It follows that there is an optimum cell spacing, which is not necessarily the minimum achievable with the available processing technology. The trench gate topology avoids this problem.

Summary

The structure of the IGBT is basically that of a conventional vertical MOSFET, with the addition of an extra p doping to form the collector. From the structure and its equivalent circuit, it can be seen that the resulting device can be regarded as a Darlington configuration, with a MOSFET driving a bipolar transistor output stage.

IGBT and MOSFET processing technologies are similar. As a consequence, any fabrication problems are also likely shared, and any problems displayed by the MOSFET due to processing difficulties are liable to be seen in the IGBT.

The IGBT's steady-state characteristics are seen to resemble the MOSFET's, but switching transients are quite different. At turn-on, the IGBT behaves like a very high voltage MOSFET, turning on quickly but with a huge on-state voltage; once the bipolar action begins, the device fills with carriers and the on-state voltage drops. At turn-off, the main current fall is similar to a slow MOSFET, but the tail current which appears is distinctly bipolar.

Packaging the IGBT for use in circuits has proved a long-standing challenge as ratings have increased; two distinct approaches have been the plastic module and the press pack, and the two approaches each have their strengths and weaknesses. The plastic module is amenable to the integration of extra functionality within the package, but has historically suffered from long-term reliability concerns. The press pack is a much longer standing solution, but cannot presently compete in terms of functionality.

Three methods for controlling an IGBT have been described: the basic on-off ('hard switched') control; the addition of a few passive components to make a simple collector voltage transient snubber (the active snubber); and a closed-loop collector voltage controller using a reference waveform. The benefits and problems associated with each in terms of the whole IGBT module can be seen from external measurements, but the consequences for the operation of the parallel chips within the module are not seen.



Current Measurement

Introduction

Large IGBT chips have always presented a processing problem. It has proved difficult to manufacture them for rated currents above approximately 100 amps if a reasonable yield is to be achieved. Novel repair techniques have increased ratings a little [39] but the maximum IGBT chip size is still an order of magnitude smaller than that of the GTO. IGBT modules rated for currents much larger than this are invariably constructed using parallel chips, but there have been few reported investigations of the consequences of doing this. Previous work has explored the parallel operation of discrete devices of most kinds, and IGBT manufacturers issue guidelines on parallel operation of both discrete devices and modules, but parallel operation within a module is largely unexplored. Anecdotal evidence suggests that early IGBT modules suffered many unexplained failures.

In order to make useful measurements of currents within the module, it must be subjected to the conditions it would experience in a typical application. A typical application for the modules available for this investigation, which are rated at 1800V, 150A, would be an induction motor drive of tens of kilowatts. This is obviously impractical to deal with in a small laboratory for various reasons, chiefly cost, noise, space and power. An alternative solution is required, the step-up converter being a suitable circuit topology. It can be implemented in a small space with a minimal power requirement, enabling current sharing during the turn-off transient to be investigated.

Measurement of the current flowing in the individual IGBT chips within a module presents another problem. Existing current measurement methods for large currents with fast switching transients are either large and expensive, or

have a significant insertion impedance, so a much more compact technique is needed. Considering the options available led to the development of a suitable method based on the magnetic field.

Switching IGBTs

A convenient method for reproducing the conditions experienced by a power switch in a real application is to use it as the switching element in a step-up (boost) converter. Here, at turn-off, the switch will experience both full current and voltage simultaneously, as would happen in a bridge converter driving an induction motor, for example. The circuit diagram of the apparatus used and typical waveforms are shown below in figure 15. The IGBT and freewheel diode could be viewed as part of one leg of a converter, with L_S as an inductive load, such as a motor, and C_{out} as the DC link supplying the converter. By suitable choice of pulse width and supply voltage, the peak current can be chosen, and then the pulse repetition rate and R_L can be adjusted to produce the required 'DC link' voltage.

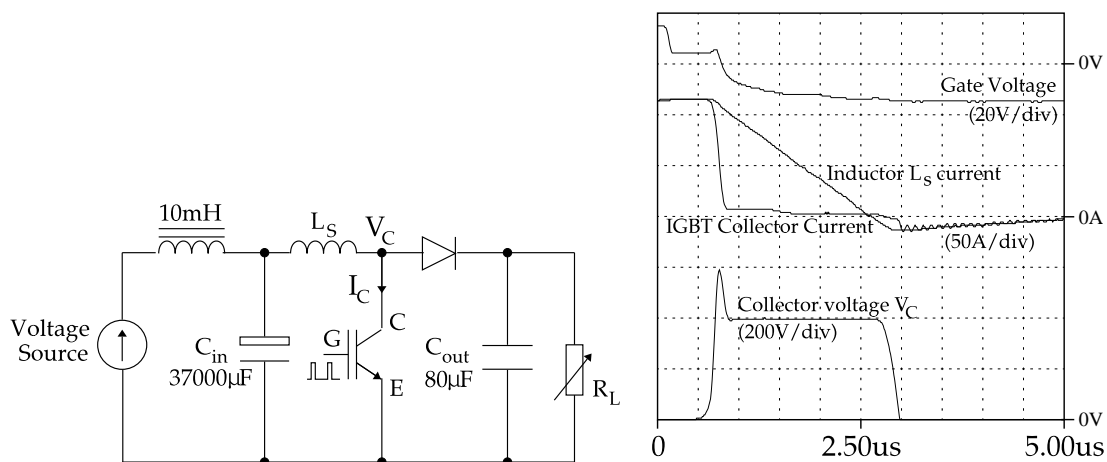


Figure 15. Step-up (boost) converter arrangement and turn-off waveforms

The investigation here focuses on device turn-off, as the turn-on of a fast power device is mainly governed by its external circuit. The collector inductance acts as a turn-on snubber, and so the collector voltage falls to a relatively low value before the collector current is high. In contrast, at turn-off, the collector voltage rises before the collector current can fall, and so the power dissipation at turn-off is much larger than at turn-on.

Initial Experiments

Initial investigations as part of a previous project had revealed the problem of current sharing between chips in a multi-chip module. An 'open' (i.e. no covering gel or plastic cover) module containing four 75A chips in parallel was modified, with separate collector connections brought out for each chip. The current in each chip could then be measured, using a conventional method, some distance away from the confined space of the module's own busbars. It was seen that, at turn-off, the chip currents diverged significantly. (one of the four chips had failed at some time previously, so there are only three collector current traces)





Figure 16. Module for initial experiments.
Note the long external collector leads
for current measurements.



Typical waveforms; collector currents
(top, 10A/div) and voltage (bottom, 200V/div)

It was obvious from these early measurements that further investigation was worthwhile. The simple approach described is unsuitable for anything more than this crude initial qualitative investigation, as significant modification of the IGBT module is required to make the measurements. For results which give a real insight into switching operation, a non-invasive method of current measurement is required. It would also be desirable to be able to measure all of the chips' individual transient currents simultaneously.

Measurement Techniques

The requirement for a non-invasive current measurement technique capable of measuring individual chips' currents within a switching IGBT module imposes a number of constraints and requirements:

(1) The module may not be modified mechanically. For example, bringing out separate power connections for current measurements, as used above, introduces a significant amount of extra inductance. With the large dI/dt during switching, any stray inductance could significantly alter results; for example, a turn-off transient of $1000A/\mu s$ will produce a voltage of 1V per nH of interconnecting inductance. An approximate value of 1nH per mm of interconnection means that these could play a significant role in the module's operation, as demonstrated in [40].

(2) The current measurement probe must have negligible insertion impedance. Obviously any measurements taken could be influenced by any extra impedance; the impedance of the whole current path through a module being of the order of $5m\Omega$. It is assumed that the lack of insulating gel in an 'open' module will make little difference to the switching properties of the module.

(3) An adequate high-frequency bandwidth is necessary, since oscillations between paralleled devices have been seen to occur at tens of MHz, at least in the case of MOSFETs [6]. However, due to the operation of the boost converter, with short current pulses at infrequent intervals (typically $20\mu s$ to $50\mu s$ repeated at 10Hz to 2kHz), the low-frequency response is unimportant.

(4) An adequate peak current capability is required. The largest IGBT chips switch up to 100 amps under normal circumstances, and may switch 1000 amps under short circuit conditions.

(5) Immunity to stray signals from other nearby conductors is vital. In the cramped confines of an IGBT module, there are many conductors carrying currents being switched at high speeds, generating unwanted magnetic fields. These must be ignored.

(6) Capacitive coupling must be minimised, as a high dV/dt is experienced during switching.

Existing Techniques

There are a number of existing current measurement techniques, which are well-proven and reliable. For various reasons, primarily size, none of them are suitable for the purpose, as shown below. However, in order to develop a suitable technique, it is necessary to look at these existing techniques and how they might be adapted for the task in hand.

Current Transformer

The current transformer is a long-established and well-respected method for measuring large currents, but the transformer itself is bulky. Essentially it is a transformer whose secondary is connected to a known impedance; the conductor being monitored provides a single-turn primary. The voltage across the known impedance is monitored to calculate the secondary current, and hence the primary current can be calculated from the ratio of primary to secondary turns. A theoretical treatment of the current transformer and its limitations, along with some modern developments can be found in e.g. [41]. As with any transformer, this type of device is unable to transform DC, and so the DC component of the current being monitored is lost. (its magnetic core will also saturate in the presence of DC currents) This is of little consequence for the step-up converter under test, as the duty cycle is very low; furthermore, it is transient performance which is of greater interest here.

Commercial devices capable of measuring sufficient current are of the order of 50mm diameter, which is obviously impractical within an IGBT module, where only a few millimetres of space are available for any current measuring probe. However, current transformers do have a role to play, in measuring a module's total current for calibration purposes, as will be seen later. A well-known type is the Pearson current monitor [42], the type 110 being suitable for this application. The insertion impedance of the Pearson type 110 used in this way is quoted as $0.2\text{m}\Omega$, with a 3dB bandwidth of quoted as 1Hz - 20MHz, a rise time of 20ns, and saturating at 0.5A.s. Whilst 20MHz is adequate for the purpose of measuring the total module current, differential oscillations between chips could be at much higher frequencies. Improvements on the current transformer's bandwidth have been attempted [43] [44], but the size of such a probe remains impractical for chip current measurements.




Figure 17. Pearson current monitor (left), LEM converter (top) and IGBT module




Figure 18. Tektronix current probe and IGBT module for size comparison

Tektronix Current Measurement System

This is another widely-used current measurement system [45], and is based on a combination of a Hall probe for DC measurements and a transformer for AC measurements. It has similar problems to the Pearson current transformer: it is large, and it is expensive. There is the additional problem of the introduction of a delay of around 50ns in its measurements. The probe required for the current levels here, the A6303, has a bandwidth quoted as 15MHz; again, it can provide a reference measurement for other techniques, but the physical size of the A6303 makes this an unattractive option. Its major advantage over most current transformers is that it is a ‘clip-on’ device; i.e. it can be removed and inserted into a circuit without needing to disturb the circuit. This would be a desirable feature for any current measurement probe used inside an IGBT module.

LEM Converter

The LEM converter [46] is based on flux compensation and a Hall probe, and can measure DC current. It does however have many of the other disadvantages of the current transformer, and additionally a very limited bandwidth of around 100kHz; its ability to measure DC is of no use here. It is intended for use as a current measurement device in an inverter’s control loop, for example, rather than for the rapid transient current measurements required here. Typical device ratings are quoted [46] as being 100A/ μ s, an order of magnitude or more slower than the currents seen here. Consequently, it is not considered further.

Resistive Shunts

Perhaps the simplest method of current measurement is to make use of Ohm's Law, and insert a known resistance in the circuit of interest, provided that the resistance is small enough to avoid altering the circuit's operation. The voltage across this resistance can then be monitored to see the current waveform, with the caveat that the resistance's parasitic inductance will limit the bandwidth available. Using a co-axial construction to achieve a minimum inductance, resistive shunts are a well-established technique for measuring large currents with a high bandwidth. Suitable items are available commercially, and it is often used as a reference measurement for evaluating new current sensing techniques.

Two major problems with this technique present themselves, notably size and invasiveness. The circuit must be opened in order to insert the known resistance, and this is unacceptable here. Furthermore, a suitable shunt for the current and frequency under consideration is large, particularly the coaxial types required for high current, wide bandwidth use. The lack of isolation from the power circuit adds a further problem, if a relatively minor one by comparison. The resistive shunt remains, however, quite suitable for measuring the total current flowing through a module or circuit. Development continues in these fields [47].

Rogowski Coil

A Rogowski coil, also known as a 'Gauss Worm', is an air-cored transformer dating from 1912 [48]. In its early years, its use was limited by the fact that its output requires integration to reproduce the waveform of the current it is monitoring. Until relatively recently this confined its use to applications where current pulses were of huge magnitude and short duration. Such current pulses could produce a signal large enough to be integrated by a passive resistor-capacitor network [49], or even by just a very low resistance [50].

The arrival of the operational amplifier made the production of a much higher performance integrator feasible, and opened the way for a much broader field of application for the Rogowski coil. Early attempts with IC integrators yielded

items with a bandwidth typically from 50Hz to 1MHz [51], which was quite adequate for investigation of circuits involving the GTO, for example, but the MOSFET and the then relatively newly-arrived IGBT demanded a wider bandwidth. Nevertheless, the Rogowski coil was beginning to challenge the resistive shunt which was well established in the field of power electronics; a particular advantage of the Rogowski coil being that, unlike the resistive shunt, it was totally isolated from the power circuit. On the other hand, it lacks the capability for DC measurements.

As experience of its use and understanding of its design and operation improved, so the ratings of devices available increased, helped by improvements in operational amplifiers [52], [53]. The trend towards higher bandwidth has continued, with parts approaching 10MHz [54] being offered, which is approaching suitability for measurements with IGBTs. Its main problem for the work undertaken here remains its size, commercial devices being offered in the range of 100mm to 2m diameter. Commercial pressures have resulted in the development of physically larger, higher current sensors with extended low frequency responses, rather than the smaller, higher bandwidth ones required here.

An interesting application described briefly in the concluding paragraphs of [55], and in more detail in [56], is an investigation of the current distribution across a GTO wafer during switching. A number of coils were situated in concentric circular slots in the GTO's cathode plate, and hence current distribution within the wafer, at least in the radial direction, could be investigated. Development continues, including a smaller (and much cheaper) implementation based on coils formed from tracks on a PCB [57].

A considerably smaller coil, described as being a little over $\frac{1}{2}$ " diameter and $\frac{1}{4}$ " high, has been reported [44], along with the mathematical theory behind its operation. Its ability (shared by all Rogowski coils) to compensate out fields from nearby conductors is noted, a feature which would certainly be useful when trying to measure a particular busbar's current amongst the densely packed busbars in a typical IGBT module. The op-amp based integrator presented to accompany this coil is less notable, in that relatively complicated arrangements are required to reset its zero offset.

More recently, another small Rogowski coil was reported [58] as being installed within an IGBT module - albeit one with copper blocks instead of IGBT chips - but regrettably no details of the device, its capabilities or its commercial availability are given in that paper. From the photograph published it appears to be of the order of 15mm diameter, which is much smaller than the previously seen commercial Rogowski coils.

Other Approaches

A number of other current measurement techniques have been proposed, each generally falling into one of three categories. Variations and developments of existing principles have been seen, including a current transformer with a non-linear amplifier to extend its dynamic range [59]. Combinations of existing techniques are also interesting, such as the combination of Rogowski coil and Hall probe [60] [61] to add DC capabilities to the Rogowski coil, but size remains a problem.

Novel approaches are far less common, a recent notable one being the use of amorphous materials as a measurement probe [62]. A conductor of suitable material is formed by very quick cooling; this leaves insufficient time during cooling to form a crystalline structure in the normal way. The resulting conductor has an electromagnetic property which is useful here: its impedance to AC signals changes in the presence of an applied magnetic field. A wire made in this manner, wrapped around the conductor carrying the current to be monitored, will have an impedance which is a function of the current being measured. To date the maximum frequency is low (40kHz), although this may be increased.

Magnetic Field Coils

Having considered the existing current measurement techniques available for the investigation, it is obvious that none is suitable as it stands. The Rogowski coil is the closest in terms of functionality, but commercially available sizes are too large for the purpose. Evidently miniaturisation is possible, and the ability to compensate out unwanted magnetic fields is vital in the cramped confines of an IGBT module. With some development, a more compact current measurement system based on the magnetic field can be envisaged. The amorphous wire

system is the other possible candidate for development into a current measurement system which would be useful for work of the nature undertaken here. However, it lacks the stray field compensation of the Rogowski coil, and the available expertise with magnetic field coil based measurement systems dictated that the field coil technique was a much more practical route to pursue.

The Rogowski coil's principle of operation is quite simple: a conductor carrying a current i produces a magnetic field \mathbf{H} around it proportional to that current. Furthermore, a conductor in a changing magnetic field will have a voltage V_{ind} induced across it proportional to that rate of change. It follows that placing a conductor near a changing current, and integrating the voltage induced along it, will enable the original current waveform to be recovered, less any DC component. It is possible to calibrate the system, either mathematically if the exact geometry of the coil is known [63], or by comparison with a reference measurement at two current levels.




Figure 19. Measurement of current via the \mathbf{B} -field

Earlier research [64] on current redistribution in single GTO wafers had resulted in the development of a current measurement system comprised of a small magnetic field coil and a high-bandwidth integrating amplifier. The coil was placed with its axis parallel to the plane of the GTO wafer under test; current redistribution within the GTO wafer could then be monitored. Although this probe was not directly related to the Rogowski coil, as it consisted of a single cylindrical coil, development of this method into one suitable for the task in hand was the obvious route to take.

As it stood, the coil was too big, and it lacked discrimination. Placing it in or near the module would produce a sizeable, unwanted, output signal. This is due to the fact that any changing current (in the diagram above) with a y or z component will induce a voltage across the coil. Furthermore, the signal is sensitive to movement of the coil within the xy plane, and to tilting around the y and z axes, all of these causing a change in output signal magnitude, making calibration difficult.

Compensated Field Probes

Initial attempts to improve the field probe were directed at miniaturisation: a coil of 34SWG enamelled copper wire was wound around a length of insulated single-strand copper wire. The coil was then covered in heatshrink sleeving, and the inner strand of wire removed. Heating the coil with a hot air gun allowed it to be formed into a 'U' shape which would fit neatly around a module bus bar. It was realised that such a probe could be vastly improved by omitting the coils around the curved part of the 'U', leaving two parallel coils, electrically connected in series. These are placed either side of the bus bar, as shown below.




Figure 20. Theory of the compensated field probe: two coils around a conductor




Figure 21. Practical realization of the compensated field probe

In this form, it provides a large degree of cancellation of unwanted signals, as depicted below, although elimination of stray pickup from the interconnections requires some care in twisting the return wires together. Currents along the x axis produce magnetic fields which do not link the coils at all; currents along the y axis produce fields which link both coils equally, but since the coils are connected in series, the voltages which are induced simply cancel out. Currents along the z axis outside the probe will produce an unwanted signal, as they link both coils, but by different amounts; however, experiments showed that judicious probe placement could render this effect negligible. Such a probe has several desirable properties, provided that it is carefully sized. It can be positioned around a bus bar, with movement in the x and y directions being impossible, and in the z direction making no difference. Tilting is only possible around the y axis. A fuller theoretical treatment of this is given in [65].

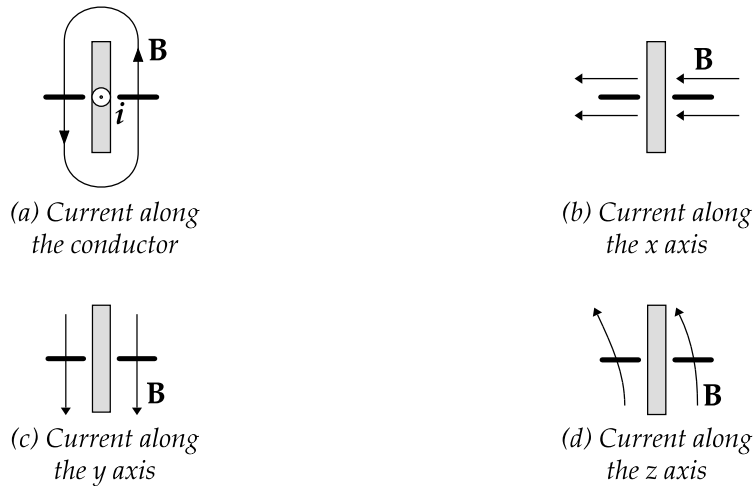


Figure 22. Field compensation (plan view of figure 20)

Integrator Design and Testing

An integrator design was already available from earlier work conducted on a similar current redistribution problem in GTOs [64], and this was adapted for the purpose. It is a combination of active and passive integrators; at lower frequencies, the AD5539 op-amp performs the integration with the precision R-C network in its feedback path. Once its output falls to unity gain, the passive R-C network at the input provides the integrating function. By suitable choice of components - the two 1kΩ resistors in the input circuit appear in parallel - the transition between the two is not noticeable, as can be seen from the Bode plot (figure 24). The circuit diagrams are included here for reference.





Figure 23. Integrator main circuit diagram (above) and supply decoupling (below)



Unwanted noise was found to be coming into the integrator through its power supply, which is a pair of isolated 5V linear supplies. The transients generated by the IGBT's rapid switching were being picked up by the mains leads, and coupled through the integrator supply into the integrators themselves. Capacitive and inductive decoupling with feed-throughs for the supply into the integrator boxes produced little improvement, so the supply leads were wound around ferrite cores. This presents a high-impedance input to any common mode noise, such as that being coupled into the integrator; this resulted in the noise pickup being substantially reduced. Battery operation may provide a better alternative, although substantial redesign would be required to the existing apparatus.





Figure 24. Bode plot of measurements of the integrator's response: magnitude response (above) and phase response (below)



As can be seen, the integrator has a satisfactory magnitude response over a range from around 10kHz to 50MHz. The resistor across the op-amp's input terminals reduces noise levels and causes the gain to level off below 10kHz as

seen; this is unlikely to be of any consequence for the work presented here. Regrettably the appropriate equipment was not available to test beyond 50MHz; it might be inferred from the Bode plot that the gain is starting to drop away towards the top of the range tested, but it is difficult to prove this without specialist test equipment as the integrator gain is very small at such high frequencies.

Calibration

The magnetic field measurements described above need little calibration to provide a qualitative estimate of the magnitude of the current redistribution problem. For more detailed work, such as predicting the difference in power dissipation (and hence temperature) of individual chips, calibration is required. Two methods were used for calibration of the probes: either with the probes outside the device, or with the probes in their final positions.

