Characterisation of the MuSIC muon beam and design of the Eu-XFEL LPD/CCC interface firmware

Samuel Leslie Cook

UCL

High Energy Physics

Submitted to UCL in fulfilment of the requirements for the award of the degree of PhD.

Declaration

I, Samuel Leslie Cook, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Samuel Leslie Cook

Abstract

As it is now known that neutrinos oscillate, this entails that lepton number is not a conserved quantity and required modification of the Standard Model. The same mechanism also allows charged leptons to violate lepton number, but at an immeasurable rate. Therefore any observation of e.g. a muon decaying to an electron in the absence of neutrinos would be a signal of physics beyond the Standard Model. A test facility, but also an experiment capable of searching for such a kinematic effect has been set-up in Osaka, Japan.

The thesis 'Initial Measurements at the MuSIC Beam-Line' documents several experiments that were carried out using the MuSIC muon beam in Osaka. The experiments focused on understanding the beam and characterising it. Three main measurements were made: total charged particle flux, muon flux and muon momentum spectrum.

Each experiment used plastic scintillators and MPPCs to detect the charged particles. Muons were identified by looking for muon decays between two scintillators. Data acquisition was performed using NIM for signal shaping and logic; and CAMAC or VME for readout via TDC and ADC modules.

In addition to direct experimental measurements, a simulation of MuSIC was made using the 'G4Beamline' and 'Geant4' packages which allowed detailed interrogation of the experiment to aid understanding of the results. It also provided a test bed upon which to refine the setup for later measurements.

These measurements allow us to confirm the performance of the novel pion capture solenoid which is an integral part of several future experiments (primarily COMET but also proposed neutrino factories).

In addition to this core work, the thesis deals with my services writing firmware for the LPD detector being built for the XFEL project in Hamburg, Germany. The firmware is an interface between the generic clock and control card (being developed at UCL) and the custom ASIC (being developed at RAL). The interface was written using VHDL and receives, translates, interprets then transmits control instructions to the detector.

Contents

Ι	\mathbf{Ch}	naracterisation of the MuSIC beam				25
1	Exe	cutive Summary				27
2	Intr	roduction				31
	2.1	Muon Sources				31
		2.1.1 Scientific Motivation for High Intensity Muon Sources				33
	2.2	MuSIC				34
		2.2.1 MuSIC design				36
		2.2.2 MuSIC: Scientific Motivation				37
	2.3	Physics				39
		2.3.1 Charged Particle Interactions in Matter				39
		2.3.2 Muon decay				42
		2.3.3 Charged Lepton Flavour Violation		•	•	44
3	\mathbf{Sim}	ulation				49
	3.1	Introduction				49
	3.2	Geant4				49
		3.2.1 Detector Constructor				50
		3.2.2 Execution model				51
	3.3	G4beamline				52
	3.4	Implementation				53
		3.4.1 G4Beamline				53
		3.4.2 Geant4				59
		3.4.3 Detector Constructor				59
		3.4.4 Primary Action Generator				64
		3.4.5 Physics List				64
		3.4.6 The Field				65
	3.5	Simulation Analysis		•	•	67
4	Cha	racterising the beam				71
	4.1	Introduction				71
		4.1.1 Scintillators Preparation				72
		4.1.2 Multi-Photon Pixel Counters				72

	4.1.0	Secondary Emission Chamber
	4.1.4	NIM, CAMAC and VME
	4.1.5	Discriminators
	4.1.6	Gates and Latches
	4.1.7	Logic units
	4.1.8	Scaler
	4.1.9	Analogue to Digital Converter
	4.1.10	Time to Digital Converter
	4.1.11	Registers
4.2	Charge	ed Particle Flux
	4.2.1	Experimental Set Up
	4.2.2	Detector efficiency
	4.2.3	Results
	4.2.4	Analysis
4.3	Muon	Lifetime
	4.3.1	Experimental Set Up
	4.3.2	Results
	4.3.3	Analysis
4.4	Momer	ntum Spectrum
	4.4.1	Experimental Set Up
	4.4.2	Data Processing
	4.4.3	Secondary calculations
	4.4.4	Results
	4.4.5	Analysis
ТС	4.4.5	analysis
IC		ind Control Interface for the Darge I fixer Detector 1
I C Inti	roductio	on 1
I C Inti 5.1	r oducti e Conver	on 1 ntions 1
Int: 5.1 5.2	r oducti Conver The Eu	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1
I C Inti 5.1 5.2	r oducti Conver The Eu 5.2.1	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1
I C Inti 5.1 5.2	roduction Convert The Eu 5.2.1 5.2.2	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1
Inti 5.1 5.2	roductie Conver The Eu 5.2.1 5.2.2 5.2.3	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1 Scientific Motivation 1
Int 5.1 5.2 5.3	roductie Conver The Eu 5.2.1 5.2.2 5.2.3 Detecte	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1 Scientific Motivation 1 ors at EuXFEL 1
I C Inti 5.1 5.2 5.3	roductic Conver The Eu 5.2.1 5.2.2 5.2.3 Detecto 5.3.1	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1 Scientific Motivation 1 ors at EuXFEL 1 The Large Pixel Detector (LPD) 1
I C Intr 5.1 5.2 5.3 5.4	roducti Conver The Eu 5.2.1 5.2.2 5.2.3 Detecto 5.3.1 DAQ a	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1 Scientific Motivation 1 ors at EuXFEL 1 The Large Pixel Detector (LPD) 1 und control systems 1
Int: 5.1 5.2 5.3 5.4	roducti Conver The Eu 5.2.1 5.2.2 5.2.3 Detecto 5.3.1 DAQ a 5.4.1	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1 Scientific Motivation 1 ors at EuXFEL 1 The Large Pixel Detector (LPD) 1 und control systems 1 Clock and Control Card (CCC) 1
Inti 5.1 5.2 5.3 5.4 5.5	roducti Conver The Eu 5.2.1 5.2.2 5.2.3 Detecte 5.3.1 DAQ a 5.4.1 Firmwa	on 1 ntions 1 uropean X-ray Free-Electron Laser: an overview 1 Synchrotrons and FELs 1 X-ray production at EuXFEL 1 Scientific Motivation 1 ors at EuXFEL 1 The Large Pixel Detector (LPD) 1 und control systems 1 clock and Control Card (CCC) 1 are 1

6.1 Requirements 133 6.2 Design Principles 134 6.3 The Design 135 7 Implementation 137 7.1 Top level 137 7.1.1 Interface 137 7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Generics 139 7.1.4 Implementation 141 7.2 Registers 139 7.1.3 Interface 142 7.2.1 Interface 142 7.2.2 Implementation 142 7.3.2 Registers 149 7.3.3 Implementation 159 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop	6	Des	ign 13	3
6.2 Design Principles 134 6.3 The Design 135 7 Implementation 137 7.1 Top level 137 7.1.1 Interface 137 7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 142 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2.1 Start, no-veto, stop 1		6.1	Requirements	3
6.3 The Design 135 7 Implementation 137 7.1 Top level 137 7.1.1 Interface 137 7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 141 7.3.1 Interface 142 7.3.2 Registers 149 7.3.3 Implementation 150 7.4.4 Transmitter 154 7.3.2 Registers 149 7.3.3 Implementation 150 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 166 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178<		6.2	Design Principles	4
7 Implementation 137 7.1 Top level 137 7.1.1 Interface 137 7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3.1 Interface 142 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA		6.3	The Design	5
7.1 Top level 137 7.1.1 Interface 137 7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2. Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3.1 Interface 142 7.3.2 Registers 144 7.3.1 Interface 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 <td>7</td> <td>Imp</td> <td>lementation 13</td> <td>7</td>	7	Imp	lementation 13	7
7.1.1 Interface 137 7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 150 7.4 Transmitter 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 150 7.4.5 Registers 155 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2.1 Start, no-veto, stop 166 8.2.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA <t< td=""><td></td><td>7.1</td><td>Top level</td><td>7</td></t<>		7.1	Top level	7
7.1.2 Generics 139 7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Receiver 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 183			7.1.1 Interface	7
7.1.3 Registers 139 7.1.4 Implementation 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.2 Interface 154 7.4.3 Registers 154 7.4.4 Implementation 150 7.4 Transmitter 154 7.4.3 Registers 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler			7.1.2 Generics	9
7.1.4 Implementation 141 7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix			7.1.3 Registers	9
7.2 Receiver 141 7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3 Veto Filter 144 7.3 Registers 149 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 1			7.1.4 Implementation \ldots 14	1
7.2.1 Interface 142 7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words		7.2	Receiver	1
7.2.2 Implementation 142 7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 155 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link in			7.2.1 Interface	2
7.3 Veto Filter 144 7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 155 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184			7.2.2 Implementation $\ldots \ldots 14$	2
7.3.1 Interface 147 7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 184 A.3 Local Link interface 184		7.3	Veto Filter	4
7.3.2 Registers 149 7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			7.3.1 Interface	7
7.3.3 Implementation 150 7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			7.3.2 Registers	9
7.4 Transmitter 154 7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			7.3.3 Implementation $\ldots \ldots 15$	0
7.4.1 Controlling the ASIC 154 7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2 Results 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184		7.4	Transmitter	4
7.4.2 Interface 155 7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 166 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184			7.4.1 Controlling the ASIC \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots 15	4
7.4.3 Registers 156 7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 165 8.2.2 Veto/No-veto 166 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			7.4.2 Interface	5
7.4.4 Implementation 159 8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 166 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			7.4.3 Registers	6
8 Testing 165 8.1 Methodology 165 8.2 Results 165 8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			7.4.4 Implementation $\dots \dots \dots$	9
8.1 Methodology 165 8.2 Results 165 8.2.1 Start, no-veto, stop 165 8.2.2 Veto/No-veto 166 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184	8	Test	ing 16	5
8.2 Results 165 8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184		8.1	Methodology	5
8.2.1 Start, no-veto, stop 166 8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184		8.2	Results	5
8.2.2 Veto/No-veto 168 8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			8.2.1 Start, no-veto, stop	6
8.2.3 Reset Command 170 8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			8.2.2 Veto/No-veto	8
8.2.4 Stop with down-scaler 173 8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			8.2.3 Reset Command	0
8.2.5 RDMA 176 8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			8.2.4 Stop with down-scaler	3
8.2.6 Local Link 178 9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			8.2.5 RDMA	6
9 Conclusion 181 A EuXFEL Appendix 183 A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184			8.2.6 Local Link	8
A EuXFEL Appendix183A.1 ASIC command words183A.2 RDMA interface184A.3 Local Link interface184	9	Con	clusion 18	1
A.1 ASIC command words 183 A.2 RDMA interface 184 A.3 Local Link interface 184	А	Euλ	(FEL Appendix 18	3
A.2 RDMA interface 184 A.3 Local Link interface 184		A.1	ASIC command words	3
A.3 Local Link interface		A.2	RDMA interface	4
		A.3	Local Link interface	4

List of Figures

1.1	Measured muon momentum spectrum for freely decaying muons at MuSIC.	
	The muon rate is normalised to a proton beam current of 1 nA. \ldots .	29
2.1	Feynman diagram showing muon production through pion decay	32
2.2	Schematic of the completed MuSIC complex. Taken from [1]	35
2.3	Schematic of the current status of MuSIC. Taken from [11]	35
2.4	Photo of MuSIC showing its 2010 status, since then it has had further concrete	
	blocks added around it to increase shielding	36
2.5	Energy loss by a muon in copper as a function of $\beta\gamma$ or momentum. The solid	
	line indicates the total stopping power, dashed lines indicate different contri-	
	butions and the vertical grey bands indicate the different phenomenological	10
	regimes. Taken from ref [2].	40
2.6	Electron energy loss per 'radiation length' (X_0^{-1}) as a function of electron	
	energy (not momentum). The material in this case is lead ($Z = 82$ and $V_{z} = 6.37 \sigma/\text{cm}^{2}$). Taken from ref [2]	/1
97	$X_0 = 0.51$ g/cm ⁻). Taken from ref [2]	41
2.1	Nuclear capture of a norative muon. This process is very hard to detect	40
2.0	directly. The net effect is to reduce the muon's apparent lifetime by removing	
	them from the population.	43
2.9	Comparison of the theoretical values (as determined using the Primakoff-	
	Goulard formula) and experimentally determined muon capture rates for a	
	range of materials, taken from [3]	44
2.10	$\mu \to e \gamma,$ as a Standard Model process. The black dot represents the conversion	
	of the muon neutrino to an electron neutrino, in the Standard Model this	
	process has a predicted branching ratio $\sim < 10^{-54}$ and has not been seen	45
2.11	Effective mass scales, Λ for expected sensitivities to muon cLFV processes for	. –
	various interaction strength ratios, κ . Figures from [22]	47
3.1	Wireframe of the G4BL simulation of MuSIC. The dark red is the iron yoke,	
	the bright red are the solenoids, the light blue is the pion capture solenoid	
	shielding and the yellow line is a simulated proton which doesn't interact in	
	this case. The co-ordinates used are defined at the end of the beam pipe	54

3.2	Distributions of the different charged particles in co-ordinates X against Y at the end of the MuSIC beam-pipe.	55
3.3	Distribution of total charged particle flux for X against Y at the end of the MuSIC beam-pipe.	56
3.4	Evolution of the particle distribution in the G4BL simulation of MuSIC. Po- sition 1 is 1 cm after the pion capture solenoid and positions 2 and 3 are after the first dipole and at the end of the beam pipe respectively	57
3.5	Position of the beam monitors used to track the particle distribution through the G4BL simulation. Note: the monitors used to create the distribution used in the final Geant4 simulation have a much smaller, 50 cm, radius compared to 2 m used in this image (as the smaller radius monitors would not be visible).	58
3.6	Schematic of the detector configuration for simulation. In the face on view the dotted lines are the scintillators, blue for the eight upstream scintillators and red for the five downstream scintillators.	60
3.7	Schematic of a downstream paddle. The upstream paddles have the same layout but are have a different width and height (0.5 mm and 30 mm respectively). The WLS fibre protrudes from the scintillator for easy attachment of the MPPCs to the support structure	61
3.8	Scintillation spectra for EJ-212. The points represent the values passed to Geant4 and are based on those read from the manufacturer's data-sheet [27].	62
3.9	Absorption and re-emission spectra for BCF-91A wavelength-shifting fibre. The points represent the values used in Geant4 to approximate the manufac- turer's curve as given in the data-sheet [28]	63
3.10	Decay time spectrum of 1×10^6 negative muons hitting a 1 mm copper stopping target in a vacuum using an unmodified version of Geant4. The decay times are fitted using $N \exp(-x/\tau)$. As can be seen the Geant4 value, 1,161 ns, does not agree with either the experimentally determined value of 163.5 ns [3] or the theoretical value of 194 ns (calculated using the equation 2.6). This bug has been fixed in more recent versions of Geant4 (versions > 4.9.6 patch 02). Only the muonic lifetime in copper is fitted as the decays are dominated by	
	captured muons.	66
3.11	Decay time spectrum of 1×10^6 negative muons hitting a 1 mm copper stopping target in a vacuum using a version of Geant4 modified to use the experimental muonic copper lifetime of 163.5 ns. The spectrum is fitted using $N \exp(-x/\tau)$. Only the muonic lifetime in copper is fitted as the decays are dominated by captured muons	67
3.12	The magnetic field strength for the region of interest. The black bar marks the position of the scintillators and stopping target.	68
4.1	Photograph of an MPPC (type S10362-11 with ceramic package) from the Hamamatsu technical report [29].	73

4.2	Oscilloscope waveform for an MPPC with a histogram of peak voltage on the left. The histogram shows the clear discrimination between number of photons available on a MPPC. This type of MPPC has $400 \times 50 \times 50 \text{ um}^2$ APDs in a	
	$1 \times 1 \text{mm}^2$ package taken from [29]	74
13	Simulated distribution of number of photons reaching one of the MPPCs at	• 1
4.0	tached wie wevelength shifting fibre to a $380 \times 50 \times 3.5$ mm ³ scintillator	74
1 1	MDDC Detern Detection Efficiency (DDE) as a function of photon wave	14
4.4	har representation of the provided and	
	for the 100U 050U and 025U models respectively). Taken from [20]	75
4 5	for the 1000, 0500 and 0250 models respectively). Taken from [29]	75
4.5	blocks indicate NIM units whilst vertical units were mounted in a CAMAC	
	crate. The two 'AND' blocks were made using NIM modules	81
4.6	Measurement of the 1D charged particle flux as given in table 4.6. Where two	
	measurements were made at a single location an offset has been applied to	
	the point to allow clearer reading of the plot, the measurements were actually	
	at the same position. The errors on the vertical position are the width of the	
	scintillator used	86
4.7	2D plot of the charged particle flux across the front of the MuSIC beam	
	pipe. The final three measurements were rescaled using the ratio of the two	
	measurements at $(0, 20)$ cm. Both the dot colour and the z-axis indicate the	
	flux	86
4.8	Comparison of simulated and measured total charged particle flux	89
4.9	Simulated results for the 2D charged particle rate normalised to a proton	
	current of 1 nA. Colour is equal to the rate as given by the z-axis. Simulated	
	using $9 \times 10^{\circ}$ in G4BL, the rate was taken to be those particles seen within	00
4.10	the area of the scintillator used.	89
4.10	Results of the first measurement of the muon lifetime with residuals from	
	the fit in the lower plot. The data was binned in 166 ns bins to remove the	
	effects of noise on the data, only entries with an ADC value of 1,100 (which	
	Excluded the pedestal) were used. The fit was applied from 600 hs to 20 μ s. The residuals show that fit gets were at earlier times although this is to	
	he expected with an exponential. Fitting with two exponentials should have	
	improved the fit at earlier times but this was found not to be the case	02
1 11	Besults of the second measurement of the much lifetime with residuals of the	52
4.11	fit in the lower plot. The data was split between bins with a width of 50 ns	
	and the data was fitted for values between 200 ns and 16 µs. The upper limit	
	of the fit was lowered from 20 μ s to 16 μ s to remove an apparent discontinuity	
	in the data. The discontinuity sits at 16 μ s and is especially clear in the	
	residual plot. The cause of the discontinuity is unknown. \dots	93
4.12	Experimental set up of the detector for MuSIC 5 (widths not to scale) see	
	table 4.8 for actual sizes. The upstream scintillators are outlined with blue	
	dashes whilst the downstream ones are red dots	95

4.13	Total muon momenta at the end of the beam pipe. The number of initial muons simulated was 95,700. The plot has been normalised to the number of	0.0
	muons per μ A of proton beam.	96
4.14	Momenta at the end of the beam pipe of muons that subsequently stopped. A stopped muon is considered to be one which is seen in the upstream scintillator and has a daughter electron seen in the downstream scintillator i.e. the muon has stopped between the scintillators. This shows the momenta that the different degraders 'select'. The number of initial muons simulated was 95,700. The plot has been normalised to the number of muons per μ A of proton beam	97
4.15	Trigger rate of each run (black) with the raw SEC count (blue). 100 samples were taken from the scaler data. Potential triggers were used to exclude any rate limiting due to dead time. The central 90 % of the data has been fitted with an order 0 polynomial with the parameters given. The SEC values are the differences of the count at the start of the time window and the end, this was done because the SEC value was not reset between measurements within a run. The SEC data suggests that the fluctuations in the trigger rate are not due to changes in the beam	101
1 16	Zoomed region (8 000 to 8 600 ns) of run 448 channel D1 showing the si	101
4.10	nusoidal component of the background the background has been fitted using	
	$N_h + N_e \sin(2\pi \frac{t-\phi}{\pi})$.	103
4.17	Per channel fitting of data for run 448, for this run there was no degrader.	105
4.18	Per channel fitting of data for run 451, for this run a 0.5 mm aluminium	
1.10	degrader was used.	106
4.19	Per channel fitting of data for run 452, for this run a 0.5 mm aluminium degrader was used.	107
4.20	Per channel fitting of data for run 455, for this run a 1 mm aluminium degrader	
1.20	was used.	108
4.21	Per channel fitting of data for run 458, for this run a 5 mm aluminium degrader	
	was used.	109
4.22	Per channel fitting of data for run 459, for this run a 5 mm aluminium degrader	
	was used.	110
4.23	Integrated number of free muon decays summed over all channels	112
4.24	Integrated number of muon decays in copper summed over all channels	112
4.25	Corrected free muon decay rate and simulated free muon decay rate. The red	
	lines are the simulated results, the blue lines the measured results including the MPPC efficiency error and the black lines are the measured values without the MPPC efficiency error	116
1 96	Corrected decay rate of muonic copper. The simulated rates are not shown	110
4.20	as the statistics for muonic copper decays were too low to extract reasonable values from. The blue line indicates the error due to systematic error, the red	
	line is the MPPC efficiency error and the black line the statistical error	116

5.1	Schematic of the beam-lines for EuXFEL, black denotes electrons whilst red corresponds to X-rays. The blue 'SASE' blocks are SASE sections whilst the 'U' blocks are spontaneous emission undulators (which don't self amplify)	
	Taken from [33].	123
5.2	Plot of peak X-ray brilliance against energy for a range of current sources as well as the predicted values for EuXFEL (here labelled 'XFEL'). The blue dots show measured peak brilliance for several energies at the existing FLASH fa- cility, 'FLASH (seeded)' is a proposed extension using a micro-bunch 'seeded' electron beam. Taken from [33].	124
5.3	The CCC RJ45 wiring diagram. Arrows on the inputs/outputs to FEE indi- cate signal direction, apart from the status line all are input to the FEE, i.e. the status line is the only line expected to send signals from the FEE to the CCC. Modified from [31]	127
5.4	The timing structure of electrons at EuXFEL and the resultant X-ray pulses. Modified from [31]	128
7.1	Top level block diagram. The clk and rst signals are supplied to each block and have been excluded from block diagrams unless otherwise noted. clk can either be an internal signal if the board is running in a standalone mode or is the clock received from the CCC. rst is an internal board reset, not to be confused with the RESET command which is intended to reset the ASIC (not the FEE).	138
7.2 7.3	Block diagram of the receiver entity	144 145
7.4	Flow diagram of the veto receiver entity. This de-serialises and splits the veto signals into the command and its payload, the Bunch ID. Note that this system also has the same flaw as the command receiver, i.e. it doesn't count bits as they are received so it can enter an undefined state if a malformed veto	
7.5	signal is sent	146
	to load the bunch pattern.	151
7.6	Flow diagram of the veto filter	152
7.7	Flow diagram for the veto decision logger.	153
7.8	Block diagram of the transmitter	160
7.9	Schematic of the control flow in the transmitter entity	162
7.10	General control logic flow for START, STOP and RESET states	163

8.1	A START, NO-VETO, STOP sequence with a single NO-VETO prior to the stop	
	signal. The blue braces indicate either input/output word sequences whilst	
	the red lines show logical sequences.	167
8.2	A selection of vetos/no-vetos.	169
8.3	The reset sequence being triggered by the cmd line from the CCC	171
8.4	The reset sequence being triggered via the control register	172
8.5	The clk_to_asic down-scaled by a factor of 100. The section delimitated	
	by the yellow guides is shown in figure 8.6. Note that the second word sent	
	(after the down-scaler trigger) is the READ_OUT_START command, the flag is	
	also down-scaled.	174
8.6	The STOP command and down-scaler start word being sent. Note the LocalL-	
	ink transfer on the 'll_' lines	175
8.7	Read back of the first 4 locations of every entity accessible via RDMA	177
8.8	readout of the veto log for 3072 bunches	179

List of Tables

2.1	Comparison of several leading muon sources. As most sources have multiple muon beam-lines the highest flux station at each source was chosen: μ E4 at PSI; port-1 of the RIKEN-RAL facility and the M20A beam-line at TRIUMF. The difference in momenta is due to both μ E4 and port-1 using surface muons whilst M20A uses cloud muons. Beam type refers to whether the beam is continuous or pulsed. P is the muon momentum (either as an average if reasonably mono-energetic or a range), I is the proton beam's current	33
2.2	Experimental values of G_1 - G_4 as determined by TRIUMF [3]	45
2.3	Current limits on the standard cLFV processes for muons. Note that several different materials have been tested for the process $\mu + X \rightarrow e + X$ but the gold measurement, made at SINDRUM II, is currently the most accurate	46
3.1	Summary of the particle distributions at the end of the MuSIC beam-pipe for 900 million initial protons. The number of particles of that type are given along with the mean and RMS of their distribution along the horizontal and vertical axis.	56
3.2	Compounds and mixtures used for the simulation. No refractive index is given for mylar as it's opaque. An additional property of EJ-212 was the light yield which was $10,000 \text{ MeV}^{-1}$	62
3.3	Dimensions of the stopping target and degraders. The degrader width of 0 mm is when the degrader is removed and replaced with air. The Z axis is the main beam direction with the X and Y axes being in the horizontal and vertical directions respectively.	64
3.4	Dimensions of the simulated paddle components. The wavelength-shifting fibres were the same for both up- and down-stream counters, as were the 'MPPCs'. The WLS fibre's cladding is said to be 3% of the total diameter so this is what is used [28].	64
3.5	Comparison of integrated numbers of muon decays to the counted number of decays in simulation of 5×10^5 initial muons. The free and copper components of the fit were considered separately. The results are for negative muons only as the positive muons only have the free muon lifetime	70

4.1	A summary of the five MuSIC beam-times with notes on the measurements made.	71
4.2	Details of the scintillators used for the two measurements of the charged par- ticle flux	79
4.3	Positions at which the charged particle flux was measured, 1D refers to the first run in which only the vertical displacement was measured, 2D refers to the second run in which horizontal measurements were also take. The distances from the beam-pipe were 6 cm for 1D measurements and 85 cm for 2D, this was due to mechanical constraints.	81
4.4	Summary of the data taken for measuring the MPPC detector efficiency. Configuration refers to which of the three MPPCs were tested. The efficiency is the simple ratio of the MPPC count to the PMT count whilst the adjusted efficiency has been scaled by the ratio of the areas of the scintillators, i.e. by $\frac{40\times40}{\pi3.5^2}$	82
4.5	Efficiencies of individual MPPCs calculated using the values from table 4.4, and equation (4.2). The two values below the line represent the total efficiency as calculated using the individual values and the measured value. The final value, $\epsilon_{\text{Ave.}}$ is the mean efficiency of the three individual efficiencies. The error on the average efficiency is the proportional difference between the measured and calculated values of the efficiency of all three MPPCs (see equation (4.3)).	84
4.6	Summary of the results from the measurement of the vertical charged particle flux. The double line indicates the division between runs with all four MPPCs working and runs with only three. The errors are statistical whilst the error on F is calculated as the errors on the un-scaled measurements at 0 cm added in quadrature. The error on the flux is calculated as the propagated errors of the other values. Counts were taken for 50 s. The SEC conversion factor, K , was 0.03408 nA. The factor was taken to be the ratio of two measurements made at 0 cm. The 0 cm measurements were used as this was the only pair of measurements that were made with the experiment in both conditions (i.e. all MPPCs functional and later, with one broken and another damaged).	85
4.7	Table of the results from the 2D measurement. The values above the double line were made with 20 mm Mg stopping target upstream whilst those below had 5 mm Cu. The scaling factor F for these sets of runs was made at $(0, 20)$ cm. The trigger count was measured over 20 s apart from the first two measurements C_{off} which were made over 11 s (these are the two measurements above the first horizontal line). The SEC measurements were made over 100 s. Counts were made using a CAMAC controlled scaler and an enable gate triggered via interrupt register. The column for the SEC bias, S_{off} , is excluded as it was constant at 9 and the SEC conversion factor, K , was 1.514 nA. As with the 1D measurement the errors are statistical. The last	
	three values also include an error due to the calculation of the scaling factor.	87

 figure 4.12	94 95 98
 subsequently stop, at the end of the beam pipe. The degrader is the thickness of the aluminium, the width and height are both 400 mm. The Geant4 simulation was used; in total 95,700 initial muons were simulated, for further details see chapter 3	95 98
 simulation was used; in total 95,700 initial muons were simulated, for further details see chapter 3. 4.10 Summary of the runs selected for this analysis. The degrader used aluminium and the stopping target 0.5 mm copper. 4.11 Dead and live times for each run along with the number of potential and good triggers. Run time is included to show the origin of the different counts. 	95 98
 4.10 Summary of the runs selected for this analysis. The degrader used aluminium and the stopping target 0.5 mm copper. 4.11 Dead and live times for each run along with the number of potential and good triggers. Run time is included to show the origin of the different counts. 	98
and the stopping target 0.5 mm copper	98
4.11 Dead and live times for each run along with the number of potential and good triggers. Run time is included to show the origin of the different counts	
	99
4.12 Summary of the fitted values for equation (4.10). These are the values used for the calculation of the integrals and, ultimately, the muon rates. The	
two metime parameters, γ_f and γ_c , were fixed at the values 2, 190.9811 \pm 0.0022 ns [2] and 163.5 ± 1 ns [3] respectively, T was fixed at 60 ns	104
4.13 Decay counts and rates, summed over all downstream scintillators. Δz is the	
thickness of the aluminium degrader	111
cays were taken as the systematic errors for each value (55.85% and 1.14% respectively).	114
4.15 Corrected rates for freely decaying muons. The error column shows the error without the contribution from the MPPC efficiency and with it as is clear	
this is the dominant source. The momentum values are the mean and RMS	
as determined by simulation (see section 4.4.1 and table 4.9)	115
4.16 Corrected rates of muonic copper decay. Errors are split into three: total error (including systematics and MPPC efficiency); error excluding systematics but including MPPC efficiency; and error excluding both systematic and efficiency	
contributions (i.e. statistical). The momentum values are the mean and RMS as determined by simulation (see section 4.4.1 and table 4.9).	115
4.17 Calculated muon production efficiencies and total fluxes at $1 \ \mu$ A. 'Component' is the simulated fraction of the total beam that stops in each configuration.	
'Rate' is the sum of the corrected rates for free and copper decays. The total flux was calculated using equation (4.17). The average total flux is $(4.2 \pm 1.1) \times 10^8$ muons s ⁻¹ where the error is the same fractional error as is on the MPPC efficiency, $\Delta \epsilon_{\text{MPPC}}^2 \sim 26 \%$. The efficiency is the total flux	
divided by the full beam power (400 W). The average efficiency of MuSIC is calculated to be $(10.4 \pm 2.7) \times 10^5$ muons W ⁻¹ .	118
5.1 Description of typographic conventions	121
5.2 Specification of the fast control signals and their payloads	129 130

7.1	Top level interface for the clock and control interface	139
7.2	A table of the generics used in the design, their type, name, where they are	
	used (R=receiver, V=veto-filter, T=transmitter, TL=top-level).	140
7.3	Summary of the top level delay registers. The default can be overridden via	
	the generics (section $7.1.2$)	141
7.4	The full set of commands received from the CCC. The payload for the START	
	is the train ID (32b), the bunch pattern ID (8b) and a check-sum (8b). The	
	(NO-)VETO payloads are the same, the 12b bunch ID.	142
7.5	Top level interface of the receiver entity.	143
7.6	Truth table for veto decisions. Count is the number of 'no-veto's already sent,	
	it is assumed to be 512 but can be any number less than this	147
7.7	Interface for the veto filter	148
7.8	Default reset values for pattern ID registers	150
7.9	Comparison of dynamic veto and reset modes	154
7.10	Interface for the transmitter	157
7.11	Control register layout. With the exception of the status register the reset	
	values of each register can be changed by setting the appropriate reset-generic,	
	these values only take effect when the rst line is asserted <i>not</i> when the RESET	
	signal is received	159
7.12	Flag information, command words that the shift-entity looks for and will flag	
	along. Default values are for 22b words, see section 7.4.3 for more details.	
	When looking for flags the sync bits (the first two bits) are ignored so do not	
	need to be set	161
7.13	Description of the shift-entity's behaviour depending on state	161
A.1	ASIC command words, see [36] for a full description and recommended use of	
	these commands.	183
A.2	Standard RDMA interface	184
A.3	Minimal local link interface as used by the veto logger	184

Acknowledgements

This thesis has been the work of too many years, with the aid by too many people to accurately and fully thank all of those who deserve to be thanked. That being said, with the hope of doing something to repay their kindness, I'd like to briefly acknowledge those who've helped this work come to fruition.

Firstly, my supervisor, without who's tireless support, corrections and ready need for caffeine this would never have happened. My lab-mates and peers at UCL, particularly Richard and Andrew who endured more hours underground with me than anyone should. Kuno-Sensei, Sato-Sensei and all my colleagues in Japan who were both excellent hosts and teachers. My supervisor at RAL, John, and all those who I worked with on the LPD/CCC interface for Eu-XFEL.

I'd like to thank my friends (particularly Antonio) and family for their forbearance over the last few years of dark mutterings when ever the subject of my studies has been brought up, it was never that bad (although maybe that tiring) and your continued support made it possible. Finally I'd like to thank my wonderful partner, Morag.

Preface

This thesis is divided into two parts. The first part covers work done at the Muon Science Innovative Channel (MuSIC) beam-line and the second part, discusses work carried out on the European X-ray Free Electron Laser (EuXFEL).

The first part, discussing MuSIC, has two main themes: simulation and measurement. An introduction (chapter 2) covers the essentials of the project and the physics behind it. Chapter 3 describes the development of the MuSIC simulation and its deployment in creating a rigorous analysis. Finally chapter 4 discusses the results of the measurements made at MuSIC.

The second part is concerned with creation of the firmware interface between the Clock and Control Card (CCC) and the Large Pixel Detector (LPD). It discusses the design, implementation and testing of the firmware ahead of use in EuXFEL.

Part I

Characterisation of the MuSIC beam

Chapter 1

Executive Summary

The Muon Science Innovative Channel (MuSIC) [1] beam line is a prototype muon production system being developed at the Research Centre for Nuclear Physics (RCNP) in Osaka, Japan. It uses a novel Pion Capture Solenoid (PCS) that captures a much larger fraction of the produced pions than other muon beam lines do. Once the pions have been captured they are directed in to the Muon Transport Solenoid (MTS) where they can decay to produce muons. MuSIC aims to achieve a muon production efficiency (per Watt of proton beam) that is far higher than is currently seen at existing beam lines.

High flux, low energy muon beams will enable the creation of dedicated muon Spin Resonance (μ SR) devices, they are a first step in making neutrino factories as well as being vital in particle physics studies such as the measurement of rare decays (e.g. $\mu \rightarrow eee$). In addition, muon beams can be used in the development of other technologies, e.g. MuSIC will be used to conduct research into Fixed Field Alternating Gradient (FFAG) storage rings and test their viability as an acceleration technology for muons.

A key use of MuSIC will be as a prototype for the Coherent Muon-to-Electron Transition (COMET) experiment being built at J-PARC in Tokai, Japan. COMET will use a similar, albeit larger, pion capture solenoid to produce a pulsed beam of 10^{11} muons s⁻¹ that will enable it to make the most precise measurements of the rare muon decay: $\mu N \rightarrow eN$. The insights gained at MuSIC, both in terms of the technology and the science, will help improve the measurements made at COMET.

In order to better understand MuSIC it was simulated using Geant4 and G4Beamline (G4BL). Two core simulations were made, one of the PCS and MTS made using G4BL and a second of the detectors made using Geant4. The G4BL simulation was used to test the initial designs of MuSIC and later to provide realistic particle distributions for the Geant4 simulation. The detector simulation, made in Geant4, was used to test detector configurations and refine analysis techniques.

Commissioning of the PCS and the first 36° of the MTS was completed in 2010. The final design of MuSIC has a planned 180° MTS but the initial section is enough to begin characterising the beam. Over two years there were five periods of beam time, each lasting between 24 and 72 hours, for which the RCNP's proton beam was available for use at MuSIC.

During the five beam times, detectors and data acquisition (DAQ) were installed, tested

and used to make three different measurements that help characterise the MuSIC beam. The three measurements were of the total charged particle flux, the muon lifetime and the muon momentum spectrum. Each experiment used simple detectors made of plastic scintillators with Multi-Photon Pixel Counters (MPPC) for read out.

The charged particle flux was measured in two separate beam times, firstly with a long strip scintillator $(1\times3\times38 \text{ cm}^3)$ and then with a disk scintillator $(3.5\times\pi\times2 \text{ cm}^3)$. The first measurement was made 6 cm from the end of the MTS at a range of heights. The peak charged particle flux was measured to be $(94.3\pm4.0)\times10^3$ particles $nA^{-1} \text{ s}^{-1}$ at 5 cm above the beam pipe's centre. The second measurement was made 85 cm from the end of the beam pipe at a range of horizontal and vertical positions. The peak flux for the second measurement was $(11.16\pm0.11)\times10^3$ particles $nA^{-1} \text{ s}^{-1}$ at (0, 20) cm, i.e. directly above the beam pipe's centre. The second most intense location was at (-17,0) cm, to the left of the beam centre (when facing it). Both these measurements show that the most intense part of the beam is above the centre and the second also shows that the this spot is slightly offset towards the left, which agrees with the simulation although not on the predicted intensity.

The muon lifetime was measured during two beam runs as a method of demonstrating the presence of muons in the beam. The measurement was made by using two scintillators on either side of a metal stopping target. By plotting the time between events in each scintillator the muon decay spectrum could be constructed and the muon's lifetime measured. The first measurement of the muon lifetime was made during the same beam time as the second charged particle measurement and recorded a lifetime of $(2,016\pm89)$ ns which is lower than the canonical value of (2196.9811 ± 0.0022) ns [2] although this is likely due to the presence of a stopping target which would have reduced the negative muon's lifetime. The second measurement attempted to account for the presence of the stopping target by fitting the muon lifetime with two exponentials and hence measured the free muon lifetime to be (2202 ± 20) ns and the lifetime of negative muons in the copper stopping target to be (195 ± 17) ns. The lifetime of muonic copper was slightly higher than the canonical value of (163.5 ± 1.0) ns measured in by Suzuki et al. [3] but is still in reasonable agreement.

The final measurement, the muon momentum spectrum, was the most complex. Six measurements were made using one of four different thickness aluminium degraders. The degraders reduced the energy of any charged particles and so filtered out different momenta to stop on the copper stopping target. By plotting the decay spectrum and integrating the area under the exponential portion it was then possible to count the number of stopped muons and hence estimate the number of muons of that momentum. The momentum spectrum measured for freely decaying muons is shown in figure 1.1. This figure also shows the reasonable agreement between the simulation and the measurement although the errors on the measurement are large.

Using the data from the momentum spectrum measurement it was possible to predict the total muon flux at MuSIC when operated at full power^{*} and this was calculated to be $(4.2 \pm 1.1) \times 10^8$ muons s⁻¹ which is on par with the world's most intense muon source, PSI,

^{*}To prevent excessive dead time in the equipment currents of <1 nA were used rather than the RCNP's maximum current of 1 μ A.



Figure 1.1: Measured muon momentum spectrum for freely decaying muons at MuSIC. The muon rate is normalised to a proton beam current of 1 nA.

which has a peak flux of 4.8×10^8 muons s⁻¹. The large errors on the total flux are due to some unavoidable assumptions made in the calculation. Firstly the calculation of the detector efficiency was made using an earlier measurement. Secondly the value is coupled strongly to the simulation of the beam that was used to estimate the ratio of stopping muons for each configuration. Despite these shortcomings the data clearly shows a high flux of muons which is comparable to that measured at PSI. The key difference between MuSIC and PSI is that PSI produces its muons using a 1.2 MW proton beam compared to the RCNP's 400 W beam, giving per-Watt yields of $(10.55 \pm 2.7) \times 10^5$ muons W⁻¹ and 292 muons W⁻¹ respectively. Therefore MuSIC is clearly the most efficient dedicated muon source in the world.

Chapter 2

Introduction

MuSIC is a new muon source housed at the Research Centre for Nuclear Physics (RCNP) in Osaka, Japan. Commissioning of the beam began in 2010 with installation of the Pion Capture System (PCS) and the first 36°, of a planned 180°, of Muon Transport Solenoid (MTS). To date there have been five periods of beam time.

Over the course of the five sessions of beam time three core measurements have been made that will be discussed in this thesis:

- 1. Charged particle flux.
- 2. Muon lifetime.
- 3. Muon momentum spectrum.

All of these experiments have had the aim of characterising the beam at MuSIC and confirming the efficiency of muon production. The measurements themselves are discussed in chapter 4.

While still incomplete, enough of MuSIC exists to test the efficacy of the PCS and existing MTS. The PCS uses a novel design in order to to maximise muon production. This new design is predicted to be several orders of magnitude more efficient than current methods.

The rest of this chapter is divided into three sections: Muon sources (section 2.1), MuSIC (section 2.2) and physics (section 2.3). 'Muon sources' discusses existing muon beams and goes on to make the case for high intensity muon sources. The section on MuSIC discusses its design and the scientific motivation for building it. Finally the physics section reviews the theory behind charged particle interactions in matter, muon decay and charged Lepton Flavour Violation (cLFV).

2.1 Muon Sources

The muon was the first second-generation particle to be discovered, it has been used to test the theory of relativity [4], measure the size of the proton [5] and as a way to detect neutrino interactions which helped lead to the discovery of their oscillations [6]. Not only is the muon an interesting particle in its own right, but a powerful and precise tool of further research.

Muons are a lepton; they are ~200 times more massive than their sibling, the electron^{*}. Unlike the electron, muons can decay and the almost always[†] muons decay via the process $\mu \rightarrow e\nu\overline{\nu}$, the muon's lifetime prior to undergoing decay is 2.1969811±0.0000022 μ s [2]. Muons interact via the electroweak forces and are, as far as we know, elementary particles.

More than any other property, the lifetime of the muon, encompass the biggest problem with muon sources. The fact that muons decay; unlike protons, neutrons[‡] and electrons, makes manipulating muons a race against time. All beams have similar stages: production, acceleration, cleaning and finally deployment; for muons this has to happen in the 2.2 μ s that the muons exist for. Despite the difficulty of creating muon beams the usefulness of a non-electron lepton beam is not to be underestimated.

As a particle that isn't common in nature muons cannot be produced in the same way as, e.g. electrons: via ionisation and subsequent acceleration. The most common method for muon production is via pion decay (see figure 2.1), this process has a branching ratio of $99.98770\pm0.00004 \% [2]$ [§]. The pion has a lifetime of only 26.033 ± 0.005 ns [2], i.e. $\sim 1\%$ of the lifetime of the muon meaning that given a suitable amount of time, the pion contamination of the muon beam can be kept low. Pion production occurs through proton-proton interactions above the energy threshold of ~ 380 MeV. The simplest method achieving proton-proton interactions above the pion production threshold is through the bombardment of a fixed target with a proton beam.



Figure 2.1: Feynman diagram showing muon production through pion decay.

Whilst all muon sources use pion decays to generate their beams there are several broad categories of beam. Muon beams are generally categorised depending on where the muons originate: 'cloud', 'surface' and 'sub-surface'. Most muon facilities will have one or more of each type of beam, as they are not incompatible and all have different properties. 'Cloud' beams collect pions (and some muons) then transport them to a 'decay pipe' in which the remainder of the pions can decay. In general cloud type beams have a larger range of momenta as the pions decay in flight. 'Surface' and 'sub-surface' beams are both generally of lower momenta (< 30 MeV/c) than cloud types as they collect muons produced by pions

^{*}The mass of the muon is $105.6583715 \pm 0.0000035$ MeV, an electron is 510.998928 ± 0.000011 keV [2]. $^{\dagger} \approx 100\%$ of the time [2].

[‡]The neutron's lifetime is 880.0 ± 0.9 s [2] long enough to make it stable in comparison to the muon.

[§]The next most common decays are $\pi \to \mu \nu_{\mu} \gamma$ and $\pi \to e \nu_e$ which have branching ratios of $(2.00 \pm 0.25) \times 10^{-4}\%$ and $(1.230 \pm 0.004) \times 10^{-4}\%$ respectively.

decaying at rest on (or just inside) the surface of the target. Surface and sub-surface beams are in high demand for muon Spin Resonance (μ SR, see section 2.2.2) experiments which require high fluxes of low momentum muons.

There are currently several dedicated muon sources^{*} around the world (see table 2.1). All of the current muon sources have their facilities acting in a parasitic mode on their main proton beam-line, with the majority of the proton beam being used for neutron sources. Most muon facilities have several separate beam-lines that can operate semi-independently, with each beam-line often specialising in a certain sub-set of muons. For example, RIKEN-RAL has 4 experimental ports: port-1 is used to investigate muon-catalyzed d-t fusion (μ CF), port-2 and port-4 are dedicated μ SR experiments; and port-3 is used for low energy muon production.

Name	Muon beam			Proton Beam		Ref
	m Flux~(muons/s)	Type	$P \; ({ m MeV/c})$	Power (W)	$I~(\mu A)$	
PSI	4.8×10^8	Cont.	28	1.2×10^6	2000	[7]
RIKEN-RAL	6×10^5	Pulse	27	1.6×10^3	200	[8]
TRIUMF	$3.5 imes 10^6$	Pulse	25 - 85	$7.0 imes 10^4$	140	[9]

Table 2.1: Comparison of several leading muon sources. As most sources have multiple muon beam-lines the highest flux station at each source was chosen: μ E4 at PSI; port-1 of the RIKEN-RAL facility and the M20A beam-line at TRIUMF. The difference in momenta is due to both μ E4 and port-1 using surface muons whilst M20A uses cloud muons. Beam type refers to whether the beam is continuous or pulsed. *P* is the muon momentum (either as an average if reasonably mono-energetic or a range), *I* is the proton beam's current.

2.1.1 Scientific Motivation for High Intensity Muon Sources

As has been noted muons are highly active area of research, both as a tool for further research and a focus of study. The activity surrounding muon beams is only increasing as new technologies become available and our ability to control these short lived beams increases.

As a focus for study, in and of themselves, muons present several interesting avenues for research. Muons are a well understood particle whose interactions are known to high precision. Because muons are so well understood they make an excellent arena to test the limits of our knowledge: by making precision measurements of the muon and its interactions we can test regions of theory that are otherwise inaccessible. For example, by looking for charged Lepton Flavour Violation (cLFV, see section 2.3.3). Through cLFV processes such as $\mu \rightarrow eee$ or $\mu N \rightarrow N^*e$, we can test the limits on many 'beyond the Standard Model'

^{*}In theory any proton source, with energy above the pion production threshold, can operate as a muon source but only a few have dedicated facilities.

theories. Obviously, in order to measure a process that has, so far, never been seen you need an ample supply of the source material so high intensity muon sources are vital for these measurements.

Currently one of the most active areas of research in particle physics is that of neutrinos. Whilst very common in nature the neutrino's tiny cross section makes study of it very difficult. In recent years this has been solved through the creation of neutrino beams such as T2K [6]. Neutrino beams are normally created using the decay products of pions but there is growing interest in building so-called 'neutrino factories'. One proposed design for a neutrino factory uses an oval storage ring to hold muons until they decay to create a mixed electron/muon-neutrino beam, obviously if many more neutrinos are wanted then significantly higher intensity muon sources are needed.

A commonly described evolution of the neutrino beam is to create a second muon beam and counter rotate them in order to produce a muon collider. A leptonic collider is highly attractive: as elementary particle collisions are cleaner and can be more finely controlled whilst the higher mass of the muon over the electron means that there is much less energy loss to synchrotron radiation. A muon collider would present many challenges but, with the International Linear Collider (which uses electrons) expected to be ~ 31 km long [10], a more compact alternative is highly desirable. If a muon collider can be built as the next stage of a neutrino factory, so much the better.

Particle physics presents just one facet of the many uses to which muon beams can be put. One of the most common uses of muon beams is probing the magnetic properties of materials through μ SR, other applications include using muons to determine the nuclear matrix elements of materials, measuring trace elements in a sample through the emission of muonic X-rays and research into μ CF.

2.2 MuSIC

MuSIC is a proof of concept beam-line that aims to produce the most muons per proton of any existing source. MuSIC aims to have a muon flux comparable to that of PSI's (the current world best) using a beam with $1/_{3,000}$ the power. This can be done by one of two naive methods: increasing the efficiency of production or using a more powerful initial beam. At both RIKEN-RAL and PSI the muon beams are produced using only a fraction of the initial proton beam. It is common that less than 10% of the initial proton beam is used to create muons, the rest continues to neutron sources. Obviously to create a high intensity muon beam using as many of the protons as possible should be the first design goal. There is a second advantage to moving away from a parasitic muon production: not only is the entire proton beam used but more of the pions can be captured. In a parasitic design the PCS has to allow the for the continuation of the proton beam, this places limits on the geometry that can be used and produced sub-optimal efficiencies. If the requirements for a continuing proton beam are removed though, the PCS geometry can be optimised purely for pion capture making further gains over merely increasing the proton current.

Figure 2.2 shows the expected complete MuSIC complex including the proposed experi-



Figure 2.2: Schematic of the completed MuSIC complex. Taken from [1].



Figure 2.3: Schematic of the current status of MuSIC. Taken from [11]



Figure 2.4: Photo of MuSIC showing its 2010 status, since then it has had further concrete blocks added around it to increase shielding.

mental stations. The schematic in figure 2.3 shows the current status, with the Pion Capture System (PCS) and the first 36° of Muon Transport Solenoid (MTS) complete. Figure 2.4 is a photo of the status of MuSIC in 2010, prior to the addition of further shielding.

2.2.1 MuSIC design

The core component of MuSIC is its PCS. The MuSIC PCS is designed to use as much of the proton beam and capture as many pions as possible. The PCS achieves this with a $2.1 \times 2.4 \times 1.6 \text{ m}^3$ [1] super-conducting magnet that has a maximum strength of 3.5 T. The PCS has an extremely large solid angle making MuSIC capable of capturing a significant portion of the pions and directing them into the muon transport solenoid (MTS).

The MTS is intended to perform three roles within MuSIC: the collection and decay of travelling pions; selection of low-momentum (30-50 MeV/c) muons and selection of muon charge. Ultimately the MTS will form a 180° semi-circle but currently only the first 36° has been constructed. The MTS's designed 10 m length provides the requirements for pion decay; its radius and magnetic field provide momentum selection whilst the application of a dipole field will select charge. Like the PCS the MTS uses super-conducting magnets to provide the 2 T main field and the 0.04 T dipole field.

The proton beam for MuSIC is provided by the RCNP's ring cyclotron. The cyclotron produces a 400 MeV proton beam (enough for pion production) with a maximum current of 1 μ A giving a maximum power of 400 W. The beam is produced in an ion source then
accelerated using two cyclotrons: the Azimuthally Varying Field (AVF) [12] cyclotron and then the ring cyclotron. AVF, can accelerate a variety of ions (up to ¹⁸O) using a voltage of 60 kV, for protons this corresponds to a final energy of ~65 MeV as a continuous beam. The ring cyclotron then accelerates the protons to ~400 MeV.

The MuSIC PCS has been optimised to capture backwards travelling pions and muons with a maximum transverse momentum (P_T^{max}) of 52.5 MeV. The two key parameters that determine P_T^{max} are the magnet's bore and the field strength, through optimisation these were determined to have values 10 cm and 3.5 T respectively. To obtain a 3.5 T field copperstabilized NbTi superconductor was used, this requires cooling to 4 K which is achieved through the use of 3 GM cryocoolers.

The target used is a graphite cylinder 20 cm long with a 2 cm radius. To maximise production the target is rotated from the PCS axis to align its major axis with the proton beam. The superconducting magnet is shielded using stainless steel that has a maximum thickness of 27 cm, the shielding tapers on either side of the target (see figure 2.3). The taper is more rapid in the backwards direction in order to capture the maximum number of pions and muons.

As already noted the muon transport solenoid (MTS) has to select muon momentum and charge whilst letting pions decay. The MTS has a designed length of 10 m which is predicted to reduce the pion contamination to less than 0.1%. Application of the dipole field can be used to select charge.

2.2.2 MuSIC: Scientific Motivation

The general arguments for high intensity muons beams have already been presented (section 2.1.1), MuSIC, though is not technically a high intensity muon beam. Assuming that it meets full design specifications it should have a comparable intensity to the muon beam-lines at PSI. To this end the scientific case for MuSIC is slightly different to that for a full high intensity muon beam. The key aim of MuSIC is in proving the PCS and MTS technology, what PSI achieves with a 1.3 MW proton beam MuSIC aims to achieve with a 400 W beam. This section will set out the other uses to which MuSIC will be put.

Prototype for COMET

The COherent Muon to Electron Transition (COMET [13]) experiment intends to place the most stringent limits in the world on the charged Lepton Flavour Violating (cLFV) decay $\mu + X \rightarrow e + X$. According to the Standard Model, including neutrino oscillation, this decay should have a branching ratio of $\sim 10^{-54}$ which is far beyond current detector capabilities, this means that should such a decay be seen it would be a so-called 'smoking gun' of new physics. Many beyond the Standard Model theories (e.g. [14]) predict this decay at branching ratios that should be detectable ($\sim 10^{-16}$). The design of COMET calls for a pulsed muon rate of $\sim 10^8$ muons/s. In order to achieve this a PCS similar to that of MuSIC is required but with a larger field (5 T) in order to cope with the more powerful muon beam. A MTS

similar to that of MuSIC's is also planned in order to select low momentum muons that can be stopped on a titanium target.

The first goal of the measurements made at MuSIC is to verify that the intensities required for COMET are feasible, they will be used to refine the simulations used at COMET allowing a more well optimised final design. Additionally the expertise and techniques developed at MuSIC should be directly applicable to aspects of COMET's running, for example in gauging expected secondary particle fluxes.

Precision Measurements of cLFV muon decays

As well as supporting COMET's search for $\mu + X \rightarrow e + X$ it is planned that MuSIC will run a complementary search for the cLFV decay $\mu \rightarrow eee$ [1]. Whereas COMET's main background is beam related and mitigated through the use of a pulsed beam, the main background to $\mu \rightarrow eee$ at MuSIC is expected to be accidental, meaning that a continuous beam poses fewer additional problems and the benefit of a higher rate.

Using a novel detector design MuSIC aims to improve on the current limit of 1.0×10^{-12} , set in 1988 by the SINDRUM I collaboration [15], by a factor of ~ 150. Exclusion at this level would limit beyond Standard Model physics to an effective mass scale of $\mathcal{O}(1,000)$ TeV (see section 2.3.3).

This improvement will be achieved through a higher stopping rate, longer running time and increased acceptance. Obviously a more intense, longer, run will increase the number of accidentals by a predicted factor of 2,000. To compensate for the increased background, improvements to the timing, vertex, energy and momentum detection will be needed but it is envisaged that these are attainable.

As discussed in section 2.2.2 detection of a cLFV event would be a huge discovery and by searching for such processes both at COMET and MuSIC a far greater region of phase space can be excluded.

Muon Spin Resonance

Muon Spin Resonance (μ SR) is a technique in which a highly polarised beam of low-energy (< 50 MeV) muons are used to detect the internal magnetic structure of a sample. The muons are stopped on the surface to be studied and the resultant positron recorded. By studying the distribution of the emitted positrons it is possible to determine the magnetic structure of the material.

 μ SR is an obviously powerful technique in determining the properties of a material and is becoming increasingly useful as the study of super-conductors increases. Used in concert with the existing, pulsed, μ SR facilities at J-PARC, MuSIC would provide a powerful tool for investigating a wide range of materials.