External calibration can be achieved by measuring a current transition with a compensating field coil and a 'reference' method (e.g. Pearson current transformer) on the same conductor, and comparing the two measurements. The probe can then be placed around any conductor in the device. This method is simple but relies on the assumption that the probe is placed in equivalent positions during calibration and use.

Alternatively, internal calibration can be performed with the probe in its final position. Each chip in the module is fitted with a probe, and the integrated outputs are summed; this sum can then be compared with the total module current, as measured outside the module with a reference device. A refinement is to calibrate the probes against each other before fitting to the module. This latter method was used to obtain the following waveforms for a two-chip module; obviously such a method is only viable for a small number of probes.

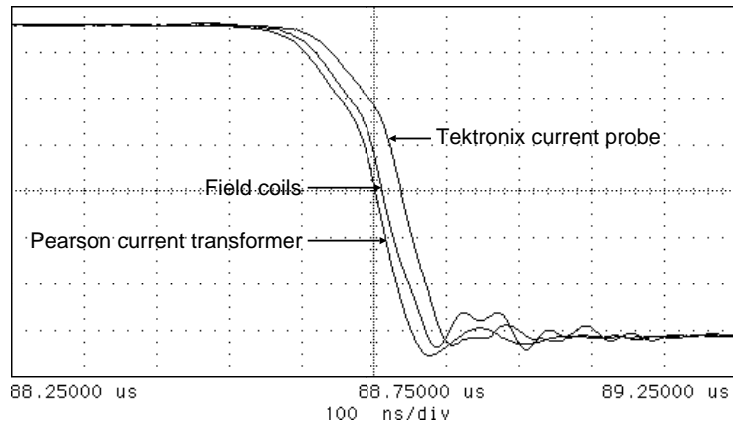


Figure 25. Calibrating probes. Left to right are Pearson current transformer, sum of two field probes, and Tektronix current probe.

As can be seen, there is good agreement between the three methods. The field coil waveform follows that of the Pearson closely, but with a 10ns delay; the Tektronix current probe lags behind by another 30 to 40ns. The extra delay introduced by the Tektronix measurement is well-known. Either of these current measurements can be used to calibrate the field coil measurement system in amps / volt.

Initial Measurements with the Field Probe

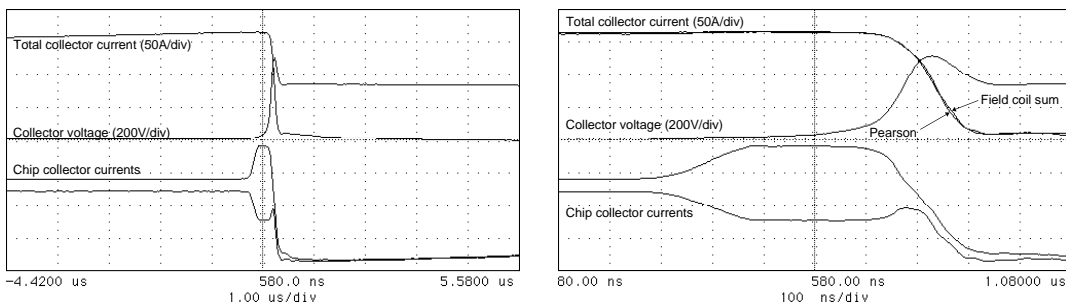


Figure 26. Early measurement result for a two-chip IGBT: overall view (left) and magnified (right)

An example measurement with the field probe is seen above; the top two traces appear at first sight to be typical IGBT switching waveforms with an inductive load. The gate voltage (not shown) begins to fall at time zero, and the IGBT begins to turn off 300ns later. The individual chip currents (the lower two traces), measured using the field probes, show that the two chip currents diverge significantly for a period of around 500ns during the voltage rise. It can also be seen that they settle at new values for around 300ns. There is no sign of oscillation between the chips.

The second of these two figures also shows the total module current as measured externally with the Pearson current transformer, which is compared with the sum of the two chip currents; it can be seen that the two traces are virtually co-incident, supporting the assertion that the measurement technique is valid, and suitable for the proposed use.

Summary

A circuit topology for testing IGBTs in realistic operating conditions has been shown. Existing commercial current measurement techniques have been seen to be unsuitable for measurements within an IGBT module, and the experimental technique developed for similar work with the GTO has been adapted and improved for the purpose. Some characterization of this system has also been undertaken.

4

IGBTs in Parallel

Introduction

In chapter 3, uneven current sharing was seen between paralleled IGBT chips during turn-off, and this is investigated. A review of the literature concerning operation of discrete devices in parallel is presented first, and then the few publications on operation of chips in parallel are also covered.

The IGBT is in its saturation region during switching, and hence its gate voltage should be controlling its collector current. Hence it is possible to deduce some information about the mechanisms involved in the current sharing of paralleled IGBTs from the behaviour of the individual chips' gate terminals.

Differences in the chips themselves, and differences in the impedances of their interconnections, are suggested as potential sources of current imbalances. The modules tested in this investigation have pairs of chips in a geometrically symmetrical arrangement, and so, for simplicity, the assumption is made that the impedances of solder joints and bond wires are identical between chips. Obviously such a symmetrical layout cannot be achieved in modules with large numbers of chips, four being the practical maximum in plastic-packaged modules.

Discrete IGBTs in Parallel

The parallel operation of many types of semiconductor switches has received attention over a long period, both as paralleled discrete devices, and as paralleled modules which themselves already contain multiple devices. Bipolar transistors [5], MOSFETs [6], GTOs [66], IGBTs [67], and even MCTs [68] have been studied.

Parallel operation of existing packaged devices is a technique which could improve the power output of a circuit without resorting to designing bigger devices or modules, provided that the current is shared evenly amongst the parallel devices during both switching and conduction. Such current sharing rarely happens by default, and various techniques have been developed to try to enforce it. Most manufacturers provide application notes on how to achieve this (e.g. [69], [70], [71], [72], [73]). Parallel operation of whole converters offers another alternative for producing higher output powers from existing modules, but this is complicated [74].

Soon after the introduction of the IGBT, then known as the Insulated Gate Transistor (IGT), parallel operation of the relatively small devices then available was studied [75]. Choosing well matched devices is the main consideration addressed in that study, although the definition of 'well matched' for transient operation is less straightforward than might be supposed. The use of a gate resistor for each device, rather than connecting all the gates together via a single resistor, is suggested. Differences in wiring inductances were subsequently tackled [40], which additionally recommends the use of a single gate resistor for all devices. The debate on the need for individual gate resistors continues, with several manufacturers' application notes [69], [70], [73] suggesting that they are necessary. Wiring inductances for a more general case, including mutual inductances, have also been studied [76]. Thermal effects are also considered [40], and thermal stability is considered to be good in all reasonable cases.

Equalisation and minimisation of stray inductances and resistances in the power and gate circuits for each device are often cited as prerequisites for good current sharing. Interestingly, the minimization of stray inductance has been questioned, with the suggestion that a certain amount of inductance in each collector circuit can be beneficial [73]; similar techniques have been used in bipolar transistor applications [77], typically employing emitter resistors. A small-signal analysis of the MOSFET ([6] and developed further in [78]) led to the suggestion of the use of ferrite beads on the gate connections to provide the required impedance. On the other hand, minimising gate inductance has been suggested as necessary to avoid oscillations [79].

In the on-state, the simplifying assumption that all devices have the same collector voltage can be made. It is straightforward to examine two devices' output characteristics (collector current vs collector voltage) and predict current sharing, given knowledge of the resistances in the power circuit. Statistics can be brought to bear to predict typical behaviour of parallel components [80], although only the on-state is readily tackled in this way. Several manufacturers include recommendations for matching on-state voltages [81] or 'diode mode voltage' [82]. A complication here is the temperature coefficient of resistance: this is positive for the MOSFET, negative for the bipolar transistor and GTO, and can be either for the IGBT, depending on the technology employed and the current being carried. Those with a negative temperature coefficient require either close thermal coupling, or a 'ballast' resistor in their power circuit to avoid thermal runaway, solutions familiar from bipolar transistor practice.

Transient operation is considerably more complicated, because one can no longer make the simplifying assumption that the collector voltages are equal, as any fast-switching device will develop sizeable voltages across any stray inductance in its power circuit. Similarly, the use of separate gate resistors for each device when MOSFETs or IGBTs are operated in parallel means that the gate voltages cannot be assumed to be the same. In the case of paralleled GTOs, the gate provides little control once switching has been initiated, whereas the IGBT, MOSFET and bipolar transistor all remain active during switching.

Beyond the schemes which merely hope to achieve the best results from the available devices, other methods have been developed in an attempt to force current sharing. Turning first to the bipolar transistor, an interesting idea (if impractical for more than two devices) is given, namely a transformer with two identical windings in opposite senses, one in the power circuit of each transistor [5]. This should resist a current increase in either transistor unless there is an identical increase in the other one. A variant on this scheme [83] extends this technique to arbitrary numbers of BJTs, although it is hardly an attractive scheme. A similar method has been proposed for GTOs [66], and the waveforms presented there suggest a fair degree of success. 'Balancing' inductors have also been suggested as a means for connecting arbitrary

numbers of switches of any given type in parallel [84], seemingly presenting an ideal solution. As with most 'ideal' solutions, there is a drawback; in this case, the balancing inductor required is of the order of some microhenries, which is a sizeable component for the current ratings under consideration here.

The gate-controlled nature of the IGBT means that paralleled devices can be forced to share current using a suitable feedback controller for each device [85], [86], [87]. Provided that the performance of the controller is sufficient, any number of IGBTs with any reasonable parameter spread can be operated in parallel, although a suitable controller is complicated and requires considerable development. Both static and dynamic current sharing can be achieved in this way.

IGBT Chips in Parallel

Shortly after the introduction of the IGBT, it was appreciated that paralleled chips would be necessary to achieve high current ratings. As early as 1989 it was reported that matching device parameters, and both matching and minimizing parasitic impedances, was necessary to assure good current sharing [88]; no results were given to support these assertions, however. At the time of writing, ten years on, chip ratings have advanced to the 200 amp range by improvements in device processing technology and by 'repair' techniques. Certain types of chip defects can be dealt with by isolating the relevant area of the chip at the end of the fabrication process, before assembly into a module. The yield of batches of larger chips can then approach that of smaller types [39]. However even these larger chips are some way from the ratings of 1000A and higher demanded from a single switching device, and so parallel operation is set to continue.

The spread of on-state voltage within a wafer lot, and across several lots, has been reported [89], but no information on other parameters is presented, nor any details of the consequences of paralleling mismatched chips. Differences in threshold voltages are touched on briefly in [90], and a technique for identifying them after module assembly described, but no mention is made of the significance of this in terms of transient current sharing. In [91] current measurements were made of current distribution between three sections of a

module, with the intention of comparing these with thermal measurements, but no significant redistribution was seen between the sections within that module. Similar measurements in another investigation [92] also revealed relatively little current imbalance during transient operation.

Non-uniform operation within a single chip, due to differing propagation delays in the gate signal to different parts of the chip, has been reported as a possible source of failures [93], [94]. Although only referring to intra-chip current distributions, it could reasonably be expected that similar effects would exist between chips. For the purposes of this investigation, individual chips are assumed to act homogeneously, i.e. no attempt is made to investigate the distribution of current within a single chip.

As detailed earlier, the steady state current distribution is straightforward to predict from measurements taken from chips before they are assembled into modules, so this investigation focuses on transient current distribution.

Gate Topology

There is a range of opinions on the importance of the various items in the gate circuit on the transient behaviour of paralleled IGBTs. The modules used in this investigation, type GP150MHB16S from Mitel, have two IGBT chips in parallel in each half of the module, with separate gate resistors, as shown in the circuit diagram below. A number of complete modules were supplied, and these were used for the early experiments; later, a number of individual chip assemblies were supplied, and these could be made up into modules with the appropriate busbars.




Figure 27. Two chips in parallel within a module, forming a single switching element (half a module)

The gate voltages of the two chips (V_{G1} and V_{G2}) were monitored individually on the bond wires between gate resistors and chips. Figure 28(a) shows the results of these measurements on a particular module: the top traces show the module's gate voltage (V_G) and the individual chip gate voltages; the chip gate

voltages are quite clearly diverging. It can be deduced from these traces that the total charge supplied to the two chips' gate terminals is quite different, suggesting that C_{iss} is different for the two chips. However, once the collector voltage begins to rise, the individual collector currents (I_{C1} and I_{C2}) converge once more, so the power difference is relatively small.




Figure 28. Individual chip gate voltage experiments, (a) with the module as supplied, and (b) with chip gates tied. The current redistribution evident in the left-hand figure is virtually eliminated.

The gate voltage traces seen here bear some resemblance to their respective collector currents, as would be expected when the IGBTs are in their active regions. It has been suggested [95] [96] that tying the two gates together (i.e. connecting the points V_{G1} and V_{G2}) might improve the current sharing. Others [40] suggest that this is unwise, as current oscillations between the two chips could result. Interestingly, this turns out not to happen, as seen above; at least, not in this case. The current sharing is substantially improved at turn-off.

As a point of interest, the gate resistors in the module were removed completely, and the gate drive's output resistance also reduced to a minimum. Figure 29 shows that the oscillations which might have been expected are just beginning, but they do not appear to affect turn-off significantly. The lower gate resistance does, however, reduce the turn-off time, and consequently E_{off} will be lower. In a later chapter, the theory behind the minimum value of R_G is investigated, with a view to speeding up switching and reducing losses if possible. It may be that oscillations occur, but do not interfere with the turn-off process, or even that there is insufficient time for them to occur at all.




Figure 29. Measurements with minimum gate resistance.
Differential oscillations are just beginning between the two chips.

Parameter Measurements

The chips used for the following tests were supplied bonded onto DBC tiles by their manufacturer, along with some data, namely collector voltage at collector currents of 1A, 50A, 75A and 150A at rated gate bias; collector leakage current at zero gate bias with collector voltages of 100V, 800V, and 1600V; and threshold voltage $V_{GE(th)}$. The latter requires some explanation: the gate and collector are connected together, and a voltage applied between gate and emitter. The threshold voltage is defined at the point where a collector current of 5mA flows. The manufacturer's data was verified before starting to investigate the effects of operating IGBTs with different characteristics in parallel.

Threshold Voltage, $V_{GE(th)}$

The threshold voltage, sometimes referred to as the 'diode mode' voltage, is fairly straightforward to measure, using nothing more than two digital multimeters (an ammeter and a voltmeter), a variable-voltage power supply, and a resistor. The collector and gate of the IGBT chip are shorted, and it is then connected to the supply via the resistor and ammeter. Convenient choices were a supply voltage of 60V and a series resistance of 10k Ω ; a little adjustment of the supply voltage to set a current of 5mA through the chip then enables $V_{GE(th)}$ to be measured using the voltmeter. It is perhaps an odd definition of threshold, but one which appears to be standard and one which is easily reproducible.

On-state Voltage, $V_{CE(sat)}$

The on-state voltage $V_{CE(sat)}$ is defined as the collector voltage when the IGBT is carrying its rated current at a specified gate voltage. The terminology is

confusing, as the IGBT is actually operating in its linear region at this point, not its saturation region. It is a term inherited from bipolar transistors, whose saturation region is encountered at low collector voltages.

Direct measurement of $V_{CE(sat)}$ can be done with a suitably large DC power supply, provided that a substantial heatsink is provided to cope with the resulting device heating; unfortunately a suitable power supply was not available at this point. Measurement during the operation of the step-up converter with an oscilloscope proved not to be a realistic option, as the input stage of the oscilloscope available could not display the (low, around 3 volts) on-state voltage accurately in the presence of the large transients (around 300 volts) at turn-off. It should also be realised that an IGBT takes a finite time to reach its minimum on-state voltage once turned on, as the carrier distribution within the device takes a while to settle, so such measurements would be of dubious quality. Instead, a small test circuit proved to be straightforward and flexible.




Figure 30. V_{CE} measurements

The device under test is permanently switched on with the desired gate voltage, typically 15V, and an external circuit ramps the device current up to the value desired and back down again. A simple pressure mount assembly was devised to allow the rapid interchange of DBC tile-mounted IGBT chips, simply requiring four bolts to be undone to release one tile ready for the next to be tested. Results from these tests generally proved consistent with data supplied with the samples tested, in many cases to within 1%.

A commercial curve tracer gave the typical 'family' of characteristic curves shown in figure 31. This shows the collector voltage versus collector current for different values of gate voltage; from these curves one can also estimate the on-state voltage for particular values of collector current and gate voltage.




Figure 31. Collector current (I_C) vs collector voltage (V_{CE}) with different gate voltages (V_{GE})

Gate Resistors

One aspect of the IGBT assemblies is the range of values of the chips' individual gate resistors. The resistors in use on the modules under test were fabricated from doped silicon, and were found to have a significant spread of values. A batch of 12 was found to have a mean of 6.4Ω and standard deviation of 0.18Ω . Obviously two chips with significantly different resistors will turn off at different rates, and redistribution will be evident.

Subsequent investigation revealed that silicon chip style resistors are known to have a significant tolerance, around $\pm 8\%$, due to doping variations during manufacture. It was also found that at least one manufacturer has moved to using metal film resistors, commonly available in $\pm 1\%$ tolerance grade for a fraction of a penny each. The press-pack style of packaging has, in at least one case, seen the disappearance of discrete chip gate resistors altogether, an extra layer of polysilicon below the gate metal pad being designed to fulfil this role instead. The author is not aware of any investigation as to the feasibility or practicality of this solution, nor of any current sharing measurements in such press-pack IGBTs. Indeed, observation of current sharing in press-pack devices has so far been limited to externally observable consequences of oscillations [6].

Matched Module

In order to test the usefulness of the existing parameter measurements, a 'matched module' was constructed. A nominal definition of matched was chosen, namely less than 0.1 V difference in both $V_{CE(sat)}$ and $V_{GE(th)}$. A number of individual chips were tested to find a closely matched pair, and these were

then connected up with busbars to mimic a module. Although lacking the heatsink of the commercial modules, switching tests can be carried out at low frequency, so that the heating effect on the chips is minimal. The electrical characteristics of the fabricated module should be very similar to those of the commercial ones, as identical busbar components were used.




Figure 32. Matched module. Good current sharing is evident.

The current sharing can be seen to be very good. The waveforms of gate voltages, measured on the chip bond wires, are also seen to be closely matched. The small variation in chip currents causes little difference in the power dissipation, as the collector voltage is low at that time.

To check the consistency of the IGBT chips' characteristics with average temperature, the pair was heated to approximately 100°C; again the chips share very well. The switching is seen to be delayed by 100ns when compared to the cold result. This is due to the increased value of the gate resistors - being heavily doped silicon, they possess a substantial positive temperature coefficient of resistance, as demonstrated later. An interesting feature is the slight divergence in the gate waveforms during the current fall.




Figure 33. Matched module, heated to approx 100°C. Sharing is still good.

Since absolutely uniform cooling is unlikely in practice, the chips were tested with a substantial temperature difference; one chip was heated to approximately 100°C whilst the other remained at approximately 30°C. The

steady-state sharing is affected, and a significant level of current redistribution occurs, leading to uneven switching losses. The current fall is delayed to a similar extent to that with both chips hot, but once again, there is good current sharing during the current fall.

At higher temperatures, the threshold voltage falls, but since $V_{GE} \gg V_{GE(th)}$, this will not directly affect the on-state voltage. However, the on-state voltage rises with increasing temperature at typical operating currents, so one would expect the hotter chip to carry less current. As seen later, the gate resistor's value increases dramatically with temperature, and so the hotter device delays turn-off of the parallel pair, with the hotter chip's current rising substantially in the period leading up to its turn-off.




Figure 34. Matched module, one chip only heated to approx 100°C, the other remaining at approx 30°C.

The tests on this matched pair indicate that differential heating affects the current redistribution via the temperature dependence of both the chip characteristics and the gate resistance. The implications of these results are complex, depending on many features of the exact circumstances of operation. For example, at a low switching frequency, with conduction losses dominating switching losses, there is a distinct negative feedback mechanism at work; conversely, at high frequency, where the switching losses dominate, a positive feedback mechanism can be envisaged. The number of chips (or modules, for that matter) in parallel and their thermal coupling will also have important roles to play. It may be possible to define a maximum switching frequency for a given situation.

Mismatched Module

In an attempt to quantify the levels of current redistribution suffered by unmatched chips, an arbitrary definition of unmatched was adopted: $V_{CE(sat)}$

could differ by 0.6V or more, or $V_{GE(th)}$ could differ by 1.2V or more; either condition being sufficient to define an unmatched condition. Accordingly, a pair of chips was identified which were unmatched on $V_{GE(th)}$, but matched on $V_{CE(sat)}$. These were tested in the same way as before, and the steady state current sharing is good, as expected. The current redistribution is quite significant at about 35% (i.e. I_C peaks at 35% above its on-state value) and, unlike the previous case, this persists into the switching part of the cycle. This has the expected effect on the distribution of the switching power losses as shown in the power curves. Integrating the power curves over time to give the switching energies of the individual chips shows that the chip taking the higher current has 25% higher switching losses than the chip shedding current.

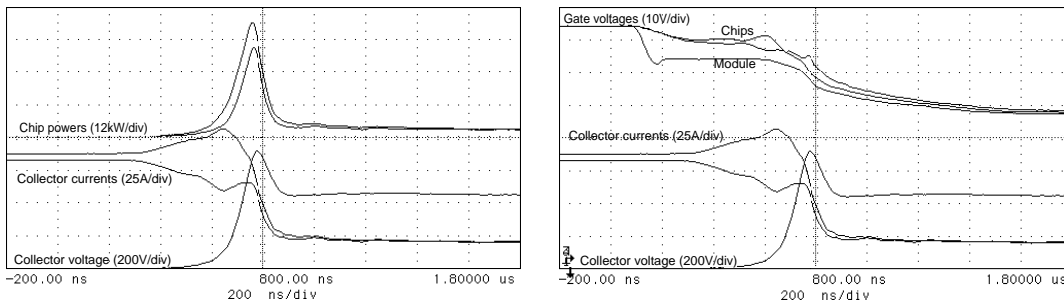


Figure 35. Mismatched module, showing difference in gate behaviour and switching losses.

Attempts to quantify the relationships between differences in the parameters and current redistribution were unsuccessful due to the complexity of the situation and the lack of suitable test samples.

Gate Drive Techniques

As described previously, there are several different techniques for controlling an IGBT, from a simple on-off control implemented in a single chip, to closed-loop control of collector voltage. The effect of the control scheme on the chips' behaviour was investigated. A particular pair of chips (half of a four chip half-bridge module, chosen at random from a production batch) was selected, and this pair was used with each of the three different gate drive schemes described earlier. The module in question was taken from production before it was fitted with its plastic cover and insulating gel, but after it was electrically complete, and was fitted with the current measurement coils as above. The module was first switched using the simple 'hard switch' gate drive to form the basis for

comparisons with other techniques. Some redistribution is evident here during the period when the collector voltage is low, but the currents reconverge as the collector voltage rises, so the power imbalance is small.




Figure 36. Driven with a simple 'hard switched' gate drive, the IGBT pair exhibits some redistribution. Moving to the other gate drive types, both cause significantly different behaviour from the simple hard switched case. The current waveforms diverge at the start of turn-off, but they do not reconverge once the collector voltage begins to rise. It can be seen that making use of the gate's active control capability can make switching loss distributions much more unbalanced. The shapes of the two current waveforms are similar between these two cases, and yet completely different to the hard switched case.

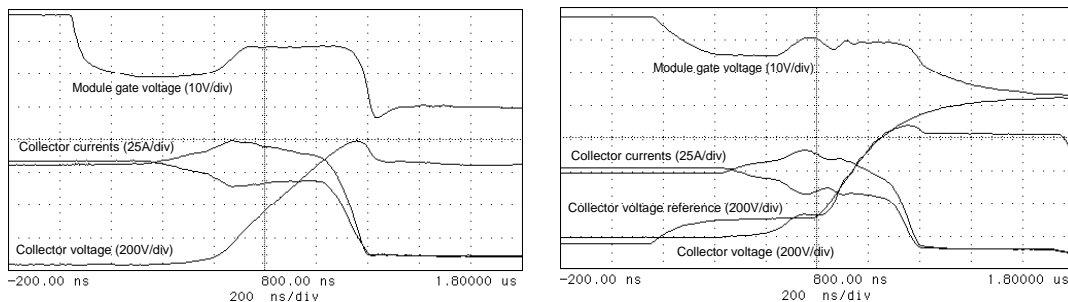


Figure 37. The same module operated with an 'active snubber' type gate drive (left) and a closed-loop collector voltage control gate drive (right). In both cases, the current imbalance lasts for much longer, and the currents do not converge during the current fall. Consequently the difference in switching losses is much greater.

Baseline Modules

A number of modules were assembled with adjacent chips from the centre of a single wafer; this should yield the closest matched chips from any given batch, unless extensive testing and selective assembly are undertaken. It is therefore expected that the transient and steady-state performance should be approaching the optimum achievable.




Figure 38. Initial measurements with a 'baseline' module with a normal hard switched gate drive
 As can be seen, the on-state sharing is good, as is the turn-off; however there is substantial redistribution during the time prior to turn-off. The chip gates do not appear to be diverging substantially at this point; linking the gates together with a short connection suggests otherwise, as shown below in figure 39. The redistribution is virtually eliminated.

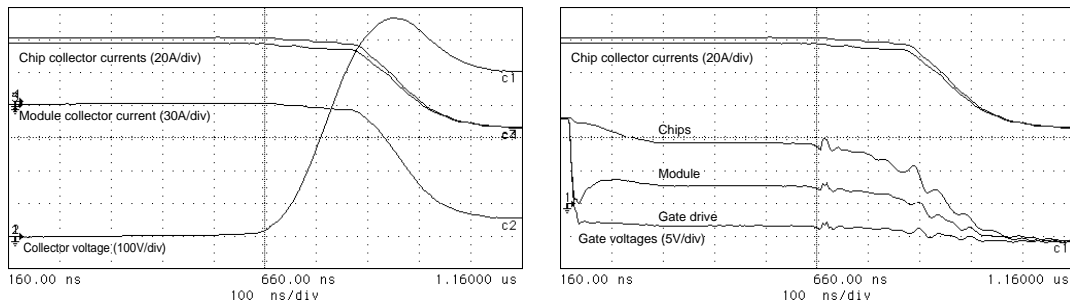


Figure 39. Chip gates tied; the redistribution has gone.

Since tying the two chip gates together virtually eliminates redistribution during turn-off in this case, the transfer characteristics of the two chips must therefore be the same, at least under these conditions; the behaviour when operated in parallel is therefore unexpected. A number of reasons for this could be envisaged, such as differing input capacitances. However such differences would probably be coupled with differences in the transfer characteristics due to the design of the IGBT's top structure.

The silicon gate resistors were then replaced with standard 1% metal film resistors; the nearest common value being 6.2Ω. A pair were selected with minimal difference, and the resulting waveforms are shown below. In this case at least, the use of identical gate resistors has removed the problem. It follows that the module's original gate resistors were not identical, and this was confirmed by measurement. Further investigation of the 12 resistors in this batch of modules showed a mean of 6.41Ω and standard deviation of 0.18Ω, giving nominal values of 6.41Ω±5.6%.




Figure 40. The baseline module in figure 38 above, but with the gate resistors replaced with two 6.2Ω metal film resistors. Current sharing during turn-off is much improved.

Gate Resistors

As noted earlier, the individual chip gate resistors are often made from silicon. This is a poor choice given their substantial positive temperature coefficient of resistance, as shown below in figure 41. A hotter chip (or module) will thus have a larger gate resistor, producing slower switching and higher switching losses, as seen earlier in figure 33, and so thermal runaway is possible. There is also the question of doping variations, with consequent variations in the (cold) resistance of the resistor chips.




Figure 41. Resistance of a gate resistor varying with temperature over a typical working range. The resistance approximately doubles between 25°C and 125°C.

From the experiments so far, it appears that the gate circuit is responsible for the turn-off problems. However, it is not clear whether this is the whole picture; i.e. if two chips matched on $V_{GE(th)}$ and $V_{CE(sat)}$ are used together with matched resistors, is redistribution still a potential problem? For an answer to that

question, we turn to another of the baseline modules supplied for this work. Again the module is switched with the gates both tied and untied, and once more the problem is reduced by tying the gates together, but in this case the problem still persists.

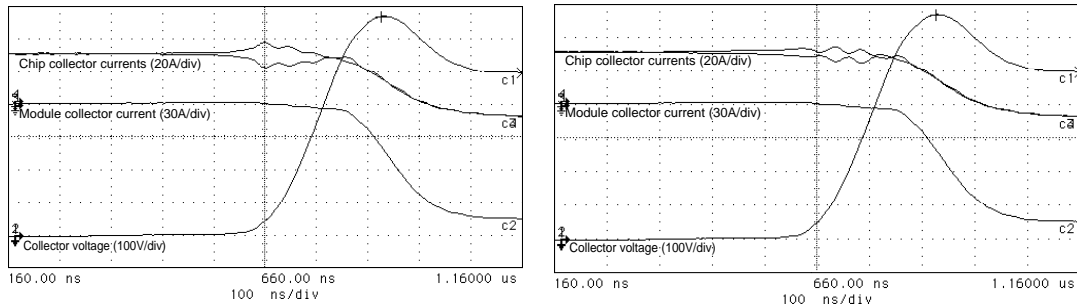


Figure 42. A different baseline module, with gates as supplied (left) and tied (right)

Thus neither matching chips on existing parameters, nor choosing chips from the same area of the same wafer, is sufficient to reliably provide ideal results under parallel operation. It therefore follows that some investigation will be required as to what might change between chips, and yet fail to be shown up by the existing static measurements. Combining the two aspects of selection by matching chips from a single wafer may produce better results, but insufficient samples were available to explore this possibility. In any case, such a solution would prove inconvenient to a manufacturer assembling modules as it could lead to large numbers of chips awaiting suitably matching partners.

Summary

Gate topology, and in particular the use of one gate resistor per chip as traditional for paralleled power MOSFETs, is seen to be a major factor in current redistribution. Gate drive techniques which rely on the IGBT gate's ability to control switching substantially worsen any existing transient current sharing problems.

Whilst well-matched chips with well-matched individual resistors usually behave well, the use of one resistor per chip allows chip differences to be reflected in their switching behaviour in parallel. The tolerances of the resistors currently used appear to be substantially worse than would be expected, and their temperature dependence is both large and undesirable; both of these problems being straightforward to solve using better materials.

The question of whether these resistors are actually necessary is raised, and no satisfactory answer can be given at this point. It seems from the experimental work here that they may not be required in the case of the hard-switched IGBT; the inherently slower nature of the IGBT may mean that the oscillations seen in the paralleled MOSFETs either do not occur at all, or do not become large enough to be problematic.

Redistribution due to chip parameter variations is also investigated, with simple practical methods for measuring the parameters quoted by manufacturers developed. It proves to be difficult to relate measured parameter variations to their consequences, and it may be that the static measurements are insufficient to characterise an individual chip in order to make predictions of behaviour when switching in parallel. It is possible to define a plethora of parameters which can be measured and compared; however, it is unclear which ones will yield useful information about parallel operation.

5

Stability Analysis

One particular aspect of the parallel operation of two IGBTs which requires study is whether such a system is prone to differential oscillations between the IGBTs. Such oscillations have been noted, primarily during turn-off at lower voltages and currents, as shown in figure 43. To ascertain the liability of paralleled IGBTs to such oscillations, and the features of module design which are likely to influence this, the stability of the IGBT is investigated. Although the analysis as stated here pertains to parallel operation of chips in a module, a similar analysis of parallel operation of modules, or indeed of cells within a chip, could be made by adapting the model.




Figure 43. Differential oscillations seen whilst switching a module at a much lower current than its rating.

The Small-Signal IGBT Model for Saturated Region Operation

The stability of the parallel operation of MOSFETs was analysed by Kassakian and Lau [6]. The IGBT can be successfully modelled for this purpose by adapting the basic small-signal MOSFET model: the addition of a parallel 'Output' capacitance C_O as shown in figure 44 accounts for the stored charge in the wide n base region [97].




Figure 44. The small-signal equivalent circuit of a single IGBT

A further difference in the case of the IGBT chips in the modules used in this investigation is the presence of the ‘Kelvin’ emitter contact, so it might be assumed that the emitter inductance coupled into the gate circuit is zero. However each chip only has one Kelvin contact per chip, so some parts of the chip will see a significant emitter impedance, and a suitably representative value is required for a single ‘lumped’ emitter impedance. The output capacitance C_O appears in parallel with C_{CE} , and so is incorporated into C_{CE} . This leads to the model shown in figure 45 for a pair of IGBTs in parallel, switching a resistive load. This is a more general case than the two situations considered previously by Kassakian & Lau [76], but the results in that investigation can be arrived at by disregarding some of the elements incorporated here as appropriate. The model is only valid for operation in the saturation region, i.e. at high collector voltage; in the linear region, the behaviour of the IGBT is quite different.




Figure 45. Two IGBTs in parallel, with parasitics and resistive load.

The concern here is with differential current oscillations, as it is redistribution of current between parallel IGBTs which is of interest. Consequently, a number of components in this system can be dispensed with, as they will play no part in any differential oscillations, namely V_G , R_L and V_{CC} ; their terminals can be considered as small-signal ground. This leaves a

circuit which is completely symmetrical, so it can be split into its two halves and one half considered on its own (figure 46); the resulting circuit analysis is relatively straightforward. If one of these half circuits is oscillatory, the other will be too, with the same natural frequency. Furthermore, they will be prone to oscillate at that frequency, the two oscillating in antiphase.




Figure 46. Differential mode circuit of a single IGBT for analysis

Characteristic Equation

The characteristic equation of this system can be derived and analysed for stability. The mathematics of this are covered in Appendix B; here it suffices to state the characteristic equation. For convenience a number of constants are defined at this point:

$$C_G = C_{GC} + C_{GE} \quad (5.1)$$

$$C_C = C_{GC} + C_{CE} \quad (5.2)$$

$$C_E = C_{CE} + C_{GE} \quad (5.3)$$

$$C_t = C_{GC}C_{GE} + C_{GC}C_{CE} + C_{GE}C_{CE} \quad (5.4)$$

The characteristic equation is then:

$$\begin{aligned} & s^4 C_t (L_G L_C + L_G L_E + L_C L_E) \\ & + s^3 (C_t (L_G R_C + R_G L_C + R_G L_E + R_E L_G + R_C L_E + R_E L_C) \\ & + (L_E/R_O)(L_C C_{GC} + L_G C_{GE} + L_G C_{GC}) + (L_C/R_O)(C_{GC} L_G + C_{GE} L_G + C_{GE} L_E) \\ & + g_m C_{GC} (L_E L_G + L_C L_G + L_C L_E)) \\ & + s^2 (C_t (R_G R_C + R_G R_E + R_C R_E) + (L_G C_G + L_E C_E + L_C C_C) \\ & + (L_E/R_O)(C_{GC} R_C + C_{GC} R_G + C_{GE} R_G) + (L_C/R_O)(C_{GC} R_G + C_{GE} R_G + C_{GE} R_E) \\ & + (R_E/R_O)(C_{GC} L_G + C_{GC} L_C + C_{GE} L_G) + (R_C/R_O)(C_{GC} L_G + C_{GE} L_G + C_{GE} L_E) \end{aligned}$$

$$\begin{aligned}
& + g_m C_{GC} (L_E R_G + L_G R_E + L_E R_C + R_E L_C + L_G R_C + L_C R_G) \\
& + s (R_G C_G + R_E C_E + R_C C_C + g_m L_E + (L_C + L_E) / R_O) \\
& + (R_E / R_O) (C_{GC} R_C + C_{GC} R_G + C_{GE} R_G) + (R_C / R_O) (C_{GC} R_G + C_{GE} R_G + C_{GE} R_E) \\
& + g_m C_{GC} (R_G R_E + R_G R_C + R_C R_E) \\
& + g_m R_E + (R_C + R_E) / R_O + 1 = 0 \tag{5.5}
\end{aligned}$$

The characteristic equation is thus a 4th order polynomial of the form

$$a_4 s^4 + a_3 s^3 + a_2 s^2 + a_1 s + a_0 = 0 \tag{5.6}$$

The Routh-Hurwitz criterion can then be used to analyze this polynomial for stability, the criteria being

$$a_1 a_2 a_3 > a_0 a_3^2 - a_1^2 a_4 \tag{5.7}$$

$$a_i > 0 \tag{5.8}$$

By finding the value of R_G which produces equality in equation (5.7), we find the minimum R_G required for stability for a given set of device and parasitic impedance parameters. This is straightforward to achieve with a mathematical package such as MathCAD [99]; the criteria in (5.8) must also be checked. These are trivially true provided that all of the circuit parameters are positive. The question of multiple solutions to (5.7) must be considered: a_0 through a_4 are all functions of R_G , so solving (5.7) could yield up to three values of R_G . Obviously the largest value of R_G satisfying this equality is the minimum value above which stability can be assured.

The usefulness of this analysis is greatly enhanced by the ability to investigate how this minimum value of R_G changes as parameters change. Some of the IGBT's electrical characteristics, particularly C_{GC} and C_{CE} , are heavily dependent on the collector voltage; meanwhile, g_m and C_{CE} are dependent on collector current, so it will be necessary to find the minimum value of R_G over the whole range of conditions to assure stability of a given module design. Design changes in both the module and the chip can also be investigated for their impact on stability.

Determining Model Parameters

Module Parasitic Impedances

A set of parameters was estimated for two 1800V 75A chips operating in a half-bridge module. Resistance R_C was measured by passing a known current through the module and measuring voltages across the various parasitic elements. R_C was calculated to be of the order of $2\text{m}\Omega$; R_E , as noted earlier, is difficult to define, and was estimated to be $0.2\text{m}\Omega$.

The inductance L_C , representing the differential collector inductance and the differential emitter impedance not coupled into the gate loop, was measured during switching transients using $V=LdI/dt$ to be approximately 20nH ; L_G was estimated as 20nH visually (this includes the inductance of the emitter Kelvin connection which is not within the power circuit), and L_E was estimated to be 2nH .

Device Parameters

The parameters g_m and R_O are easily found from data sheet or measurement data. For the 1800V chip used in the later modelling exercises, g_m is approximately 80 at rated current (75A), falling away to 10 at low current (10A). R_O can be determined from the slope of the I_C vs V_{CE} graph at a particular gate voltage: at rated current and moderate collector voltages, R_O is calculated to be approximately 100Ω .

The terminal capacitances of a module are usually given on manufacturers' datasheets as functions of collector-emitter voltage, and the chip capacitances can be deduced from them. However, it is standard practice to quote these at zero gate voltage, and hence zero collector current, for a single test frequency (usually 1MHz). Stability investigations using these values give values of R_G which are unrealistically small. However, at turn-off, the capacitances C_{CE} and C_{GC} are much larger due to the charge stored in the device's wide drift region during the conduction period.

Silvaco's device simulation suite [100] offers two methods for determining the device capacitances. The simpler (and much quicker) approach is to use a SPICE-like small signal AC analysis, which is described in the manual as being

based on a linearisation of the device model, but little else is given by way of description. Qualitatively the results are plausible for the zero collector current situation, but less so for the situation here. The alternative is to perform a full transient simulation containing the IGBT. Two DC voltage sources (gate and collector) bias the device to the operating point of interest, and then a small AC source is applied in series with one of its terminals. By examining the resulting current in the AC source, the capacitance at that terminal can be deduced.




Figure 47. IGBT circuit with bias and AC source (left) and small-signal equivalent circuit (right) to determine C_{ISS}




Figure 48. IGBT circuit with bias and AC source (left) and small-signal equivalent circuit (right) to determine C_{OSS}




Figure 49. IGBT circuit with bias and AC source (left) and small-signal equivalent circuit (right) to determine C_{RSS}

This approach yields C_{ISS} and C_{OSS} , as shown in figures 47 and 48, although some care is needed with the latter when measuring I_C as some of the collector current flows through R_O . C_E presents a problem (figure 49): the small-signal voltage V_{AC} not only produces a current in C_{GE} and C_{CE} as desired, but also changes the current source by $g_m V_{AC}$. Consequently, examining the change in emitter current in response to the small-signal AC source will not yield the value for $(C_{GE} + C_{CE})$ required due to the influence of the current source. The problem

does not arise in the case of the C_{ISS} measurement, as the current source is shorted; for the C_{OSS} measurement, the change in V_{GE} is zero. However, by examining figure 48 more closely, it can be seen that consideration of I_G and I_E separately enables C_{GC} and C_{CE} to be distinguished.

Hence the three measurements available are C_{GC} , C_{CE} and C_{ISS} , from which the third device capacitance C_{GE} can be deduced. C_{GC} and C_{CE} in particular are seen to be heavily voltage dependent; furthermore, both are substantially larger when the IGBT is active, as noted earlier.

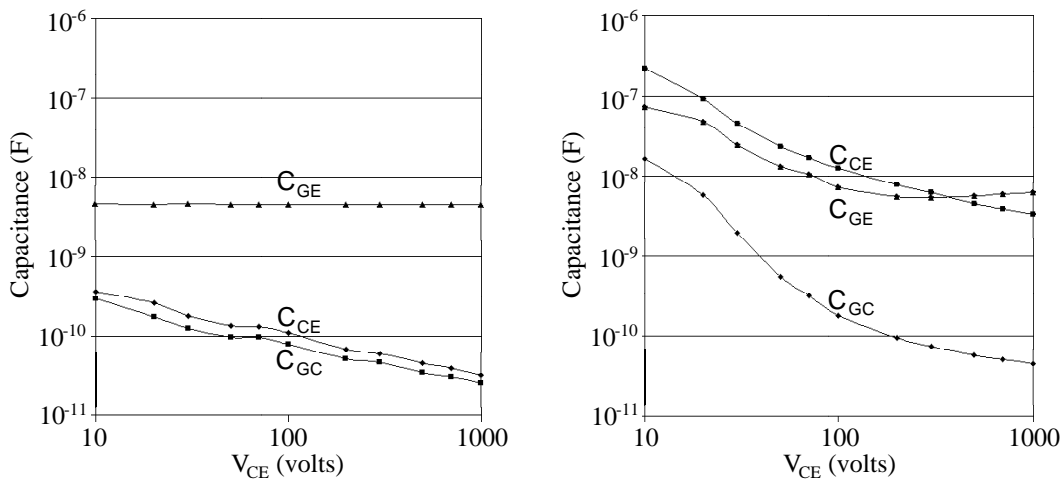


Figure 50. Device capacitances as a function of collector voltage, with $V_{GE}=0V$ (left) and $V_{GE}=7V$ (right)

As expected from manufacturers' datasheets, at zero gate voltage, C_{GE} is approximately constant over the range of collector voltages; C_{CE} and C_{GC} are of similar magnitude to each other, and voltage dependent. The value of the ratio $C_{GC} : C_{CE}$ is approximately constant at around 0.75.

With a gate voltage of 7V, and corresponding collector current of approximately the chip's rated value (75A), a different picture emerges. All three capacitances are both voltage-dependent, and much larger. C_{GC} is seen to have risen to over 10nF, and this is responsible for the slow voltage rise during the early stages of device turn-off. C_{GE} has risen to around 75nF at low collector voltage, and C_{CE} is a massive 220nF - very different from the sub 1nF values which might be expected from the manufacturers' data.

Results

With so many potential variables, it is essential to fix all but one or two, which can then be varied, producing a curve or family of curves of R_{Gmin} .

Ultimately the aim is to map $R_{G_{min}}$ as a function of many other variables, but mapping the whole spread of possible circumstances is clearly impractical. The approach taken to the problem here is to choose a point of interest, and investigate the area around that point; for example, one can see the effect of altering stray emitter impedance in an existing module design. An alternative approach is to look at quite different points of interest, changing many variables at once; in this case one could investigate the consequences of moving from a plastic module to a 'press-pack' type of construction, for example.

Values used for an initial calculation were $R_C=2m\Omega$; $R_E=0.2m\Omega$; $g_m=80$; $L_C=20nH$; $L_E=2nH$; $L_G=20nH$; $C_{GC}=95pF$; $C_{CE}=7.8nF$; $C_{GE}=5.5nF$; $R_O=100\Omega$. The capacitances were chosen at $V_{CE}=200V$, $I_C=75A$ as representative values which would be expected during the switching transition. Inserting these values into the characteristic equation gives a baseline value for R_G of 2.18Ω , which is not incompatible with the 6Ω to 7Ω used in practice.

Variation in Collector Voltage

As shown in figure 50, the IGBT's terminal capacitances change significantly with collector voltage. The consequences of this for the required gate resistor are shown in figure 51; the gate voltage is constant, and hence the collector current is approximately constant at the IGBT's rated level. It can be seen that the required gate resistance increases from almost zero at low collector voltage to 4Ω at $1000V$.