Fixed-Field Alternating Gradient research

One of the biggest obstacles to the development of a muon-collider is the acceleration of the beam. In a traditional electron accelerator a large synchrotron ring can be used to store and accelerate a beam over a long time period, unfortunately, with their short lifetime this technique is not viable for muons. One proposed solution to this problem is to use a Fixed Field Alternating Gradient (FFAG) storage ring. An FFAG uses a fixed magnetic field that has a non-uniform gradient to accelerate and store a beam.

FFAGs are considered as a potential storage ring for muons as they can rapidly accelerate a particle without adjustment of the magnets. In this way many designs for muon colliders plan to use FFAGs to accelerate a collimated ('cooled') muon beam prior to collision. The RCNP facility already hosts a six-segment FFAG that has been used to explore phase-rotation of an alpha-particle beam. Once MuSIC is completed it is intended that these phase-rotation and further acceleration experiments be carried out on a muon beam.

2.3 Physics

This chapter covers the two primary physics processes of concern in characterising the beam at MuSIC: particle interactions with matter and muon decays. There is also a brief discussion of cLFV which is one of the main scientific motivations for MuSIC.

2.3.1 Charged Particle Interactions in Matter

All charged particles will ionise matter they pass through, the amount that any particular particle ionises any particular material is highly complex and beyond the scope of this discussion, a more thorough discussion can be found in many reviews (e.g. [2]). The general properties of ionisation are described by the Bethe formula:

$$-\left\langle \frac{dE}{dx}\right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2}\right]$$
(2.1)

where:

- $-\left\langle \frac{dE}{dx}\right\rangle$ is the average energy loss in MeVg⁻¹cm²
- $K = 4\pi N_A r_e^2 m_e c^2$, (N_A is Avagadro's number and r_e is the classical electron radius)
- m_e the electron mass.
- c the speed of light.
- z the incident particle's charge (i.e. for pions and muons, 1).
- Z the atomic number of the target atom.
- A the atomic mass.
- β is the incident particle's speed as a fraction of c (i.e. $\frac{v}{c}$).
- γ is the particle's gamma factor (i.e. $\frac{1}{\sqrt{1-\beta^2}}$).

- T_{max} is the maximum energy that can be transferred to a free electron in a single collision with the incident particle.
- *I* is the mean excitation energy (in eV).
- $\delta(\beta\gamma)$ is the correction to the energy loss due to the 'density effect'

The Bethe formula is generally applicable to any charged particle but the values for the density effect and mean excitation are not: their values are different for electrons when compared to heavier particles (e.g. muons and pions). Another important note is that, for electrons at low energies, whilst ionisation is the dominant effect there are other processes that need to be taken into consideration (Bhabha scattering, e^+ annihilation, Møller scattering). This difference can be seen in figures 2.5 and 2.6 for muons and electrons respectively. As can be seen for electrons there is no plateau between ionisation and bremsstrahlung regime, as there is for muons and pions.



Figure 2.5: Energy loss by a muon in copper as a function of $\beta\gamma$ or momentum. The solid line indicates the total stopping power, dashed lines indicate different contributions and the vertical grey bands indicate the different phenomenological regimes. Taken from ref [2].

At MuSIC the loss of energy by charged particles to matter is very important as it is the main method for detecting particles, and hence, characterising the beam. When a particle deposits energy in certain materials through ionisation they will produce light, this light can then be detected. The exact amount of light produced is dependent on the material and the



Figure 2.6: Electron energy loss per 'radiation length' (X_0^{-1}) as a function of electron energy (not momentum). The material in this case is lead (Z = 82 and $X_0 = 6.37 \text{g/cm}^2$). Taken from ref [2].

amount of energy deposited. A common empirical approximation to the light yield is Birks' Law:

$$\frac{d\mathcal{L}}{dx} = \mathcal{L}_0 \frac{dE/dx}{1 + k_B dE/dx}$$
(2.2)

where \mathcal{L} is the luminosity, \mathcal{L}_0 is the luminosity at low specific ionisation density, k_B is Birks' constant which has to be measured for the scintillator. Birks' law describes two main regimes: linear and quenched. Initially the light yield is roughly linear but as the amount of energy deposited increases the material becomes saturated and the amount of additional light produced declines. For PVT (the scintillator used at MuSIC) the light yield remains linear for rates of energy deposition of less than ~0.5 MeV mm⁻¹ [16], i.e. muons with momenta between 56 MeV/c and 1 TeV/c [17].

2.3.2 Muon decay

As has been noted free muons decay $\approx 100\%$ of the time through the process: $\mu \to e\nu\bar{\nu}$. This process is the same for both positive and negative muons, although there is a sign change and the neutrino/anti-neutrino flavours will swap. The biggest difference between positive and negative muons is that negative muons can become bound to a nucleus in the same way as electrons. This section will start with a general discussion of the important parameters of muon decay and then continue with a discussion of how this property of negative muons changes their decay.

The decay of a muon, like any other decay, can be modelled using the distribution:

$$P(t) = N e^{-t/\tau} \tag{2.3}$$

where P(t) is the survival probability of a particle after time, t; N is a scaling factor to the size of the initial population and τ is the population's lifetime. For a free muon in vacuum the lifetime is $\sim 2.2 \ \mu$ s. As figure 2.7 shows the muon decay is mediated by a charged W boson, the presence of the two neutrinos mean that the only practically detectable result of the muon decay is the electron. The muon's (comparatively) long lifetime makes identification of them possible through plotting decay times; this is discussed in greater detail in chapter 4.

In addition to freely decaying, negative muons can undergo additional interactions with matter, unlike positive muons. As a negatively charged lepton they can be become bound to a nucleus, forming a muonic atom. A muon, becoming bound to an nucleus will rapidly transition to the un-occupied 1s state, this normally takes less than 10^{-13} s [18]. In the 1s state the muon will be able to interact weakly with the nucleus. The main interaction, known as 'nuclear capture', is $\mu + p \rightarrow \nu_{\mu} + n$ and can be seen in figure 2.8. Whilst this process is undetectable reduces the apparent lifetime of the muon by reducing the total number of muons that can decay detectably. As well as nuclear capture, negative muons that are bound to atoms can undergo normal decay. This 'decay in orbit' occurs in the 1s shell and tends slightly increase lifetime, due to the reduced energy available for the process. The total



Figure 2.7: Feynman diagram showing muon decay.

lifetime of a negative muon is then reciprocal of the rates of these two processes:

$$\tau_{\mu^{-}} = (\Delta_t)^{-1} \tag{2.4}$$

$$\Delta_t = \Delta_c + Q\Delta_d \tag{2.5}$$

where Δ_t is the total interaction rate; Δ_c is the rate of nuclear muon capture; and $Q\Delta_d$ is the rate of decay in orbit. Q is the 'Huff-factor' [19], a theoretical value which accounts for the reduction in energy available to bound muons, Δ_d is the rate of the positive muon decay (i.e. $\tau_{\mu^+} = (\Delta_d)^{-1}$). The positive muon decay rate is used because measurements of the negative muon's lifetime are less accurate due to the discussed nuclear effects.



Figure 2.8: Nuclear capture of a negative muon. This process is very hard to detect directly. The net effect is to reduce the muon's apparent lifetime by removing them from the population.

A theoretical approximation to Δ_c , as a function of atomic mass, A, and atomic number, Z, was devised by Primakoff and Goulard [20]:

$$\Delta_c(A,Z) = Z_{eff}^4 G_1 \left[1 + G_2 \frac{A}{2Z} - G_3 \frac{A - 2Z}{2Z} - G_4 \left[\frac{A - Z}{2A} + \frac{A - 2Z}{8AZ} \right] \right]$$
(2.6)

where the values G_{1-4} are parameters to be fitted and Z_{eff} is the effective value for the atomic number, as calculated by Ford and Wells [21], this accounts for the muon's 1s radius being much smaller than an electrons. The difference between the predicted capture rate and the measured rate is given in figure 2.9 with the fitted values of G shown in table 2.2. As can be seen there are still disagreements between theory and experiment, for this reason in later simulations the measured values of capture rates are preferred.



Figure 2.9: Comparison of the theoretical values (as determined using the Primakoff-Goulard formula) and experimentally determined muon capture rates for a range of materials, taken from [3].

2.3.3 Charged Lepton Flavour Violation

The discovery that neutrinos have mass and that they undergo flavour mixing shows that lepton flavour is not conserved, making charged Lepton Flavour Violation (cLFV) a possibility.

In the Standard Model of particle physics the simplest diagram of cLFV is when the emitted muon neutrino is virtual (see figure 2.10), for this to be true it must undergo conversion before being absorbed by the electron; the rate of this is governed by the neutrino mixing angles and masses. In the Standard Model this processes has a very small branching

	Value
Number of data	30
G_1	260.
G_2	-0.040
G_3	-0.26
G_4	3.24
(Expt Fit)/Expt. (%)	4.1

Table 2.2: Experimental values of G_1 - G_4 as determined by TRIUMF [3].

ratio [22]. E.g. for $\mu \to e\gamma$, the branching ratio is:

$$\operatorname{Br}(\mu \to e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_{i=2,3} U_{\mu i}^* U_{ei} \frac{\Delta m_{i1}^2}{M_W^2} \right|^2$$
(2.7)

$$< 10^{-54}$$
 (2.8)

where $U_{\nu i}^{(*)}$ are elements of the neutrino mixing matrix, Δm_{i1}^2 is difference in neutrino masses and M_W is the mass of the W-boson.

This extremely low branching ratio makes cLFV a 'smoking gun' for new physics, should such a process be observed then it clearly indicates physics beyond the Standard Model. As a cLFV process has never been seen all current experiments have only been able to set upper limits on the branching ratio, the current world best is a limit of 7×10^{-13} (90 % C.L.) set by SINDRUM II in measuring $\mu + Au \rightarrow e + Au$ [23].*



Figure 2.10: $\mu \to e\gamma$, as a Standard Model process. The black dot represents the conversion of the muon neutrino to an electron neutrino, in the Standard Model this process has a predicted branching ratio $\sim < 10^{-54}$ and has not been seen.

There are three core cLFV processes of interest: $\mu \to e\gamma$, $\mu \to eee$ and $\mu + X \to e + X$ (where X is an atomic nucleus), the current limits of these processes are given in table 2.3. At higher energies there are further processes involving taus (e.g. $\tau \to \mu\gamma$) and hadronic systems (e.g. $K_L \to e\mu$) but for brevity they will not be discussed here.

^{*}This is an effective branching ratio in order to make direct comparisons with over cLFV processes easier.

Process	Limit (90 % C.L.)	Experiment	Ref.
$\mu \to e\gamma$	$< 2.4 \times 10^{-12}$	MEG	[24]
$\mu \to eee$	$< 1 \times 10^{-12}$	SINDRUM	[15]
$\mu + Au \to e + Au$	$< 7 \times 10^{-13}$	SINDRUM II	[23]

Table 2.3: Current limits on the standard cLFV processes for muons. Note that several different materials have been tested for the process $\mu + X \rightarrow e + X$ but the gold measurement, made at SINDRUM II, is currently the most accurate.

Comparisons of the limits in table 2.3 with the branching ratio in the Standard Model (equation (2.8)) show there is still large regions of phase space that has yet to be excluded. Many beyond the Standard Model theories predict higher branching ratios for cLFV processes through the addition of loop diagrams; by placing limits on the rates of the these processes the mass scale of the model can be constrained. If cLFV is observed in multiple processes then it becomes possible to constrain the model further as the branching ratios for the different processes are linked via the Lagrangian that describes them.

To study the general features of a model that includes measurable cLFV an effective Lagrangian can be constructed, one such Lagrangian, taken from [22], is:

$$\mathcal{L}_{cLFV} = \frac{m_{\mu}}{(\kappa+1)\Lambda^2} \overline{\mu}_R \sigma_{\mu\nu} e_L F_{\mu\nu} + \frac{\kappa}{(1+\kappa)\Lambda^2} \overline{\mu}_L \gamma_{\mu} e_L(\overline{e}\gamma^{\mu}e)$$
(2.9)

where L and R are chirality; $F^{\mu\nu}$ is the photon field strength and m_{μ} is the muon mass. The operators are parameterised using two constants: Λ , the effective mass scale, and κ , a dimensionless parameter that determines the relative strength of the Lagrangians' operators. In effect κ predicts the relative branching ratios between the different cLFV processes whilst Λ predicts at what mass new physics will occur. The first term is the term governing the rate of loop interactions (e.g. figure 2.10) whilst the second term determines the rate of the tree level interactions (i.e. the exchange of a heavy particle), because of the mass scale these are treated as contact interactions. It's important to note that other effective Lagrangians exist but this is a good illustrative example. Figure 2.11 shows the relationship between κ and Λ , as can be seen mass scales up to 300 TeV have already been excluded. Should MuSIC improve on SINDRUM I's limit for $\mu \to eee$ by the predicted factor of 100 then the region up to 1,000 TeV will also be excluded.



(a) $\mu \rightarrow e$ in Ti and $\mu \rightarrow e \gamma$ experiments

(b) $\mu \to eee$ and $\mu \to e\gamma$. The target sensitivity for $\mu \to eee$ measurements at MuSIC is shown in the solid black line.

Figure 2.11: Effective mass scales, Λ for expected sensitivities to muon cLFV processes for various interaction strength ratios, κ . Figures from [22].

Chapter 3

Simulation

3.1 Introduction

MuSIC was simulated to aid design of the detectors as well as to improve analysis of the data produced. Two core simulations were created, one using G4Beamline (G4BL) and the other using Geant4.

The G4BL simulation was mainly used to simulate the bulk of MuSIC and the hadron interactions. A more detailed and configurable Geant4 simulation was used to simulate the detectors. This section will discuss the general features of Geant4 and G4BL before a more detailed discussion of the actual implementation. The discussion will focus on the work done on the Geant4 simulation for the fifth MuSIC beam time as this was the most complete simulation performed. Simulations were used in earlier beam times but not to the same extent, although the G4BL simulation was used to interpret the total charged particle flux measurements. Less simulation work was done for the muon lifetime measurement as it was simpler to compare the measured lifetimes with precise values made by other experiments.

3.2 Geant4

Geant4 is described as 'a toolkit for the simulation of the passage of particles through matter' [25]. Rather than a complete program it is a collection of pre-compiled libraries that comprise the basics required to simulate various interactions. It has additional libraries that provide further functionality such as visualisation.

Geant4 is written in C++ using an object orientated approach that allows the user to either use Geant4's implementation of a class or write their own. This flexibility also means that all the functionality of Geant4 is explicitly opt-in, making the resultant programs much faster than they would otherwise be (as only the components that are selected are used). Geant4's speed does come at a cost which is that a larger amount of configuration is required than would be needed for less flexible solutions (e.g. G4BL).

Even with its increased flexibility Geant4 has certain requirements that must be met in order to work. These requirements embody the minimum information needed for Geant4 to create a simulation. The requirements are a definition of what to simulate, where to simulate it and the processes to simulate. Geant4's base requirement is that there is an instance of one of each of three classes^{*}:

- **The detector constructor** a description of the space in which the particle is to be simulated.
- The physics list which physical processes to simulate.

The primary action generator – a description of the initial particle to simulate.

With these three classes defined a simple program can be linked with the Geant4 libraries and run. The libraries use the primary action generator to create the initial particle, its interactions are described by the physics list and the world is specified by the detector constructor. In addition to these required components a simulation can be further refined with custom classes that replace or expand those supplied by Geant4.

The rest of this section will discuss how Geant4 uses the above three classes, and several others, to simulate a particle. This discussion will be split into two: discussion of the detector constructor and discussion of Geant4's execution model. The detector constructor will be treated separately as it is a primarily static description of a space. The execution model will cover the dynamic aspects of the simulation, e.g. how physical processes are implemented.

3.2.1 Detector Constructor

The detector constructor has two key jobs: defining the materials and defining the volumes of the detector. The materials describe the physical properties of various elements, compounds and mixtures. The volumes describe the position, rotation, material and logical attributes of the detector.

Material Properties

The material properties in Geant4 can be thought of as members of one of two sets: referred to here as 'core' and 'extended'. The core properties cover those values that are required for reasonable calculations of energy loss in the material, the extended properties cover more detailed calculations and properties that are more specialised.

The core properties for an element are the atomic number, the mass number and the molar density. Compounds and mixtures are described as ratios of their component elements at a specific density. Extended properties can be any property of the material that the user is interested in, for our purposes these were the optical properties of the materials and are discussed below.

Due to the computationally complex nature of interactions between light and material Geant4 doesn't implement them by default. When enabled Geant4 uses single values or

^{*}Whilst there is a technical distinction between an instance of a class (an object) and the class itself for the sake of clarity the word 'class' is used throughout.

tables of values to describe the various optical processes that it can simulate. The tables it uses are defined in terms of photon energy and missing values are generated via interpolation.

Detector Volumes

The detector volumes describe the physical pieces of the detector. Each volume to be simulated has its description split between three classes: the geometry, the logical volume and the physical volume. The geometry describes the shape and size of the volume. The logical volume describes how the volume relates to the larger simulation. Finally the physical volume places and orientates the volume with respect to the other volumes.

The volume's geometry is normally described by either basic polyhedra (cubes, tubes, trapezoids etc.) or compound shapes made from basic polyhedra. Geant4 provides basic tools for creating compounds made of the intersection, difference and addition of other geometries (either basic polyhedra or other compound polyhedra). Using these tools it's possible to create suitably complex constructions in code. Geant4 does offer more powerful tools (e.g. geometry importing) but these were not used for this simulation.

The logical volume stores the specific attributes of the component. The main information that the logical volume maintains is what material a component is (this can be changed using messengers) and what shape it is. A single logical volume can be associated with many physical volumes to make repetitive constructs easy to define.

As has been noted the physical volume positions the component. Additionally it stores which logical volume this component is in (as well as which logical volume describes it). Storing the parent volume means that a single logical volume can contain many sub-components to make bulk manipulation easier, as all placements are done with respect to the parent volume's origin.

3.2.2 Execution model

Geant4 manages simulation by splitting it into several layers with each layer representing a level of complexity within the simulation. The top-most level is the *run*, this represents a collection of *events* and all the information that is static between them (e.g. the detector configuration). An event is a collection of particle *tracks*. Each event starts with an initial particle (as defined by the primary action generator) the particle's track is split into *steps* with each step representing a period of time in which nothing 'significant' happens. As significant processes occur (e.g. particle decay) new tracks (to represent the new particles) are added to the event and they too stepped through their simulation. A track finishes when either the particle decays (and produces new tracks), it leaves the simulated volume or comes to rest. Once all the tracks finish, the event is over and the next event begins. Normally the user will specify the number of events that any run should contain, although other criteria can be used (e.g. a time out).

A 'significant' occurrence is one which will make further calculations of properties of the particle difficult and therefore limit the length of the particle's step. Determination of this step length is split between multiple classes within Geant4 but they can be broadly separated into four categories: detector considerations, user imposed limits, continuous physical processes and discrete physical processes. In each category the method is broadly the same, any class that can have an effect on the particle calculates a distance from the current position at which the effect will occur, e.g. a detector may report that the particle will leave the current volume in 10 mm whilst a decay process may report that the particle will likely decay in 5 mm. Once all the proposed step lengths have been determined the shortest is found and the particle is transported that distance and properties (position, time, energy, particle type(s)* etc.) are updated.

As has been stated the main consideration for the step length with respect to the detector are the volume boundaries which generally delimitate a change in material that will obviously determine which processes occur and how. Other detector considerations include the effect of a field on the particle's transport. Continuous processes generally only limit a particle by reducing its energy to zero (obviously some processes can increase it but these rarely limit a step length). Discrete processes obviously limit a particle by changing it in some way, for example by having a particle decay. Once a step length has been determined any continuous effects are applied to the particle, if it's at a volume boundary it moves into the next volume, any limiting discrete processes are applied and the next step length is calculated.

The processes (discrete or continuous) are supplied by the physics list. The physics list registers which processes can occur to which particles and then assigns a class to handle that eventuality. In this way physical processes can be enabled or disabled at will. It is also possible to register several classes to a single particle-process pair in order to describe different situations. The most common use of this is to apply different models depending on particle energy. Obviously the full range of physical processes that need to be described to accurately simulate the Standard Model is quite large, luckily Geant4 has extensive precompiled lists that cover the major types of particle interaction as well as several specialised lists that cover specific situations e.g. low energy regimes which are often ignored in high energy simulations as they can be computationally expensive.

3.3 G4beamline

G4beamline (G4BL) is a program written by Muons, Inc. [26]; it is a particle tracking simulation program that uses Geant4. Unlike Geant4, G4BL is a pre-compiled program where the user writes simple scripts that define the parameters of the simulation. In general creating simulations in G4BL is much quicker than Geant4 as a significant amount of 'boilerplate' code is removed at the cost of flexibility. The use of Geant4, rather than G4BL, for the detector simulation was mainly driven by the finer control Geant4 offered.

G4BL simulates the major interactions that a particle can undergo but cannot simulate optical processes, detectors and has more limited output options than Geant4. The physics used in G4BL is less well optimised for low energy than Geant4 can be made to be. It also does not have the same flexibility as Geant4 does for use with macros.

^{*}Should the particle decay it is replaced at the end of the step with the appropriate daughter particles.

3.4 Implementation

As has been discussed the simulation was split between G4BL and Geant4. This section will focus on the implementation of the Geant4 detector simulation as this constituted the bulk of the work carried out.

3.4.1 G4Beamline

G4BL was used to describe the bulk of MuSIC: the initial proton beam, the pion capture solenoid: target, capture solenoid, the shielding and return yoke; and the transport solenoids: bending magnets and pipe sections. A virtual detector, with radius 50 cm, was placed at the end of the beam pipe and recorded the particles that pass through it. A wireframe rendering of the set up can be seen in figure 3.1 and the co-ordinates used for the simulation and analysis are defined on it: X is horizontal across the beam pipe, Y is vertical across the beam pipe and Z is in the direction of the beam.

Using this script G4BL, simulated the computationally expensive hadronic interactions at the production target. The output of the G4BL simulation could then be reused for multiple detector set ups by importing the particle distribution into Geant4.

G4BL was used to simulate 900 million initial protons, the resulting beam at the end of the MTS was used as input for the Geant4 simulation. Figure 3.2 shows the flux for each particle type whilst figure 3.3 shows the total charged particle distribution at the end of the beam-pipe. Table 3.1 shows the number of each particle type and the mean position. As figure 3.2 shows the positive and negative particles split with this most clearly seen in the heavier pions and muons. Obviously the total particle flux is dominated by the electrons with the muons spread across the beam. The majority of the electrons, negative muons and pions originate in target interactions whilst positrons, positive muons and pions are produced both at the target in later interactions.



Figure 3.1: Wireframe of the G4BL simulation of MuSIC. The dark red is the iron yoke, the bright red are the solenoids, the light blue is the pion capture solenoid shielding and the yellow line is a simulated proton which doesn't interact in this case. The co-ordinates used are defined at the end of the beam pipe.





Particlo	Count	Horizonal position (mm)		Vertical position (mm)		
1 ai ticle	Count	Mean	RMS	Mean	RMS	
e ⁻	$1,\!362,\!161$	$-61.55 {\pm} 0.05$	$62.29 {\pm} 0.04$	$98.73 {\pm} 0.04$	44.62 ± 0.03	
e^+	183,074	$-37.10{\pm}0.16$	$67.91 {\pm} 0.11$	50.20 ± 0.15	$66.03 {\pm} 0.11$	
μ^-	8,940	$-45.10{\pm}0.75$	$70.97 {\pm} 0.53$	$65.35 {\pm} 0.69$	$65.54{\pm}0.49$	
μ^+	86,480	$-26.12{\pm}0.27$	$79.55 {\pm} 0.19$	-1.75 ± 0.27	$79.16 {\pm} 0.19$	
π^{-}	4,169	$-57.50{\pm}0.99$	$63.99 {\pm} 0.70$	$84.65 {\pm} 0.87$	$56.27 {\pm} 0.62$	
π^+	$53,\!406$	-44.19 ± 0.32	$72.79 {\pm} 0.22$	-28.50 ± 0.32	$74.04{\pm}0.23$	
р	$60,\!589$	-43.37 ± 0.30	$73.15 {\pm} 0.21$	-26.28 ± 0.31	$76.80{\pm}0.22$	
$\overline{\mathrm{p}}$	0	n/a	n/a	n/a	n/a	
n	$88,\!553$	-8.31 ± 0.33	$98.84{\pm}0.23$	0.10 ± 0.31	$92.78 {\pm} 0.22$	

Table 3.1: Summary of the particle distributions at the end of the MuSIC beam-pipe for 900 million initial protons. The number of particles of that type are given along with the mean and RMS of their distribution along the horizontal and vertical axis.



Figure 3.3: Distribution of total charged particle flux for X against Y at the end of the MuSIC beam-pipe.

The evolution of the final distribution is shown in figure 3.4 with the position of the monitors used shown in figure 3.5. The conversion of positive pions into positive muons can be seen with the increase of the latter at the second position.



Particle counts at various positions in the MTS

Figure 3.4: Evolution of the particle distribution in the G4BL simulation of MuSIC. Position 1 is 1 cm after the pion capture solenoid and positions 2 and 3 are after the first dipole and at the end of the beam pipe respectively.



Figure 3.5: Position of the beam monitors used to track the particle distribution through the G4BL simulation. Note: the monitors used to create the distribution used in the final Geant4 simulation have a much smaller, 50 cm, radius compared to 2 m used in this image (as the smaller radius monitors would not be visible).

3.4.2 Geant4

The Geant4 simulation was split into several key classes in addition to the detector constructor, primary action generator and physics list previously discussed. These classes were:

The magnetic field which gave the field strength within the detector.

- **The stepping action** which analysed every step of a particle's track and if it passed certain criteria, passed the track to ROOT for recording.
- **ROOT** which read and wrote ROOT files that were used to import the G4BL distribution and record events from the stepping action.
- The messengers a set of classes for configuring the simulation without re-compiling the program.

The implementation of each of these classes will now be discussed with the exception of the ROOT and messenger classes which primarily implemented utilities and had no impact upon the simulation or the data it produced.

The messengers were used to specify parameters to be used in the following run. These parameters are supplied via configuration files called macros. Four aspects of the simulation were controlled via macros: the primary action generator, the physics list, the magnetic field and the detector constructor. It was then possible to write further scripts that generated macros to test a range of configurations, e.g. a script was used to generate several macros that ran the simulation with a range of different detector setups.

3.4.3 Detector Constructor

The detector constructor is by far the most complex of the classes used in the simulation. The implementation of the detector construction was split over several functions, this discussion will follow similar lines. The materials will be discussed, then the optical properties and finally the geometry of the detector. Prior to these more detailed discussions a more holistic view of the detector will be given.

The detector simulated is the same as that used in the final measurements of the muon momentum spectrum. A schematic of the detector is shown in figure 3.6. The detector has four core components: the degrader, two counters and a stopping target. The degrader acts as a filter of entering particles' energy by reducing it; any particles with a momentum that is too low will stop in the degrader. The stopping target works in the same way but because there are counters on either side we can detect when a particle enters or exits. If a particle doesn't exit the stopping target within a certain time it can be assumed to have stopped. If the stopped particle is a muon we can then look for its daughter electron and construct a decay time spectrum. By selecting the thicknesses of the degrader and stopping target we can change the distribution of momenta that stop in the detector.

The layout of the components in the detector is: the degrader; the thin, upstream counter; the stopping target; then finally the thicker downstream counter. The degrader and stopping

target are of plates of metal (aluminium and copper respectively). The counters are each composed of several strips of scintillator (8 in the upstream counter and 5 in the downstream) that are spread across the height of the beampipe. Each scintillator strip has a Wave-Length Shifting (WLS) fibre mounted on the back of it which (in the experiment) was connected to a Multi-Pixel Photon Counter (MPPC). A schematic of a strip can be seen in figure 3.7.