Figure 51. Minimum value of R_G required at different collector voltages.

Variations in Device Parameters

The characteristic equation derived above can be compared with the two special cases in [76] to see the effects of neglecting some of the parasitic elements. C_{CE} and C_{GC} are used as the two variables, with the other parameters being fixed at the values given above. It can be seen from the figures (below) that the two special cases do not give good results here. In the first special case in [76], it is assumed that $R_C = R_E = L_E = 0$, and in this case C_{CE} has little effect on stability. The values of R_G obtained are not unreasonable. In the second special case, it is assumed that $L_C = R_E = 0$, and it is seen that for small C_{CE} and larger values of C_{GC} , no gate resistor at all is needed. The third calculation includes all of the stray impedances, and it can be seen that the values of R_G are significantly different to the other two cases.

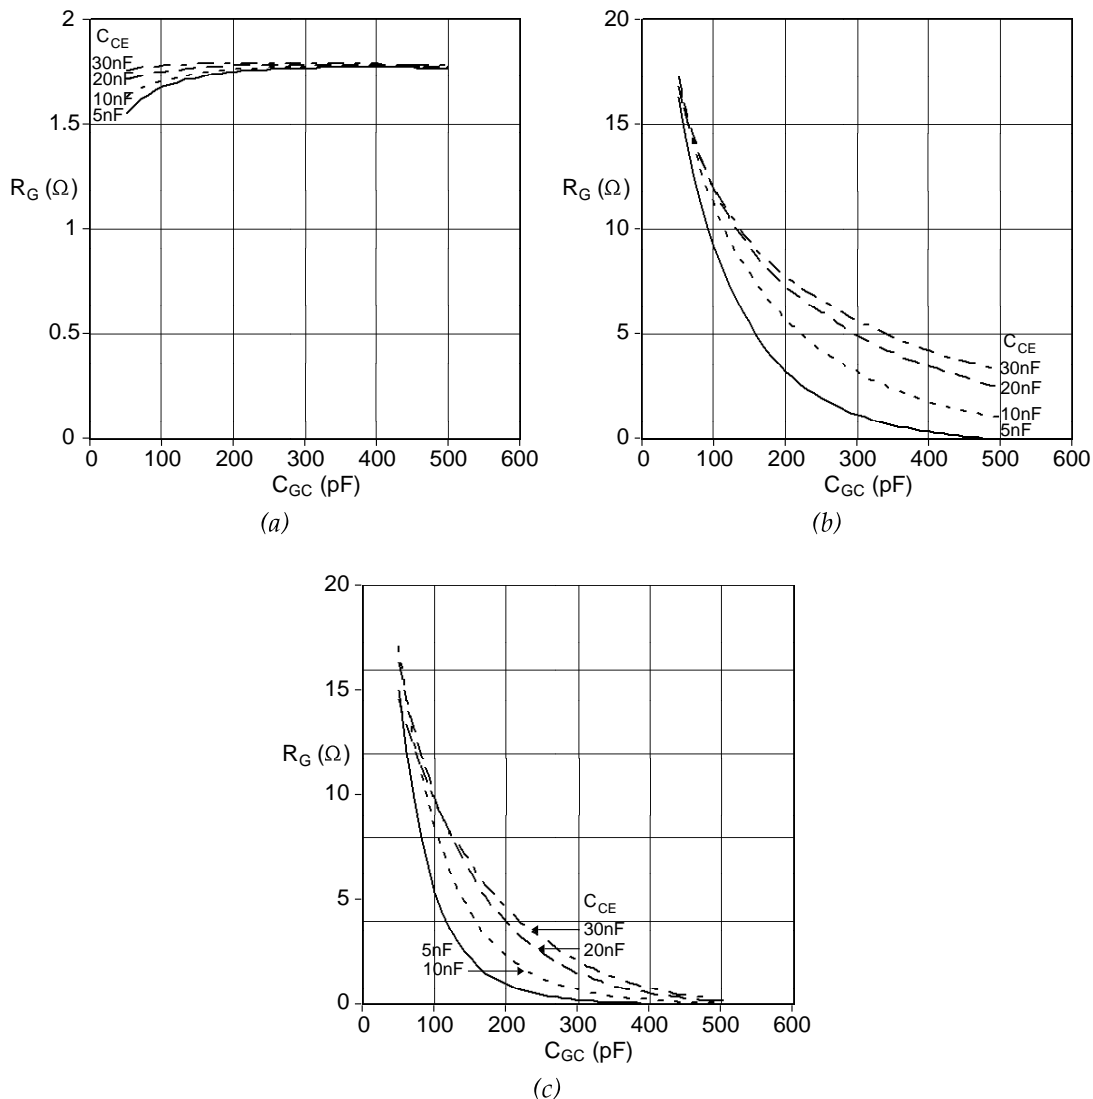


Figure 52. Minimum value of R_G required for different values of C_{CE} with varying C_{GC} calculated with Kassakian & Lau's simplified models (a) and (b), and compared with full model (c)

The effects of varying C_{GE} and $g_{m'}$ two more parameters which are dependent on the IGBT itself rather than on their electrical interconnections, are shown in figure 53. Once again, the three models are compared, and again the simplest model is significantly different from the other two in both form and magnitude. The second gives approximately the right form, but the magnitude of R_G is rather large. The models of Kassakian & Lau are thus insufficiently detailed for use in the IGBT analysis undertaken here.

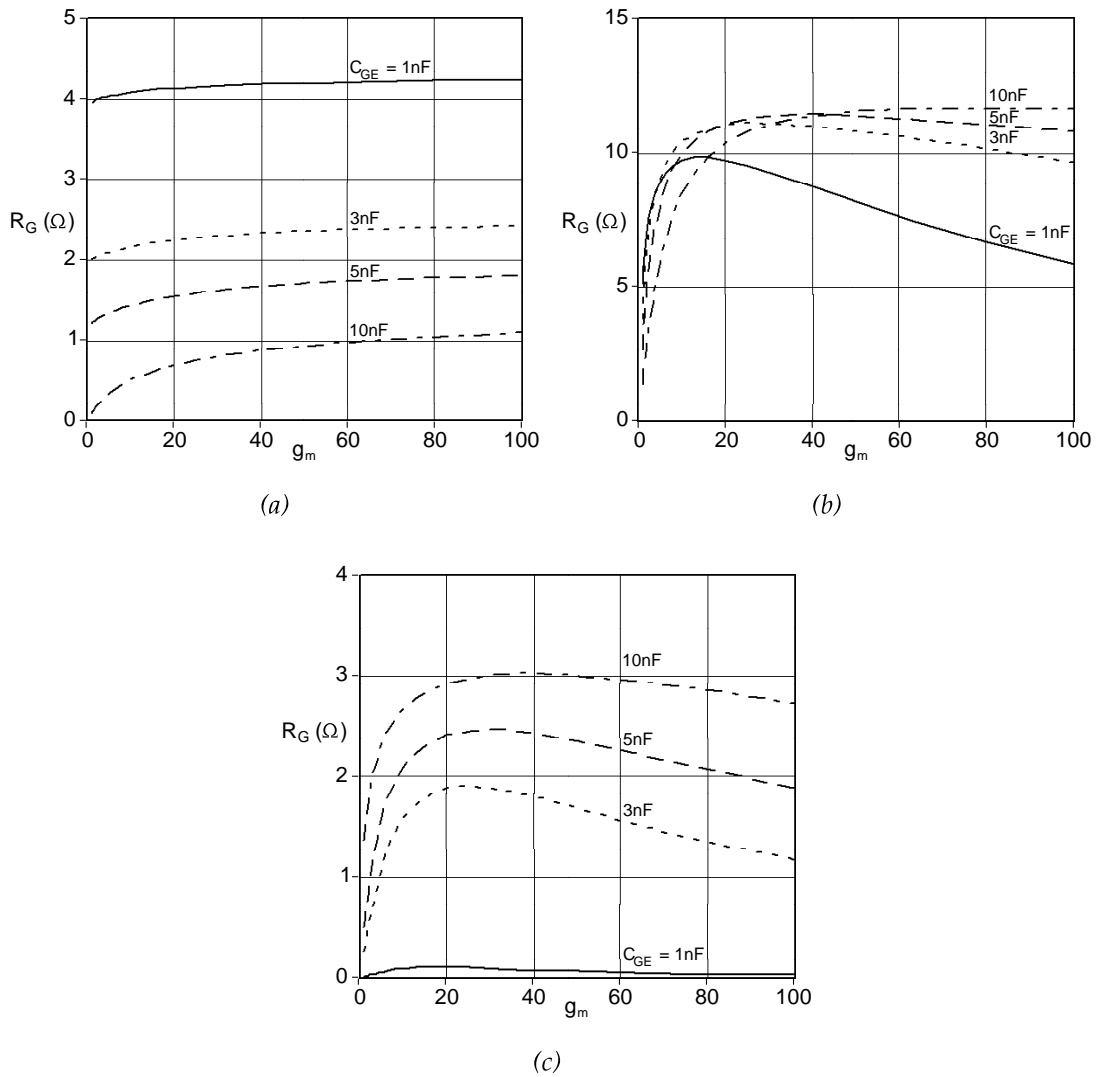


Figure 53. Minimum value of R_G required for different values of C_{GE} with varying $g_{m'}$ calculated with Kassakian & Lau's simplified models (a) and (b), and compared with full model (c)

It appears that, in terms of the IGBT chips themselves, large values of C_{GC} and small values of C_{CE} and C_{GE} are desirable for stability. It is also seen that the effect of variations in $g_{m'}$ is relatively smaller over most of the working current range being considered here.

Variations in Stray Impedances

The effects of variations in four of the stray impedances are shown in figure 54. Emitter inductance in the gate loop is clearly undesirable; for the parallel IGBT chip situation, even a few nanohenries has a significant impact on stability. Similarly, inductance in the gate connection is unwanted. On the other hand, R_E (and R_C , not shown here) can be neglected in this topology. It follows that module design can have a considerable effect on stability, and more importantly, that some aspects of module design demand much closer consideration than others.




Figure 54. Minimum value of R_G required for different values of L_E with varying L_G , and different values of R_E with varying L_C .

Alternative Module Topologies

Whilst it has been seen which of the device and module parameters are important for stability for the plastic module type package, it is also instructive to consider the 'press pack' type of module in the same way. Although no test information is available, some representative values can be tried to gain some insight into the likely implications of a move to this type of package, assuming that similar chips were used. Initial estimates were $L_E=5\text{nH}$, $L_C=2\text{nH}$, $L_G=10\text{nH}$; $R_E=0.5\text{m}\Omega$, $R_C=0.2\text{m}\Omega$. It appears that, given these values, press-pack style devices require a substantial gate resistor of around 7.5Ω . This indicates that a thorough analysis of the stray impedances of press-pack devices is necessary.

Stability Analysis With Circuit Simulation

It has been shown by analysis of the characteristic equation of the IGBT model that both chip characteristics and module topology have a significant influence on the choice of a gate resistor to assure stability of a particular IGBT module. As noted previously, this analysis of the characteristic equation only provides the minimum value of gate resistor required for stability given a set of parameters, and how this will change with changes in those parameters. It does not, however, reveal anything about the features of a stable system. For this it is necessary to investigate the IGBT model in more detail; a circuit simulator such as PSPICE [98] provides a convenient method for doing this. Furthermore, the complexity of analysing a circuit to find its characteristic equation is avoided, so (for example) parallel devices can be examined with a degree of asymmetry in their interconnection, a situation which would be unwieldy to analyse mathematically. Naturally, the limits of validity of the small-signal IGBT model itself still apply.

A closed-loop circuit, such as that shown in figure 55(a), can be analysed for stability by opening its feedback path and investigating the transfer function of the resulting open-loop circuit, as shown in figure 55(b). In the case of the equivalent circuit of the IGBT, the resulting open-loop circuit is shown in figure 56. The same parameter values are used as for the initial calculation with the characteristic equation above.





Figure 55. (a) Closed-loop circuit



(b) Loop opened and AC signal source added

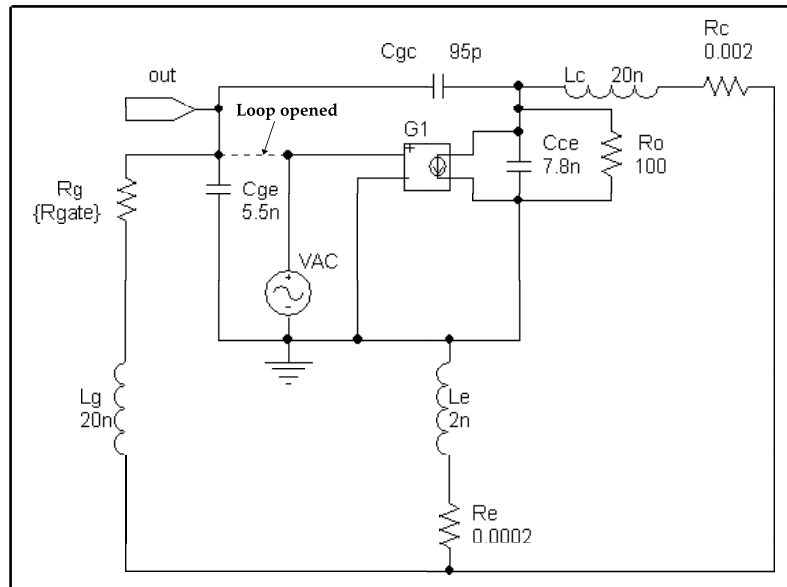


Figure 56. Single IGBT equivalent circuit, as used in the MathCAD analysis, but split for analysis of transfer function.

With PSPICE it is straightforward to perform the required analysis, sweeping the frequency of the AC source over a range for each of a number of values of R_G . This yields a family of Bode plots, shown in figure 57. Three distinct resonances are shown, produced by two poles and a zero. At approximately 12MHz, C_{CE} resonates with L_C and L_E ; increasing R_G has little effect on this resonance, as the gate circuit is not involved. C_{GE} resonates with L_E and L_G at a slightly higher frequency of approximately 15MHz, and R_G is seen to damp this resonance. Both of these frequencies can be calculated from $\omega=(LC)^{-1/2}$, agreeing well with observations from the Bode plot. The third feature, the zero at 33MHz, is less straightforward: C_{GC} , L_C , L_G and L_E are all involved, and again it is damped by R_G . Each of these three resonances contributes a 180° phase shift between V_{in} and V_{out} , whilst R_E and L_E contribute 90° at 16kHz.




Figure 57. Bode plot of the response of the circuit in figure 56 for different values of R_G .

However, the Bode plot provides insufficient information to determine whether the circuit is stable in this case. The Nyquist diagram must be drawn, and then the stability of the circuit can then be determined from the trajectory's encirclements of the +1 point. The Nyquist diagram corresponding to figure 57 is shown in figure 58; below a certain value of R_G , the circuit will be unstable.





Figure 58. Nyquist diagram corresponding to figure 57 (above) and close-up showing +1 point (below)



It is useful to relate this result back to the Bode plot, which is much more straightforward to produce and interpret in relation to the circuit. In this particular case, the part of the trajectory which could encircle the +1 point on the Nyquist diagram corresponds to frequencies between approximately 15MHz and 35MHz: the crossings of the 360° line on the Bode diagram are

equivalently the crossings of the real axis on the Nyquist diagram. Consequently, by investigating the loop gain at these two 360° crossings on the Bode diagram, the required information can be deduced. If the loop gain is greater than unity at both these points, or less than unity at both of these points, or it does not cross the 360° line at all, the +1 point is not encircled, and so the circuit is stable; otherwise, the circuit is unstable.

Returning to figure 57, the circuit is clearly stable for $R_G=3.2\Omega$ (the 360° line is not crossed), and unstable for $R_G=1.6\Omega$ (the 360° line is crossed once with gain of 20dB and once with a gain of -6dB), so the critical value is somewhere between these values. Refining the selection of R_G values leads to the conclusion that the critical value of R_G is 2.18Ω . This agrees with the baseline value calculated with MathCAD using the Routh-Hurwitz stability criterion. It is also possible to plot Nyquist diagrams using MathCAD as the transfer function can be found in this single-chip case, and once again the same value for R_G is obtained.

It is interesting to note that for $2.18\Omega < R_G < 2.5\Omega$, the Bode diagram indicates two points with zero phase and gain of greater than unity, which could logically be regarded as a sufficient condition to produce oscillation. In this case the circuit remains stable as already noted; this can be confirmed by closing the loop of the circuit in figure 56, applying an AC excitation, and performing a time domain (transient) simulation using PSPICE: exponentially growing oscillations are found only for values of R_G below 2.18Ω . The Nyquist diagram correctly predicts this as shown in figure 59.




Figure 59. Nyquist diagram for $R_G=2.0, 2.2, 2.4$ and 2.6Ω showing +1 point

Stabilization

The Bode plot is a useful aid to understanding the behaviour of the circuit, and provides an insight as to how it may be stabilised without the need to introduce large gate resistors. The three resonances can be moved around within the frequency domain by adjusting the parasitic inductances, which could in reality be achieved by changes in module design. By making suitable choices for these inductances, it is possible to move the zero between the two poles, as shown in figure 60; it can be seen from the phase response that the circuit should now never oscillate. In this case, L_G has been increased to 200nH; increasing L_C to 150nH has a similar effect.




Figure 60. Bode plot of the response of the circuit in figure 56, with $L_G=200\text{nH}$.

Additional collector inductance is a long-standing remedy for oscillation dating from the bipolar transistor [77], and was shown to be beneficial in figure 54(b); however gate inductance appeared to be undesirable in figure 54(a). In the light of the result presented in figure 60, the result of figure 54(a) is re-examined for a much larger range of gate inductance in figure 61. It can be seen that substantially increasing L_G may be useful as suggested in previous work [6], [78].




Figure 61. Minimum value of R_G required for different values of L_E , with varying L_G .

Stabilisation measures which depend on the introduction of substantial extra inductance are not without their drawbacks: increasing L_C will increase the peak collector voltage during switching, and increasing L_G will slow the IGBT's gate control. A substantial reduction of L_C , moving one of the poles beyond the zero to achieve the same result, appears to be a more attractive possibility, but this is impractical.

Another method for improving stability is raising C_{CC} : at 400pF, the circuit is stable for any $R_G > 0.1\Omega$, or for 600pF stability is achieved for any positive R_G . This could be achieved by use of a small external capacitor, as in the active snubber style of gate drive. Increasing C_{CC} in this way will limit the peak rate of collector voltage change achievable during switching, which may be undesirable if the fastest possible switching is required; against this, the reduction in gate drive resistance will reduce the switching delay, and so it is not possible to make a general statement as to whether switching will be faster or slower.

It can be seen that investigation of the open-loop transfer function of the IGBT's equivalent circuit with Bode and Nyquist diagrams yields the minimum R_G for stability for a given set of circuit and device parameters. It is also seen that the value given is the same as that obtained from applying the Routh-Hurwitz stability criterion. Consequently the need to derive the characteristic equation for each circuit variation can be dispensed with, a crucial step towards making investigation of more complicated systems feasible.

Asymmetric IGBT Modules

Having avoided the need to analyse each circuit configuration mathematically, more complex systems may be studied by adding a second IGBT equivalent circuit. Thus two chips within a module can be investigated, and in particular, the consequences of any asymmetry in their connections can be seen. However, one can no longer make all of the simplifying assumptions which led to the investigation of a single chip with no load or power supply, depicted in figure 46. Instead, there are several different situations arising during a switching cycle. For example, during the tail current phase (phase 4 in figure 9), the collector voltage is constant, but the collector current varies. The oscillations

seen in figure 43 start during the voltage rise at turn-off (phase 2 in figure 9), and so this situation is considered further here. There is a constant current flowing through the IGBT module during this phase, and as a current source can be considered as a small-signal open circuit, the power supply branch can be disregarded once again. Paralleling two of the small-signal IGBT models thus leads to the circuit of figure 62. A small imbalance is introduced in the form of different emitter inductances.

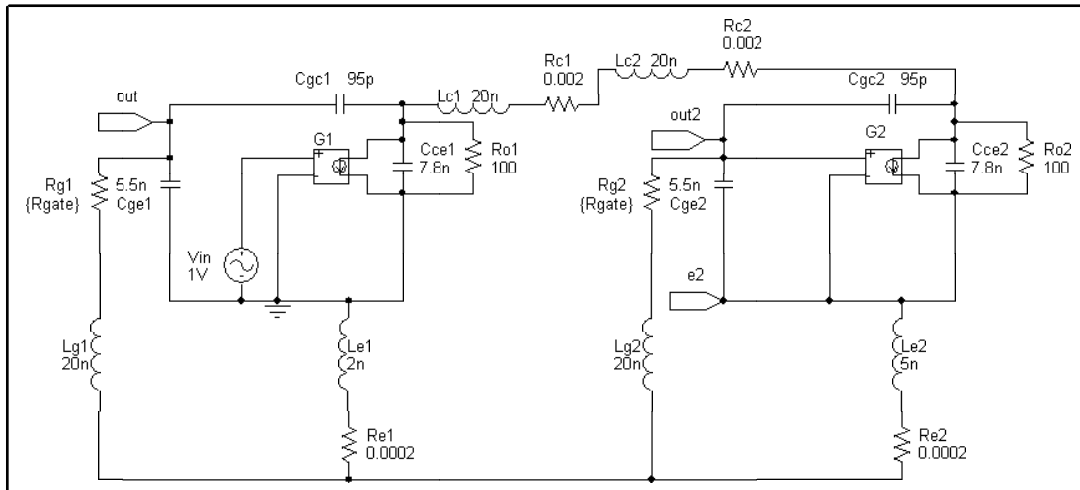


Figure 62. Small-signal model for two parallel IGBTs with uneven emitter inductance

The analysis of this situation is not as straightforward as was the case for figure 56: there is no longer a single oscillator whose feedback loop can be opened. Opening the feedback loop of one oscillator and investigating its frequency response is insufficient to reveal whether the circuit is stable, as shown in figure 63(a). Here the feedback loop of the left-hand IGBT is opened and the response plotted; it would appear that $R_G=3.2\Omega$ produces oscillation, but not $R_G=1.6\Omega$. However, the output of the right-hand IGBT when the left-hand IGBT is excited is shown in figure 63(b): all the values of $R_G<3.2\Omega$ cause the right-hand IGBT to oscillate. Closer investigation reveals that the circuit oscillates in one or other mode for all $R_G<3.5\Omega$, compared with 2.18Ω and 4.51Ω for the two individual oscillators. Opening the loop of the right-hand IGBT instead produces the same result, and this is verified by transient simulations.