Figure 3.6: Schematic of the detector configuration for simulation. In the face on view the dotted lines are the scintillators, blue for the eight upstream scintillators and red for the five downstream scintillators.

Materials

To simulate the detectors at MuSIC a fairly simple selection of materials were required, the elemental metals used for the stopping target and degrader; and several plastics that made up the scintillators. Geant4 treats materials as either pure elements or as mixtures (even if they're a compound). For the materials used at MuSIC only six elements were needed (H, C, N, O, Al and Cu), these had their standard values for mass number, atomic number and, where needed, density. The majority of these were used as components of plastics and



Figure 3.7: Schematic of a downstream paddle. The upstream paddles have the same layout but are have a different width and height (0.5 mm and 30 mm respectively). The WLS fibre protrudes from the scintillator for easy attachment of the MPPCs to the support structure.

for approximating the composition of air but copper and aluminium were used in their pure form (hence the inclusion of their density).

The compounds and mixtures (table 3.2) were, with the exception of air, components of the scintillators. Mylar was used to wrap the scintillators; EJ-212 [27] (a trade name for NE-102A, based on polyvinyltoluene) is the core scintillating material; and BCF-91A [28] is the (WLS) fibre, which has a polystyrene core with a PMMA cladding. The optical properties that were included in the simulation were:

- 1. The refractive index, n.
- 2. The emission spectrum that described the scintillation light produced by ionising particles.
- 3. Absorption spectrum, used in concert with an additional emission spectrum this described the WLS fibres.
- 4. Scintillation yield, for EJ-212 this was $10,000 \text{ MeV}^{-1}$ as stated by the manufacturer and assumes a linear response.

These properties were only implemented for the materials involved in the optical system. The emission and absorption spectra for the scintillator and WLS fibres were approximated to the curves given in the material data sheets ([27] and [28]), the values used are plotted in figures 3.8 and 3.9 respectively.

$\operatorname{Compound}/$	Components	Batio	Density	Refractive	Rof
Element	Components	Itatio	(g/cm^3)	Index	Iter
Air	N, O	4:1	1.290	1.00	—
Mylar	Н, С, О	4:5:2	1.40	_	[17]
EJ-212	Н, С	5.17:4.69	1.023	1.58	[27]
BCF-91A (core)	H, C	4.82:4.85	1.05	1.60	[28]
BCF-91A (clad)	Н, С, О	8:5:2	1.2	1.49	[28]

Table 3.2: Compounds and mixtures used for the simulation. No refractive index is given for mylar as it's opaque. An additional property of EJ-212 was the light yield which was $10,000 \text{ MeV}^{-1}$.



Figure 3.8: Scintillation spectra for EJ-212. The points represent the values passed to Geant4 and are based on those read from the manufacturer's data-sheet [27].

Geant4 simulated BCF-91A absorption/emission spectra



Figure 3.9: Absorption and re-emission spectra for BCF-91A wavelength-shifting fibre. The points represent the values used in Geant4 to approximate the manufacturer's curve as given in the data-sheet [28].

Detector geometry

MuSIC was simulated as an $8 \times 8 \times 8$ m³ volume of air containing the detector. The size of the volume was chosen to simplify transformations of the detector with respect to the magnetic field. The detector was implemented as four placed volumes inside a single logical volume that allowed all of them to be correctly positioned with respect to the magnetic field. The four sub-volumes were the upstream counter, the downstream counter, the degrader and the stopping target. This was the set-up used for the muon lifetime measurements and the final momentum spectrum measurements made during runs 4 and 5.

The degrader and stopping target were configured so that their material and thickness could be changed using macros. This made testing configurations prior to beam-time much quicker and easier. This also meant that the degrader could be removed by creating a volume of air. The materials used for the degrader and stopping target were simple elemental metals (aluminium and copper respectively). Both components were implemented as simple sheets with the dimensions given in table 3.3.

The counters were more complex. Each counter was split into a number of 'strips' (eight upstream, five downstream). A strip was composed of a scintillator, a 135° section of pipe (also made of scintillator material) that mimicked the optical cement and a length of wavelength-shifting fibre (made of a core and shell). Mylar was then used to define an optical surface that existed at the boundary between scintillator and air. An optical surface is a special construct in Geant4 that is not a full volume but a set of physical effects that

Component	Matorial	Dimension (mm)		
Component	material	Х	Y	Ζ
Stopping Target	Copper	370	310	0.5
Degrader	Aluminium	400	400	0, 0.5, 1, 5

Table 3.3: Dimensions of the stopping target and degraders. The degrader width of 0 mm is when the degrader is removed and replaced with air. The Z axis is the main beam direction with the X and Y axes being in the horizontal and vertical directions respectively.

Component	Material	Number	Dimensions (mm^3)
Upstream	EJ-212	8	$380 \times 30 \times 0.5$
Downstream	EJ-212	5	$380 \times 50 \times 3.5$
WLS fibre (core)	BCF-91A (core)	1/paddle	$500 \times (r - 0.5)$
WLS fibre (cladding)	BCF-91A (clad)	1/paddle	$500 \times (7 - 0.5)$
MPPC	Air	2/paddle	$1 \times 1 \times 1$

Table 3.4: Dimensions of the simulated paddle components. The wavelength-shifting fibres were the same for both up- and down-stream counters, as were the 'MPPCs'. The WLS fibre's cladding is said to be 3% of the total diameter so this is what is used [28].

occur when optical photons move from one volume to the next, in this way the reflective effect of mylar wrapping was created. A $1 \times 1 \times 1 \text{ mm}^3$ volume of air is placed at both ends of each wavelength-shifting fibre to act as a Multi-Pixel Photon Counter (MPPC) detector region. The dimensions used for the paddles are are given in table 3.4 the entire layout can be seen in figure 3.6.

3.4.4 Primary Action Generator

The primary action generator is a relatively simple class that specifies the properties of the initial particle to be simulated. This means, for all implementations, defining the particle's type, position within the volume and initial momentum.

At MuSIC the key task of the primary action generator was to either import the particle distribution from the G4BL simulation or create particles with a similar distribution. The source and type of particle could be configured using messengers.

3.4.5 Physics List

Geant4 provides the majority of physics processes as precompiled lists that cover the major aspects of a simulation. For MuSIC the key attribute of each list was the accuracy at low energies (< 20 MeV). In this respect the most important lists were those dealing with

electromagnetic processes. To cover this 'option3' and the 'extra' lists were used which have improved low energy simulations and more detailed processes. Most importantly the 'extra' list implements muon capture at rest. The few hadronic processes were simulated with the 'HP' QGSP-BERT and the elastic physics models, both of which have better methods for simulating low energy neutrons.

One problem with the default supplied physics lists is that the implementation of muon capture is wrong and doesn't accurately model the muon capture rate for copper. This is obviously a problem as copper is the main stopping target used at MuSIC. The apparent lifetime for negative muons in copper produced using unmodified Geant4 was ~ 1.1 μ s compared to the experimental value of 163.5 ns [3] or the theoretical value of 194 ns (calculated using equation 2.6). This difference can be seen in figure 3.10 and figure 3.11 which show the muon decay spectra of negative muons interacting with a 1 mm copper stopping target using the unmodified and modified versions of Geant4 respectively. As these are illustrative plots only negative muons were simulated and the only material was copper.

The modification used was to the 'G4StopElementSelector' class where the function 'Get-MuonCaptureRate' had an additional element added to its list experimentally determined muon capture rates. The added value corresponded to the muon capture rate in copper, which is $5.72 \pm 0.05 \times 10^6$ s⁻¹. This capture rate corresponds to 93.5 % of the negative muons being captured by the nucleus and the remaining 6.5 % decaying in orbit. This was acknowledged as a bug by the Geant4 development team and has been fixed in more recent versions of the framework (i.e. more later than 4.9.6 patch 02).

3.4.6 The Field

Field maps for the entire detector region were available but to improve efficiency only the region directly surrounding the detector was simulated, this is shown in figure 3.12. Geant4 uses interpolation to find values for the magnetic field that weren't present in the file. The file listed position and the magnetic field at each position. To save space the file relied on the vertical symmetry to account for the lower half of the detector.

Stepping Action

Geant4's stepping action was used to record the status of the simulation. The stepping action is called for every step a particle makes along its track. Using this, steps of interest were written to a ROOT file for later analysis. Two sets of data were recorded: the 'truth' data and the 'MPPC' data. The truth data recorded every charged particle in any of the core detector components: the degrader, the scintillators, and the stopping target. The data recorded was:

- Particle type
- Position
- Momentum



muon lifetime in 1mm Cu (un-modified Geant4)

Figure 3.10: Decay time spectrum of 1×10^6 negative muons hitting a 1 mm copper stopping target in a vacuum using an unmodified version of Geant4. The decay times are fitted using $N \exp(-x/\tau)$. As can be seen the Geant4 value, 1,161 ns, does not agree with either the experimentally determined value of 163.5 ns [3] or the theoretical value of 194 ns (calculated using the equation 2.6). This bug has been fixed in more recent versions of Geant4 (versions > 4.9.6 patch 02). Only the muonic lifetime in copper is fitted as the decays are dominated by captured muons.



muon lifetime in 1mm Cu (modified Geant4)

Figure 3.11: Decay time spectrum of 1×10^6 negative muons hitting a 1 mm copper stopping target in a vacuum using a version of Geant4 modified to use the experimental muonic copper lifetime of 163.5 ns. The spectrum is fitted using $N \exp(-x/\tau)$. Only the muonic lifetime in copper is fitted as the decays are dominated by captured muons.

- Time since the start of the event
- Which detector component the particle was in
- Track ID
- Parent particle's track ID

The MPPC data attempted to more closely approximate the actual data produced by the detector. This was done by logging photons that were found in an MPPC volume, unlike the truth data where particles were recorded if they were in the scintillator at all (i.e. could conceivably be detected). To reduce the overhead of running the simulation only the time and position of the photon were recorded.

3.5 Simulation Analysis

As well as providing a method to test the effects of different detector constructions the Geant4 simulation was also used to verify the analysis technique used in the final measurements. To test the analysis technique a very large numbers of muons were produced using a Gaussian approximation of the G4BL muon distribution. This approximation was useful for verification of the analysis technique but could not be used with actual data as the distributions they produced were poor fits to the G4BL output.



Magnetic field at MuSIC detector

Figure 3.12: The magnetic field strength for the region of interest. The black bar marks the position of the scintillators and stopping target.

The core analysis technique used was to create muon lifetime distributions, fit them with a function and then use the integral of the appropriate portion of the function to calculate the number of muon decays that had occurred. This method could be used because no other constituent of the beam would have a lifetime comparable to that of the muons: electrons and protons are stable whilst charged pions have a lifetime of 26.003 ± 0.005 ns [2] which is suitably short, even compared to muonic copper, to have little impact on the muon decay spectrum (neutral pions have a lifetime of 8.52×10^{-8} ns and don't produce charged daughters).

The first step of the protocol was to create histograms of downstream-upstream hit times. In the measurements this was done by an item of hardware but for the simulation this was done by looking for muons in the upstream counter and their daughter-electrons in the downstream counter. The difference in times was then binned, any differences of less than 50 ns were ignored as this was a region that we were blind to in the experiment.

During verification a variety of bin widths were tested in an attempt to minimise the error on any one bin whilst aiming to avoid loss of signal. The was especially important for fitting the copper component of the decay spectra as this is significant contribution for only the first few bins. Ultimately bin widths of 16 ns were settled on.

The simulation was then fitted using a double exponential:

$$N_{cu}\exp(-t/\tau_{cu}) + N_f \exp(-t/\tau_f). \tag{3.1}$$

as the simulated data was pure the background terms used in the final analysis (see section 4.4.4) were not included. Using the same method as the experimental technique the values of τ_{cu} and τ_f were fixed at their measured values (163.5 ns and 2,197 ns respectively) for the fitting. Unlike the previously fitted data (figure 3.11) the large number of positive muons overwhelm the effect of negative muons captured by air. The fitted values for equation (3.1) were extracted into individual exponentials and integrated over the region 50 ns to 20 μ s.

The integrated approximation to the number of decays could then be compared directly to the counted number of muon decays in the simulation. The results of the integration method compared to the counted number of decays are given in table 3.5. As can be seen from the table there is good agreement between the counted number of decays and the number produced by the integration method. This verifies the use of the integration method which is the simplest method of estimating the number of muon decays and the only method that can be used with this data.

Degrader	Thickness	Free component		Copper of	component
Material	(mm)	Count	Integral	Count	Integral
Air	5	11722 ± 108	11613 ± 119	881 ± 30	965 ± 58
	0.5	11827 ± 109	11789 ± 119	837 ± 29	852 ± 57
Δ1	1	11723 ± 108	11739 ± 118	891 ± 30	852 ± 56
111	5	7078 ± 84	7079 ± 92	453 ± 21	441 ± 42

Table 3.5: Comparison of integrated numbers of muon decays to the counted number of decays in simulation of 5×10^5 initial muons. The free and copper components of the fit were considered separately. The results are for negative muons only as the positive muons only have the free muon lifetime.

Chapter 4

Characterising the beam

4.1 Introduction

Characterisation of the MuSIC beam was carried out over two years and through a number of experiments. Over the course of two years there were five runs making three significant measurements to characterise the beam: the total charged particle flux; the muon flux (by measuring the muon lifetime); and the muon momentum spectrum. The details of the runs are covered in table 4.1 along with which experiments were carried out. Several other experiments were also carried out by other groups but these are not discussed here (e.g. neutron flux measurements and Mo muon-bombardment experiment).

Dates		Measurements
Start	Stop	
29 July 2010	31 July 2010	Charged particle flux.
13 February 2011	16 February 2011	Charged particle flux.
15 February 2011		Muon lifetime.
19 July 2011	21 July 2011	Muon lifetime.
		Muon yield (via muonic X-rays).
22 October 2011	23 October 2011	Neutron flux.
		Muon momentum spectrum (via lifetime).
18 June 2012	22 June 2012	Muon yield (via muonic X-rays).
		Mo muon-bombardment.

Table 4.1: A summary of the five MuSIC beam-times with notes on the measurements made.

This chapter has been split into four sections: the rest of this section will be an introduction to the equipment used for the measurements, with each of the three measurements then treated individually. Each measurement section will cover the set up, the results and analysis of the data.

4.1.1 Scintillators Preparation

As discussed in chapter 2, scintillating materials produce light when charged particles travel through them. There are several categories of scintillator, those used for our measurements were plastic as they are easy to work with, inexpensive and available in a range of sizes. There are two main considerations in preparing a plastic scintillator: maximising light collection and preventing external light contamination. The standard approach to deal with both of these considerations is to polish and wrap the scintillator.

The scintillator surface is highly vulnerable to damage either through scratches or grease. Damage to the surface of the scintillator is problematic as it inhibits light collected through total internal reflection (TIR). Anything that reduces the amount of TIR increases the amount of light absorbed outside of the scintillator, where it cannot be detected. In order to prevent scratches, scintillators must be handled carefully and protective films are only removed at the last moment. As well as careful handling, the scintillators are cleaned to remove dirt and dust by polishing them with iso-propanol, this removes grease that can also degrade the surface.

Scintillators are generally wrapped with two thin layers: an inner reflective material and an outer blackout material. The inner layer is normally aluminium-mylar or foil, this is wrapped loosely in order to leave a small air gap which increases the amount of TIR (and hence light collection). The outer layer is normally a black plastic wrap. This is carefully sealed with tape to prevent light leakage (another aid to this is obviously turning any lights off in the experimental area). An important consideration is to keep the layers as thin as possible in order to minimise energy deposition; to this end overlaps are kept to a minimum.

In order to attach sections of the scintillators together and to attach detectors to the scintillator optical cement and grease were used. Optical cement was used when a permanent fixing was required while grease was used if the components would need to be separated. To use the cement the components were cleaned using iso-propanol, then the two-part glue mixed and applied. A foam jig was used to hold the pieces in place whilst the cement set. When non-permanent joins were required, e.g. to attach MPPCs to WLS-fibres, optical grease was used to form the connection. Optical cement and grease have refractive indices close to (if not the same as) that of the scintillators to maximise transmission either between sections of the scintillator (e.g. from the scintillator bulk to the WLS-fibre) or from the scintillator to the MPPC.

4.1.2 Multi-Photon Pixel Counters

Multi-Photon Pixel Counters (MPPCs, see figure 4.1) are highly sensitive devices able to accurately count photons over a wide range of intensities. An MPPC is a small device (normally $\mathcal{O}(1 \times 1) \text{ mm}^2$) made from many Avalanche PhotoDiodes (APD) operated in 'Geiger-mode'. Each individual APD is a single pixel within the MPPC. The APDs are connected together so that the MPPC's output is the sum of the outputs of the individual pixels. As can be seen in figure 4.2 there is clear banding that corresponds to the number of incident photons making counts, accurate to the single photon, obtainable when operating
within the MPPC's limitations (see below).



Figure 4.1: Photograph of an MPPC (type S10362-11 with ceramic package) from the Hamamatsu technical report [29].

APDs make use of the photoelectric effect and large voltages to produce an 'avalanche' of electrons when triggered by an incident photon. A limitation to this system is that an individual APD's output is roughly constant regardless of the number of incident photons. Should the number of photons become comparable to the number of pixels then the MPPC becomes saturated and produces a constant signal as all the pixels fire, rather than clear bands as pixels fire in isolation. For the energy range predicted at MuSIC saturation is not considered to be a significant problem as the number of photons is predicted to be below the one hundred pixels used in our MPPCs (see figure 4.3).

An MPPC has several key attributes: Photon-Detection Efficiency (PDE), peak sensitivity, time resolution, operating voltage (V_0) , gain (M) and dark current. The PDE and peak sensitivity respectively describe the likelihood of detection of a photon of given wavelength (see figure 4.4) and the wavelength that the MPPC is most sensitive to (typically ~440 nm). Time resolution for MPPCs is generally very good, normally 200–300 ps for FWHM at the single photon level which is more than accurate enough for our purposes. The operating voltage is the potential required to make the MPPC work, due to variance in manufacture this is given individually for each MPPC and has to has to be set correctly for optimum performance, Hamamatsu's MPPCs have $V_0 = 70 \pm 10$ V. The gain of the MPPC indicates the strength of the signal response to a photon, typical values are between 10⁵ and 10⁶, this corresponds to signals of $\mathcal{O}(1-10) \ \mu V$ which must be amplified for accurate measurement. The dark current is a measure of how noisy a particular MPPC is; it is the rate of false signals that have an effective strength equivalent to 0.5 photo-electron (p.e.). Typical values for the dark current are in the range 100–500 kcps (×10³ counts per second) these can be accounted for by correct setting of triggers.

As can be seen, whilst MPPCs have many useful features they do have several features that need careful treatment to make them useful: mainly amplification and noise reduction. These two problems are closely related and in fact, when treated properly one will often help with the other. Given the size of the signal from an MPPC it is obvious that amplification is



Time

Figure 4.2: Oscilloscope waveform for an MPPC with a histogram of peak voltage on the left. The histogram shows the clear discrimination between number of photons available on a MPPC. This type of MPPC has $400 \times 50 \times 50 \mu \text{m}^2$ APDs in a $1 \times 1 \text{mm}^2$ package taken from [29].



Figure 4.3: Simulated distribution of number of photons reaching one of the MPPCs attached, via wavelength-shifting fibre, to a $380 \times 50 \times 3.5$ mm³ scintillator.



* Photon detection efficiency includes effects of crosstalk and afterpulses

Figure 4.4: MPPC Photon Detection Efficiency (PDE) as a function of photon wavelength. A range of pixel densities are shown (100, 400, 1600 pixels/ 1×1 mm² for the 100U, 050U and 025U models respectively). Taken from [29].

required to make the signal useable, in fact, as will be discussed below, an unfiltered MPPC signal is too small for detection by most standard equipment. Amplification is done using linear amps which are applied as soon as possible to reduce noise due to cabling as well as attenuation.

4.1.3 Secondary Emission Chamber

As well as studying the resultant beam, knowledge of the initial proton beam is essential for normalisation. This is done using a Secondary Emission Chamber (SEC). A SEC uses thin foils of gold placed within a strong electrical potential. When protons pass through the foils a small number of electrons are produced as part of ionisation, these can then be collected and counted. The SEC only measures a fraction of the beam and even then, only indirectly. In order to calibrate the SEC, runs are performed in which the entire proton beam is absorbed with a copper block downstream of the SEC. Measurement of the total current produced by the block can then be used to determine the fraction absorbed by the SEC, and hence the conversion factor. Measurement of the SEC was carried out using scalers (see below) that counted the cumulative charge passing through the SEC.

4.1.4 NIM, CAMAC and VME

Rather than develop custom data acquisition hardware, a modular crate system was used. A crate system supplies power and mechanical fixings for a range of modules, these modules can then be connected together to produce a system. Three crate systems were used: Nuclear Instrument Module (NIM), Computer Automated Measurement And Control (CAMAC) and

Versa Module Eurocard bus (VMEbus or VME). NIM supplies power only, whilst CAMAC and VME provide a 'backplane' through which data can be transferred to a control card.

NIM modules were primarily used to provide signal processing and data acquisition logic. The CAMAC and VME crates were used for data readout. Ultimate control and storage was carried out using a PC which recorded the data for later analysis. The rest of this section will discuss the basic modules used in NIM, CAMAC and VME.

4.1.5 Discriminators

A discriminator works by producing a digital signal when an analogue input signal exceeds some preset threshold. Converting the analogue signal to digital makes it much easier to manipulate as many other modules can only use a digital signal. The discriminators were powered by NIM modules and had their thresholds set manually.

Discriminators were used to signal when a certain number of photons were detected by an MPPC i.e. when a particle had passed through the scintillator. As the discriminators have a minimum threshold of ~ 25 mV the MPPC signals were amplified to make them useable. As every MPPC has a different gain the thresholds for the discriminators had to be set individually. To normalise between MPPCs the discriminator thresholds were set based on the number of photons the analogue signal corresponded to (see figure 4.2). The trigger level (in p.e.) was the same for all MPPCs and was chosen to maximise the signal to noise ratio, a typical value was ~ 2.5 p.e. which suppressed most of the dark current without compromising sensitivity to charged particles. An oscilloscope could then used to translate the photon (p.e.) threshold into a voltage threshold (in mV) that can be used by the discriminator.

4.1.6 Gates and Latches

Gates are a set of modules that generally have several related functions. Depending on the mode they can change the length of a signal, delay it, 'latch' it or start a clock signal. The first two modes are generally used for creating 'gates' (on or off signals) for other modules: e.g. turning on a module to record the shape of a signal or indicating to the PC that there is data waiting to be read. A 'latch' is a signal that remains on until another signal switches it off, these are often used when it's unclear how long something will take e.g. sending data to the PC. Clock signals were generally used as calibration information for other measurements e.g. recording the SEC.

4.1.7 Logic units

Logic units provide boolean logic ('AND', 'OR', 'NOT') for processing digital signals. The limiting factor is the number of inputs that can be combined and the complexity of the combinations. Normally signals are combined in a block with a single module having several discrete blocks. Each block will perform a single logical operation (OR or AND) on its inputs.

The inputs of a block can either be normal or negated. A block (or sometimes the entire module) will generally have a veto signal that will stop outputs until the veto is cleared.

Logic units were generally used for simple tests such as checking that all MPPCs on a scintillator had triggered, this helped reduce the number of false positives in addition to the application of a discriminator. Logic units were also used to prevent attempts to process two events at once: if a second event arrived before the previous event was fully processed then the second event was ignored. This leads to 'dead time' but is unavoidable.

4.1.8 Scaler

A scaler is a counting module. It receives a digital input, that when asserted, increments its counter. The most important factor of a scaler is its maximum speed (the maximum input frequency). Events that occur more rapidly than the input frequency will not be properly counted (either not being counted at all or counted as a single event), a typical maximum input frequency is 100 MHz. Most scalers can be daisy-chained together in case their maximum value is exceeded and some allow multiple inputs that can be counted independently. NIM scalers will generally be readout by eye (e.g. using a display) whilst VME/CAMAC modules will be readout using the crate's data-bus.

4.1.9 Analogue to Digital Converter

There are two primary attributes of an MPPC signal that were measured: its size and its separation from other signals. In order to measure the size of the MPPC signal, an Analogue to Digital Converter (ADC) was used. ADCs work in a variety of ways depending on the exact type of measurement they are making, two common versions used at MuSIC were Peak Sensing ADC (PS-ADC) and Charge-integrating ADCs (QDC). Both types of ADC measure a component of the input analogue signal, the PS type measures the peak voltage whilst the QDC measures the integrated charge, both make these measurements only when enabled by the ADC's 'gate'. The ADC's gate is normally set to be slightly longer than the expected signal from an MPPC, i.e. \sim 50 ns.

Several factors define the ADC: the number of bits it is able to read out, the range of input signals that it can measure, its linearity, digitisation/read time ('dead time') and number of channels. The number of bits the ADC has, the range and its linearity work together to determine the accuracy of the ADC; the number of bits determines the resolution between any two values, the range determines the values that the ADC can measure whilst the linearity maps measured values to actual voltages. Dead time measures how long is required to measure the analogue signal and how long it takes to then send the digitised values to the controlling PC i.e. for how long the module is inoperative, 'dead'. The number of channels on an ADC is a measure of how many inputs it can measure simultaneously, normally one is ascribed to each MPPC so that the triggering signal can be recorded.

A key feature of using an ADC is the pedestal, obviously any analogue signal is going to have some noise that represents its zero level, in an ADC this manifests as a large peak in the lower bins of the ADC, removal of the pedestal is often done by the ADC itself and through the DAQ but it does sometimes have to be removed from the data as well.

4.1.10 Time to Digital Converter

A Time to Digital Converter (TDC) measures the time difference between signals. The core parameters for a TDC are the number of bits of precision on each channel, the maximum duration the TDC can record, its linearity, its dead time and number of channels of input. A TDC measures time between a 'start' and a 'stop' signal. Normally a single start signal is provided to the entire TDC with each channel receiving a stop signal. The measured times between the global start signal and the individual stop signals are then readout independently.

An alternative TDC used in later experiments is the Multi-Hit TDC (MH-TDC). Rather than measure a single start/stop pair for each channel the MH-TDC has a ring-buffer to record data for each channel. A ring-buffer has a fixed length (the 'window'), typically $\mathcal{O}(50) \ \mu$ s, when a signal (a 'hit') is received its position within the window is marked. The ring-buffer records every hit it receives, regardless of the start signal. When the start signal is received the buffer is readout, this can be done immediately or after a delay. By using a delay signals both before and after the start signal can be recorded.

4.1.11 Registers

Registers are modules that allow simple logical signals to be sent to or from a PC. In our systems they have two key uses: either indicating to the controlling PC that there is data to be read (an 'interrupt') or for the PC to indicate that all the data has been read and that the DAQ should be reset. The interrupt time had to be carefully delayed such that there was long enough for the digitisation to complete whilst minimising dead time. The DAQ reset signal was sent by the PC once it had read all the modules to indicate that data taking could resume.

4.2 Charged Particle Flux

The first experiment carried out at MuSIC aimed to measure the total flux of charged particles. As well as this measurement it was also a test-bed for the technology and techniques used in later experiments.

Two separate measurements of the charged particle flux were made: one using a long strip scintillator that measured the flux at a range of vertical positions and a second that used a smaller disk that measured the flux at a number of points across the face of the disk.

The DAQ used for these experiments was designed to have a lot greater functionality than was ultimately used. The original design aimed to use the difference of arrival times at the various MPPCs for greater positional accuracy whilst the ADC measurements were intended to measure the energy of the particle; neither of these techniques were successful but the full DAQ is included for completeness.

4.2.1 Experimental Set Up

The experimental design for both measurements was broadly the same: the scintillators both had four MPPCs attached. In the first measurement a long strip scintillator was used and the MPPCs were distributed evenly across the short ends, for the second measurement the disk had the four MPPCs placed evenly around the circumference (see table 4.2). NIM crates housed the modules for the triggering logic, digitisation was performed via CAMAC modules and the read out data was recorded on a PC to be analysed afterwards. The scintillators were both wrapped with aluminium foil and black wrap and the MPPCs attached via optical cement. The scintillators were mounted on aluminium supports at the end of the beam-pipe.

Shape	Volume (mm^3)	MPPCs	MPPC positions
Strip	$380 \times 30 \times 10$	4	± 5 mm on each end.
Disk	$35^2 \times \pi \times 20$	4	Radially, every 90°.

Table 4.2: Details of the scintillators used for the two measurements of the charged particle flux.

The aim of these early experiments was to make three measurements: the total trigger rate, the difference in arrival times of light at the MPPCs and the distribution of the number of firing pixels at each MPPC. The total trigger count is proportional to the flux of charged particles through the scintillator. The timing and photon counts were hoped to give information on positioning and the amount of energy deposited but this was found to be incorrect. In the case of positioning information it was hoped that the left/right correlation of arrival times could be used to position the interaction within the scintillator but it was found that the variance on arrival times was too large compared to the difference between them so accurate resolution was impossible. Measurement of the particles' energy didn't succeed due to lack of accurate calibration data.

Ultimately only the measurement of the trigger rate was successful. The trigger rate is proportional to the charged particle flux, assuming negligible dead time. To ensure a negligible dead time, as discussed below, a simplified DAQ system was used and MuSIC was run with a reduced beam current (< 1 nA instead of the MuSIC design current of 1 μ A). Another concern for measuring the charged particle flux are minimally ionising particles (MIPs). MIPs, by definition produce a minimal amount of scintillation light, making their detection difficult. The thickness of the scintillators, especially the second (20 mm) should still be sensitive to MIPs.

The DAQ used to measure the trigger rate was a subset of the full DAQ, used to measure the time and energy:

- 1. Amplify the MPPC signal to a detectable level.
- 2. Remove dark current using a discriminator.
- 3. Create a trigger and veto from the coincidence of all MPPCs: an 'AND' operation applied to the output of the discriminators.
- 4. Record the analogue MPPC signal using an ADC.
- 5. Use the trigger and the first MPPC's signal to start the TDC.
- 6. Delay three of the remaining signals to form the stops to the TDC (the delay can be removed to regain the accurate time).
- 7. Remove the veto on the system once readout is complete.

This can be seen as a schematic in figure 4.5.

To make the final measurement of the trigger rate, the majority of the DAQ was ignored and the four-fold coincidence counted using a scaler. By counting the raw triggers, effects of dead time could be negated and a truer value for the number of interactions obtained. One down side of this design is that by using a four-fold coincidence no cross-checking between channels could be carried out to find relative efficiencies, unfortunately this deficiency was not noticed until too late. For the 1D experiment, measurements were made over 50 s with control done via hand, this obviously had limitations.

The positions of the two experiments are given in table 4.3 with respect to the centre of the beam-pipe. 'By-eye' readout of the 1D measurement was done using a display mounted on the scaler used to count triggers. For the 2D measurement a CAMAC module was used that was read out by computer. It's important to note that not only were the detectors very different but due to experimental constraints the 2D measurement was made significantly further from the end of the beam pipe than the 1D experiment (85 cm compared to 6 cm).



Figure 4.5: Schematic of the DAQ used for measuring the charged particle flux. Horizontal blocks indicate NIM units whilst vertical units were mounted in a CAMAC crate. The two 'AND' blocks were made using NIM modules.

Measurement	Distan	ce from centre of be	Scaler	Readout	
	Horizontal	Vertical	Longitudinal	Time (s)	
1D	0	-5, -15, 0, 15, 5	6	50	By-eye
2D	-17, 0, 17	-16, 0, 20, 25	85	20	CAMAC

Table 4.3: Positions at which the charged particle flux was measured, 1D refers to the first run in which only the vertical displacement was measured, 2D refers to the second run in which horizontal measurements were also take. The distances from the beam-pipe were 6 cm for 1D measurements and 85 cm for 2D, this was due to mechanical constraints.

Configuration	Run Time	MPPC	PMT	Fficioner	Adjusted
Configuration	$(\times 10^{3} {\rm s})$	Count	Count	Enciency	Efficiency
123	86.0	1651	$1,\!161,\!165$	0.00142	0.0591 ± 0.0015
12	92.3	6055	$1,\!229,\!389$	0.00493	0.2048 ± 0.0026
13	87.1	5653	$1,\!179,\!116$	0.00479	0.1993 ± 0.0027
23	84.4	3999	$1,\!129,\!137$	0.00354	0.1472 ± 0.0023

Table 4.4: Summary of the data taken for measuring the MPPC detector efficiency. Configuration refers to which of the three MPPCs were tested. The efficiency is the simple ratio of the MPPC count to the PMT count whilst the adjusted efficiency has been scaled by the ratio of the areas of the scintillators, i.e. by $\frac{40 \times 40}{\pi^3 5^2}$.

4.2.2 Detector efficiency

As well as making measurements of the charged particle flux another important measurement, made soon after the second beam time, was measurement of the efficiency of the circular detector. This was done using three large $(40 \times 40 \times 1 \text{ cm}^3)$ scintillator paddles. Light from the scintillators was detected using Photo-Multiplier Tubes (PMT). The three large paddles were positioned sandwiching the circular scintillator, two above and one below.

The circular detector's efficiency was calculated using the ratio of detections by the PMTs to detection by the MPPCs (once size considerations had been taken into account). It was assumed that the rate of detection by the circular scintillator was proportional to the detection rate in the three large scintillators. The first test was the efficiency of all three remaining MPPCs (one had become irrevocably damaged during beam time) then the efficiency of pairs of MPPCs was tested. Testing consisted of counting the occurrence of three-fold coincidence on the PMTs and either two or three fold coincidence on the MPPCs. Data was taken for approximately one day for each configuration. The MPPCs were not tested individually to ensure that dark current events were not included, although these were reduced by using a discriminator to produce the signals unfortunately this introduced the same bias as discussed in section 4.2.1. The results of this measurement are shown in table 4.4.

Assuming that the total efficiency (ϵ_t) is the product of the individual efficiencies $(\epsilon_{1,2,3})$ then it can be expressed as:

$$\epsilon_t = \epsilon_1 \epsilon_2 \epsilon_3 \tag{4.1}$$

with the measured efficiencies of the pairs of MPPCs we can calculate the individual effi-

ciencies by solving:

$$\epsilon_{12} = \epsilon_1 \epsilon_2 \qquad \implies \epsilon_1 = \frac{\epsilon_{12}}{\epsilon_2}$$

$$\epsilon_{13} = \epsilon_1 \epsilon_3 \qquad \implies \epsilon_2 = \frac{\epsilon_{12} \epsilon_3}{\epsilon_{13}}$$

$$\epsilon_{23} = \epsilon_2 \epsilon_3 \qquad \implies \epsilon_3 = \sqrt{\frac{\epsilon_{23} \epsilon_{13}}{\epsilon_{12}}}$$

This can be generalised to:

$$\epsilon_i = \sqrt{\frac{\epsilon_{ij}\epsilon_{ik}}{\epsilon_{jk}}} \tag{4.2}$$

where ϵ_i is a single efficiency we want to calculate and ϵ_{jk} is the combined measurement of the other two efficiencies. Applying this to table 4.4 we get the efficiencies for the MPPCs shown in table 4.5. It's important to note that these efficiencies represent more than just the individual MPPC's quantum efficiency but the entire gestalt of systematic effects that contribute to the efficiency of the specific MPPC including scintillator acceptance (both geometric and energetic), light collection and transmission. As can be seen the estimated total efficiency ends up being larger than what was actually measured this is likely due to the bias introduced by requiring coincidence. Using these values the average efficiency of a single MPPC can be calculated and given an error as seen in the final line of the table. The error on the average efficiency was taken to be the proportional difference between the calculated and measured efficiencies, i.e.:

$$\Delta_{\epsilon} = \epsilon_{\text{Ave.}} \times \frac{123_{Calc} - 123_{Meas}}{123_{Meas}} \tag{4.3}$$

as well as the discussed bias it is important to emphasise that this efficiency is strongly coupled to both the threshold used with for the discriminators and the physical set up of the detector. Its use in later calculations (specifically those in section 4.4) is done with the understanding that is a better approximation to the efficiency that none but not a replacement for an actual measurement which could not be carried out.

4.2.3 Results

Prior to the main experiments the conversion factors from SEC count to proton current was determined, these values were determined to be, respectively for the 1D and 2D measurement: 0.03408 nA and 1.514 nA. The results from the measurements are presented in table 4.6, for the 1D case, and table 4.7 in the 2D case. The counts were converted to rates and the flux calculated using:

$$j = \frac{F(C - C_{off})}{I_p} \tag{4.4}$$

$$I_p = K(S - S_{off}) \tag{4.5}$$

MPPC	Efficiency
1	0.5265 ± 0.0064
3	0.3786 ± 0.0046
2	0.3889 ± 0.0048
123_{Calc}	0.0775 ± 0.0016
123_{Meas}	0.0591 ± 0.0015
$\epsilon_{\text{Ave.}}$	0.431 ± 0.134

Table 4.5: Efficiencies of individual MPPCs calculated using the values from table 4.4, and equation (4.2). The two values below the line represent the total efficiency as calculated using the individual values and the measured value. The final value, $\epsilon_{Ave.}$ is the mean efficiency of the three individual efficiencies. The error on the average efficiency is the proportional difference between the measured and calculated values of the efficiency of all three MPPCs (see equation (4.3)).

where j is the flux, F is a scaling factor, C is the trigger count with the beam on, C_{off} is the background trigger rate (i.e. with the beam off), I_p is the proton current, S is the SEC count, S_{off} is the SEC count with the beam off and K is the SEC to current conversion factor. The scaling factor, F, was used to account for different conditions between runs. F was taken to be the ratio of the measurements made at the same location. I.e. the scaling factor for the 1D case was the ratio of the two measurements made at 0 cm. For the 1D measurement there was a difference between the first four and last five measurements due to damage to the MPPCs (one broke and another became detached from the scintillator). The 2D measurements were made at a larger distance than the 1D measurements as there was another experiment upstream that prevented closer positioning. The upstream experiment was used to make the muon lifetime measurement discussed below, it consisted of two scintillators on either side of stopping target. The stopping target was initially 20 mm of magnesium that was changed to 5 mm of copper for the final 3 runs.

4.2.4 Analysis

As can be seen from the figures 4.6 and 4.7 the beam spot is positioned slightly up and to the left of the centre of the pipe. Where there were multiple measurements at the same position there is good agreement which suggests that the measurements are correct.

The agreement between the measurements and the simulation is less good as shown by figure 4.8 where clearly the simulated rate and the measured rate are off by a significant amount (although both have the same shape). Comparing the simulation of the 2D case (figure 4.9) with the measurement (figure 4.7) we see that the simulation predicts a much more intense beam spot than what we see. This is likely due to the G4BL simulation having a much simpler model of the magnetic field at the position of the 2D measurement. The fluxes (with the exception of the peak at (0 cm, 20 cm)) are lower in the simulation than

Height	Trigger (Count	SEC Count		SEC Count		Factor	Flux (j)
(cm)	(C)	(C_{off})	(S)	(S_{off})	(F)	$(nA^{-1} s^{-1})$		
0	$2,\!151,\!736$		1,255	381	1.0 ± 0.0	$72,200 \pm 3,300$		
-5	$1,\!438,\!685$	30	1,286			$46,600 \pm 2,100$		
-15	446,302		1,212			$15,770 \pm 760$		
-15	502,596		1,208			$17,830 \pm 860$		
5	1,663,702		1,298	398		$91,300 \pm 4,000$		
5	1,420,836	83	1,142			$94,300 \pm 4,800$		
15	1,080,170	00	1,336		1.684 ± 0.034	$56,900 \pm 2,400$		
15	$1,\!185,\!051$		$1,\!371$			$60,200 \pm 2,500$		
0	1,307,015	34	1,298	404		$72,200 \pm 3,200$		

Table 4.6: Summary of the results from the measurement of the vertical charged particle flux. The double line indicates the division between runs with all four MPPCs working and runs with only three. The errors are statistical whilst the error on F is calculated as the errors on the un-scaled measurements at 0 cm added in quadrature. The error on the flux is calculated as the propagated errors of the other values. Counts were taken for 50 s. The SEC conversion factor, K, was 0.03408 nA. The factor was taken to be the ratio of two measurements made at 0 cm. The 0 cm measurements were used as this was the only pair of measurements that were made with the experiment in both conditions (i.e. all MPPCs functional and later, with one broken and another damaged).



Figure 4.6: Measurement of the 1D charged particle flux as given in table 4.6. Where two measurements were made at a single location an offset has been applied to the point to allow clearer reading of the plot, the measurements were actually at the same position. The errors on the vertical position are the width of the scintillator used.



Figure 4.7: 2D plot of the charged particle flux across the front of the MuSIC beam pipe. The final three measurements were rescaled using the ratio of the two measurements at (0, 20) cm. Both the dot colour and the z-axis indicate the flux.

Х	Υ	Trigger	Count	SEC Count	Factor	Flux (j)	Simulated (j)
(cm)	(cm)	(C)	(C_{off})	(S)	(F)	$(nA^{-1} s^{-1})$	$(nA^{-1} s^{-1})$
0	0	68,761	58	51	1.0 ± 0.0	$5,400 \pm 1,000$	$4,287 \pm 65$
-17	20	$117,\!947$	58	51	1.0 ± 0.0	$9,300 \pm 1,700$	$1,347 \pm 37$
-17	-16	39,528	48	68		$2,210 \pm 330$	492 ± 22
-17	0	$97,\!190$	78	73		$5,010 \pm 710$	$6,952 \pm 83$
17	0	$63,\!160$	91	76	10100	$3,110 \pm 430$	$1,174 \pm 34$
17	20	105,097	88	71		$5,590 \pm 810$	902 ± 30
17	-16	27,010	55	67	1.0 ± 0.0	$1,540 \pm 230$	975 ± 31
0	-16	43,634	97	69		$2,400 \pm 360$	$1,752 \pm 42$
0	20	$142,\!536$	73	64		$8,600 \pm 1,300$	$11,160 \pm 110$
0	0	$75,\!376$	60	66		$4,360 \pm 670$	$4,287 \pm 65$
0	25	45,950	35	46		$6,800 \pm 2,100$	$1,216 \pm 35$
-17	25	$69,\!582$	119	47	1.65 ± 0.37	$10,000 \pm 3,000$	914 ± 30
0	20	80,088	238	60		$8,600 \pm 2,400$	$11,157 \pm 106$

Table 4.7: Table of the results from the 2D measurement. The values above the double line were made with 20 mm Mg stopping target upstream whilst those below had 5 mm Cu. The scaling factor F for these sets of runs was made at (0, 20) cm. The trigger count was measured over 20 s apart from the first two measurements C_{off} which were made over 11 s (these are the two measurements above the first horizontal line). The SEC measurements were made over 100 s. Counts were made using a CAMAC controlled scaler and an enable gate triggered via interrupt register. The column for the SEC bias, S_{off} , is excluded as it was constant at 9 and the SEC conversion factor, K, was 1.514 nA. As with the 1D measurement the errors are statistical. The last three values also include an error due to the calculation of the scaling factor.

measured but not to the same degree as in the 1D case.

As figure 4.9 shows both measurements agree on the general position of the beam spot. The measurements are not directly comparable due to the differences in both scintillator volume and position with respect to the end of the beam-pipe. Obviously the flux for the 2D measurement is generally lower due to smaller scintillator and greater distance from the beam-pipe.

The detector efficiency measurement for the 2D apparatus allows a more accurate calculation of the flux for this case whilst also providing a bench-mark efficiency for future uses of the scintillator/MPPC configuration (where it wasn't always possible to make an analogous measurement).

The biggest problem, and one which was a cause of many issues, was the fragility of the MPPCs. Both the MPPC itself and the optical cement that bonded them to the scintillator were liable to break and this caused many problems that could not be easily fixed during the time available.

A significant issue with the 1D measurement is the lack of a confirmation measurement that allows verification of the scaling factor applied to the later results. Even if the re-scaled points are ignored then they confirm that the beam spot is in the upper half of the pipe, a measurement later confirmed by the 2D measurement.

Comparing the measured and simulated rates yields some interesting features. There is broad agreement on the general shape of the distribution but, using normalised plots, it appears that the measured charged particle rate for the 1D measurement is consistently higher than the simulated value by a factor of between 2 and 10. No obvious explanation of this presents itself as the measured value, should be, a lower bound whilst the simulated value should be closer to the 'true' value due to lack of detector effects and similar. Whilst this suggests that further development of the G4BL simulation is required it is encouraging that MuSIC is performing above expectations, it also indicates that with improved experiments these sorts of measurements should provide useful data for tuning these models.



Figure 4.8: Comparison of simulated and measured total charged particle flux.



Figure 4.9: Simulated results for the 2D charged particle rate normalised to a proton current of 1 nA. Colour is equal to the rate as given by the z-axis. Simulated using 9×10^8 in G4BL, the rate was taken to be those particles seen within the area of the scintillator used.

4.3 Muon Lifetime

Measurement of the muon lifetime was done to confirm the presence of muons in the beam. As has been discussed muons have a relatively long lifetime (2.197 μ s) that is reasonably distinct from any others; this acts as a useful method of identification. To make the measurement two scintillators are placed on either side of a stopping target. The entire detector is placed perpendicular to the muon's direction of travel. The muons lose energy as they traverse the stopping target and some number of them will decay emitting an electron. By measuring the time differences between a detection at the upstream scintillator (the muon) and a detection at the downstream scintillator (the decay electron) we can produce a spectrum of the characteristic exponential decay and measure the muon lifetime. The curve produced is the sum of the many exponentials, each corresponding to the material in which negative muons were captured or the free decay of positive muons. As negative muons are less common at MuSIC than positive ones and because in many materials negative muons have similar lifetimes only two exponentials are typically used: one for the positive muons and another for negative muons captured by the stopping target which has a measurably different lifetime.

In total three measurements of the muon lifetime were taken: one concurrently with the 2D charged particle flux measurement, a high-statistics run and then measurements that form the basis of the muon momentum spectrum measurement. The first two of these measurements will be discussed here and the final one will be discussed later.

4.3.1 Experimental Set Up

The basic experimental set up for muon lifetime measurement has already been discussed: two scintillators sandwiching a stopping target. Both measurements were made using the same scintillator and stopping target: two $380 \times 50 \times 3.5$ mm³ scintillators and a $370 \times 80 \times$ 6 mm³ pure copper target. An MPPC was mounted at either end of both scintillators (4 MPPCs in total) to provide readout.

The DAQ used was a simple extension of the previous versions. This time, rather than measuring the time differences between the different individual MPPCs the TDC was used to measure the time between a hit on the upstream scintillator and then a hit on the downstream. A 'hit' was considered to be a co-incident signal on both MPPCs of that scintillator. As a measure to reduce spurious signals due to other charged particles, e.g. electrons, a veto window of 50 ns was used. The veto window was a period following the upstream hit in which no hits were allowed in the downstream scintillator. The veto-window is used to reject particles that didn't decay between the scintillators or which decay but aren't muons (e.g. pions, which have a 26 ns lifetime).

4.3.2 Results

Figures 4.10 and 4.11 show the TDC measurements that were made to record the muon lifetime. The TDC bin values are converted to ns using:

$$t = K(T - T_0) (4.6)$$

where t is the real time, in ns, K is the calibration constant (0.025, measured by manufacturer), T is the value recorded in the MH-TDC and T_0 is the time when the start signal for the TDC was sent, as recorded by the MH-TDC on a dedicated channel.

The data is then fitted with either one or two exponential decay functions and a flat background. Whether a single or double exponential is used depends heavily on the amount of data that was taken. A single function was used to fit the free decay of muons (neglecting the captured negative muons) whereas a second exponential was used to model the decay of captured negative muons inside the stopping target. Fitting the copper-decay exponential requires much more data as a significant proportion of the events from this are ignored because of the veto-window. In the first 50 ns ~26 % of all copper decays will have already occurred in comparison to only ~2% of free muon decays, this is made worse as the number of positive muons is significantly higher than the number of negative muons making the captured muon signal much lower.

4.3.3 Analysis

As can be seen in figures 4.10 and 4.11 the exponential decays are good fits to the data with reasonable chi-squared values (192.6/114 and 516/231 respectively). The second muon lifetime is in agreement with the canonical value of $(2.1969811 \pm 0.0000022) \ \mu\text{s}$ [2] whilst the first measurement for the first run is low, $(2,016\pm89)$ ns. This low measurement is likely due to the fit not accounting for any of material in the detector or the stopping target itself. The plastic detector is mostly carbon in which muons have a lifetime of $(2,026.3\pm1.5)$ ns [3] whilst the copper lifetime is (163.5 ± 1.0) ns [3]. Both measurements have a large background but the core fact remains that large numbers of muons were detected and could be identified by their lifetime.

The second measurement produces a slightly high value for the lifetime of negative muons in copper, (194.8±17.1) ns but this discrepancy is small ($< 2\sigma$), there are a number of possible explanations for this. Firstly there is an obvious problem with the data gathering as seen in the discontinuity at 16 μ s, whilst this is excluded from the fit it may be that this is having some more subtle effects that aren't accounted for. A further possible cause of error is the noise seen in the residuals which seems to have a periodic component to it, this is further explored in section 4.4.



Figure 4.10: Results of the first measurement of the muon lifetime with residuals from the fit in the lower plot. The data was binned in 166 ns bins to remove the effects of noise on the data, only entries with an ADC value of 1,100 (which excluded the pedestal) were used. The fit was applied from 600 ns to 20 μ s. The residuals show that fit gets worse at earlier times although this is to be expected with an exponential. Fitting with two exponentials should have improved the fit at earlier times but this was found not to be the case.



Figure 4.11: Results of the second measurement of the muon lifetime with residuals of the fit in the lower plot. The data was split between bins with a width of 50 ns and the data was fitted for values between 200 ns and 16 μ s. The upper limit of the fit was lowered from 20 μ s to 16 μ s to remove an apparent discontinuity in the data. The discontinuity sits at 16 μ s, and is especially clear in the residual plot. The cause of the discontinuity is unknown.

Component	Matorial	Din	nensio	Number	
Component	Material	Х	Υ	Ζ	number
Degrader	Al	400	400	0.5, 1, 5	1
Upstream scintillator	EJ-212	380	30	0.5	8
Stopping Target	Cu	370	310	0.5	1
Downstream scintillator	EJ-212	380	50	3.5	5

Table 4.8: Dimensions of the various detector components. The 'Z' is the beam direction, 'Y' is vertical and 'X' is horizontal. The layout of the detector can be seen in figure 4.12.

4.4 Momentum Spectrum

Measurement of the muon momentum spectrum was carried out using an extension of the muon lifetime measurement technique. The number of scintillators was increased to improve coverage of the beam-pipe. To make the momentum measurement, a range of degraders were used to select different (initial) ranges of muon momenta that would then stop. To determine the relation between degrader and stopped muon momentum the simulation was used.

4.4.1 Experimental Set Up

Detector

The detector consisted of an aluminium degrader whose thickness could be varied to select different momentum ranges, a thin (0.5 mm) upstream counter and a thicker (3.5 mm) downstream counter on either side of a 0.5 mm copper stopping target, table 4.8 gives details of the dimensions used. A schematic of the detector can be seen in figure 4.12. Based on simulation (see chapter 3) we predicted the mean initial momentum of muons that decay for different degrader thicknesses, the initial muon momentum distribution is given in figure 4.13 and the momentum distributions for the muons that stop in the detector can be seen in figure 4.14, the mean momenta are also tabulated in table 4.9. The origin of the high-end tail seen figure 4.14 is uncertain.

The counters used were mylar wrapped scintillators with readout performed by MPPCs mounted on either end of a wavelength-shifting fibre bounded along the long axis of the scintillator using optical cement. The signals from each pair of MPPCs are combined and amplified before being passed to the data acquisition system (DAQ) for processing.

Two important considerations were made with the choice of scintillators: the upstream scintillator had to minimise perturbation of the beam whilst providing a suitable trigger from muons and the downstream scintillator had to be effective at detecting electrons. To this end a thin (0.5 mm) upstream scintillator and thicker (3.5 mm) downstream scintillator were chosen. One downside of this set up is that in using a thin upstream scintillator the efficiency for detecting minimally ionising particles is reduced (the simulated light yield for a



Figure 4.12: Experimental set up of the detector for MuSIC 5 (widths not to scale), see table 4.8 for actual sizes. The upstream scintillators are outlined with blue dashes whilst the downstream ones are red dots.

Degrader	Momentum (MeV/c)			
(mm)	Mean	RMS		
0.0	45.27 ± 0.24	$20.99 {\pm} 0.17$		
0.5	$49.62 {\pm} 0.23$	$18.89 {\pm} 0.16$		
1.0	$52.60 {\pm} 0.22$	$17.71 {\pm} 0.16$		
5.0	66.09 ± 0.22	$14.80 {\pm} 0.16$		

Table 4.9: Simulated mean and Root Mean Squared (RMS) momenta of muons, that subsequently stop, at the end of the beam pipe. The degrader is the thickness of the aluminium, the width and height are both 400 mm. The Geant4 simulation was used; in total 95,700 initial muons were simulated, for further details see chapter 3.



Figure 4.13: Total muon momenta at the end of the beam pipe. The number of initial muons simulated was 95,700. The plot has been normalised to the number of muons per μA of proton beam.

MIP is only ~ 3 photons per MPPC) but as muons that subsequently stop are desired (which deposit 3–4 times more energy) this has a negligible effect.

Data Acquisition

The DAQ system used was a simple evolution of that used in the previous experiments: a discriminator unit with paired logic block formed a trigger on a suitably strong upstream signal with no downstream signal within a 50 ns veto window. The triggering signal was recorded using QDC (although it is not analysed here) and a MH-TDC was used to record hit times at all scintillators $\pm 20 \ \mu$ s of the trigger. To increase the accuracy of the MH-TDC the trigger time is recorded on a separate channel to the individual MPPC signals called 'TDC0' which can then be removed later.

For a signal to be recorded or used as a trigger the discriminator was set to exclude events with fewer than eight photons in the upstream scintillator and fewer than ten in the downstream. The reason for these high values is that each analogue signal, as passed to the discriminator, is the combination of the signals from both the MPPCs on a scintillator. The thinner upstream scintillator has a lower threshold as it's aimed at selecting slow muons that will stop, unfortunately this means it cannot detect mips but these are unlikely to stop. Conversely the downstream scintillator has been set up to select electrons which will deposit more energy and the higher threshold will help reduce background.

A scaler was used to record system diagnostics through regular polling. The values



Figure 4.14: Momenta at the end of the beam pipe of muons that subsequently stopped. A stopped muon is considered to be one which is seen in the upstream scintillator and has a daughter electron seen in the downstream scintillator i.e. the muon has stopped between the scintillators. This shows the momenta that the different degraders 'select'. The number of initial muons simulated was 95,700. The plot has been normalised to the number of muons per μ A of proton beam.

Run ID	Time (sec)	Current (pA)	Degrader (mm)
448	9,221	15	0.0
451	1,001	15	0.5
452	4,924	13	0.5
455	6,307	13	1.0
458	5,144	14	5.0
459	2,452	12	5.0

Table 4.10: Summary of the runs selected for this analysis. The degrader used aluminium and the stopping target 0.5 mm copper.

recorded were:

- 1. SEC output.
- 2. Triggers (i.e. U AND \overline{D} AND \overline{BUSY}^*).
- 3. Potential triggers (i.e. U AND \overline{D}).
- 4. A clock (for accurate run measurement).