Figure 63. Frequency response of the circuit of figure 62 with the left-hand IGBT excited.
 (a) Output of the left-hand IGBT (above)
 (b) Output of the right-hand IGBT (below)



Discussion

From the comparison of the model presented here with the simplified ones used by Kassakian & Lau, it appears that the latter are unsuitable for use in investigating the issues involved in parallel operation of IGBTs in a module. This is due to certain stray elements being disregarded which become important in IGBTs, as the chips are large in comparison to MOSFETs. In particular, disregarding L_E is unwise in this case, as the inductance of the emitter bond wires plays a significant part in the module's behaviour. Furthermore, the importance of the single emitter Kelvin connection has become apparent in the sensitivity of the minimum R_G value to the value of L_E .

Some previous work on parallel operation of MOSFETs [101] recommended the use of ferrite beads around gate connections to avoid oscillations. The wisdom of applying this to IGBTs remains difficult to elicit; in the case studied here, a certain amount of gate inductance significantly increases the minimum gate resistor required, but further increases prove beneficial. The reasons for this can be seen from the movement of the various resonances on the Bode plot of figure 57. Collector inductance (which includes emitter inductance outside of the gate loop) does appear to be beneficial for stability. However the choice of suitable values for these parameters to achieve maximum stability is complicated, as the IGBT's full range of operating conditions must be considered. It is unlikely that there is a set of parameters which can guarantee stability without a gate resistor over all conditions.

Perhaps the most important point is that the stability analysis tells us little about the transient performance of parallel IGBTs, other than the possibility of oscillation. Under fast switching conditions, it is possible that oscillations will not appear, because a finite time is needed for them to become established. It is also crucial to realise that conditions during turn-off are radically different from those during turn-on due to the charge built up in the device, and, as has been seen, due consideration must be given to simulating or estimating the device's capacitances.

The analysis presented provides a grasp of which elements are important in oscillations, and as such is important in looking at current sharing. It has also highlighted the limitations of the modelling work done previously.



Device Simulation

Introduction

The practical investigation in the previous chapters suggests that reviews of device design and characterization are necessary, with three potentially useful outcomes sought. Firstly, improvements in the IGBT design or manufacturing process which mitigate the current sharing problems would be a desirable result; this would facilitate better module ratings without complicating the chip selection process. Secondly, improved characterization of the chips to enable better matching for selective assembly is possible, if less desirable, as it may lead to large numbers of chips being held 'in stock' awaiting suitably matched partners to make up large modules. Thirdly, it may be possible to improve module design so that the current sharing problems can be reduced without needing to change the IGBT chip design or manufacture. This could lead to improvements at minimal extra cost.

Improving the IGBT's design by production of a number of wafer lots with subtle differences, and then investigating the behaviour of each of these, is not a practical way forward in the absence of any indication of which aspects of device design are likely to be important. Not only is the fabrication of a suitable range of wafers liable to be expensive and to have a long lead time, but producing the closely controlled experimental conditions required to achieve meaningful results when testing the devices produced is not straightforward and may be impossible. Consequently, attention turns instead to computer simulation of IGBTs operating in parallel. Here a single parameter can be changed, and its effects on behaviour investigated in a tightly controlled environment. Investigation of identical devices with different stray impedances, or different devices with identical stray impedances, becomes possible. Obviously these results must ultimately be related back to module and chip design to derive any benefit, but simulation should at least provide a starting direction for investigation.

Device Simulation Software

The work carried out here was undertaken with a physics-based numerical simulation suite produced by Silvaco [100]. The top (emitter) section of a typical IGBT structure file is shown below in figure 64. The solid grey areas are the device terminals, and the graduated grey shading in the area below them represents the silicon and its net doping; more highly doped areas are lighter shades. The solid black lines within the silicon indicate the positions of the junctions J2 and J3 shown in figure 1, and the white lines connect the mesh points at which the physical equations governing the device's behaviour are solved by the simulator.




Figure 64. The emitter end of a typical IGBT model

Suitably meshed semiconductor devices can be investigated in isolation, using the two-dimensional device simulator (ATLAS) alone, and their static internal behaviour and resulting terminal characteristics found. However, of more use for this investigation is MIXEDMODE, a SPICE-like time and frequency domain circuit simulator, which additionally allows the user to incorporate two-dimensional devices simulated by ATLAS. This method is very suitable for investigation of current sharing issues in IGBT modules, as the circuit can include both the interconnections within a module, and a realistic external circuit. Changes in either the circuit, or one of the IGBTs, can be investigated.

For the sake of computation time, only two devices are used; a typical simulation run then takes of the order of five hours on the available computing hardware, unless oscillations develop.

This investigation is primarily aimed at externally observable consequences of changes in device and circuit details. Considering the case of a circuit change (different gate resistors, for example), the different power losses in the two chips can be directly attributed to that change in the external circuit. Most details of the device's internal behaviour are of relatively little direct interest; however, some may indicate the areas to consider if the device design is to be altered to accommodate the different circuit conditions. Similarly, in the case of a device parameter change (different gate oxide thickness, for example), it is not feasible to directly measure that parameter for each chip before assembly into a module; one has to rely on electrical characteristics, such as on-state voltage and threshold voltage. Here it is clearly more desirable to relate the different power losses to these electrical characteristics.

The model of the IGBT itself can be produced in two distinct ways. The first way is to use ATHENA, another component of the software suite, which allows the simulation of the production process. An IGBT model is yielded from details of the various processes applied to a bare silicon wafer. The second alternative is to use another program, DEVEDIT, to directly specify the dimensions of the device, along with its doping profiles. The latter approach is adopted here, using a model produced by the manufacturers for their 1800V, 75A IGBT chips; these had been supplied in the 'baseline' modules used for much of the practical work. The p-base doping profile was measured, and combined with knowledge of the device dimensions and processing steps to produce the model used. Although this may yield a model which does not exactly match the devices used in the tests, it is effects obtained by parallel operation which are of most interest here, and it is anticipated that any reasonably accurate IGBT model should show similar effects. Previous work with an earlier 1600V 75A chip model showed qualitatively similar effects, although numerical measurements were only undertaken with the later model.

Simulation Circuit Design

On inspection of a typical four-chip half-bridge module, it appears that there is the following arrangement for each chip (below). Obviously some of these elements are equivalent and can be combined together to simplify the circuit; for example, impedances in the collector and emitter terminals can be combined, where they are not also coupled into the gate circuit. The parasitic capacitances between the tracks on the tile and the heatsink are ignored, as they are much smaller (typically 10pF) than the IGBT chip's own capacitances. The IGBTs' anti-parallel diodes are neglected, as they will never enter conduction during the boost converter's operation.

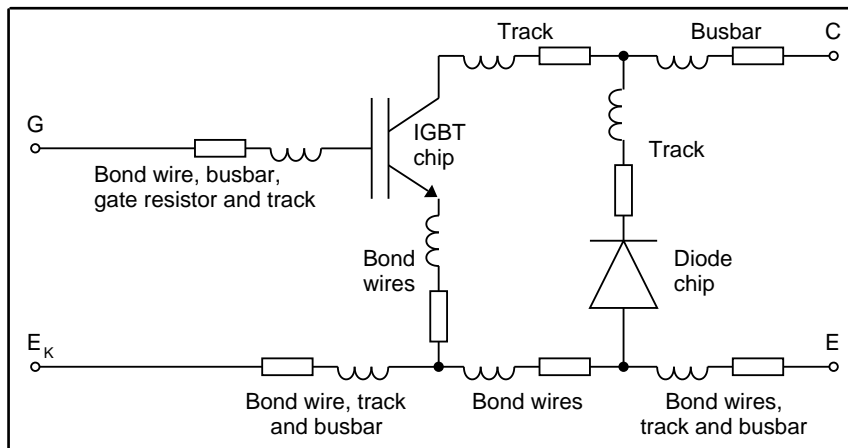


Figure 65. Schematic diagram of one IGBT chip and its associated parasitic elements

Two of these models can then be paralleled by connecting the relevant terminals together, and additionally considering parasitic elements common to both chips. Adding in the remaining components of the boost converter used for the experimental work produces the model shown below in figure 66. The power circuit is slightly amended from that used for experimental work, which was shown in figure 15; the DC supply and its associated large smoothing capacitors and inductor have been replaced by a single current source I_{DC} , and the output resistor / capacitor combination replaced with a voltage source. These amendments provide much more convenient control of the simulation's conditions; for example, the DC rail voltage can be altered with one simple change, whereas the experimental rig required some alterations of switching frequency, pulse width, supply voltage and load resistance to achieve the same.

If two of the chip models depicted in figure 65 are connected in parallel

without further alteration, the emitter Kelvin connections could carry a redistribution current in the event of uneven current sharing. Depending on the relative impedances of the main emitter connection and the Kelvin connection, some of the load current could flow out of the Kelvin connection from one chip, back through the other, and then leave via the second chip's emitter bond wires, an issue raised in [102]. However, the circuit topology depicted in figure 66 was adopted as sufficiently sophisticated for the required task.




Figure 66. The circuit used for simulations with the Silvaco software suite

All of the inductances and resistances, with the exception of L_{GG} , L_{DC} , R_{GG} , R_{G1} and R_{G2} , represent stray impedances within the module and its associated power circuit. R_{G1} and R_{G2} are the two gate resistors within the module, and R_{GG} is the output resistance of the gate drive. V_{GG} is the voltage at the gate terminal of the module. The remainder of the resistances and inductances represent the bond wires and busbars within the module. The lack of resistance in the collector circuit is due to the voltages developed across them being negligible compared with the DC link voltage; the emitter resistances are retained as they are coupled into the gate circuit, which is obviously much more sensitive to small changes in voltage.

Model Parameter Measurements

The foregoing provides a circuit model to explore, once the values of the impedances depicted above are found; as described in chapter 5, this is non-trivial due to their small magnitudes. Resistances of the order of fractions of a

milliohm, and inductances of tens of nanohenries or less, are notoriously difficult to determine reliably. Whilst the absolute values were relatively unimportant in chapter 5, as a range of values was investigated in each case, here the values need to be refined.

Two approaches present themselves: one is to calculate the required impedances from the module's geometry. Resistances can be estimated from dimensions of the bond wires and tracks and their resistivities, and estimates for contact resistances added in; the inductances can be calculated from the Biot-Savart law and the module's geometry. This is not particularly attractive, given the complicated geometry of the IGBT module.

The alternative approach is experimentally, as chosen here, being more in keeping with the aims of this work. For components in the power circuit, a large dI/dt is achievable by simply switching the IGBT module on and off. dI/dt can be measured with field coils, and voltage probes and an oscilloscope can be used to measure voltages, yielding the inductances required using $V=L.dI/dt$. Similarly, resistances in the power circuit can be measured by switching the module on and applying a known current. The required voltages can be measured with an accurate voltmeter, and the resistances deduced from $V=I.R$.

The gate resistors are large enough to measure with conventional means, but the inductances in the gate circuit prove a little more difficult, as dI/dt is usually small in normal operation. There is also the problem of interaction with the power circuit if this is undertaken whilst the module is switching a load current. Estimates were made by introducing a gate-emitter short on the chip's bond wires and applying a substantial dI/dt .

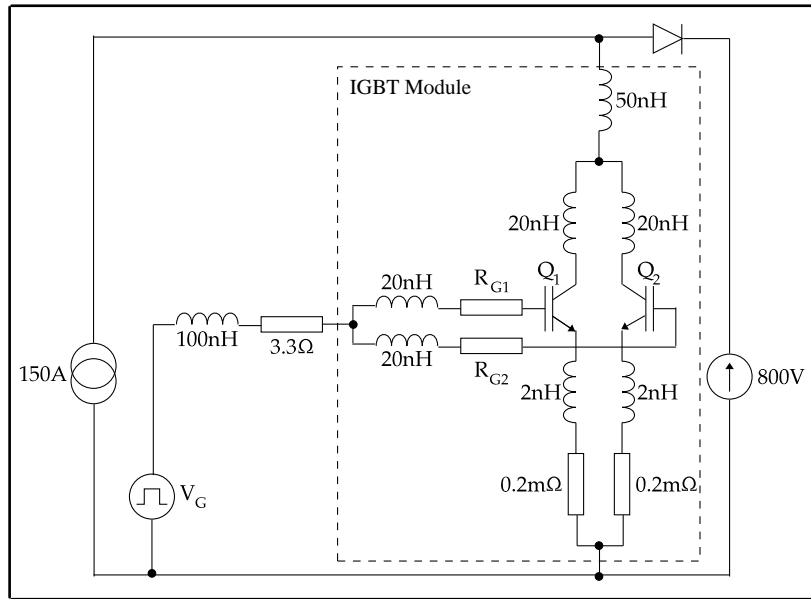


Figure 67. The simulated circuit, with measured and estimated parasitic components

Gate Resistors

In the chapter 4, the importance of equal gate resistors became clear, and so some simulation of the impact of mismatched gate resistors is undertaken. Based on the spread of gate resistors determined experimentally, two values were chosen to represent a realistic scenario, namely 6.2Ω and 6.5Ω . These were used in the circuit shown in figure 67, which is otherwise symmetrical.

Uneven resistivity of the polysilicon gate tracks has been cited as the cause of inhomogeneous turn-off of a single IGBT chip [94]: this is effectively the same problem, on a smaller scale. Experiments using discrete IGBTs with different gate resistors were presented in that paper, and it would be interesting to compare the results obtained, but the differences in gate resistors used are much larger, and so comparison is of limited value.

At a simple level, it could be expected that the IGBT with the smaller gate resistor would begin to turn off earlier than its partner, which would consequently carry progressively more current until it too began to turn off. However, the turn-off process is rather more complex than this, as was described in chapter 2, and so the waveforms produced are not as would be expected from this simple explanation.

Figure 68 shows the simulated turn-off waveforms, divided into five phases for the purposes of explanation. The two chip currents diverge once the chip gate voltages reach their plateau levels, the divergence being slow during phase

2, but then accelerating during phase 3. The divergence peaks distinctly, and then the collector currents reconverge during phase 4, the latter part of the collector voltage rise. The gate voltages are virtually indistinguishable until phase 3, when the collector current divergence is large; even then, there is no similarity between gate voltage and collector current waveforms. It is thus suggested that the latter stages of turn-off are not directly under gate control, but instead are merely dependent on the device and circuit conditions resulting from the earlier stages of turn-off.




Figure 68. Simulation of two identical chips in parallel, but different gate resistors (6.2Ω and 6.5Ω). Similar results have been observed by others for mismatched gate resistors [102], and also during the experimental work, shown in figure 69. The collector currents show the divergence and reconvergence expected from the simulation result above, although the distinction between phases 2 and 3 is not seen. The gate voltage waveforms are of relatively little use; it appears that the measurements of the chip gate voltages have suffered from the inclusion of a little stray emitter inductance, despite efforts to measure them as closely as possible to the chips.




Figure 69. Measured turn-off of parallel chips with different gate resistors

The simulated collector current waveforms seen in figure 68 require some explanation. First, a distinction must be drawn between two causes of current flow through the collector and emitter terminals of the IGBT. During conduction, around 75% of the emitter current flows through the MOS channel. This passes on and out of the collector terminal; meanwhile, holes are being injected at the collector end of the device. This hole current flows to the emitter terminal, although it does not pass through the MOS channel. During switching, current continues to flow out of the device terminals; the electron current continues to leave via the collector, and hole current via the emitter. However, the MOS channel is constricted, and eventually completely cut off, and so the continuing flow of current through the device empties the device of the excess carriers injected during conduction, causing the depletion region to form and the collector-emitter voltage to rise. In other words, the current flowing through the device's power terminals can have two different origins, but this is not directly visible from outside the device.

With the aid of the device simulator ATLAS and its visualisation tools, the evolution of the MOS channel current during switching can be observed within the device. As an aid to understanding the processes involved, the simulation is re-run, but with a much larger difference between the gate resistors, namely 6.2Ω and 12.4Ω. The results are shown in figure 70, with the same five phases indicated.




Figure 70. Switching waveforms for two chips in parallel with a large difference in gate resistors

Throughout phase 1, the two chips carry the same current; figure 71 shows the electron current densities in the two chips at the beginning of this phase, which are identical as expected. The first chip's gate voltage falls to its plateau level during this phase, and the second chip's gate voltage falls more slowly.




Figure 71. Electron currents in the two IGBTs during normal conduction. In this sequence of figures, white represents a current density of 200A/cm². The IGBT with the larger gate resistor is on the right.

During phase 2, the gate voltage of the IGBT with the smaller gate resistor continues to fall, and its MOS channel is cut off. Figures 72 and 73 show the two IGBTs' electron current densities at $t=2.5\mu\text{s}$ and $t=2.7\mu\text{s}$, which are part way through phase 2 and the end of phase 2 respectively. Unlike the single device turn-off situation, where an inductive load maintains a constant collector current, here one IGBT's collector current can fall at this stage: the external circuit only forces the sum of the two collector currents to be constant. It can be

seen that by the end of this phase, the first IGBT's channel is completely off, and its collector current has fallen to zero, the second IGBT's channel carrying the extra current. The gate voltage of the second IGBT rises again as its collector current increases.




Figure 72. Electron currents at $t=2.5\mu\text{s}$, during phase 2




Figure 73. Electron currents at $t=2.7\mu\text{s}$, the end of phase 2

During phase 3, the second IGBT begins to turn off, its gate capacitance being slowly discharged by its gate current. By the end of phase 3, the accumulation layer under the gate oxide has gone and the channel is becoming constricted, as seen in figure 74. The first IGBT's gate voltage remains low. Little change is seen in the collector currents of the two chips.




Figure 74. Electron current densities at $t=2.9\mu\text{s}$, the end of phase 3.

Figure 75 shows the second IGBT's electron supply being much reduced during phase 4. The collector voltage rise accelerates, and the gate terminal now has little influence, as only a small fraction of each device's current is coming through its MOS channel. Consequently the two devices are left to share current between themselves. The total current is still constant, and the two collector voltages approximately equal, so the progress of the depletion regions in the two devices determines the current sharing between them. Since they were carrying the same current during conduction, and the carrier lifetimes are long, the carrier distributions should be the similar, and so the current sharing will then tend towards a balance, as seen.




Figure 75. Electron current densities at the end of phase 4: both MOS channels are virtually off.

Finally, in phase 5, the last stages of turn-off proceed normally; the collector voltage reaches the DC link voltage, and then the two collector currents fall together. As the two IGBTs are identical, and no longer under gate control, the current sharing during the remainder of the turn-off is good. The collector currents seen in the measurements of figure 69 bear out this result, and in fact remain equal for a brief period at this point, supporting the theory that the channel current has ceased.

The detailed description of IGBT turn-off presented above illustrates an important point: under hard switching, the MOS gate controls only the first part of turn-off, whereas the action of the wide open-base bipolar transistor forming the bulk of the device governs the remainder. This effect can be seen more clearly from figure 76, where the electron and hole currents are plotted separately; the sequence of events described above can be seen clearly. It is also notable that, since the gate terminal can only control the MOS channel current, control of dV/dt or dI/dt during switching must be done with the MOS channel still conducting.




Figure 76. Hole and electron currents plotted separately for the same turn-off as figures 63 through 67.

Returning to figure 68, where a 0.3Ω difference in gate resistors produced a substantial current imbalance during turn-off, it appears that the simulated current imbalance is exaggerated compared to what might be expected from the real measurements; the real IGBT pair had gate resistors differing by 0.7Ω . This could be due to a difference between the device model and the actual devices. It

may also be a consequence of the necessarily simplified view of the module itself; stray capacitances are assumed to be dominated by those of the IGBT itself, and inductances and resistances are difficult to measure reliably. Magnetic coupling between conductors, discussed in [76], can further complicate the issue, but these couplings are not modelled here.

Perhaps the most important result here is the loss distribution, as this is what will ultimately govern the heat distribution, and consequently the operating limits and long-term reliability, of a module. The total turn-off losses per chip (E_{OFF}) are 42.25mJ and 43.06mJ, a difference of just 1.92%, despite the fact that one chip's collector current is briefly more than 50% higher than that of the other. Closer examination shows that the collector voltage at this point is low whilst the collector currents are diverging, but as it rises the collector currents fall together. During the tail time (i.e. the time when collector current remains after the main current fall), there is still a significant collector current remaining, whilst the collector voltage is at the DC rail voltage; the tail time makes a significant contribution to switching losses, but again the current sharing is good. Thus the loss imbalance is less than might be first thought from inspection of these waveforms.

Gate Oxide Thickness

The thickness of the gate oxide layer is known to be difficult to control; whilst manufacturers are reluctant to disclose exact information, variations of the order of 10% are expected. These obviously affect the threshold voltage: as a certain electric field strength is required to form the inversion layer for conduction, a thicker gate oxide layer than the nominal 0.100 micron will mean that a higher voltage is required on the gate terminal for this to occur. Standard MOS theory suggests that the relationship should be linear (e.g. [103] page 283), and this is borne out by simulations, as shown in the graph below. As would be expected, the on-state voltage $V_{CE(sat)}$ increases slightly with increasing threshold voltage; if the gate voltage is maintained at the same value, with a higher threshold voltage, the inversion layer is shallower. The MOS channel will then contribute a little more to the on-state resistance of the device, as seen in figure 78.





Figure 77. Effect of gate oxide thickness on threshold voltage $V_{GE(th)}$ (above) and on-state voltage $V_{CE(sat)}$ (below)



Here a 1% variation in thickness is taken as a nominal starting point. It assumed for the purposes of these simulations that the oxide thickness is uniform over a chip; as detailed earlier, this is not necessarily the case. The threshold voltages ($V_{GE(th)}$) of these two chips are 4.45V and 4.50V, approximately 1% different, and the on-state voltages ($V_{CE(sat)}$) at 75A are 2.121V and 2.123V, approximately 0.1% different.

The results when the pair are switched in parallel are seen in figure 78. The on-state sharing is good, as expected from the close matching of $V_{CE(sat)}$; however there is a substantial collector current divergence prior to turn-off. Inspection of the gate waveforms shows a little divergence between the two individual gates.

As in the previous simulation, the collector current variations during switching are somewhat larger than would be expected for a real life situation, at least from qualitative comparisons with the waveforms obtained to date with the 'baseline' modules. Clearly with the expected 10% variation in oxide thickness the collector current redistribution problem will be much worse.




Figure 78. Switching waveforms for two chips in parallel with different oxide thickness

As in the case of different gate resistors, the loss distribution is relatively good, E_{OFF} being 42.43mJ and 42.86mJ for the two chips; increasing the thickness difference to 4% increases the difference in losses, the two values being 41.87mJ and 43.48mJ. The differing shapes of the two gate voltage waveforms are now more apparent, and can be seen to cross over at $t=2.6\mu s$; the collector current waveforms are very similar to those seen in figure 68.




Figure 79. Switching waveforms with the oxide thickness increased to 4%

P-base Doping

Device doping is notoriously difficult to control accurately over a single chip, let alone a wafer or batch of wafers. For GTO wafers, it was found that p-base doping variations of 15% between two cells on a single wafer were not inconceivable [104], and the consequences of variations of this size were investigated. The IGBT's requirement for more closely controlled processing technology has reduced that variation substantially, a figure of 2% being quoted anecdotally, although manufacturers are unwilling to divulge exact figures. The effects of doping variations of this order of magnitude were investigated in the vicinity of the supplied model's original value of $5 \times 10^{17} / \text{cm}^3$. As with the gate oxide thickness, the p-base doping affects the threshold voltage $V_{\text{GE(th)}}$; in this case, an increase in the doping will mean that a higher electric field is required to produce the inversion layer required for conduction. The expected rise in $V_{\text{GE(th)}}$ for rising p-base doping is seen below, $V_{\text{GE(th)}}$ being approximately proportional to the square-root of the p-base doping ([103] page 283), and this is not incompatible with the graph below, if difficult to visualise. $V_{\text{CE(sat)}}$ also rises slightly for reasons outlined previously.





Figure 80. Effect of p-base doping on threshold voltage ($V_{\text{GE(th)}}$) and on-state voltage ($V_{\text{CE(sat)}}$) (over)



Increasing the p-base doping of the baseline chip by 2.5% produces a similar threshold voltage shift to that obtained by increasing its gate oxide thickness by 1%. Paralleling two chips with this p-base difference yields sharing during the on-state better than 1%, and yet during turn-off 10% of the current from one chip is transferred to the other as seen in figure 81. The losses are very similar to those seen in the case of the 1% oxide thickness variation, at 42.45mJ and 42.85mJ. The waveforms seen here bear a close resemblance to those obtained with oxide thickness variations and shown in figure 78, suggesting that the two non-idealities affect the chips in similar ways, namely by altering the threshold voltage.




Figure 81. Different p-base doping. The collector current waveforms are similar to those in figure 78.

Gate Oxide Thickness and P-Base Doping Combined

Since both of these device variations alter $V_{GE(th)}$, the question arises as to whether the resulting changes in switching behaviour can be related directly to $V_{GE(th)}$, without requiring knowledge of the details of the device changes. This would clearly be desirable if chip selection is to be used in order to fully exploit the current ratings of parallel chips.

A pair of chips with the same $V_{GE(th)}$ were simulated in parallel. A 4% decrease in gate oxide thickness, coupled with an 11.1% increase in p-base doping, produced two chips with virtually identical threshold voltages (4.486V and 4.484V) and on-state voltage (2.134V and 2.139V). Whilst it is accepted that this doping variation is larger than would normally be expected, the switching characteristics of these two chips in parallel are interesting.

The on-state sharing is seen in figure 82 to be good, as would be expected from their terminal characteristics, if slightly worse than previously. The two gate voltages diverge during the early stages of turn-off, as do the collector currents, but later on the gate voltages are close together, despite which the two collector currents remain different. An unusual feature of this particular turn-off is that the collector current waveforms cross over twice during the collector voltage rise, at $t=2.35\mu\text{s}$ and $t=2.85\mu\text{s}$. The loss imbalance is once again small, the two IGBT turn-off losses being 42.55mJ and 42.70mJ, a difference of 0.35%.




Figure 82. Switching waveforms with different p-base doping and oxide thickness: collector current and voltage (left) and gate voltage (right)

The turn-off characteristics of the pair of chips can be radically altered by coupling their gates, a solution suggested in [95] and [96]. Figure 83 shows that the collector currents no longer cross over prior to turn-off, and they have a smaller peak divergence: the general shape of the waveforms during the storage time are completely different. The losses for the two chips are now 42.51mJ and 42.73mJ, a difference of 0.52%. In both linked and unlinked cases, the collector currents fall together.




Figure 83. Collector current and voltage waveforms (left) and gate waveforms (right) for the same chips used in figure 82, but with the chip gates tied with a 1Ω resistor

It seems that another distinguishing characteristic is required if chips are to be ideally matched. The results here suggest that one possibility is collector current at a particular collector and gate voltage, so the devices' transfer characteristics were investigated.

The graph in figure 84 shows a small but significant difference between the two chips. Assuming that the two gate voltages are identical, as in the case of linked gates, one can predict values for the two chip collector currents from these transfer characteristics. Closer inspection shows that at a V_{GE} of a fraction over 7V and a total current of 150A, the two chip currents are approximately 73.1A and 76.9A; this compares with values during the transient simulation of 73.2A and 76.8A at this point. It would appear that this is a reasonable method for predicting current distribution during switching with tied gates. Such predictions remain of limited value when considering the more normal topology, however, due to the decoupled operation of the two chips if the gates are not tied.




Figure 84. Simulated transfer characteristics (V_{GE} vs I_C at fixed $V_{CE}=100V$) of the two chips

Gate Drive Techniques

It was seen during the experimental work that using the gate's ability to control the collector voltage during turn-off was beneficial in some circumstances, such as operation of devices in series, but had consequences for the collector current distribution during turn-off. The active snubber style of control is eminently suitable for simulation with the Silvaco software, as it comprises only two passive components; however, the closed-loop voltage control type of gate drive is far less convenient. A simplified model of this could be used, but it was felt that the results would add little to those obtained with the simple active snubber.

Here, the simulations shown in figures 82 and 83 were re-run with the addition of the active snubber, producing the waveforms shown in figures 85 and 86. It is clear that the extended turn-off has prolonged the period for which the collector currents are unbalanced. The loss imbalances have increased from 0.35% to 2.34% and from 0.52% to 3.03% respectively compared to the hard switched case. Evidently the use of the controllability comes at a price; however the module as a whole behaves extremely well, despite its internal problems.




Figure 85. Switching waveforms for chips with different oxide thickness and p-base doping, and active snubber.




Figure 86. Switching waveforms for chips with different oxide thickness and p-base doping, gates linked with a 1Ω resistor, and an active snubber.

Intra-chip Issues

A complication arises from the fact that the emitter Kelvin connection (the bond wire which returns the gate current to the gate drive) is typically connected at one corner of each chip. The emitter metallisation on the top of the chips under investigation is split into eight segments, and these are connected together with the emitter bond wires, as shown in figure 87 below. The gate signal is carried along thin tracks between the emitter segments, which means that different parts of the chip will have different impedances in their gate connections.




Figure 87. Single IGBT chip and bond wires, showing asymmetry in gate current path arrangements

The circuit diagram in figure 88 shows a circuit model of the four of the segments in one chip. There are effectively four equivalent pairs of IGBTs, apart from the Kelvin connection to one pair (A and B). The return paths for the gate current from the different areas of the chip are significantly different, as for seven of the eight segments, the gate return current must flow along the bond wires carrying the main emitter current. The di/dt produced by switching transients will cause voltage spikes in the gate control loop for these parts of the chip. The bond wires involved in the gate current return path for the top two segments of the chip are highlighted. Clearly the impedance in this path is significantly different to that in the path for the bottom left segment, and this could lead to inhomogeneity in switching of the separate areas of the chip.

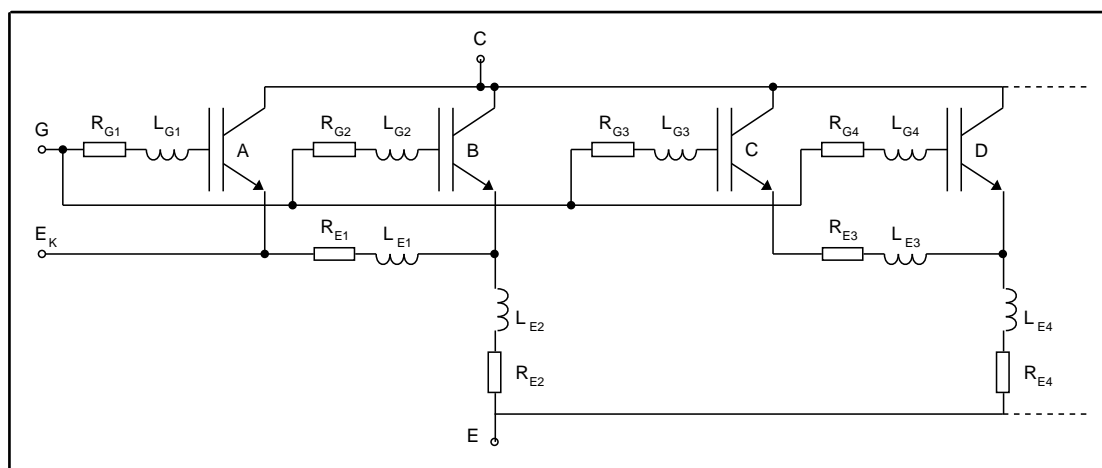


Figure 88. Equivalent circuit for single IGBT chip shown in figure 87

It is considered that the two problems, i.e. inter-chip and intra-chip redistribution, are effectively the same problem of parallel IGBTs, but with different interconnecting impedances in the emitter circuit, notably in the Kelvin connection arrangement. For the time being the potential intra-chip problems are

ignored and the chip assumed to act as a single IGBT; the validity of this assumption should be explored in due course. It can reasonably be assumed that a parallel IGBT simulation which works for the case of parallel chips should be relatively straightforward to modify to investigate the intra-chip case, assuming that the interconnecting impedances can be determined.

Summary

A representation of the IGBT module has been developed and implemented in Silvaco's device simulation software suite, and the effects of a number of possible changes in the IGBT and its circuit investigated. The case of different gate resistors has been examined in some detail, as this is a situation which arises in production modules using silicon resistors. It appears that, as might be expected in this situation, one IGBT turns off before the other; more specifically, one IGBT's MOS channel turns off before the other. The remainder of turn-off is largely governed by the devices themselves and the external circuit, rather than by gate control.

The effects of differing gate oxide thickness and p-base doping have also been examined, and superficially at least appear to be very similar; both alter the threshold voltage of the IGBT. Once again, one IGBT's MOS channel turns off before the other, and so the waveforms are similar to those observed for the gate resistor case. The picture is somewhat complicated by the fact that these device changes do not only affect the threshold voltage: these two effects can be balanced against each other to produce two chips with identical threshold voltages, and yet they will exhibit redistribution during parallel turn-off. In this case, the gate capacitances of the two chips are different, and so the result is similar to the case of differing gate resistors: the R-C time constants of the two gate circuits are different.

The active snubber circuit provides an insight into the internal operation of IGBT modules when the gate's controlling ability is used to control turn-off more closely, for example when the rate of voltage rise is controlled for series operation. In this case, the current unbalance during the collector voltage rise is seen to stabilise; the IGBTs remain in their active region, and the current sharing is dictated by the saturation region transfer characteristics of the two chips. The MOS channels in the IGBTs remain on for an extended period.



Discussion

Current Measurements

The field coil technique, whose development was outlined in chapter 3, has facilitated reliable measurement of individual chip currents as required, and provides the basis for the practical aspects of this work. Applying this technique to a module used as the switching element in a step-up converter enables the current redistribution within an IGBT module to be studied under a full range of conditions, including different collector currents, collector voltages, and gate drive techniques.

The turn-off transient was studied extensively, whereas the turn-on transient has received very little attention. Consideration of the turn-on process of a fast semiconductor switch with a collector inductance shows that the power dissipation at turn-on is relatively low. The collector inductance acts as a turn-on snubber, and the device voltage falls rapidly, whereas the device current increases more slowly. Thus the consequences for the loss distribution are not very significant, whereas the turn-off transient has been seen to produce significant loss imbalances.

Should further investigation of turn-on be required, the step-up converter circuit of figure 15 can be modified to facilitate investigation of turn-on, using the topology depicted in figure 89. A second power switching device (*B*) is placed in parallel with the device under test (*A*). Device *B* can be any fully controllable semiconductor device capable of supporting the required voltage and carrying the required current. Device *B* is turned on until the supply current reaches the test value, and then turned off; the supply current then commutates to the diode. After a suitable interval to ensure that device *B* is turned off fully, the device under test (*A*) can be turned on under conditions virtually identical to those seen in a real inverter.

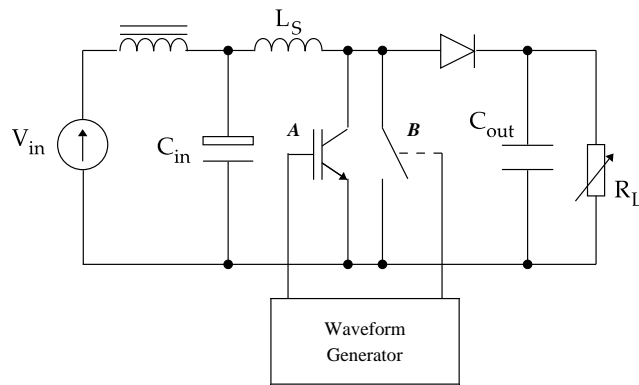


Figure 89. Proposed IGBT turn-on rig

To date the field probes have only seen use in the laboratory environment, but they have many more potential applications. In particular, they could be incorporated into a module operating in a full-scale inverter and used to monitor the long-term evolution of the module's behaviour. Such an investigation would yield useful information about the ageing characteristics of IGBT modules, and possibly the processes involved. The long-term reliability of modules could then be analysed.

Characterization of the current measurement system remains a desirable objective, but this was not achieved fully. It was seen in figure 43 that the whole system can resolve oscillations occurring at 20MHz; exploration of the system's response at frequencies beyond this has been attempted using a rig designed to generate extremely fast current changes. This was frustrated by the impedance mismatches between the probe and connecting cable, and between connecting cable and integrator, as these cause ringing which masks the true current transient. A substantial redesign of the integrator and careful choice of interconnecting cable will reduce this ringing. It is also noted that the apparatus as it stands lacks a reference current measurement capable of reproducing such transients; this could be remedied with a coaxial resistive shunt.

However, the integrator has been characterised using a signal generator, the results of which were seen in figure 24, and it was seen that its magnitude response is correct to at least 30MHz. With the available equipment it proved impossible to make useful measurements beyond this; the output signal amplitude diminishes with increasing frequency due to the integrating action, eventually becoming too small to measure reliably even with the largest

available input signal. The phase response is less ideal and, whilst this will not mask differential oscillations such as those seen in figure 43, it could cause distortion of transient current distributions during switching.

Most modern plastic modules are inconvenient for current measurement, and often have no possibility of using the existing field probe. Such modules typically have several chips soldered onto a single DBC tile, as indicated in figure 90. Clearly such an arrangement means that measurement of individual chip currents using the existing current probe around a busbar cannot yield individual chip currents.




Figure 90. Diagrammatic representation of a typical three-chip tile

One option is to persist with an existing current measurement system, i.e. either the field coil and integrator, or Rogowski coils. A smaller field probe could be produced which would fit around the bond wires; however the space available is very limited, and such a probe would not be well suited to the existing methods of hand production with a minimum of facilities. Depending

on the exact layout of the module, maintaining the compensating nature of the field coils could be difficult. Miniature Rogowski coils may provide a solution.

The advent of the press-pack style of packaging presents another challenge. The use of any sort of measurement technique within the packaging will be extremely difficult due to the confined space available. However it is probable that a Rogowski coil can be fabricated to fit around the individual poles of the emitter pole piece, in a similar fashion to earlier work with GTOs [56], but even then the problem of making external connections remains. The press-pack module cannot simply be operated without its lid in the same manner as a plastic module, as the lid also forms one of the main power terminals. The ceramic body of the package could be modified to allow the necessary external connections, or a different external ceramic casing may be feasible. A more realistic alternative given the available facilities would be to use the magnetic field mapping techniques developed for similar work with the GTO [64], although the results are far less straightforward to obtain and then require considerable interpretation.

Redistribution within a single chip has not been investigated in detail. The different gate current return paths for different parts of an individual chip were highlighted in chapter 6, but to date no suitable current measurement method for investigating this has been produced. The compensated field coils are difficult to use here as they would need to be placed around the emitter bond wires: even if they could be miniaturized to fit around individual bond wires, their compensating ability would be seriously hampered by the adjacent bond wires. As shown in figure 22, the presence of other conductors parallel with, and immediately adjacent to, the one being monitored will reduce the accuracy of the compensation. As with the press-pack style of module, it may be necessary to resort to magnetic field mapping.

Another option is adapting the flux-gate magnetometer for such a task, an approach presented very recently [105]. At the time of writing this technique is only suitable for static current measurements, and no others have been seen which can approach the bandwidth and stray noise rejection of the magnetic field coil / integrator system for measuring current redistribution during switching transients.

Stability

Early research into paralleled IGBTs [40] [75] suggested that separate gate resistors were required, and this is the stance taken by most manufacturers. Investigation of the IGBT's small-signal model in chapter 5 agrees that an amount of gate resistance is required to avoid differential oscillations, and these have been seen experimentally under certain conditions. However, it has been shown here that individual resistors cause transient current distributions to be significantly worsened if there are any differences between chips.

From the small-signal model, it was shown that changes to module or chip design could reduce the gate resistance required, which would be beneficial in terms of speed of operation and transient current distribution. However the required changes are not necessarily feasible. For example, stability can be assured for the particular IGBT under the conditions investigated in chapter 5, even with no gate resistor, if the collector inductance is increased to 150nH. This could be achieved by adding approximately 10cm of plain conductor to each chip's collector connection, but clearly this is not straightforward to accommodate within the confines of an IGBT module. The introduction of magnetic materials could achieve the same inductance in a more realistic space. However, in both cases, each chip would need to be mounted on its own tile to enable the inclusion of the collector inductance, which would reduce the number of chips which could be accommodated in a given size of module. The added inductance would also have undesirable consequences for the electrical properties of the module. Turn-on would be slowed, and the collector voltage overshoot at turn-off would be increased by approximately 100V.

At least one manufacturer has dispensed with discrete per-chip gate resistors for their press-pack devices, using instead the resistance of the device's gate polysilicon layer, whereas another has retained the conventional discrete resistors [39]. Oscillations were discovered in the latter style of module during the later stages of turn-off. It was found that a certain amount of emitter inductance would suppress these, although the extra inductance was unwelcome, and this provides an interesting contrast with figure 54(a), where emitter inductance is seen to be undesirable. An alternative solution, in the

shape of permalloy rings around each emitter connection, provided the same suppression without the extra inductance. The permalloy rings offer a frequency-dependent impedance, and if suitably chosen, their impedance at the oscillation frequency is sufficient for suppression purposes, but normal switching transitions are not unduly affected.

Perhaps the most interesting aspect of these particular oscillations is the fact that they occur during the 'tail time', i.e. after the device's main turn-off (phase 4 in figure 9), in contrast with the result of figure 43, where oscillations were seen during the voltage rise. It has been suggested that these oscillations could be caused by the gate terminal appearing to have negative capacitance at high collector voltages [107], [108]. During the tail time, the IGBT current is low and the collector voltage high, whereas the investigation in chapter 5 focussed on rated collector current and a collector voltage of 200V. The small-signal model used for that analysis, shown in figure 44, is only relevant to the saturation region, and could not be used for analysis during the tail time, as the IGBT is no longer in its saturation region.

The limits of the mathematical analysis must be recognised, and in particular, that this is a small-signal analysis. It has already been seen that device capacitances and gain (g_m) change substantially with current; this may explain the fact that differential current oscillations were not observed to grow exponentially, but rather to find a certain limit, as seen in figure 43. A similar result was found under short circuit conditions [109]. It must be also be recognised that the IGBT's parameters used in the analysis are only valid at one particular device voltage and current, and so stabilisation measures cannot be relied on unless a thorough investigation under all the various possible operating conditions is carried out. In the first instance an exhaustive search of all possible collector currents and DC link voltages will be necessary, including short circuit operation. This may lead to a better understanding of which particular operating regimes require the most stabilisation, and hence reduce the scope of the investigation required for subsequent IGBT chips. Radical changes such as a move to the press-pack type package would clearly require a full investigation.

Transient Current Distribution

It was hoped at the outset that a number of externally measurable electrical parameters of IGBT chips could be used to predict their behaviour when operated in parallel. It was also hoped that maximum tolerances on these electrical parameters could be specified to guarantee a certain maximum loss imbalance between chips operated under given conditions. Experience with device simulation has shown that the situation is considerably more complex than initially anticipated.

The results presented in chapter 6 suggest that some measure of plateau voltage could provide a useful extra measure which could be used to improve the matching of IGBT chips for parallel operation. Such a measurement could be achieved in practice by applying a given collector voltage, and ramping the gate voltage until the collector current reaches its rated value. Combined with the existing threshold voltage measurement, the device transconductance is effectively matched.

The waveforms seen in figures 61 and 63 bear some comparison with those shown in [94]. Whilst the latter uses more substantial differences between gate resistors than the results presented here, the comparison is nonetheless interesting, as some aspects of the waveforms are similar. With a ratio between the two gate resistors of 5:1, it is seen in [94 figure 8b] that the chip with the smaller resistor turns off first, but recommences conduction at a later stage. The two collector currents do not, however, reconverge during the main current fall and tail time, as was the case in figure 70. This may be due to the two devices being mismatched, as suggested by divergence during the tail time when the devices were switched with identical gate resistors [94 figure 8a].

The differing effects of variations in some of the device parameters were seen in chapter 6. For example, a 1% variation in gate oxide thickness gives similar turn-off results to a 2.7% variation in p-base doping. On the other hand, small variations in carrier lifetime and anode doping produce little direct effect on switching performance. It can be seen that switching performance is largely governed by the behaviour of the MOS channel at the emitter end of the device, whereas the static distribution comes mainly from the bipolar behaviour and

can be predicted statistically [80]. Such information can be combined with knowledge of device processing technology to deduce which aspects of device processing are most critical to ensuring optimum paralleled device performance. Of the device parameter variations investigated here, the gate oxide thickness requires the tightest control, although the effect of variability of the silicon gate resistors is significantly larger.