Obviously the clock and SEC were used for normalisation whilst the triggers and potential triggers were used for calculation of dead time (see section 4.4.3).

4.4.2 Data Processing

There were three stages in preparing the data for analysis: reading from VME, calibration and binning. At the end of these processes the data was analysed and the final measurements made. Six data sets were chosen for this analysis, the run conditions are listed in table 4.10.

The data was read from VME using the MIDAS program ('a general purpose data acquisition system for small and medium scale experiments' [30]), this dealt with the low level interface to VME and allowed basic plotting of the information.

Once the data was saved, it was calibrated to convert it to real world measurements. The calibration formula for the MH-TDC was:

$$t' = 0.024414(t - \text{TDC0}) \tag{4.7}$$

where t' is the calibrated time, t is the MH-TDC value and TDC0 is the time at which the trigger occurred. The calibration co-efficient of 0.024414 was determined using a known

^{*}Where 'U' is a suitable upstream signal, 'D' is a signal in the downstream scintillator and 'BUSY' indicates that the system is busy. Barred values indicate negation, i.e. \overline{D} indicates no downstream signal and \overline{BUSY} indicates that the system isn't busy.

Run	Degrader	Degrader Length Trigger		$s(\times 10^3)$	Live time	Dead time
	(mm)	(s)	Good	Potential	(%)	(%)
448	0.0	9221	9653.0 ± 3.1	15678.8 ± 4.0	61.57 ± 0.03	38.43 ± 0.03
451	0.5	1001	767.32 ± 0.88	1070.4 ± 1.0	71.69 ± 0.11	28.31 ± 0.11
452	0.5	4944	3459.9 ± 1.9	4667.8 ± 2.2	74.12 ± 0.05	25.88 ± 0.05
455	1.0	6307	4090.2 ± 2.0	5372.4 ± 2.3	76.13 ± 0.05	23.87 ± 0.05
458	5.0	5144	2049.1 ± 1.4	2455.0 ± 1.6	83.47 ± 0.08	16.53 ± 0.08
459	5.0	2452	915.31 ± 0.96	1077.5 ± 1.0	84.95 ± 0.12	15.05 ± 0.12

Table 4.11: Dead and live times for each run along with the number of potential and good triggers. Run time is included to show the origin of the different counts.

clock as input to the MH-TDC and then detecting a set number of ticks that could then be read off as MH-TDC values.

The final stage of data preparation was to bin the data, this was done to make later analysis quicker. Each channel had its MH-TDC times combined in 1 ns bins. This left suitable accuracy whilst vastly reducing the data set's size.

4.4.3 Secondary calculations

Using the scaler information two checks could be made: the dead time and the Periodic Trigger Rate (PTR). The dead time of the system is a measure of how many events are lost because the system is busy. The PTR is a check that the rate of triggering events is reasonably constant, deviations from a constant rate would indicate possible problems with the data.

Dead time

As was stated in section 4.4.1 the scaler was used to record diagnostics on the DAQ. To calculate the dead time, the trigger count and the potential trigger count are needed which can be used thus:

$$\text{Live time} = \frac{\text{Good Triggers}}{\text{Potential Triggers}}$$
(4.8)

Dead time
$$= 1 - \text{Live Time}$$
 (4.9)

where potential triggers are those with a signal in the upstream scintillator and no corresponding signal downstream within the 50 ns veto window. A good trigger is a potential trigger that occurs without the system being busy. The results of this calculation are given in table 4.11.

Trigger Rate

To ensure that there were no problems during data taking the trigger rate was checked. The trigger rate gives a measure of how consistent the detector was as well as highlighting any possible problems in the run. Each run was split into 100 equal sections and the trigger rate determined using the values recorded by the scaler. The results are plotted in figure 4.15 and fitted with a flat function. As can be seen as the degrader's thickness increases the trigger rate goes down. There are also some peaks in the first run (448), it is unclear what caused this. The SEC data (also plotted) does not suggest that the fluctuations in the rate are due to changes in the beam.



Figure 4.15: Trigger rate of each run (black) with the raw SEC count (blue). 100 samples were taken from the scaler data. Potential triggers were used to exclude any rate limiting due to dead time. The central 90 % of the data has been fitted with an order 0 polynomial with the parameters given. The SEC values are the differences of the count at the start of the time window and the end, this was done because the SEC value was not reset between measurements within a run. The SEC data suggests that the fluctuations in the trigger rate are not due to changes in the beam.

4.4.4 Results

As has been discussed there are six data sets that were used for this analysis, the details of them are given in table 4.10. To do this analysis only stop times from the downstream scintillators were used. The upstream start (T0) times were then subtracted from these to form the muon decay spectra. Each run was analysed as discussed in section 4.4.2. The binned data was then fitted using:

$$N_f \exp\left(\frac{-t}{\tau_f}\right) + N_c \exp\left(\frac{-t}{\tau_c}\right) + N_s \sin\left(2\pi \frac{t-\phi}{T}\right) + N_b \tag{4.10}$$

t is the time in ns. The four terms each have their own subscript based on what they describe: free muon decay (f); decay of negative muons captured by the stopping target (c); a sinusoidal background (s) and a flat background (b). The two muon decay terms are parameterised with a scale factor (N) and a lifetime (τ) . The sinusoidal term has a period (T) and phase, relative to the trigger time, (ϕ) .

Originally only a flat background was used but this was found to be insufficient to describe the data and the sinusoidal component was introduced. The sinusoidal effect is thought to be minor bunching of the protons due to the AVF acceleration system which has a stated frequency of 6–18 MHz [12] (or a period of 167–56 ns respectively). The flat background is made by travelling particles that don't form a trigger but are recorded by the MH-TDC.

In order to have good statistics whilst maintaining sensitivity to the relatively fast copper component of the decays a bin width of 16 ns was used. A large bin width (e.g. 100 ns) could have been used to smooth the sinusoidal noise but this would result in very few bins in which the copper component was significant. Whilst the ideal solution would be to use binning such that the sinusoidal component was cancelled out this was not feasible due to limitations of the fitting program. In the end a bin width was chosen that resulted in a small χ^2 whilst remaining sensitive to copper.

Fitting all the parameters could not easily be done freely due to the range of values to which it was being fitted; some features (e.g. the periodic noise) had a scale of $\mathcal{O}(60)$ ns compared to the free decay $\mathcal{O}(2,000)$ ns and the entire fit range $\mathcal{O}(20,000)$ ns. To mitigate the problem of scale the known values, τ_f and τ_c , were fixed at their given values of 2, 196.9811 ± 0.0022 ns [2] and 163.5 ± 1 ns [3] respectively. The period, T, was constrained to 60 ± 5 ns which was determined to be a good value from fitting the noise in isolation (see figure 4.16).

Once the data was fitted the f and c terms were integrated individually to calculate the number of free and copper decays the muons underwent. The fitted histograms are shown in figures 4.17 to 4.22 and the results are summarised in table 4.12. The fitting shows that the bottom channels in the downstream scintillator (D4 and D5, counting from the top) have the worst fit results, often chi-squared values of > 10,000/1,239 (ndf), this is compared to values of $\mathcal{O}(2,000)/1,239$ for the other channels. Of the other channels D1 and D3 are generally best but D2 is consistently good enough to be useable (unlike channels D4 and D5 which are removed from later analysis). All channels exhibit a slight exponential component to the data in the negative region, this is most clear in channels D4 and D5, the source of



Figure 4.16: Zoomed region (8,000 to 8,600 ns) of run 448, channel D1, showing the sinusoidal component of the background, the background has been fitted using $N_b + N_s \sin(2\pi \frac{t-\phi}{T})$.

this is unclear but suggests that there is a problem in that may be introducing bias to the result.

Run	Ch.	N_b	N_s	ϕ	N_c	N_f
	D1	576.85 ± 0.80	121.10 ± 1.15	$2.19~\pm~0.21$	549 ± 31	1377.2 ± 5.7
	D2	914.39 ± 1.0	197.28 ± 1.44	10.33 ± 0.16	1524 ± 38	1937.2 ± 6.9
448	D3	164.22 ± 0.43	40.68 ± 0.63	18.22 ± 0.36	592 ± 21	620.3 ± 3.4
	D4	32.07 ± 0.19	$6.89 \hspace{0.2cm} \pm \hspace{0.2cm} 0.28$	57.02 ± 0.96	193 ± 11	172.8 ± 1.7
	D5	27.79 ± 0.18	$6.37 \hspace{.1in} \pm \hspace{.1in} 0.26$	$4.31~\pm~0.94$	24.8 ± 7.8	91.1 ± 1.4
	D1	39.55 ± 0.21	2.38 ± 0.31	25.3 ± 2.4	29.4 ± 8.1	101.4 ± 1.5
	D2	62.27 ± 0.26	17.74 ± 0.38	13.28 ± 0.50	153 ± 11	178.2 ± 1.9
451	D3	11.52 ± 0.12	3.60 ± 0.17	20.4 ± 1.2	85.6 ± 6.5	57.76 ± 0.99
	D4	$2.18 \hspace{0.2cm} \pm \hspace{0.2cm} 0.05$	0.61 ± 0.08	6.5 ± 3.0	25.9 ± 3.5	16.59 ± 0.49
	D5	1.97 ± 0.05	0.65 ± 0.07	$8.7 \hspace{0.2cm} \pm \hspace{0.2cm} 2.6$	7.3 ± 2.6	$8.78~\pm~0.40$
	D1	161.25 ± 0.43	43.56 ± 0.61	$4.08~\pm~0.31$	208 ± 17	435.5 ± 3.1
	D2	252.14 ± 0.54	71.73 ± 0.76	12.39 ± 0.25	706 ± 23	789.9 ± 4.0
452	D3	46.29 ± 0.23	15.70 ± 0.34	19.82 ± 0.55	335 ± 13	266.7 ± 2.1
	D4	8.95 ± 0.11	$2.49 \hspace{0.2cm} \pm \hspace{0.2cm} 0.16$	$0.00~\pm~1.64$	128.2 ± 7.5	76.6 ± 1.1
	D5	8.01 ± 0.097	2.37 ± 0.14	$7.54~\pm~1.43$	25.9 ± 5.2	38.71 ± 0.82
	D1	187.67 ± 0.46	46.99 ± 0.66	4.38 ± 0.32	265 ± 19	539.1 ± 3.4
	D2	293.99 ± 0.58	73.76 ± 0.83	12.41 ± 0.26	896 ± 26	967.9 ± 4.4
455	D3	54.02 ± 0.25	17.18 ± 0.37	20.27 ± 0.53	438 ± 15	326.4 ± 2.3
	D4	10.37 ± 0.11	2.86 ± 0.16	$0.00~\pm~0.36$	195.1 ± 8.5	87.4 ± 1.1
	D5	$9.30 \hspace{0.2cm} \pm \hspace{0.2cm} 0.10$	2.63 ± 0.15	4.8 ± 1.3	41.9 ± 5.8	47.00 ± 0.90
	D1	73.73 ± 0.29	21.53 ± 0.43	0.000 ± 0.037	219 ± 15	350.2 ± 2.4
	D2	114.37 ± 0.37	34.08 ± 0.54	5.32 ± 0.38	799 ± 20	562.4 ± 3.1
458	D3	21.47 ± 0.16	8.40 ± 0.24	9.89 ± 0.71	236 ± 11	178.1 ± 1.6
	D4	3.805 ± 0.070	1.25 ± 0.10	$0.00 \ \pm \ 0.19$	72.6 ± 5.8	49.23 ± 0.80
	D5	$3.728~\pm~0.067$	1.436 ± 0.098	$0.00 \ \pm \ 0.76$	19.1 ± 4.2	28.11 ± 0.65
	D1	30.67 ± 0.19	8.86 ± 0.28	56.68 ± 0.72	115.3 ± 9.8	155.0 ± 1.6
	D2	48.07 ± 0.24	14.91 ± 0.35	$4.36~\pm~0.56$	369 ± 13	254.0 ± 2.1
459	D3	$8.87 \ \pm \ 0.11$	3.65 ± 0.16	11.3 ± 1.0	110.1 ± 7.3	81.0 ± 1.1
	D4	$1.631~\pm~0.046$	$0.579~\pm~0.068$	48.2 ± 3.1	39.2 ± 4.9	22.01 ± 0.54
	D5	1.496 ± 0.043	$0.615~\pm~0.063$	0.0 ± 1.0	8.0 ± 2.8	12.73 ± 0.43

Table 4.12: Summary of the fitted values for equation (4.10). These are the values used for the calculation of the integrals and, ultimately, the muon rates. The two lifetime parameters, τ_f and τ_c , were fixed at the values 2,196.9811±0.0022 ns [2] and 163.5±1 ns [3] respectively, T was fixed at 60 ns.
























Dun	Δz	I_p	D	Γ	Copper Decays		Free Decays	
nun	(mm)	(nA)	(s)		Count	Rate (nA^{-1})	Count	Rate (nA^{-1})
448	0.0	0.01534	9221	0.6157	19895 ± 400	228.5 ± 4.6	528076 ± 1283	6065 ± 15
451	0.5	0.01546	1001	0.7169	1999 ± 114	180.2 ± 10.3	45269 ± 355	4080 ± 33
452	0.5	0.01313	4944	0.7412	9319 ± 240	$193.6 {\pm} 5.0$	200246 ± 732	4161 ± 15
455	1.0	0.01332	6307	0.7613	11937 ± 265	$186.6 {\pm} 4.1$	246059 ± 803	3847 ± 13
458	5.0	0.01363	5144	0.8347	$9365 {\pm} 202$	160.1 ± 3.5	$146384{\pm}570$	2502 ± 10
459	5.0	0.01238	2452	0.8495	4440 ± 136	172.1 ± 5.3	65781 ± 378	$2550{\pm}15$

Table 4.13: Decay counts and rates, summed over all downstream scintillators. Δz is the thickness of the aluminium degrader.

Once calculated via integration the number of muons can be summed for all the channels in a run. The summed value can then be normalised between runs to the proton beam current using:

$$R = \frac{N_{\mu}}{I_p DL} \tag{4.11}$$

where R is the normalised rate of muon decays (in nA⁻¹), N_{μ} is the number of muon decays (either free or from muonic copper), I_p is the proton current in nA, D is the duration of the run in seconds and L is the live time (as given in table 4.11). The rates and number of muon decays, summed over all channels, are given in table 4.13. The results back up the earlier measurements made on the beam flux (section 4.2) as the larger values of N_f , N_c and N_b are seen in channels D1, D2 and D3, i.e. the beam position is in the upper portion of the detector.

Figures 4.23 and 4.24 show the rates of decays from free muons and muonic copper respectively. These plots are the 'pure' rates based solely on the measured values. Between the pairs of runs (e.g. runs 451/452 and 458/459) there is reasonable agreement and the general shape suggests that (unsurprisingly) as larger degraders are used there are fewer decays of stopped muons.

As can be seen in the fit histograms, the χ^2 values for the fitting of channels D4 and D5 are generally poor and because of this they were removed from the final, summed, rate calculation. The main cause of the poor χ^2 value is likely the low statistics compounded by the fact that, as the counters share a common trigger those scintillators that are further away from the beam spot will record a greater fraction of noise.



Figure 4.23: Integrated number of free muon decays summed over all channels.



Figure 4.24: Integrated number of muon decays in copper summed over all channels.

Once the muon rates (figure 4.23) have been calculated they can be corrected to account for other effects. This is shown in figure 4.25 where additional terms have been incorporated to account for the photon acceptance, the MPPC efficiency and the systematic errors due to the fitting procedure.

The photon acceptance was calculated using the optical simulation and a simple muon gun particle source. The photon acceptance was defined to be the ratio of muons that decay in a detectable manner ('potential decays') to those that produce photons that are detected ('detected decays'):

$$A = \frac{\text{Detected decays}}{\text{Potential decays}} \tag{4.12}$$

a detectable muon was considered to be one that traverses the upstream detector and has a corresponding daughter electron traverse the downstream detector. The requirement for detected photons are that the number of photons seen exceed the threshold (eight and ten photons for the up and downstream MPPCs respectively, see section 4.4.1). Putting the values from the simulation into equation (4.12):

$$A = \frac{8362}{10500} \tag{4.13}$$

$$= (79.6 \pm 1.2)\% \tag{4.14}$$

obviously this is a value based purely on simulation of a complex system so introduces potentially large errors into the final results which cannot be easily quantified.

The efficiency of any single detector was taken to be the square of the average efficiency of a MPPC, as measured in section 4.2.2:

$$\epsilon_{\rm MPPC}^2 = (18.6 \pm 11.6)\% \tag{4.15}$$

whilst this value was for a different detector configuration with different thresholds it is hoped that it's a reasonable first-order approximation to the actual efficiency of the MPPCs although, again as with A, it is hard to quantify the total error introduced. This is especially true for the efficiency as it was measured under different circumstance, with different discriminator thresholds on a scintillator with different mounting system and geometry.

The systematic errors were largely assumed to be due to the inaccuracies in the method of fitting and integrating the data. Sources of this have already been discussed but to quantify it the key values were varied to show the effects of changing the parameters of the fit itself, rather than the parameters that were fitted (e.g. τ_f). The key parameters were determined to be the point from which the fit was applied (the 'lower bound') and the width of the bins used to fit the data. The bin width was varied to 8 and 32 ns whilst the lower bound could only be increased so was changed to 75 ns (there is no data before 50 ns). Using these values all the data was re-fit and the difference from measured value taken. The results are given in table 4.14. The maximal fractional differences from the standard measured using 16 ns bins and a 50 ns lower bound were used as the systematic error. The systematic errors were 55.85% and 1.14% for copper and free decay counts respectively. As is clear, the systematic

Bun	Bin	Lower	Copper				Free	
num	bound (ns)	width (ns)	Decays	Difference	Fraction	Decays	Difference	Fraction
448			228.66	0.17	0.0007	6064.08	0.69	0.0001
451			179.62	0.54	0.0030	4080.08	0.27	0.0001
452	8	50	193.81	0.17	0.0009	4159.70	1.28	0.0003
455	0	50	186.84	0.22	0.0012	3845.32	1.36	0.0004
458			160.44	0.35	0.0022	2502.20	0.11	0.0000
459			172.77	0.64	0.0037	2549.44	0.79	0.0003
448			184.25	44.24	0.1936	6097.10	32.33	0.0053
451		50	139.95	40.21	0.2232	4103.83	24.03	0.0059
452	30		148.37	45.27	0.2338	4191.79	30.81	0.0074
455	02	50	141.78	44.84	0.2403	3874.43	27.75	0.0072
458				110.46	49.63	0.3100	2530.72	28.42
459			126.63	45.50	0.2643	2577.57	27.34	0.0107
448			132.05	96.44	0.4221	6047.05	17.72	0.0029
451			93.76	86.39	0.4796	4078.37	1.43	0.0004
452	16	75	105.99	87.66	0.4527	4157.79	3.20	0.0008
455		10	99.67	86.95	0.4659	3846.58	0.10	0.0000
458			70.68	89.41	0.5585	2517.28	14.97	0.0060
459			84.21	87.92	0.5108	2561.92	11.69	0.0046

Table 4.14: Results of the systematics fits. The maximal values for copper and free decays were taken as the systematic errors for each value (55.85% and 1.14% respectively).

error on the copper fit is much higher but this is what we would expect. The systematic error on the copper count is included in the corrected rate but not for the free rate as the error on the efficiency is the dominant term then.

The corrected muon decay rate for material (or freely), R_a , is then:

$$R_a = \frac{N_{\mu}}{I_p L D A \epsilon_{\text{MPPC}}^2} \tag{4.16}$$

where the symbols I_p , L and D have the same meaning as in equation (4.11). The symbols: ϵ_{MPPC} and A are the average efficiency of an MPPC and the photon acceptance respectively. Table 4.15 shows the results of this calculation and figure 4.25 compares it to simulation. As can be seen there is good agreement between the corrected rate of muon decay and what is expected from simulation.

The same protocol can be applied to the muonic copper decays as well as adding the systematic error to get an corrected rate that is shown in figure 4.26 and table 4.16. Unfortunately due to the limitations of the simulation not enough negative muons are produced for accurate comparison and so only the measured values are shown here.

Momentum	Buns	Corrected Rate	Error	(nA^{-1})
(MeV/c)	itano	(nA^{-1})	(incl. eff.)	(excl. eff.)
45 ± 21	448	41,000	$25,\!000$	610
50 ± 19	451, 452	28,000	$17,\!000$	430
52 ± 18	455	26,000	$16,\!000$	390
66 ± 15	458, 459	17,100	$11,\!000$	260

Table 4.15: Corrected rates for freely decaying muons. The error column shows the error without the contribution from the MPPC efficiency and with it, as is clear this is the dominant source. The momentum values are the mean and RMS as determined by simulation (see section 4.4.1 and table 4.9).

Momentum	Dung	Corrected Rate	Error (nA^{-1})			
$({ m MeV/c})$	TUIIS	(nA^{-1})	Total	(excl. sys.)	(excl. eff.)	
45 ± 21	448	1,540	1,700	960	38	
50 ± 19	451, 452	1,260	1,400	790	43	
52 ± 18	455	1,260	1,400	780	34	
66 ± 15	458, 459	1,120	1,200	700	27	

Table 4.16: Corrected rates of muonic copper decay. Errors are split into three: total error (including systematics and MPPC efficiency); error excluding systematics but including MPPC efficiency; and error excluding both systematic and efficiency contributions (i.e. statistical). The momentum values are the mean and RMS as determined by simulation (see section 4.4.1 and table 4.9).



Figure 4.25: Corrected free muon decay rate and simulated free muon decay rate. The red lines are the simulated results, the blue lines the measured results including the MPPC efficiency error and the black lines are the measured values without the MPPC efficiency error.



Figure 4.26: Corrected decay rate of muonic copper. The simulated rates are not shown as the statistics for muonic copper decays were too low to extract reasonable values from. The blue line indicates the error due to systematic error, the red line is the MPPC efficiency error and the black line the statistical error.

4.4.5 Analysis

As can be seen from figure 4.25 there is a good agreement between simulation and experiment. The largest source of errors for this measurement was the lack of accurate information on the detector efficiency and in further measurements this should be corrected. Despite the large errors the data shows a clear reduction in the number of stopped muons as the degrader thickness increases and that this reduction is stable between the repeat measurements.

Throughout the entire experiment there was a gradual decline in the total trigger rate (even between matched configurations) as is seen in figure 4.15. The cause of the reduction in trigger rate is unclear, its effect on the results is corrected for through the dead time calculation.

A point of interest in the data is the difference between the first and subsequent runs. In the first run the simulation predicts a lower rate whilst in all other runs the simulation predicts a higher rate. This suggests that there may be some other effect that needs to be accounted for in either the simulation or the analysis. Another area that needs more accurate modelling is the inclusion of beam effects in order to better understand the sinusoidal noise seen in all data and account for it.

One area of tension in this analysis is the overestimation by the simulation of the muon rate, this is especially puzzling in light of the total charged particle rate where the rate was underestimated by up to a factor of two by the simulation. There is no clear source for this tension but it may well be due to particles other than muons. This seems an attractive argument given the higher than simulated rate when no degrader was used obviously to resolve this more refined measurements must be taken to verify or reject the simulation.

Based on the simulation the stopped muons account for a small portion of the beam, no more than 8 % of the total muon flux. Using the simulation the fraction of the beam which is stopped for each degrader can be estimated. Using these ratios we can calculate the total muon flux for the RCNP's maximum 1 μ A beam:

$$J = 1,000 \times \frac{R}{C} \tag{4.17}$$

where J is the total muon flux per second, R is the measured muon rate per nA and C is the fraction of the beam that was simulated to stop in that configuration. The factor of 1,000 converts from nA to μ A and hence to s⁻¹ for a 1 μ A proton beam. The calculated values for the total muon flux are shown in table 4.17 along with the per Watt efficiency, calculated using the RCNP's 400 W beam at 1 μ A.

Two things are clear from table 4.17: at full power MuSIC should be capable of producing the design goal of > 10⁸ muons s⁻¹ and it should be one of the most intense muon beams in the world. The average intensity, $(4.2 \pm 1.1) \times 10^8$ muons s⁻¹, is comparable PSI's most intense muon beam $(4.8 \times 10^8 \text{ muons s}^{-1} [7])$ but achieved with a proton beam significantly less powerful. This last point is best highlighted by comparing the number of muons produced per Watt of proton beam: PSI produces 292 muons W⁻¹ whilst MuSIC is predicted to produce $(10.4 \pm 2.7) \times 10^5$ muons W⁻¹. The largest source of errors in this analysis is the estimation of the MPPC efficiency. That being said if the flux calculations are carried out using the un-corrected rates then fluxes of $\mathcal{O}(10^7)$ muons s⁻¹ are still obtained.

Dun	Component	Rate	Total flux	Efficiency
Run	(%)	$(\times 10^3 \text{ nA}^{-1})$	$(\times 10^8 \text{ s}^{-1})$	$(\times 10^5 \text{ muons W}^{-1})$
448	$8.158 {\pm} 0.096$	43 ± 27	5.2 ± 3.2	13.0 ± 8.1
451	$7.356 {\pm} 0.091$	29 ± 18	3.9 ± 2.4	$9.8{\pm}6.1$
452	$7.356 {\pm} 0.091$	29 ± 18	$4.0{\pm}2.5$	$10.0 {\pm} 6.2$
455	$6.742 {\pm} 0.087$	27 ± 17	$4.0{\pm}2.5$	10.1 ± 6.3
458	$4.665 {\pm} 0.071$	18 ± 11	3.9 ± 2.4	$9.6{\pm}6.0$
459	$4.665 {\pm} 0.071$	18 ± 11	$3.9 {\pm} 2.5$	$9.9{\pm}6.1$

Table 4.17: Calculated muon production efficiencies and total fluxes at 1 μ A. 'Component' is the simulated fraction of the total beam that stops in each configuration. 'Rate' is the sum of the corrected rates for free and copper decays. The total flux was calculated using equation (4.17). The average total flux is $(4.2 \pm 1.1) \times 10^8$ muons s⁻¹ where the error is the same fractional error as is on the MPPC efficiency, $\Delta \epsilon_{\text{MPPC}}^2 \sim 26$ %. The efficiency is the total flux divided by the full beam power (400 W). The average efficiency of MuSIC is calculated to be $(10.4 \pm 2.7) \times 10^5$ muons W⁻¹.

The other key problem with the final flux measurements are the bias introduced towards the simulation. Obviously this has a large effect on the final flux and efficiency measurements but unless the simulation is off by quiet a large amount it seems unlikely that MuSIC's intensity will be less than 10^8 muons s⁻¹

Part II

Clock and Control interface for the Large Pixel Detector

Chapter 5

Introduction

5.1 Conventions

Throughout this document there are several typographical conventions that are observed (see table 5.1).

Type or prefix	Meaning
lower mono-spaced font	Port or signal names.
lower normal font	Port or signal names (tables only).
CAPS	Names of generics.
BOLD CAPS	Named states of a state-machine.
MONO-SPACED CAPS	Command words.
CAPS	Command words (tables only).
0xYYYY	Hexadecimal number (YYYY in base 16)
0bYYYY	Binary number (YYYY in base 2)

Table 5.1: Description of typographic conventions.

5.2 The European X-ray Free-Electron Laser: an overview

This is a discussion of the work carried out designing and implementing the firmware for the Clock and Control Card (CCC) [31] interface of the Large-Pixel Detector (LPD) [32] for use at the European X-ray Free-Electron Laser (EuXFEL) [33]. There will be a brief discussion of EuXFEL, its aims, the detectors and control systems then a more in-depth look at the design, implementation and testing of the interface.

EuXFEL is a 3.4 km Free-Electron Laser (FEL) being constructed below Hamburg, Germany. The project is scheduled to begin operation in 2016 with commissioning beginning in 2015. EuXFEL is built upon expertise and concepts prototyped at the Free-electron Laser in Hamburg (FLASH) [34] which is operated by DESY although EuXFEL will be operated as an independent research facility. The aim of EuXFEL is to produce a coherent X-ray beam with peak brilliance of 10^{33} photons/s/mm²/mrad²/0.1% BW; a pulse duration of ~100 fs and a wavelength down to ~0.1 nm.

5.2.1 Synchrotrons and FELs

X-rays are produced from electrons via synchrotron radiation. When an electron is accelerated through a magnetic field, that causes it to curve and lose energy, this energy is lost as X-rays photons. Synchrotron sources use Linear Accelerators (Linacs) to accelerate electrons before injecting them into a storage ring where they are passed through bending magnets to produce X-rays. The original synchrotron sources used the natural curvature of their storage ring to produce X-rays where more modern designs use special sets of magnets (called 'undulators') that very rapidly change the electron's course to produce X-rays. By changing the configuration of the undulators, the X-ray properties can also be adjusted.

At a FEL very long undulators are used. The long undulators are the key difference that separates FELs from synchrotrons: FEL undulators are tuned so that the electrons undergo 'micro-bunching'. Micro-bunching is a process in which the electron bunch interacts with the X-rays it has emitted. As the electrons travel those that lag will receive a boost from the X-rays emitted behind them. Meanwhile, those electrons that lead tend to lose energy via emitted X-rays. As the faster electrons lose energy and the slower electrons gain it the bunch breaks into many smaller bunches. These micro bunches are in phase and produce coherent X-rays. The process of micro-bunching produces X-ray light which, like a laser's, is coherent but it also arrives in very short pulses ($\mathcal{O}(fs)$). Whilst lasers are produced through light-amplified stimulated-emission, micro-bunching produces X-rays through Self-Amplified Stimulated-Emission (SASE)*.

5.2.2 X-ray production at EuXFEL

At EuXFEL electrons are accelerated to 17.5 GeV [33] using a superconducting linac and the X-rays are produced in one of three SASE undulators which can be combined with conventional undulators to produce photon energies from 25 keV to 0.26 keV which correspond to wavelengths of between 0.05 nm and 4.7 nm. The proposed layout can be seen in figure 5.1. EuXFEL is expected to supply 27,000 pulses of X-ray light ('bunches') per second, with each pulse having a duration of ~100 fs. The bunches are split over 10 'trains' per second, each containing up to 2,700 bunches and lasting 600 μ s. The predicted cumulation of all this is a peak brilliance of 10³³ photons/s/mm²/mrad²/0.1%BW. For comparison, the energy and peak brilliance of several existing sources (FEL and synchrotron) is given in figure 5.2. As can be seen FELs (FLASH, LCLS) have a much higher peak brilliance although, generally, with a reduced range in energy.

^{*}Technically 'FEL' is a misnomer, it should be FES as it is not a laser.



Figure 5.1: Schematic of the beam-lines for EuXFEL, black denotes electrons whilst red corresponds to X-rays. The blue 'SASE' blocks are SASE sections whilst the 'U' blocks are spontaneous emission undulators (which don't self-amplify). Taken from [33].

5.2.3 Scientific Motivation

There are two primary problems with more traditional synchrotrons: incoherent light and pulse length. As the light is produced in a long bunch of electrons it has no overall phase. This means that only samples that are crystalline (or can be crystallised, i.e. grown into a repeating pattern) can be imaged. Many structures form only poor crystals or can't form them at all. The X-ray pulse length that synchrotrons produce, whilst under normal operating conditions, are generally of order 10–100 ps [35]. Obviously this places a lower limit on the speed of things that you can 'film', again limiting the range of experiments that can be carried out. FELs solve both of these problems by producing coherent light that can have a very short pulse length, additionally because of the SASE process the peak brilliance of a FEL is vastly increased (see figure 5.2).

The primary aim of EuXFEL is to study conditions previously unseen in a laboratory setting. This aim is achieved through three core properties of EuXFEL: 'coherence, ultrahigh brilliance and time structure' [33] the combination of these gives access to three broad areas of study: the tiny, the fast and the extreme. Because of the limitations of incoherent light, pulse length and brilliance none of these regimes are easily studied at synchrotrons.

The imaging of the tiny relies on the wavelength of the light used being comparable to the scale of the structure to be imaged. At EuXFEL as well as having X-ray wavelengths sufficient to image molecules, due to the coherent nature of the light non-repeating structures can also be imaged unlike at traditional synchrotrons. Whilst the brilliance of the beam will destroy most samples very quickly, tests at FLASH show that enough time remains to produce a detailed image of the sample, even if it has not been crystallised. This means that larger structures can be imaged at an atomic scale (e.g. entire viruses) or structures that won't crystallise or only form very small, low quality crystals (e.g. protein membranes).



Figure 5.2: Plot of peak X-ray brilliance against energy for a range of current sources as well as the predicted values for EuXFEL (here labelled 'XFEL'). The blue dots show measured peak brilliance for several energies at the existing FLASH facility, 'FLASH (seeded)' is a proposed extension using a micro-bunch 'seeded' electron beam. Taken from [33].

EuXFEL's time-structure, provides the potential of the second regime, speed. As each individual flash lasts less than ~ 100 fs and each train comprising of 2,700 flashes it's possible to 'film' processes as they occur. This will make it possible for researchers to understand what happens during a phase transition or when a material reverses its magnetisation by watching it happen in high detail and without suffering the motion-blur of slower systems.

The final regime, the extreme, is driven by EuXFEL's brilliance. Able to recreate intense temperatures and pressures EuXFEL can use this to create environments not normally seen on Earth. For example: the propagation of shockwaves through a plasma or to image the stresses on a component under extreme magnetic fields.

5.3 Detectors at EuXFEL

In order to achieve EuXFEL's scientific program a broad range of detectors are required. For almost all planned experiments there is a need to image the beam's interaction with the target, the standard solution to this is a 2D pixel detector. This type of detector consists of an array of light sensitive pixels that give both position and intensity information about the incident light, with minor reconstruction an image of the target can then be formed. The basic requirements for the 2D pixel detectors at EuXFEL are [33]:

- Swiftness EuXFEL produces 27,000 X-ray pulses per second, the detector needs to be able to record a large number of these.
- **Dynamic range** In any one flash the number of photons received by any portion of the detector can vary massively (between 1 and 10^5 photons [36]) this information needs to be preserved with a good signal to noise ratio by the detector.
- **Radiation resistance** When fully operational EuXFEL is intended to be used nearly continuously, so obviously any detector used has to be able to survive the harsh environment at the end of the beam-line.

There are currently three 2D pixel detectors being built for use at XFEL: Adaptive Gain Integrating Pixel Detector (AGIPD) [37], DEPFET Sensor with Signal Compression (DSSC) [38] and Large Pixel Detector (LPD) [32]. All three satisfy the above requirements through a variety of technologies.

The main differences between the three detectors are in their approach to the dynamic range: LPD and AGIPD both have three separate gain levels giving them the required range, whilst DSSC uses the non-linearity of its DEPFET to achieve a similar outcome. There are a few other significant differences: DSSC has hexagonal pixels (AGIPD and LPD have square); AGIPD uses dynamic switching to select the appropriate gain for each pixel before storing it in a single pipeline and LPD has an entire pipeline for each gain level (this means that when a narrower gain is required it can be set and all three pipelines used for storage).

5.3.1 The Large Pixel Detector (LPD)

LPD is a 2D, 1 Mega-pixel detector designed and build by a collaboration of the Rutherford Appleton Laboratory and Glasgow University in the UK. The detector is designed to be modular with a full 1 Megapixel being made up of 16 'supermodules', each supermodule contains a single Front End Module (FEM) that controls and reads out the 128 Application Specific Integrated Circuits (ASICs) each of which has 512 individual pixels i.e. each supermodule is 65,536 pixels divided between 128 ASICs and 1 FEM. It is the FEM that then communicates with the rest of EuXFEL via the Clock and Control Card (CCC) and the Train Builder (TB) [39] (which acts as a common readout card).

There are two lines specified for controlling the ASIC during operation: the system clock (clk) and the control (cmd). The FEM's CCC-interface is responsible for receiving the generic signals from the CCC and converting them to those expected by the ASIC. There are a large number of commands that the ASIC expects in order to function (a full list is given in appendix A.1). These commands fall into a few general groups: starts, (no-)vetos^{*}, stops, resets, and testing. In each of these sets of commands there are generally two or three individual command words that effect a specific component of the ASIC (e.g. reset the write pointer, start the trigger pointer).

Each FEM uses a Xilinx Virtex-5 Field Programmable Gate Array (FPGA) and two Xilinx Spartan-3's for control and fan-out, the Virtex-5 has two softcore PowerPC440 processors that manage the resources on the FEM (e.g. configuring control registers). As well as the CCC-interface, firmware manages readout of the ASIC; the ASICs' configuration and communication with the TB. The two Spartan-3 FPGAs are used primarily to co-ordinate fan-out of the signals to the ASICs.

5.4 DAQ and control systems

In a project the scale of EuXFEL there are a large number of different subsystems that need to communicate flawlessly in order to operate. Not only does a single common clock need to be distributed between all systems but it needs to compensate for the time to transmit a signal between the different components and how this latency may change depending on local conditions (e.g. the temperature). To maintain synchronicity there are several layers of timing and control system used at EuXFEL. The top-most layer is the master clock from which all other timing signals are derived. The master clock signal is distributed, along with global information about the machine's status (e.g. the next bunch-train's ID), to the Timing Receiver (TR) cards. The 2D detectors use an additional layer (the CCC) to simplify their interface to the TR cards. The CCC removes information that is not needed by the 2D detectors and provides additional information that is received directly from other subsystems (e.g. vetos). In addition to the CCC the 2D detectors share a second common interface, the TB that provides a common method of storing and ordering data from each

^{*}The ASIC actually uses a 'trigger' rather than a 'veto', but for consistency with the EuXFEL documentation we will use 'no-veto' and 'veto' to refer to TRIGGER_FLAG_SET and NOP respectively

train.

The CCC has four primary functions in EuXFEL:

- 1. Distribution of the clocks required to maintain synchronicity with the rest of the machine.
- 2. Control of the attached FEMS.
- 3. Providing of veto information for each bunch.
- 4. Collection of status information from the FEMs.

In addition to these requirements the CCC can also operate in standalone mode (detached from a TR board) in order to facilitate testing and for use at other locations (e.g. LCLS).

The CCC communicates with FEMs via RJ45 (see figure 5.3), the four paired wires carry signals: clock, fast control, veto and status. The clock signal is a \sim 99 MHz derived from the TR and synchronised to the, \sim 4.5 MHz, bunch clock, the fast control line carries information about each train whilst the veto line carries information on whether a bunch should be kept, the status line returns the received clock to indicate good connections.

5.4.1 Clock and Control Card (CCC)



Figure 5.3: The CCC RJ45 wiring diagram. Arrows on the inputs/outputs to FEE indicate signal direction, apart from the status line all are input to the FEE, i.e. the status line is the only line expected to send signals from the FEE to the CCC. Modified from [31].

Clock signal

Rather than a simple monotonic time structure, EuXFEL has two superimposed patterns: the bunch trains and the bunches. The trains arrive at a rate of 10 Hz with each lasting only 600 μ s but containing up to 2,700×100 fs bunches, each separated from the next by ~220 ns (i.e. a rate of ~4.5 MHz), figure 5.4 shows this.



Figure 5.4: The timing structure of electrons at EuXFEL and the resultant X-ray pulses. Modified from [31]

This timing structure forces the detectors to use the time between bunches for data readout while during the bunch train they are limited to just storing data. This means that each detector is limited by the length of its on-ASIC pipeline with regards to how much data it can store (for LPD this is either 512 frames if using all three gain levels or 1536 if only using one). This structure also means there are two cycles that the detectors need to be synchronised to: the bunch clock (\sim 4.5 MHz) and secondly the bunch-train clock (\sim 10 Hz). Finally, in addition to these two machine-wide clocks there is the common CCC 'fast clock' that is used for transmitting commands from the CCC which has a frequency of \sim 99 MHz and is the clock that the ASIC works to.

Note: there is some vagueness as to the exact clock speeds used as, at time of writing a definitive value has yet to be decided on, the bunch clock is expected to remain between 4 and 5 MHz with the fast clock expected to be a simple divisor of this that results in a rate of roughly 100 MHz.

Fast control

The fast control is used to convey information about each train. This information comes in four parts: when the next train will start, when it will stop, what its ID is and what bunch pattern should be used (see 'Veto signals', below). There is also a reset command to indicate that the ASIC and FEM should be reset to a known state.

The expected use of the fast signal is shown in table 5.2, a START signal arrives with attached train ID and bunch pattern ID, after some number of vetos the STOP signal is received. Rarely, either when there is a fault or if, for example, the detector's been switched off, the RESET signal will restore the FEM and ASIC to a prepared state.

Command	Bits	Payload	Description
START	0b1100	Train ID (32b), bunch pattern ID (8b), checksum (8b)Start of the train	
SIANI	001100		
STOP	0b1010	none	End of the train
RESET	0b1001	none	Reset the FEM and ASIC
reserved	0b1111	n/a	n/a

Table 5.2: Specification of the fast control signals and their payloads.

Veto signals

Given the previous discussion of the 2D detectors at EuXFEL (section 5.3) it is obvious that it is unfortunately impossible for them to record all 2,700 bunches of the data. Equally with a single linac being divided between five, and later ten, experimental stations not every detector will be receiving all of the bunches for every train anyway. To account for this there are two veto systems that allow the detectors to select which bunches they should record for processing: the 'bunch pattern' and the 'online veto'. Either of these two sources may veto a bunch so it is only recorded if *neither* vetos it.

The bunch pattern veto is derived from the global configuration of the machine: if, for example, the first half of the electron beam is being sent to another experimental station the detector will receive a pattern that tells it to veto that portion of the train. The patterns act as masks: for each bunch in a train the pattern states whether it should be vetoed or not. The bunch patterns are decided ahead of time, and a selection of patterns^{*} are loaded when the FEM is configured, the bunch pattern to be used for each train is included in that train's header information as part of the start signal sent via the command line.

The online vetos are mainly situational, if the beam doesn't produce any X-rays or there is a fault then there is rarely any point taking data, in which case those bunches should be vetoed. These online vetos can have any source and the signals are supplied to a dedicated veto unit that is external to the CCC. The CCC will in turn pass on the veto or no-veto

^{*}Predicted to be fewer than 10.

signal to the FEM of the detector, either with an attached bunch ID or with a fixed latency from the bunch in question depending on the specifications of the detector. Online vetos are what is received via the veto line and they have a format given in table 5.3, currently LPD makes no use of the bunch ID. The CCC promises that for every bunch either a VETO or a NO-VETO signal will be sent, if neither is received then the interface will assume a NO-VETO.

Command	Bits	Payload	Notes
VETO	0b110	Bunch ID (8h)	Veto this bunch
NO-VETO	0b101		Record this bunch
reserved	0b111	n/a	n/a

Table 5.3: Veto signal specification.

5.5 Firmware

Firmware describes a broad range of technologies that bridge the divide between hardware (the physical chip and wires) and software (a program intended to run on a processor). Whilst firmware can often be used to refer to quite complex programs run on embedded systems, in this document it is used to refer to the specific logic loaded onto a programmable chip to configure its operation.

As discussed, the LPD FEM uses Virtex-5 FPGAs to run its firmware, FPGAs are made up of 'slices' of logic that can be configured in order to create powerful systems. The general layout of a slice is a block of configurable logic attached to a Look Up Table (LUT) this combination provides basic logical manipulations followed by a brute force 'if A then B' method of implementing the design. It's important to note that most modern FPGAs have additional dedicated slices that allows them to implement more specialised functions, examples include: digital signal processing, the previously discussed softcore processors, dedicated 'Block RAM' (BRAM) etc. The specialised slices mean that software can be used to control configuration settings for the firmware natively which then can run without support from an operating system or creating a custom chip whilst also making use of large optimised structures like BRAM for storage.

5.5.1 VHDL

There are a variety of languages for writing firmware, the one used for LPD is VHSIC^{*} Hardware Description Language (VHDL). VHDL works by describing the expected operation of various discrete blocks within the firmware, these descriptions are then translated

^{*}Very High Speed Integrated Circuit

(synthesised) into bit-code that, in the case of FPGAs will tell the chip how to configure itself.

Whilst a detailed description of the language is beyond the scope of this document there are several attributes of it that are required to understand the following discussions. The key difference is that VHDL is *description* language. This means that the code exists to describe the mapping from inputs to outputs, how this is actually implemented is left to the compiler. The upshot of this system is that depending on the target being compiled for (e.g. FPGA, ASIC, etc.) you can get very different implementations of the same code.

The slice-based architecture of FPGAs means that many designs are synthesised as 'apply logic to signals' then 'look up value of signals in table' and finally 'output value stored in table'. The result of this system is that most designs are split into two groups of components the 'state-machine' and the 'memory'. A state-machine is a set of states with attached conditions, the inputs to the state-machine determine which state should be selected and then that state determines what the output should be, the memory stores any persistent information needed by the state-machine either as input, or output. Throughout CCCinterface firmware there are examples of this where a state-machine implements the logic and a BRAM provides large scale storage.

In VHDL, blocks of code (called 'entities') are defined by their name, ports and generics; entities represent a cohesive unit of logic, for example a state-machine. Ports represent in or out-bound signals^{*} to that entity, these can either be single or grouped into vectors. The generics of an entity describe constant values associated with it; they can be changed on a per-synthesis scale but not by the firmware itself. Generics are mainly used in the design to specify reset values for registers and values that shouldn't be changeable at run time, but may need to change between systems (e.g. delays).

VHDL specifies two broad classes of value that can be used for ports: signals and vectors, a signal is a single bit of information whilst a vector is a collection of bits in some order. VHDL's basic signal type is called 'std_logic' (sl). std_logic is generally used for transfer of the boolean values ('1' or '0') but it can also take several other values[†] that better describe the ultimately analogue reality of hardware signals e.g. the value 'L' specifies a weak signal that should probably be low. The vector form of std_logic is a 'std_logic_vector' (slv) that is ordered 'X (up)to Y' or 'Y downto X', if the slv is converted to a numeric type (e.g. an integer) then this ordering determines the 'endedness'. E.g. if an slv (3 downto 0) is 0b1000 then it has an integer value of 8, the same value but with ordered reversed (i.e. 0 to 3) has an integer value of 1. Throughout this document order is denoted using parenthesis e.g. slv (3:0) is a std_logic_vector (3 downto 0) while slv (0:3) is (0 to 3).

^{*}VHDL also specifies 'inout' as a bi-directional port but they are not used here. ''L', 'H', 'U', 'W', 'X', 'Z' and '-'.

Chapter 6

Design

Design is a two stage process: initial design and designing whilst implementing. Very few large scale projects end up looking exactly like the early designs as use cases change and flaws emerge. To this end it's better to design from a few basic principles and use those to guide the implementation rather than a fixed plan which may, as the implementation progresses, turn out to be wrong.

This chapter discusses the uses cases that were considered for the LPD-CCC interface firmware, the more general principles that guided decisions and finally this chapter discusses the overall design as it was implemented.

6.1 Requirements

The requirements for the LPD-CCC interface can be split into several groups: those requirements made by EuXFEL/CCC, those that were made by the LPD group and those that emerged as a result of the technology. Whilst most of the requirements of one can be equally seen as requirements of the other by splitting the design based on the source certain design decisions became clearer. This section now discusses these requirements before further discussion of the design.

The first and most obvious requirement is that the interface must correctly interpret the commands received via the four lines that make up the CCC interface and respond appropriately. The firmware also had to remain synchronised with the ~ 4.5 MHz bunch clock to ensure the ASIC recorded data at the correct time. To facilitate this the interface had to respond with a fixed latency to all commands it received.

LPD requires that the interface can talk to the ASIC via the clk and cmd lines and send the correct word in response to the signal received from the CCC. In order to account for the wrapping nature of the LPD pipeline the FEM was required to log the veto decision for each bunch; without this reconstructing, the time-order of the received images would be a tricky task. To facilitate testing of the LPD ahead of EuXFEL's completion it was also decided that the interface would have to be able to run without an attached CCC in a 'single shot' style configuration controlled directly via the softcore processors (which became known as 'reset-mode', see section 7.4). A final requirement was that the clock sent to the ASIC had to be adjustable as the v 1.0 of the ASIC required a much slower clock during readout (by a factor of \sim 100, i.e. 1 MHz).

6.2 Design Principles

The main design principles were that the design be modular, flexible and simple. These criteria have many advantages and a few disadvantages, they are also strongly entwined with the pursuit of one often resulting in the development of another. By using these principles the result is a design that not only fulfils the current requirements, but is easy to adapt to future requirements, more maintainable and more understandable.

A modular system is one in which the problem is divided into smaller units. Each unit should solve one, and only one, problem. Each problem should be made as simple as possible. By creating units in this way they can be developed and tested in isolation which reduces the number of inputs and outputs to be understood. Particularly useful units can then be re-used with less testing and development. The main cost of using modules is that the total solution may not be as efficient as one made from a single entity although developing and testing such a solution may be nearly impossible.

A good example of this modularity was the control registers. Whilst each version was customised for the specific use case, each one was built from the same core design and once that was established further versions were quick and easy to develop and also had a much lower chance of failure.

The flexibility of a design is strongly coupled to how modular it is. Smaller, simpler entities are generally more re-useable than larger ones. Obviously at some point the small modules have to be combined and there are certain things that can't be achieved without a certain, minimum amount of complexity. To address this as many aspects of the design as possible were made using VHDLs generics. This meant that the same solution could be used in a number of different situations.

For example, given the number of different clock speeds that might be used with the system, the divisor to convert the fast clock into the bunch clock was set using a generic. By using a generic to describe the ratio of one clock to the other as long as the ratio remains within certain bounds the system can be easily re-configured.

The final principle, simplicity, should be key in making the firmware a long-lived piece of work. Rather than an inflexible, monolithic design that has to be thrown away at the slightest change, hopefully the solution should be easy enough to work with that even if large portions need to be updated the individual changes should be small and easy to make with few (ideally no) unforeseen consequences or dependencies. An example of this is that sections of the receiver entity are already being used to test the prototype CCC.

6.3 The Design

As discussed above each entity should have a clear task and as much as possible a simple interface to achieve it. The overall design process was to break the problem down into ever smaller entities until an entity that encompassed the most basic behaviour was encountered e.g. the veto receiver entity de-serialises the veto information and checks it for commands. Certain entities couldn't be fully decomposed without introducing non-uniform latencies, avoiding these situations was considered paramount. The main reason for the strict requirement for simple latencies is to simplify fan-out of the signals, which is carried out by dedicated FPGAs. This was especially true for implementing the veto logic in the transmitter entity, in such cases the core requirements won out over design principles.

Ultimately the design was split into three sections: receiver, transmitter and the veto filter. The receiver entity is designed to read the signals from the CCC, process them and flag which words it had received. The veto filter combines the bunch pattern, maximum number of writes the ASIC can make and veto signals whilst logging the decisions. Finally the transmitter interprets the flags set by the other two entities and send the correct word to the ASIC.

Each section used a state-machine for the logic. The veto-filter and transmitter also used BRAM and registers for large scale storage and configuration. In each entity, the number of signals in or out, was minimised to avoid confusion. Softcore access to the BRAMs and registers was done via a standardised 32b Remote Direct Memory Access (RDMA) interface. Communication with the other firmware units was implemented using std_logic flags apart from the transfer of header information via LocalLink to the readout entity.

Chapter 7 Implementation

As discussed in chapter 6 the top level design was split into three entities: receiver, veto-filter and transmitter. Each entity has its own section (below) split into an introduction on the aims of the entity; a description of its interface including generics; what registers it uses and any format information; and finally discussion of how the entity was implemented.

The interface of the entity describes in and out ports as well as any generics that it uses, 'LocalLink' and 'RDMA' refer to the collection of ports that make up a pre-defined interface (see appendices A.3 and A.2 respectively). The five attributes ('in', 'out', 'generic', 'LocalLink' and 'RDMA') are grouped together under the description 'direction' in the following tables.*

The registers section of an entity describes any memory locations that can be accessed via the softcore, what is stored in them and what (if any) format that data is expected to be in (e.g. some control registers will consist of multiple bits used to flag specific use cases). The implementation sections discuss the design decisions and intended use of each entity as well any possible limitations (e.g. some BRAMs can overflow causing undefined behaviour if incorrectly set).

7.1 Top level

The top level of the CCC interface consists of three core and two ancillary entities. The receiver, the transmitter and the veto filter form the core and a delay unit with associated configuration register form the ancillary entities. A schematic of this can be seen in figure 7.1.

7.1.1 Interface

The top level interface is given in table 7.1, the generics for the top level are discussed in section 7.1.2. As can be seen, the interface is fairly simple: the received clock, an asynchronous reset, then the two in-ports from the CCC (command and veto), two out-ports to the ASIC, four flags for use within the FEM, a set of RDMA interfaces to the softcore and a LocalLink

^{*}Except in the top level description where the generics are in a separate table.



Figure 7.1: Top level block diagram. The clk and rst signals are supplied to each block and have been excluded from block diagrams unless otherwise noted. clk can either be an internal signal if the board is running in a standalone mode or is the clock received from the CCC. rst is an internal board reset, not to be confused with the RESET command which is intended to reset the ASIC (not the FEE).

Name	Direction	Type	Description
clk		sl	Fast clock, generally from CCC.
rst	in	sl	FEE internal reset signal.
$\mathrm{cmd}_\mathrm{from}_\mathrm{ccc}$	111	sl	Serial CMD line from CCC.
veto_from_ccc		sl	Serial VETO line form CCC.
serial_to_asic	out	sl	AKA 'asic_in' at ASIC.
clk_to_asic	Out	sl	AKA 'clk_in' at ASIC.
$rsync_sent_flag$		sl	See section 7.4.2
$readout_sent_flag$		sl	"
$downscaler_start_sent_flag$	out	sl	"
${\rm downscaler_stop_sent_flag}$		sl	"
nvetos_sent		slv $(8:0)$	"
11		LocalLink	Access to the veto log.
delay_reg		RDMA	Set the internal delays, see
	Interface		section 7.1.3
tx_cmd_bram	Interface	RDMA	See section 7.4.3.
tx_ctrl_reg		RDMA	"
pattern_bram		RDMA	See section 7.3.2.
$pattern_id_reg$		RDMA	"

combined with the nvetos_sent bus to the readout entity in order to form the packet to send to the TB.

Table 7.1: Top level interface for the clock and control interface.

7.1.2 Generics

To provide maximum flexibility (as well as minimum repetition) all generics are available at the top level and are automatically propagated to all entities that use them. There are four generics that are specific to the top level: **START_**, **STOP_**, **RESET_**, and **VETO_START_** (all suffixed with '**DELAY_RST**', see table 7.2 for a full list of generics). These generics set the *reset* values for the delay registers (see table 7.3) i.e. should the **rst** signal be asserted then the values that register will have afterwards are specified by these generics. A detailed description of what these registers do is given below in section 7.1.3.