An alternative approach to the slow and laborious device simulations is circuit-level simulation [97], where a simplified one-dimensional model of the IGBT is implemented in a circuit simulator such as PSPICE. The IGBT can then be easily incorporated into, and simulated within, its environment. Such an approach offers vastly reduced simulation times, and a more accessible user-interface. The models are often readily available from the manufacturers, but they contain little relating to the physical structure of the device, in contrast to the two-dimensional physical model used in chapter 6. Clearly the accuracy of such simulation techniques is much lower; thus these techniques are not yet sufficiently sophisticated to deal with the intricacies of device operation to provide useful results here, although better models are under development.

A number of questions remain unanswered, notably the fact that less substantial redistribution is measured in practice than is predicted by the device simulations for the same situation. For example, a 5% difference in gate resistors produces a simulated current transfer from one chip to the other of 50%, whereas the corresponding measurement indicates a transfer of the order of 25%. It is possible that the device model used is not sufficiently accurate for detailed modelling purposes; for example, the split emitter metallisation has been shown to have a significant effect under short circuit conditions [109], and may also affect normal transient operation, but this issue has not been addressed here.

Module Design Issues

In larger IGBT modules employing several chips in parallel, maintaining symmetry becomes more difficult very rapidly as the number of chips increases beyond four: some degree of asymmetry is almost inevitable. The consequences of the asymmetry in the power circuit are primarily seen during conduction,

and the effect on transient current redistribution appears to be less severe than might be first thought. Thus the most important aspects of the module's design are in the gate circuit, and in particular, the emitter inductance and the topology of the gate connection.

It has been seen, both here and elsewhere [95] [96] that 'balancing resistors' between gates are useful in improving transient current sharing, and the practical work of chapter 4 and the simulations in chapter 6 have borne this out. The differential oscillations predicted by the theoretical treatment are rarely seen.

The gate resistor has two conflicting requirements on its value. A low value is desired for fast switching, whereas a high value is desired for stability. Compromises are never particularly satisfactory as neither requirement is completely met. However it may be possible to refine the two requirements so that they conflict less: the current oscillations predicted are shown to be at relatively high frequencies of 20MHz or more. The switching transients of large modules typically take 200ns or more, and so it can be seen that a frequency-dependent impedance such as a ferrite bead may offer a better solution. This is not a new idea [78] [101] but is worth revisiting. Discrimination between switching transients and oscillations will be made easier if the oscillation frequency can be increased by changes to module or chip design.

The material of the gate resistor itself is seen to be important, and the continued use of silicon is not recommended due to its high temperature coefficient. Other alternatives such as metal film are readily available with much better tolerances. It was shown in chapter 6 that small differences in gate resistor values produce significantly uneven loss distributions, which is clearly undesirable.

The gate current return (emitter Kelvin) connection is difficult to implement satisfactorily, particularly in plastic modules. If it is taken from one IGBT, the gate current from other chips is forced to return via the emitter bond wires which carry the main switching transients. If it is taken from each IGBT, it also provides an alternative path to the power circuit. It has been suggested in the case of parallel modules [102] that splitting the gate resistance between gate and

return connections could reduce this effect, and this could be applied to the case of parallel chips.

Clearly there is a significant amount of design work required to produce the correct characteristics for the chip, module and gate impedance. The potential benefit is a module which switches more quickly and more homogeneously, reducing switching losses and improving reliability. It may be that the module design can be altered to achieve improvements in performance without costly changes to chip manufacturing processes.

Future IGBTs

A number of new types of IGBT are emerging, offering different compromises between the various desirable characteristics such as switching speed, on-state voltage and maximum DC link voltage. The terrace gate IGBT offers a lower gate capacitance and hence easier gate drive, along with reducing the turn-off delays associated with the accumulation layer; C_{GC} and C_{GE} are both reduced, which will have consequences for stability when operated in parallel. The trench gate IGBT removes the 'JFET effect' and hence reduces the on-state voltage drop, and once again the gate capacitance will change.

Additionally, larger area devices are slowly being developed through improving manufacturing processes, coupled with 'repair' techniques which allow faulty sections of an individual IGBT chip to be disabled during post-manufacture testing. This will lead to a change in the emphasis from inter-chip to intra-chip current distribution issues.



Concluding Remarks and Further Work

The current probe developed to operate with the existing magnetic field based current measurement system has facilitated non-invasive measurement of individual chip currents within multi-chip IGBT modules for the first time. Being a non-invasive measurement, the modules may be operated under realistic conditions whilst measurements are made. A detailed comparison has been made with existing current measurement methods, concentrating on their performance and suitability for use within IGBT modules. The magnetic field probe system has been successfully calibrated against benchmark conventional methods, within the limits of the benchmark systems.

The measurement system has enabled the identification of two undesirable properties of the gate resistors used in the modules tested, namely their variability and temperature coefficient. Both of these resulted in poor transient current sharing within the modules. The use of metal film resistors, which have better characteristics, was proposed and tested. Current sharing during transient operation was improved.

IGBT chips, even those from the same wafer, show externally measurable differences in their electrical characteristics due to non-uniformities in device processing techniques. The measurement system was also used to develop an understanding of the behaviour of such unmatched chips when operating in parallel. It was shown that linking the chip gate connections can offer improvements to transient current sharing under such circumstances. Similarly, it was shown that reducing the individual gate resistors improves current sharing and also speeds switching, reducing switching losses. The extended use of the IGBT's active region during controlled switching was shown to worsen the distribution of switching losses, although the controllability of the module as a whole is maintained.

Paralleled devices may suffer from differential current oscillations. The MOSFET had been investigated elsewhere in this respect, using analysis of the characteristic equation of its small-signal model and a limited number of parasitic elements. The Routh-Hurwitz stability criterion was then used to determine the minimum value of gate resistor required to assure stability. This analysis was extended here to the IGBT, with a comprehensive set of parasitic elements included. The IGBT's current and voltage dependent capacitances were characterised for the analysis using the physics-based finite element device simulator ATLAS.

Certain aspects of module design were found to be particularly influential in avoiding these oscillations, and the stability of a module can be improved by suitable choice of interconnecting impedances. The Bode diagram can be used to indicate which impedances are particularly important in a given situation. As the terminal capacitances of the IGBT chip itself under switching conditions are heavily current and voltage dependent, investigation of a full range of both current and voltage conditions would be required to assure the stability of a particular combination of module and chip designs. The extension of the small-signal analysis to cover asymmetrical situations for the first time has further widened the applicability of this model.

The differential current oscillations predicted when the gate resistor used is insufficient to ensure stability have been in practice seen under certain switching conditions, although the module as a whole still switches correctly. However such oscillations do not necessarily occur when instability is predicted, as the fast switching of a hard-driven IGBT leaves insufficient time for oscillations to develop.

The two-dimensional physics-based finite element simulator ATLAS was used to investigate current sharing issues in more detail. A single switching element composed of two chips and interconnecting impedances was simulated, and the results compared with measurements. Changes to the doping or geometry of one IGBT chip of an otherwise identical pair were investigated, and the resulting changes in loss distribution quantified. The variability of the gate resistors and gate oxide thickness were found to be particularly important in

producing uneven current distributions. The influence of the gate topology on current sharing was also demonstrated, confirming the mathematical analysis.

Current redistribution was seen between parallel chips whose basic DC electrical characteristics, measured using existing methods, were identical. It was shown that these electrical characteristics are not simply related to single device features, and so multiple device changes can cancel out in terms of electrical characteristics. An additional externally measurable electrical characteristic is proposed, namely the gate plateau voltage, to enable better matching of chips for module assembly. Chips matched using this test in addition to existing measures may exhibit better current sharing.

Recommendations for Further work

A smaller current probe may facilitate measurements within single chips, and a thorough characterization of the current measurement system would be desirable. The integrator could usefully be updated with a self-zero circuit, and its impedance matched to that of the field coil. Alternative current measurement techniques will be required to investigate current distribution issues in press-pack modules; the radial and tangential field measurement technique developed elsewhere for GTOs may be suitable.

The consequences of the split emitter metallisation on the IGBT chip's upper surface should be investigated; this will require more accurate quantification of the module's stray impedances, and most particularly, those of the emitter bond wires. It may be possible to reduce the differential emitter inductance between areas of the chip, either by bonding a machined block to the top (emitter) side of the chip, or by adding an extra layer of metallisation. The bond wires would then be affixed to these. Bringing the gate connection out at one corner of the chip, as is seen in some press-pack modules, would make this easier. Inverting the chip and soldering it emitter side down may seem a simpler way to achieve this, but arranging the IGBT's forward voltage blocking under such circumstances would be difficult.

The stability analysis requires extension to cover the whole range of device operating conditions, and consideration needs to be given to analysis of the tail current time. The latter will require the production of a different small-signal

model for the IGBT. Once this has been achieved, the optimisation of the IGBT module design for stability can be tackled.

The transient operation of the IGBT module has been investigated with respect to theoretical chip variations. A practical counterpart to this could verify the conclusions reached. It may be possible to make measurements on production chips to verify the expected variations across a wafer; in combination with finite element simulation, future improvements in production methodology and control can be targetted for maximum gains in device uniformity.



References

- [1] E. D. Wolley: "Gate Turn-Off in P-N-P-N Devices", IEEE Transactions on Electron Devices, July 1966, pages 590 to 597.
- [2] J. P. Russel et al.: "The COMFET - A new high conductance MOS-gated device", IEEE Electron Device Letters, March 1983, pages 63 to 65.
- [3] B. J. Baliga: "Fast-Switching Insulated Gate Transistors", IEEE Electron Device Letters, December 1983, pages 452 to 454.
- [4] S. Duong et al.: "Investigation of the IGBT Case Explosion in Short Circuit Operation", EPE 1997, volume 2 pages 16 to 21.
- [5] P. Aloisi: "Paralleling of High Voltage Bipolar Power Darlington", EPE 1985, volume 1, pages 7 to 13.
- [6] J. G. Kassakian, D. Lau: "An Analysis and Experimental Verification of Parasitic Oscillations in Paralleled Power MOSFETs", IEEE Transactions on Electron Devices, July 1984, pages 52 to 56.
- [7] L. Lorenz et al.: "CoolMOS™ - a new milestone in high voltage Power MOS", ISPSD 1999, pages 3 to 10.
- [8] H.-R. Chang, B. J. Baliga: "500V n-Channel Insulated-Gate Bipolar Transistor with a Trench Gate Structure", IEEE Transactions on Electron Devices, September 1989, pages 1824 to 1829.
- [9] I. Omura et al.: "Carrier injection enhancement effect of high voltage MOS devices - Device physics and design concept", ISPSD 1997, pages 217 to 220.
- [10] V. A. K. Temple: "MOS-Controlled Thyristors - A New Class of Power Devices", IEEE Transactions on Electron Devices, October 1986, pages 1609 to 1618.

-
- [11] K. K. Afridi, J. G. Kassakian: "Turn-off Failures in Individual and Paralleled MCTs", ISPSD 1993, pages 60 to 65.
- [12] P. R. Palmer et al.: "The MOS and Bipolar Gated Thyristor: A Thyristor with IGBT Switching Characteristics", IEE Proceedings: Circuits, Devices and Systems, April 1998, pages 105 to 110.
- [13] T. Goto et al.: "Design Concept and Performance Consideration for Fast High Power Semiconductor Switching for High Repetition Rate and High Power Excimer Laser", Review of Scientific Instruments, July 1997, pages 2658 to 2665.
- [14] H. E. Gruening, A. Zuckerberger: "Hard Drive of High Power GTOs: Better Switching Capability obtained through Improved Gate-Units", IAS 1996, pages 1474 to 1480.
- [15] H. E. Gruening, B. Ødegard: "High Performance Low Cost MVA Inverters Realised with Integrated Gate Commutated Thyristors (IGCT)", EPE 1997, volume 2, pages 60 to 65.
- [16] B. J. Baliga: "The MOS-Gated Emitter Switched Thyristor", IEEE Electron Device Letters, February 1990, pages 75 to 77.
- [17] S. Sridhar and B. Baliga: "The Dual Gate Emitter Switched Thyristor (DG EST)", IEEE Electron Device Letters, January 1996, pages 25 to 27.
- [18] B. H. Stark: "Multiple-Mode Power Semiconductor Devices", Ph.D. thesis, Cambridge University UK, October 1999.
- [19] S. M. Sze: "Semiconductor Devices: Physics and Technology", John Wiley & Sons, 1985.
- [20] C. R. M. Grovenor: "Microelectronic Materials", Adam Hilger / IOP Publishing Ltd, 1989.
- [21] S. M. Sze (Editor): "Modern Semiconductor Device Physics", John Wiley & Sons, 1998.
- [22] M. Hiyoshi et al.: "A 1000A 2500V Pressure Mount RC-IGBT", EPE 1995, volume 1, pages 51 to 55.
- [23] T. Heinzl: "IPM for Low Power Drives", PCIM Europe, Issue 4/1999, pages 8 to 13.

-
- [24] N. Iwamuro et al.: "A New Vertical IGBT Structure with a Monolithic Over-Current, Over-Voltage, and Over-Temperature Sensing and Protecting Circuit", IEEE Electron Device Letters, September 1995, pages 399 to 401.
- [25] J.-M. Thébaud et al.: "Thermal Fatigue Resistance Evaluation of Solder Joints in IGBT Power Modules for Traction Applications", PESC 2000, pages 1285 to 1290.
- [26] A. Hamidi et al.: "Reliability of High Power IGBT Modules Testing on Thermal Fatigue Effects due to Traction Cycles", EPE 1997, volume 3, pages 118 to 122.
- [27] S. Large, P. Walker: "Rapid Cycle Testing of High Current IGBT Powerswitch Modules", PEVD 2000, pages 235 to 240.
- [28] M. A. E. Andersen: "Comparison of Three IGBT Inverters, One Hard-Switched and Two with Snubber Circuits Using a Minimum Number of Components", EPE 1993, pages 306 to 311.
- [29] L. Abraham et al.: "Investigation on IGBT Switching Process with Variable Gate Charge Current", EPE 1993, pages 323 to 327.
- [30] H.-G. Eckel, L. Sack: "Optimization of the Turn-Off Performance of IGBT at Overcurrent and Short-Circuit Current", EPE 1993, pages 317 to 322.
- [31] A. Galluzzo et al.: "Switching Characteristic Improvement of Modern Gate Controlled Devices", EPE 1993, pages 374 to 379.
- [32] S. Gediga et al.: "High Power IGBT Converters with New Gate Drive and Protection Circuit", EPE 1995, volume 1 pages 66 to 70.
- [33] C. Gerster et al.: "Gate-control Strategies for Snubberless Operation of Series Connected IGBTs", PESC 1996, pages 1739 to 1742.
- [34] P. R. Palmer, A. N. Githiari: "The Series Connection of IGBTs with Optimized Voltage Sharing in the Switching Transient", PESC 1995, pages 44 to 49.
- [35] C. Licitra et al.: "A New Driving Circuit for IGBT Devices", IEEE Transactions on Power Electronics, May 1995, pages 373 to 378.
- [36] A. Brambilla, E. Dallago: "A Driver Circuit for High Power GTO Devices", EPE 1991, volume 2, pages 448 to 451.

-
- [37] P. R. Palmer et al.: "A Comparison of IGBT Technologies for use in the Series Connection", PEVD 1996, pages 236 to 241.
- [38] F. Auerbach et al.: "6.5kV IGBT Modules", IAS 1999.
- [39] Y. Takahashi et al.: "2.5kV - 1000A A Power Pack IGBT (High Power Flat-Packaged NPT Type RC-IGBT)", IEEE Transactions on Electron Devices, January 1999, pages 245 to 250.
- [40] R. Letor: "Static and Dynamic Behaviour of Paralleled IGBTs", IEEE Transactions on Industry Applications, March/April 1992, pages 395 to 402.
- [41] B. V. Cordingley: "Wideband Terminated Current Transformers for Power Electronic Measurements", PEVD 1998, pages 433 to 436.
- [42] Pearson Electronics: "Pearson Wide Band Current Monitors", <http://www.pearsonelectronics.com/>
- [43] L. Ghislanzoni, J. A. Carrasco: "A DC Current Transformer for Large Bandwidth and High Common-Mode Rejection", IEEE Transactions on Industrial Electronics, June 1999, pages 631 to 636.
- [44] F. Costa et al.: "Wide Bandwidth, Large AC Current Probe for Power Electronics and EMI Measurements", IEEE Transactions on Industrial Electronics, August 1997, pages 502 to 511.
- [45] Tektronix: "Current Measurement Systems", <http://www.tek.com/>
- [46] J. P. Etter, J. S. Soffer: "Current and Voltage LEM Module", PEVD 1986, pages 164 to 167.
- [47] R. Bakran et al.: "A Novel System for Current and Voltage Sensing in Traction Converters", EPE 1997, volume 3 pages 294 to 299.
- [48] W. Rogowski and W. Steinhaus: "Die Messung der Magnetischen Spannung", arch. Elektrotech 1, pages 141 to 150, 1912.
- [49] R. T. Thomas: "High impulse voltage and current measurement", IEEE Transactions volume IM-19, no. 2, pages 112 to 117, 1970.
- [50] J. Cooper: "On the High Frequency Response of a Rogowski Coil", Journal of Nuclear Energy Part C volume 5 pages 285 to 289, 1973.

-
- [51] W. F. Ray and K. D. Murray: "The Use of Rogowski Coils for Current Waveform Measurement in Power Electronic Circuits", EPE 1991, volume 3, pages 379 to 383.
- [52] W. F. Ray and R. M. Davis: "Wide Bandwidth Rogowski Current Transducers: Part 1 - The Rogowski Coil", EPE Journal volume 3 no. 1, March 1993, pages 51 to 59.
- [53] W. F. Ray and R. M. Davis: "Wide Bandwidth Rogowski Current Transducers: Part 2 - The Integrator", EPE Journal volume 3 no. 2, June 1993, pages 116 to 122.
- [54] W. F. Ray and R. M. Davis: "High Frequency Improvements in Wide Bandwidth Rogowski Current Transducers", EPE 1999.
- [55] W. F. Ray and R. M. Davis: "Developments in Rogowski Current Transducers", EPE 1997, volume 3 pages 308 to 312.
- [56] D. Metzner et al.: "Non-Destructive Evaluation of the Circuit-Dependent Turn-Off Limits of Gate Turn-Off Thyristors by the Detection of the Local Current Density", ISPSD 1997, pages 85 to 88.
- [57] N. Karrer and P. Nofer-Noser: "PCB Rogowski Coils for High di/dt Current Measurement", PESC 2000, pages 1296 to 1301.
- [58] T. Ohi et al.: "Analysis and Measurement of Chip Current Imbalances Caused by the Structure of Bus Bars in an IGBT Module", IAS 1999.
- [59] F. J. Wakeman: "Non-Linear Current Transformer for Accurate Measurement of GTO Thyristor Turn-off Characteristics", EPE 1991, volume 2, pages 428 to 433.
- [60] N. Karrer et al.: "A New Current Probe with a Wide Bandwidth", EPE 1999.
- [61] N. Karrer et al.: "HOKA: A New Isolated Current Measuring Principle and its Features", IAS 1999.
- [62] M. J. Prieto et al.: "Design of Current Sensors using Amorphous Wires", EPE 1997, volume 3, pages 300 to 304.

-
- [63] A. Radun: "An Alternative Low-Cost Current-Sensing Scheme for High-Current Power Electronics Circuits", IEEE Transactions on Industrial Electronics, February 1995, pages 78 to 84.
- [64] P. R. Palmer, C. M. Johnson: "Measurement of the Redistribution of Current in GTO Thyristors During Turn Off", EPE 1989, pages 1621 to 1625.
- [65] P. R. Palmer et al.: "Non-Invasive Measurement of Chip Currents in IGBT Modules", PESC 1997, pages 166 to 171.
- [66] B. Detemmerman: "Parallel and Series Connection of GTOs in Traction Applications", EPE 1985, volume 1, pages 73 to 79.
- [67] Y. Tadros et al.: "Design Aspects of High Power PWM Inverters with IGBT", EPE 1991, volume 2, pages 83 to 88.
- [68] A. Merazga et al.: "Series and Parallel Operating of MOS Controlled Thyristor", EPE 1997, volume 4, pages 18 to 23.
- [69] Siemens: "Connecting IGBTs in Parallel (Fundamentals)", IGBT Application Hints, <http://www.infineon.com>.
- [70] International Rectifier: "Application Characterization of IGBTs", IR technical note 990.
- [71] International Rectifier: "A Chopper for Motor Speed Control Using Parallel Connected Power HEXFETs", IR technical note 941B.
- [72] Dynex Semiconductor: "Parallel Application of IGBTs", Application note AN4508-3.0, January 2000.
- [73] R. Locher: "Parallel Operation of IGBT Discrete Device", IXYS corporation tip 10-14-96, <http://www.ixys.net/t101496.html>
- [74] Ch. Keller, Y. Tadros: "Are Paralleled IGBT Modules or Paralleled IGBT Inverters the Better Choice?", EPE 1993, pages 1 to 6.
- [75] S. R. Korn: "Parallel Operation of the Insulated Gate Transistor in Switching Operations", PCIM 1986, pages 218 to 234.
- [76] P. O. Jeannin et al.: "Influence of Stray Inductances on Current Sharing during Switching Transitions in Paralleled Semiconductors", EPE 1999.

[77] Thomson-CSF: "The Power Transistor in its Environment", October 1978, chapter 9.

[78] Siliconix Inc.: "Parallel Operation of Power MOSFETs", Siliconix Inc. application note TA 84-5, June 1984.

[79] T. Ogura et al.: "High Turn-off Current Capability of Parallel-connected 4.5kV Trench-IEGTs", ISPSD 1998, pages 47 to 50.

[80] C. Wong, K. Priest: "Current Sharing for Paralleled IGBTs Using Statistics Method", IAS 1996, pages 1418 to 1424.

[81] Mitsubishi Semiconductors: "Using IGBT Modules", September 1998, pages 21 to 26.

http://www.mitsubishichips.com/data/datasheets/power/pdfs/powermos4_0.pdf

[82] D. A. Dapkus II: "Parallel Operation of IGBTs", Texas Instruments DT 94-6.

[83] M. Alyas et al.: "Fast High Current Switch for Higher Power Switching Power Supplies", PEVD 1988, pages 17 to 20.

[84] H. A. C. Braga, I. Barbi: "A New Technique for Parallel Connection of Commutation Cells: Analysis, Design and Experimentation", PESC 1995, pages 81 to 87.

[85] P. Hofer et al.: "Paralleling Intelligent IGBT Power Modules with Active Gate-Controlled Current Balancing", PESC 1996, pages 1312 to 1316.

[86] P. Hofer-Noser, N. Karrer: "Modelling of Power Converters Using Paralleled Intelligent IGBT Power Modules", EPE 1997, volume 2, pages 256 to 261.

[87] P. Hofer-Noser, N. Karrer: "Monitoring of Paralleled IGBT / Diode Modules", IEEE Transactions on Power Electronics, May 1999, pages 438 to 444.

[88] W. Bösterling et al.: "IGBT Modules: Concept - Gate Drive - Fault Protection", EPE 1989, pages 599 to 604.

[89] K. Sommer et al.: "Multichip High Power IGBT Modules for Traction and Industrial Application", EPE 1997, volume 1, pages 112 to 116.

-
- [90] A. Camera et al.: "Electrical and Thermal Characterization of High Power Press-Packed IGBT", EPE 1999.
- [91] T. Franke et al.: "Current and Temperature Distribution in Multi-Chip Modules under Inverter Operation", EPE 1999.
- [92] M. Mori, R. Saitou, T. Yatsuo: "A High Power IGBT Module for Traction Motor Drive", ISPSD 1993, pages 287 to 291.
- [93] S. Januszewski et al.: "Causes and Mechanisms of Semiconductor Device Failures in Power Converter Service Conditions", EPE 1995, volume 1, pages 625 to 630.
- [94] J. Yamashita et al.: "A Study on the IGBT's Turn-off Failure and Inhomogeneous Operation", ISPSD 1994, pages 45 to 50.
- [95] M. Tabata et al.: "Control Methods of Current Balancing for Parallel Connected IGBTs", ISPSD 1998, pages 101 to 104.
- [96] H. Miyazaki et al.: "Neutral-Point-Clamped Inverter with Parallel Driving of IGBT's for Industrial Applications", IEEE Transactions on Industry Applications, January 2000, pages 146 to 151.
- [97] A. R. Hefner: "An Improved Understanding for the Transient Operation of the Power Insulated Gate Bipolar Transistor (IGBT)", PESC 1989, pages 303 to 313.
- [98] Orcad Inc.: "Orcad Capture", <http://www.orcad.com/>
- [99] Mathsoft Inc.: "Mathcad", <http://www.mathcad.com/mathcad>
- [100] Silvaco Inc.: "Virtual Wafer Fab", <http://www.silvaco.com/>
- [101] Hexfet databook p. 5
- [102] A. Mauder, W. Scholz: "Investigation of the Static and Dynamic Current Distribution in Paralleled IGBT Modules", Power Conversion, June 1997, pages 275 to 284.
- [103] B. J. Baliga: "Modern Power Devices", John Wiley & Sons, 1987.
- [104] C. M. Johnson, P. R. Palmer: "Simulation of Wafer-Scale GTO Thyristors in Circuits", IEEE Transactions on Power Electronics, April 1991, pages 308 to 313.

[105] G. Busatto, G. P. Pepe, A. Ruosi, M. Valentino: "Non-Destructive Detection of Current in Power Modules by Means of Magnetic Measurements", IAS 2000.

[106] IGBT databook - 'Reliability - Extract from Semiconductor Reliability Handbook (Discrete Device) 1993', Toshiba, 1995.


[107] I. Omura, W. Fichtner, H. Ohashi: "Oscillation Effects in IGBT's related to Negative Capacitance Phenomena", IEEE Transactions on Electron Devices, January 1999, pages 237 to 244.

[108] I. Omura et al.: "IEGT Design Concept Against Operation Instability and its Impact to Application", ISPSD 2000, pages 25 to 28.

[109] P. R. Palmer, H. S. Rajamani, J. C. Joyce: "Behaviour of IGBT Modules Under Short Circuit Conditions", IAS 2000.

B

Derivation of the IGBT's Characteristic Equation



Single IGBT equivalent circuit, with parasitic inductance and resistance in each connection

We start by writing down the voltages and currents at each of the device's terminals. Let $Z_G = R_G + sL_G$, and similarly for Z_C and Z_E ; the current in a capacitor $I_C = sCV_C$; the voltage across an inductor $V_L = sLI_L$. The collector current due to the gate voltage is $g_m(V_G - V_E)$. Capacitance C_{ij} represents the capacitance between two terminals; this can also include any stray capacitances due to the external circuit. Thus we start with:

$$V_G = I_G Z_G \quad (B.1)$$

$$V_E = I_E Z_E \quad (B.2)$$

$$V_C = I_C Z_C \quad (B.3)$$

$$I_G = C_{GC} s(V_C - V_G) + C_{GE} s(V_E - V_G) \quad (B.4)$$

$$I_C = C_{GC} s(V_G - V_C) + C_{CE} s(V_E - V_C) - g_m(V_G - V_E) - (V_C - V_E)/R_O \quad (B.5)$$

$$I_G + I_C + I_E = 0 \quad (B.6)$$

Substitute (B.1), (B.2), (B.3) into (B.4), (B.5) to eliminate V_C, V_E, V_G :

$$I_G = C_{GC} s(I_C Z_C - I_G Z_G) + C_{GE} s(I_E Z_E - I_G Z_G) \quad (B.7)$$

$$I_C = C_{GC} s(I_G Z_G - I_C Z_C) + C_{CE} s(I_E Z_E - I_C Z_C) - g_m(I_G Z_G - I_E Z_E) - (I_C Z_C - I_E Z_E)/R_O \quad (B.8)$$

Substitute (B.6) into (B.7), (B.8) to eliminate I_E :

$$I_G = C_{GC}s(I_C Z_C - I_G Z_G) - C_{GE}s((I_C + I_G)Z_E + I_G Z_G) \quad (B.9)$$

$$I_C = C_{GC}s(I_G Z_G - I_C Z_C) - C_{CE}s((I_C + I_G)Z_E + I_C Z_C) - g_m(I_G Z_G + (I_C + I_G)Z_E) - (I_C Z_C + (I_C + I_G)Z_E)/R_O \quad (B.10)$$

Rearrange (B.9):

$$I_G = I_C s(C_{GC}Z_C - C_{GE}Z_E) - I_G s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E) \\ I_G(1 + s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E)) = I_C s(C_{GC}Z_C - C_{GE}Z_E) \quad (B.11)$$

Rearrange (B.10):

$$I_C = I_G(s(C_{GC}Z_G - C_{CE}Z_E) - g_m(Z_G + Z_E) - Z_E/R_O) - I_C(s(C_{GC}Z_C + C_{CE}(Z_C + Z_E)) + g_m Z_E + (Z_C + Z_E)/R_O) \\ I_G(s(C_{GC}Z_G - C_{CE}Z_E) - g_m(Z_G + Z_E) - Z_E/R_O) = I_C(1 + s(C_{GC}Z_C + C_{CE}(Z_C + Z_E)) + g_m Z_E + (Z_C + Z_E)/R_O) \quad (B.12)$$

Divide (B.11) by (B.12) and multiply out:

$$(1 + s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E))(1 + s(C_{GC}Z_C + C_{CE}(Z_C + Z_E)) + g_m Z_E + (Z_C + Z_E)/R_O) = (s(C_{GC}Z_G - C_{CE}Z_E) - g_m(Z_G + Z_E) - Z_E/R_O)(s(C_{GC}Z_C - C_{GE}Z_E)) \quad (B.13)$$

Rearrange (B.13) and factorize:

$$s^2((Z_G Z_C + Z_G Z_E + Z_C Z_E)(C_{GC}C_{GE} + C_{GC}C_{CE} + C_{GE}C_{CE})) + s(Z_G(C_{GC} + C_{GE}) + Z_E(C_{CE} + C_{GE}) + Z_C(C_{CE} + C_{CG})) + (g_m Z_G + Z_E/R_O)(C_{GC}Z_C - C_{GE}Z_E) + g_m Z_E(C_{GC}Z_G + C_{GE}Z_G + C_{GC}Z_C) + ((Z_C + Z_E)/R_O)(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E) + g_m Z_E + (Z_C + Z_E)/R_O + 1 = 0 \quad (B.14)$$

Define constants:

$$C_G = C_{GC} + C_{GE}$$

$$C_C = C_{GC} + C_{CE}$$

$$C_E = C_{CE} + C_{GE}$$

$$C_t = C_{GC}C_{GE} + C_{GC}C_{CE} + C_{GE}C_{CE}$$

and substitute in (B.14) to yield:

$$s^2 C_t (Z_G Z_C + Z_G Z_E + Z_C Z_E) + s(Z_G C_G + Z_E C_E + Z_C C_C + (g_m Z_G + Z_E/R_O)(C_{GC}Z_C - C_{GE}Z_E))$$

$$\begin{aligned}
& + g_m Z_E (C_{GC} Z_G + C_{GE} Z_G + C_{GC} Z_C) + ((Z_C + Z_E)/R_O)(C_{GC} Z_G + C_{GE} Z_G + C_{GE} Z_E) \\
& + g_m Z_E + (Z_C + Z_E)/R_O + 1 = 0 \tag{B.15}
\end{aligned}$$

Now replace Z_G with $R_G + sL_G$ (et al.) and multiply out:

$$\begin{aligned}
& s^2 C_t ((R_G + sL_G)(R_C + sL_C) + (R_G + sL_G)(R_E + sL_E) + (R_C + sL_C)(R_E + sL_E)) \\
& + s((R_G + sL_G)C_G + (R_E + sL_E)C_E + (R_C + sL_C)C_C \\
& + (g_m R_G + s g_m L_G + R_E/R_O + s L_E/R_O)(C_{GC}(R_C + sL_C) - C_{GE}(R_E + sL_E)) \\
& + (g_m(R_E + sL_E))(C_{GC}(R_G + sL_G) + C_{GE}(R_G + sL_G) + C_{GC}(R_C + sL_C)) \\
& + ((R_C + R_E + s(L_C + L_E))/R_O)(C_{GC}(R_G + sL_G) + C_{GE}(R_G + sL_G) + C_{GE}(R_E + sL_E))) \\
& + g_m(R_E + sL_E) + (R_C + R_E + s(L_C + L_E))/R_O + 1 = 0 \tag{B.16}
\end{aligned}$$

Group together in s :

$$\begin{aligned}
& s^4 C_t (L_G L_C + L_G L_E + L_C L_E) \\
& + s^3 (C_t (L_G R_C + R_G L_C + R_G L_E + R_E L_G + R_C L_E + R_E L_C) \\
& + (L_E/R_O)(L_C C_{GC} + L_G C_{GE} + L_G C_{GC}) + (L_C/R_O)(C_{GC} L_G + C_{GE} L_G + C_{GE} L_E) \\
& + g_m C_{GC}(L_E L_G + L_C L_G + L_C L_E)) \\
& + s^2 (C_t (R_G R_C + R_G R_E + R_C R_E) + (L_G C_G + L_E C_E + L_C C_C) \\
& + (L_E/R_O)(C_{GC} R_C + C_{GC} R_G + C_{GE} R_G) + (L_C/R_O)(C_{GC} R_G + C_{GE} R_G + C_{GE} R_E) \\
& + (R_E/R_O)(C_{GC} L_G + C_{GC} L_C + C_{GE} L_G) + (R_C/R_O)(C_{GC} L_G + C_{GE} L_G + C_{GE} L_E) \\
& + g_m C_{GC}(L_E R_G + L_G R_E + L_E R_C + R_E L_C + L_G R_C + L_C R_G) \\
& + s (R_G C_G + R_E C_E + R_C C_C + g_m L_E + (L_C + L_E)/R_O \\
& + (R_E/R_O)(C_{GC} R_C + C_{GC} R_G + C_{GE} R_G) + (R_C/R_O)(C_{GC} R_G + C_{GE} R_G + C_{GE} R_E) \\
& + g_m C_{GC}(R_G R_E + R_G R_C + R_C R_E)) \\
& + g_m R_E + (R_C + R_E)/R_O + 1 = 0 \tag{B.17}
\end{aligned}$$

This is the characteristic equation, which holds in all cases covered by the initial circuit; however the number of variables renders it rather complicated for further analysis. A number of special cases can be considered to determine whether this equation is consistent with previously published results. Setting $R_O = \infty$ and $g_m = 0$ will leave a circuit with the three terminals symmetrical; the equation does then indeed reduce to one symmetrical in C , E and G :

$$s^4 C_t (L_G L_C + L_G L_E + L_C L_E) + s^3 (C_t (L_G R_C + R_G L_C + R_G L_E + R_E L_G + R_C L_E + R_E L_C))$$

$$\begin{aligned}
& + s^2 (C_t (R_G R_C + R_G R_E + R_C R_E) + (L_G C_G + L_E C_E + L_C C_C)) + s (R_G C_G + R_E C_E + R_C C_C) \\
& + 1 = 0
\end{aligned} \tag{B.18}$$

One can also try Kassakian & Lau's special cases; in both cases $R_O = \infty$. Firstly, $R_E = L_E = R_C = 0$:

$$\begin{aligned}
& s^4 (C_t L_G L_C) + s^3 (C_t R_G L_C + g_m L_G L_C C_{GC}) + s^2 (L_G C_G + L_C C_C + g_m R_G C_{GC} L_C) \\
& + s (R_G C_G) + 1 = 0
\end{aligned} \tag{B.19}$$

which is in agreement. Secondly, for $L_C = R_E = 0$:

$$\begin{aligned}
& s^4 (C_t L_G L_E) + s^3 (C_t (L_G R_C + R_G L_E + R_C L_E) + g_m L_E C_{GC} L_G) \\
& + s^2 (C_t R_G R_C + L_G C_G + L_E C_E + g_m C_{GC} (L_G R_C + L_E R_G + L_E R_C)) \\
& + s (R_G C_G + R_C C_C + g_m (L_E + R_G C_{GC} R_C)) + 1 = 0
\end{aligned} \tag{B.20}$$

c.f. (expressed in terms of C, E, G instead of D, S, G)

$$\begin{aligned}
& s^4 (L_G L_E C_t) + s^3 (C_t (R_C L_E + R_G L_E + R_C L_G) + g_m L_E L_G C_{GC}) \\
& + s^2 (R_G R_C C_t + L_G C_G + L_E C_E + g_m C_{GC} (R_C L_E + R_G L_E + R_C L_G)) \\
& + s (R_G C_G + R_C C_E + g_m (L_E + R_G R_C C_{GC})) + 1 = 0
\end{aligned} \tag{B.21}$$

A discrepancy is noted in the 's' term of the latter case, viz. $R_C C_C$ or $R_C C_E$.

Transfer Function

The transfer function enables Bode and Nyquist diagrams to be plotted for the IGBT when operating open-loop. The derivation is similar and presented here for interest. Equations (B.1) through (B.4) plus (B.6) are retained, and equation (B.5) is replaced to suit the open-loop situation:

$$I_C = C_{GC} s (V_G - V_C) + C_{CE} s (V_E - V_C) - g_m V_{in} - (V_C - V_E) / R_O \tag{B.22}$$

Substitute (B.1), (B.2), (B.3) into (B.4), (B.22) to eliminate V_C , V_E , V_G :

$$I_G = C_{GC} s (I_C Z_C - I_G Z_G) + C_{GE} s (I_E Z_E - I_G Z_G) \tag{B.23}$$

$$I_C = C_{GC} s (I_G Z_G - I_C Z_C) + C_{CE} s (I_E Z_E - I_C Z_C) - g_m V_{in} - (I_C Z_C - I_E Z_E) / R_O \tag{B.24}$$

Substitute (B.6) into (B.23), (B.24) to eliminate I_E :

$$I_G = C_{GC} s (I_C Z_C - I_G Z_G) - C_{GE} s ((I_C + I_G) Z_E + I_G Z_G) \tag{B.25}$$

$$\begin{aligned}
I_C = C_{GC} s (I_G Z_G - I_C Z_C) - C_{CE} s ((I_C + I_G) Z_E + I_C Z_C) \\
- g_m V_{in} - (I_C Z_C + (I_C + I_G) Z_E) / R_O
\end{aligned} \tag{B.26}$$

Rearrange (B.25):

$$\begin{aligned}
I_G &= I_C s(C_{GC}Z_C - C_{GE}Z_E) - I_G s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E) \\
I_G &= I_C (s(C_{GC}Z_C - C_{GE}Z_E) / (1 + s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E))) \\
&= AI_C \tag{B.27}
\end{aligned}$$

Rearrange (B.26):

$$\begin{aligned}
I_C &= I_G (s(C_{GC}Z_G - C_{CE}Z_E) - Z_E / R_O) \\
&\quad - I_C (s(C_{GC}Z_C + C_{CE}(Z_C + Z_E)) + (Z_C + Z_E) / R_O) - g_m V_{in} \\
I_G (s(C_{GC}Z_G - C_{CE}Z_E) - Z_E / R_O) \\
&= (1 + s(C_{GC}Z_C + C_{CE}(Z_C + Z_E)) + (Z_C + Z_E) / R_O) I_C - g_m V_{in} \\
BI_G &= CI_C - g_m V_{in} \tag{B.28}
\end{aligned}$$

Substitute (B.27) into (B.28) to eliminate I_G :

$$\begin{aligned}
ABI_C &= CI_C - g_m V_{in} \\
V_{in} &= ((C - AB) I_C) / g_m \tag{B.29}
\end{aligned}$$

Define $V_{out} = V_G - V_E$

$$\begin{aligned}
V_{out} &= I_G Z_G + (I_G + I_C) Z_E \\
&= (AZ_G + AZ_E + Z_E) I_C \tag{B.30}
\end{aligned}$$

Divide (B.30) by (B.29):

$$\begin{aligned}
V_{out} / V_{in} &= g_m (A(Z_G + Z_E) + Z_E) / (C - AB) \\
&= g_m ((s(C_{GC}Z_C - C_{GE}Z_E)(Z_G + Z_E) + Z_E) + (1 + s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E))Z_E) \\
&\quad / ((1 + s(C_{GC}Z_C + C_{CE}(Z_C + Z_E)) + (Z_C + Z_E) / R_O) (1 + s(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E)) \\
&\quad - s(C_{GC}Z_C - C_{GE}Z_E) (s(C_{GC}Z_G - C_{CE}Z_E) - Z_E / R_O)) \\
&= g_m (sC_{GC}(Z_C Z_G + Z_C Z_E + Z_G Z_E) + Z_E) \\
&\quad / s^2 ((Z_G Z_C + Z_G Z_E + Z_C Z_E)(C_{GC}C_{GE} + C_{GC}C_{CE} + C_{GE}C_{CE})) \\
&\quad + s(Z_G C_G + Z_E C_E + Z_C C_C + (Z_E / R_O)(C_{GC}Z_C - C_{GE}Z_E) \\
&\quad + ((Z_C + Z_E) / R_O)(C_{GC}Z_G + C_{GE}Z_G + C_{GE}Z_E)) \\
&\quad + (Z_C + Z_E) / R_O + 1 \tag{B.31}
\end{aligned}$$

Expand Z_C , Z_E , Z_G , multiply out and rearrange:

$$\begin{aligned}
V_{in}/V_{out} = & g_m(B \cdot (s^3 C_{GC}(L_C L_G + L_C L_E + L_G L_E) \\
& + s^2 C_{GC}(R_C L_G + R_G L_C + R_C L_E + R_E L_C + R_G L_E + R_E L_G) \\
& + s(C_{GC}(R_C R_G + R_C R_E + R_G R_E) + L_E) \\
& + R_E) \\
& / (B \cdot s^4 C_t(L_G L_C + L_G L_E + L_C L_E) \\
& + s^3 (C_t(L_G R_C + R_G L_C + R_G L_E + R_E L_G + R_C L_E + R_E L_C) \\
& + (L_E/R_O)(L_C C_{GC} + L_G C_{GE} + L_G C_{GC}) + (L_C/R_O)(C_{GC} L_G + C_{GE} L_G + C_{GE} L_E) \\
& + s^2 (C_t(R_G R_C + R_G R_E + R_C R_E) + (L_G C_G + L_E C_E + L_C C_C) \\
& + (L_E/R_O)(C_{GC} R_C + C_{GC} R_G + C_{GE} R_G) + (L_C/R_O)(C_{GC} R_G + C_{GE} R_G + C_{GE} R_E) \\
& + (R_E/R_O)(C_{GC} L_G + C_{GC} L_C + C_{GE} L_G) + (R_C/R_O)(C_{GC} L_G + C_{GE} L_G + C_{GE} L_E)) \\
& + s (R_C C_G + R_E C_E + R_C C_C + (L_C + L_E)/R_O \\
& + (R_E/R_O)(C_{GC} R_C + C_{GC} R_G + C_{GE} R_G) + (R_C/R_O)(C_{GC} R_G + C_{GE} R_G + C_{GE} R_E) \\
& + (R_C + R_E)/R_O + 1) \tag{B.32}
\end{aligned}$$

It can be shown that setting $V_{in}/V_{out}=1$ in (B.32) (i.e. closing the loop) leads to the characteristic equation (B.17), as expected.



Publications

Parts of this research have been presented in six conference papers. The first, in the **PESC 1997** conference, covered the theory and development of the magnetic field coil current measurement technique, and some initial results. In the **EPE 1997** conference, practical measurements on matched and mismatched modules were shown, and the effect of temperature investigated. The focus moved to gate drive techniques and their consequences for the **PEVD 1998** paper, and thermal imaging was used to verify the results seen with the field coils. The simulation work with Atlas was discussed at **ISPSD 1999**, and **EPE 1999** saw the stability analysis presented. Some aspects of module behaviour during short circuit conditions were presented at **IAS 2000**, including current sharing measurements and simulations under these short circuit conditions.

Palmer P.R., Stark B.H., Joyce J.C.: *Non-Invasive Measurement of Chip Currents in IGBT Modules.* Power Electronics Specialists' Conference, St Louis, 1997.

Palmer P.R., Joyce J.C., Stark B.H.: *Measurement of Chip Currents in IGBT Modules.* 7th European Conference on Power Electronics and Applications, Trondheim, 1997.

Palmer P.R., Joyce J.C.: *Current redistribution in multi-chip IGBT modules under various gate drive conditions.* Seventh International Conference on Power Electronics and Variable Speed Drives, London, 1998.

Joyce J.C., Palmer P.R.: *Some causes of current redistribution in IGBT modules.* 11th International Symposium on Power Semiconductor Devices, Toronto 1999.

Palmer P.R., Joyce J.C.: *Causes of parasitic current oscillations in IGBT modules at turn-off.* 8th European Conference on Power Electronics and Applications, Lausanne, 1999.

Palmer P.R., Rajamani H.S., Joyce J.C.: *Behaviour of IGBT modules under short circuit conditions.* IEEE Industrial Applications Society Annual Meeting, Rome, 2000.