7.1.3 Registers

There is only one set of registers in the top level that is accessible; it is used for control of delays between the receiver and other entities. The 32b register provides delays of up to

Name	Type	Entities	Notes
WORD_LENGTH	integer	TL, V, T	Length of ASIC command
			words, default:22.
MAX_NVETOS	integer	TL, V, T	Maximum of no-vetos to
			accept, default:512.
N_BUNCHES	integer	TL, V, T	Number of bunches in a
	- ()		train, default:3072.
START_DELAY_RST	slv (31:0)	TL	(default: 0x00000001)
STOP_DELAY_RST	slv(31:0)	TL	(default: 0x00000001)
RESET_DELAY_RST	slv (31:0)	TL	(default: 0x00000001)
VETO_START_DELAY_RST	slv(31:0)	TL	(default: 0x0000055)
START_WORD	slv $(3:0)$	R	From CCC, (default: 1100)
STOP_WORD	slv $(3:0)$	R	From CCC, (default: 1010)
RESET_WORD	slv $(3:0)$	R	From CCC, (default: 1001)
VETO_WORD	slv $(2:0)$	R	From CCC, (default: 110)
NO_VETO_WORD	slv $(2:0)$	R	From CCC, (default: 101)
BUNCH_ID_LENGTH	integer	R	Bunch ID length (default:12)
$TRAIN_ID_LENGTH$	integer	R	Train ID length (default:32)
CHECKSUM_LENGTH	integer	R	Checksum length (default:8)
BUNCH_PATTERN_LENGTH	integer	R	Bunch pattern ID length
			(default:8)
$PATTERN_REG_(0:9)_RESET$	slv (31:0)	V	Reset values for the pattern
			register.
$\operatorname{REG}_\operatorname{RESET}_(9:0)$	slv (31:0)	Т	Register resets $(0-9)$.
SYNC_RESET_SIG	slv (31:0)	Т	Flag that the SYNC_RESET
			command has been sent.
READOUT_SIG	slv (31:0)	Т	Flag that the READOUT
			command has been sent.
DOWNSCALE_SIG	slv (31:0)	Т	Flag to start the down-scaler
			(either internal or external).
DOWNSCALER_STOP_SIG	slv (31:0)	Т	Flag to stop the down-scaler.
DOWNSCALE_FACTOR	integer	Т	Factor for the internal
			down-scaler, default: 100.

Table 7.2: A table of the generics used in the design, their type, name, where they are used (R=receiver, V=veto-filter, T=transmitter, TL=top-level).

 $2^{32}-1$ clocks for the three signals sent to the transmitter entity (i.e. start, stop and reset) and the delaying of the start signal to the veto filter, i.e. that forms the veto_start. The appropriate delay for 'veto_start' is given by:

$$VETO_START_DELAY = START_NWORDS \times WORD LENGTH + START DELAY - 2$$
(7.1)

where START_NWORDS is the number of words (as set in the transmitter control register) to be sent to the ASIC in response to the START command. This must include any NOPS, 'non-operations' a word that has no effect (normally all 0's), required to delay the actual start. WORD_LENGTH is the length in bits of each of those words (set via generic) and START_DELAY is any further delay added between the receiver and the transmitter, the '-2' accounts for the internal delay of the veto filter.

There is no delay register to control the veto signals themselves as these are assumed to be wanted with minimum latency.

Name	Address	Default
START_DELAY	0x1	0x00000001
STOP_DELAY	0x2	0x00000001
RESET_DELAY	0x3	0x00000001
VETO_START_DELAY	0x4	0x00000055

Table 7.3: Summary of the top level delay registers. The default can be overridden via the generics (section 7.1.2).

7.1.4 Implementation

A schematic of the entities at the top level is shown in figure 7.1. There are two broad paths through the system: command and veto. The command path is delayed to allow offsets to be accounted for whilst the veto path is kept as fast as possible. It's also worth noting that the veto filter receives two start signals, one to indicate that the train ID and bunch pattern ID have been received and a second to indicate when the first veto will arrive, both of these signals are derived from the same start signal but the later has to be carefully timed to coincide with the beginning of vetos.

7.2 Receiver

As stated in the introduction there are five signals that are sent by the CCC, these are transmitted via two ports: the cmd_from_ccc and the veto_from_ccc. The CCC clk carries the derived fast clock (normally 99 MHz) which is synchronised to the slow machine clock (normally 4.5 MHz) which is distributed by the timing receiver. The full list of commands

and payloads can be seen in table 7.4. As the receiver block is expected to be static during operation (i.e. not need any configuration) only generics are used and there are no externally accessible registers.

Name	Value	Payload	Notes
START	0b1100	48b	Start of the train
STOP	0b1010	n/a	End of the train
RESET	0b1001	n/a	Reset the FEM and ASIC
VETO	0b110	12b	Veto this bunch
NO-VETO	0b101		Record this bunch

Table 7.4: The full set of commands received from the CCC. The payload for the START is the train ID (32b), the bunch pattern ID (8b) and a check-sum (8b). The (NO-)VETO payloads are the same, the 12b bunch ID.

7.2.1 Interface

The top level interface for the receiver entity is shown in table 7.5. Of the generics, only the word definitions should be changed and only then for words of the same length (e.g. '110' changed for '101'). The payload values can be made shorter but not longer.

The in ports are essentially the same as the CCC specification with the addition of an internal asynchronous **rst** line which will clear any buffers and return the state-machines to IDLE.

The outputs give a single line for each command as well as buses that hold payload values until overwritten. It's important to note that the command lines go high for only a single clock whilst the payload buffers will retain their value until the next command starts, this means that, for example, the train_id will remain available until the STOP command begins to be received.

7.2.2 Implementation

The receiver module is split into two state machines, one for each of the two command lines, as seen in figure 7.2. The entities have similar designs although the command receiver entity is slightly more complex to cope with the non-constant, multi-part payloads.

The designs uses strobes to signal which type of command has been received and buses to hold the ID information. The ID-buses (train, bunch pattern, checksum and bunch) are cleared at the start of the next fast control (not veto commands), this means that the data is available until it gets replaced.

The two entities use simple state-machines designs (figures 7.3 and 7.4) that can be summarised as: 'IDLE \rightarrow LOG_COMMAND \rightarrow STROBE_COMMAND \rightarrow [LOG_PAYLOAD \rightarrow] IDLE'. For both command and veto lines the state machine is triggered by the line going

Name	Direction	Type	Description
START_WORD		slv $(3:0)$	Default serial command: 1100
STOP_WORD		slv $(3:0)$	Default serial command: 1010
RESET_WORD		slv (3:0)	Default serial command: 1001
VETO_WORD		slv $(2:0)$	Default serial command: 110
NO_VETO_WORD		slv $(2:0)$	Default serial command: 101
BUNCH_ID_LENGTH		integer	Expected bunch ID length
	Generic		(default:12)
TRAIN_ID_LENGTH		integer	Expected train ID length
			(default:32)
CHECKSUM_LENGTH		integer	Expected checksum length
BUNCH_PATTERN_LENGTH			(default:8)
		integer	Expected bunch pattern ID
			length (default:8)
clk		sl	CCC clock
rst	in	sl	FEE internal reset
cmd_i		sl	Fast command line from CCC
veto_i		sl	Fast veto line from CCC
start_o		\mathbf{sl}	Start signal to transmitter
stop_o	out	\mathbf{sl}	Stop signal to transmitter
rst_o		sl	Reset signal for transmitter
veto_o		sl	Veto to trigger veto filter
no_veto_o		sl	No veto to trigger veto filter
bunch_p_o		slv(7:0)	Bunch pattern ID to veto
			filter
bunch_id_o		slv $(7:0)$	Bunch ID to veto filter
train_o		slv (31:0)	Train ID to veto filter

Table 7.5: Top level interface of the receiver entity.



Figure 7.2: Block diagram of the receiver entity.

high. The stream is de-serialised by passing the serial commands to a shift register. Once a command is matched then the appropriate strobe is set. If the command has an attached payload then that is passed to a shift register until the appropriate number of payload bits have been received when they are written to a bus and the machine returns to 'IDLE'.

As every command on the **veto** port has the same format:

$$< \text{COMMAND2}: 0 > < \text{BUNCH ID11}: 0 >$$
 (7.2)

the veto receiver (figure 7.4) only implements two states: CMD and BUNCH_ID with the command (either VETO or NO-VETO) being flagged once it's received and the Bunch ID being recorded regardless. In the command receiver (figure 7.3) there are five states: IDLE, CMD, TRAIN, BUNCH_PATTERN and CHECKSUM. The TRAIN, BUNCH_PATTERN and CHECKSUM are used to record the START command's payloads. If the STOP or RESET commands are received, the flow returns to IDLE to await the next command.

7.3 Veto Filter

As discussed in section 5.4.1 there are two sources of vetos that needed be combined: the bunch pattern and the online sources. It is assumed that 10 bunch patterns will be sufficient for general operation. The LPD, due to its fixed pipeline size, has a third source of vetos: only 512 bunches can be recorded.* A general truth table is given in table 7.6, as can be

^{*}Unless running with a single gain setting.


Figure 7.3: Flow diagram of the command receiver. The serial signals are read and flags set as required. The 'START' payload i.e. train ID, bunch pattern ID and checksum) are stored and remain available until cmd_i next goes high i.e. the start of the next command). This system has a weakness in that if malformed commands are sent and then the loop will attempt to 'execute' what ever follows (e.g. the train ID.) and ultimately entering an unknown state.



Figure 7.4: Flow diagram of the veto receiver entity. This de-serialises and splits the veto signals into the command and its payload, the Bunch ID. Note that this system also has the same flaw as the command receiver, i.e. it doesn't count bits as they are received so it can enter an undefined state if a malformed veto signal is sent.

seen a NO-VETO is only sent if fewer than the maximum NO-VETOs have been sent and both the pattern and the veto line agree to specify that a NO-VETO should be sent.

Line	Pattern	Count	Outcome
veto	veto	> Max	veto
nveto	veto	> Max	veto
veto	nveto	> Max	veto
nveto	nveto	> Max	veto
veto	veto	< Max	veto
nveto	veto	< Max	veto
veto	nveto	< Max	veto
nveto	nveto	< Max	no-veto

Table 7.6: Truth table for veto decisions. Count is the number of 'no-veto's already sent, it is assumed to be 512 but can be any number less than this.

7.3.1 Interface

The interface for the veto filter is given in table 7.7. There are ten generics to be set that specify the reset values for the pattern ID register (see section 7.3.2).

The in ports are mainly concerned with signals from the receiver module, these are used to determine when vetos should be expected as they need to be synchronised with the bunch clock. There are two start signals that are of concern: the veto_start and start_i, the former of these indicates when the first veto should be expected whilst the latter indicates that the START command, the train ID and the bunch pattern ID have been received. The train and bunch pattern IDs are required before the first veto is received in order to correctly load the veto pattern as well as create the header for the veto logger. The bunch ID is received but, currently, nothing is done with it. The stop_i signal is used as an alternative to the maximum number of bunches to stop the veto filter.

The out ports are very simple, one is the final veto decision which is sent to the transmitter whilst the other is a count of the number of NO-VETO commands sent to the ASIC. The number of no-vetos is required for the readout entity to know how many words to expect once readout starts.

There are 3 data interfaces which are discussed in section 7.4.3. Briefly: the two RDMA interfaces allow external access to the pattern BRAM and the pattern ID register whilst the LocalLink interface is used by the readout entity to access the veto-decision-log which is required for data reconstruction.

Name	Direction	Type	Description
WORD_LENGTH		integer	Length of ASIC command word (default: 22).
MAX_NVETOS		integer	Number of n_vetos we can send (default: 512).
N_BUNCHES		integer	Maximum number of
			bunches in a train (default:
	generic	. (3072).
PAT ^T TERN_REG_(0:9)_RESET		slv(31:0)	Reset values for the pattern
		1	register (default: see (.3.2).
CIK		SI	L de la FEF
rst		SI -1	Internal FEE reset.
veto_1		SI	
nveto_1		SI	
start_1	in	si	
stop_1		si	Delayed start signal to
veto_start		51	coincide with vetos
bunch id		slv(11.0)	Not actually used
train id		slv(31.0)	Added to the local link
l'am_ra			header.
bunch pattern id		slv (7:0)	Used to select the veto
			pattern.
veto_to_tx		sl	Combined pattern and veto
	out		decision.
nvetos_sent		slv $(8:0)$	Need to know how much to
			read from the ASIC.
$pattern_bram_rdma$		RDMA	Interface to the pattern
			entity RAM, mask:
pottom id nog ndmo	interface	DDMA	Interface to the pattern
pattern_id_reg_rama		NDMA	ID/BRAM offset registers
			mask: 0x0000000F.
11		LocalLink	Interface to the veto log
			FIFO, uses a 256b data bus.

Table 7.7: Interface for the veto filter.

7.3.2 Registers

In the veto filter there are two memory blocks: the pattern ID registers and pattern BRAM, both accessed via a dedicated RDMA interface and the veto log which is accessed via local link (see appendices A.2 and A.3).

Pattern BRAM

The pattern BRAM holds the pre-defined veto patterns. These are combined with the fast veto line to decide whether or not to veto a bunch i.e. a bunch can be vetoed by either the pattern *or* the veto line). Each pattern consists of a bit for each bunch in a train (it is assumed that there will be <3,072 bunches in each train). If the corresponding bit is '1' then the bunch will be vetoed, otherwise it is dependent on the veto line.

It is assumed that no more than ten patterns will be needed in normal operation so the BRAM is configured to hold $1,024 \times 32b$ words. This means that for a 3,072 bunch train 96 words need to be defined for each pattern. The pattern is assumed to be spread across monotonically increasing addresses (e.g. address 0x0 to 0x60), the veto bits are applied 0 to 31.

Pattern ID registers

The pattern ID registers map IDs to their appropriate offset within the BRAM. There is no specification for the number of patterns needed during normal operation but 10 registers are assumed to be enough. The value of each register must be of the format:

$$< PATTERN_ID 31 : 24 > \ldots < BRAM offset 9 : 0 >$$

$$(7.3)$$

default values for the register are given in table 7.8, the reset values are set via generics at the top level and the current value can be changed via RDMA. The defaults assume that the patterns are all unique with no overlap, as such each delimitates $96 \times 32b$ words i.e. 3,072 bunches. If there is an intersection between the end of one pattern and the start of another then setting the offset to the beginning of the common section will work as expected.

Veto log

The veto log is a 13×256 b FIFO which logs the veto decision for each bunch. This information is formed into a LocalLink frame which uses 256b data words. The 256b header is defined as:

$$< \text{Train ID } 255: 224 > < \text{Bunch Pattern ID } 223: 216 > \dots$$
$$\dots < \text{Number of no-vetos sent } 191: 182 > \dots$$
(7.4)

where '...' represent padding (1s). The rest of the frame consists of the veto log which consists of a bit per bunch indicating either a veto ('1') or a no-veto ('0'). The entities are ordered MSB first, e.g. bit 255 of frame 0 indicates the first veto decision whilst bit 0 of frame 12 indicates the last decision.

Register	Value	Pattern ID	Offset
0x1	0x01000000	0x1	0x000
0x2	0x02000060	0x2	0x060
0x3	0x030000C0	0x3	0x0C0
0x4	0x04000120	0x4	0x120
0x5	0x05000180	0x5	0x180
0x6	0x060001E0	0x6	0x1E0
0x7	0x07000240	0x7	0x240
0x8	0x080002A0	0x8	0x2A0
0x9	0x09000300	0x9	0x300
0xA	0x0A000360	0xA	0x360

Table 7.8: Default reset values for pattern ID registers.

7.3.3 Implementation

The veto filter is split into two distinct processes: filtering and logging. Filtering is a requirement of the specification. The logging is required because the ASIC records no identifying information about each bunch and due to the ASIC's write pointer wrapping through memory, images are not guaranteed to be read out in chronological order; a block diagram of this solution is shown in figure 7.5.

The $1,024 \times 32b$ BRAM is used to store the veto patterns to be combined with the veto signal in order to form the veto decisions. To construct the veto decision and control access into the BRAM a state machine implementing figure 7.6 was used. The veto decisions are made at the beginning of each word to reduce latency whilst BRAM and state changes are made at the end. The final entity, the pattern ID register, exists to specify the mapping from bunch pattern ID to BRAM offsets.

The logging is performed using a simple state machine to feed in words to a 32b in, 256b out first in, first out BRAM (FIFO), this state machine (figure 7.7) also constructs the LocalLink frame header (the train ID, the bunch ID then '0' padding). The header requires 8 clocks between start_i being asserted and veto_start, to be properly formed. The FIFO is a simple First Word Fall Through (FWFT) entity with the optional empty and almost_empty signals enabled. The read out of the FIFO is performed via a 256b-word LocalLink interface. Once either stop_i is asserted or N_BUNCHES have been counted any remaining bits of the 256b word are filled with '1's and the LocalLink's src_rdy is asserted.

The LocalLink interface is implemented by wrapping the FIFO's 256b data_out port in some simple logic to make it conform to the local link standard. This mainly involves using empty to monitor the state of the FIFO, almost_empty as eof and an internal flag from the logging state-machine for the sof. The sop (start-of-payload) signal is not used to differentiate the header from the payload as the entire frame will ultimately form part of the meta-data in the header of the readout.



Figure 7.5: Block diagram of the veto filter. Note: veto_start is the start signal synchronised to the vetos, start is the START signal from the Rx block and used to load the bunch pattern.



Figure 7.6: Flow diagram of the veto filter.



Figure 7.7: Flow diagram for the veto decision logger.

7.4 Transmitter

The transmitter entity's main task is to act as the interpreter for the signals received from the CCC, translating them into command sequences that the ASIC can understand and act upon. As has been discussed above there are four signals that need to be translated for the ASIC: START, STOP, RESET and (NO-)VETO. The first three command signals all require an arbitrary number of words be sent to the ASIC whilst the veto signals are responded to by one of two words.

The transmitter can run in one of two modes; 'dynamic veto mode' and 'reset mode'. Dynamic veto mode is intended to be the normal mode of operation for XFEL whilst reset mode is intended for static runs with simple veto patterns (e.g. for testing). A comparison of the two can be seen in table 7.9. Reset mode is enabled by strobing bit 0 of the main control register (see section 7.4.3 for details). The transmitter also has an optional minor mode called 'down-scaler mode'; this is primarily for v 1.0 of the ASIC that requires a slower clock during readout. When enabled it uses a multiplier (set by generic) to slow the clock, after a set number of cycles normal speed is resumed.

	Mo	ode
	Dynamic- veto	Reset
Standalone operation	X	1
Dynamic veto decisions	1	×
	START	
Signal response	STOP	Register flag
Signar response	RESET	rtegister nag
	N/VETO	

Table 7.9: Comparison of dynamic veto and reset modes.

7.4.1 Controlling the ASIC

The LPD ASIC manual [36] specifies the command sequences required for each stage of operation of the ASIC. For the veto commands a simple selection between NOP (veto) and TRIGGER_FLAG_SET (no-veto) is required. The other commands have more complex sequence sets that can change based on operation; a brief summary will be given here and full details can be found in the LPD ASIC manual [36].

The standard response to the START command is to set the ASIC in a state ready to write to its memory, this means resetting the write and trigger points and clearing the skip register. Once these commands have been carried out the ASIC will start incrementing the write pointer through the memory locations. After additionally starting the trigger pointer, the ASIC is ready to write the contents of the pixels when a TRIGGER_FLAG_SET (i.e. NO-VETO) is received.

The command sequence for the STOP command is more complex and split over two stages: the pre and post readout sections. To start readout the ASIC is set in a stable mode thus:

- 1. The ASIC is put in power saving mode.
- 2. Read out of the data begins.
- 3. Disable on-chip resets.
- 4. Manually reset the gain, pre-amp, write and trigger pointers.

The data read back consists of one 36b word for each no-veto sent, once all of these have been sent the ASIC is returned to ready mode for the next train:

- 1. Power up the ASIC, i.e. return bias currents to full.
- 2. Re-synchronise the ASIC to the bunch (~ 4.5 MHz) clock.
- 3. Re-enable on chip resets.
- 4. Stop read out.

Note: the above constitute a basic over view of the STOP process, there are some important considerations with regards to synchronisation of these commands that are beyond the scope of this and covered in full in the LPD ASIC manual [36].

The ASIC command words have the format:

$$<$$
 SYNC 19 $><$ X 18 $><$ CMD 17 : 10 $><$ PADDING ZEROS 9 : 0 $>$ (7.5)

if a word length greater than 20b is being used (e.g. 22b, the current default) then the zeropadding is extended as required. The 'SYNC' bit is a '1' and 'X', by convention, a '0'; these are automatically prepended by the state machine. The only exception to this is the 20b command SYNC_RESET (currently set to be 0x5A5A5) which overrides the above format.

7.4.2 Interface

The transmitter has three 'sets' of generics: the control register reset values, the flag words (ending in _SIG) and the down-scaler factor. The register resets set the default value to be written to the assorted registers if rst is asserted. The flag words specify certain command words that the transmitter should scan for in order to flag them for use elsewhere. The SYNC_RESET flag is intended for use with any entities (e.g. the slow command line) that have to be synchronised to the ASIC's bunch clock. The READOUT is intended for use by the readout entity in order to prepare for readout. The DOWNSCALER_SIG is used to begin either the internal (if it's enabled) or an external down-scaler clock. The DOWNSCALER_STOP_SIG is only

for use by an external source (the internal version counts clocks). The DOWNSCALER_FACTOR specifies the factor to use for the internal down-scaler e.g. a factor of 100 means a scale of 100 fast clock cycles to each slow clock.

The in ports to the transmitter are all flags from the receiver; with the standard exception of **rst** which is the FEM internal signal (cf. **reset** which is the flag from the receiver).

The out ports fall into 3 categories: signals to the ASIC (ASIC_in and clk_in), flags (marked _sent) and start_nwords_o. The ASIC commands are defined in the specification. The flags have been discussed above with regards to their generic-defined trigger words. The final out port is a wrapper to the value of the start_nwords register (see section 7.4.3) which is used in setting the delay for veto_start.

The two RDMA interfaces are used for access to the command sequence BRAM and the control registers (section 7.4.3).

7.4.3 Registers

There are two externally accessible sets of registers in the transmitter entity, a control register and the command BRAM both are accessed via RDMA (see appendix A.2). The control registers direct the flow of the state machine whilst the BRAM stores the instruction sets to be sent to the ASIC. There are three commands that the BRAM has instruction sets for: START, STOP and RESET. Obviously these are normally received via the receiver module but the RESET command can also be sent using a toggle in the control register.

Control register

The control register specifies 10 registers:

- 1 The general state-machine control register that toggles various modes i.e. whether to use the down-scaler module and manual RESET) as well as how many bits to send downscaled if that mode is enabled.
- **2-7** three pairs of registers that store the number of words (shortened to 'nw') and offsets for the 3 different instruction sets i.e. one -nw and one -offset for each of START, STOP and RESET).
- 8 The word to be sent in case of a veto.
- **9** The word to be sent in case of a no-veto.
- 10 The status register, indicates what state the state-machine is in.

default values and addresses are given in table 7.11.

The state-machine control register takes values of the following format:

< DOWNSCALER_ENABLE $32 > \ldots <$ DOWNSCALER_BITS $20:5 > \ldots$ $\ldots < \text{RESET}_MODE_EN 0 >$ (7.6)

Name	Direction	Type	Description
$REG_RESET_(9:0)$		slv (31:0)	Register resets (0-9).
$SYNC_RESET_SIG$		slv (31:0)	Flag that the SYNC_RESET
			command has been sent.
READOUT_SIG		slv (31:0)	Flag that the READOUT command
	C		has been sent.
DOWNSCALE_SIG	Generic	slv (31:0)	Flag to start the down-scaler
			(either internal or external).
DOWNSCALER_STOP_SIG		slv $(31:0)$	Flag to stop the down-scaler.
DOWNSCALE_FACTOR		integer	Factor for the internal
			down-scaler, default: 100.
clk		sl	The CCC clock.
rst		sl	FEE reset.
start	in	sl	From the receiver entity.
stop		sl	"
reset		sl	"
veto		sl	From the veto filter.
asic in		sl	Fast serial commands to the
_			ASIC.
clk_in		sl	Clock to the ASIC.
$readout_sent$	out	sl	readout flag (for ASIC receiver
			entity).
rsync_sent		sl	SYNC_RESET flag (for slow
			control entity).
${\rm downscaler_start_sent}$		sl	Downscaler start flag.
$downscaler_stop_sent$		sl	Downscaler stop flag.
$start_nwords_o$		slv (31:0)	Number of words used for
			'START' to set veto_start delay.
bram_rdma		RDMA	Interface to the ASIC command
	Interface		word BRAM. Mask:
	Interface		0x000003FF.
$ctrl_rdma$		RDMA	Interface to the control register.
			Mask: 0x0000000F.

Table 7.10: Interface for the transmitter.

reset_mode_en is a flag to manually start the state-machine in reset mode i.e. reset_nwords worth of commands from reset_offset in the BRAM will be sent. This mode can be used instead of dynamically determining the veto to be sent in a 'fire and forget' manner.

DOWNSCALER_ENABLE indicates that the internal down-scaler should be used to send the next **DOWNSCALER_BITS**, this means that the maximum number of bits that can be sent at the down-scaled rate is $2^{16} - 1$, i.e. 65,535 or $2,978 \times 22b$ words. Obviously if down-scaling is being handled by an external clock this restriction doesn't apply.

The BRAM offsets and command sequence lengths (-offset and -nw registers respectively) have maximum values determined by the size of the BRAM i.e. 1,024. These values are set using the Least Significant Bits (LSB) of the word i.e. 9 downto 0. Obviously setting an -offset to 1,024 will result in only 1 word and unless the corresponding -nw is set to 1 this will result in an address overflow and undefined behaviour. By the same token setting an -nw register to 1,024 is allowed but will mean that any other commands to be sent must be either specified as subsets of this command sequence or not used^{*}.

The veto/no-veto registers specify the appropriate word to send when in dynamic-veto mode, the defaults are NOP and TRIGGER_FLAG_SET respectively. When run in this way the state-machine progresses START \rightarrow vetos \rightarrow STOP with vetos being determined by the veto filter.

The status register is a read-only register that logs which state the state-machine is in and has the following format:

$$<$$
 state_machine_enabled 31 > ...
... $<$ RESET 3 > $<$ STOP 2 > $<$ DYNAMIC VETO 1 > $<$ START 0 > (7.7)

Instruction set BRAM

The BRAM used to store the instruction sets is $1,024 \times 32b$ words. Words not defined by the ASIC are ignored so internal flag triggers (e.g. for DOWNSCALER_START) need not be defined ASIC commands.

The BRAM expects words to have the format:

$$<$$
 N_NOPS 31 : WORD_LENGTH $><$ COMMAND (WORD_LENGTH -1) : 0 > (7.8)

the COMMAND is expected to be correctly padded e.g. if 22b words are being used for a normal command the last 12b should be '0'. The SYNC bit doesn't need to be set. N_NOPS indicates how many NOP commands should follow the COMMAND, a NOP is pre-defined to be a SYNC-bit followed by the appropriate number of 0s. It is important to note that NOPs contribute to the number of words sent for any state and a NOP sent as a member of N NOPS must not be the final command.

^{*}This could be useful if using the reset-mode or in dynamic mode if RESET is not expected to be used.

Address	Description	Reset generic	Default value
0x1	SM-Control	CTRL_REG_RESET	0x00000000
0x2	Start: offset	$START_OFF_RESET$	0x000000000
0x3	Start: number of words	START_NW_RESET	0x00000006
0x4	Stop: offset	STOP_OFF_RESET	0x00000006
0x5	Stop: number of words	STOP_NW_RESET	0x00000007
0x6	Reset: offset	$RESET_OFF_RESET$	0x0000000D
0x7	Reset: number of words	$RESET_NW_RESET$	0x0000000F
0x8	Veto word	VETO_WORD_RESET	0x00210000
0x9	No-veto word	NVETO_WORD_RESET	0x00200000
0xA	Status	n/a	0x00000000

Table 7.11: Control register layout. With the exception of the status register the reset values of each register can be changed by setting the appropriate reset-generic, these values only take effect when the **rst** line is asserted *not* when the **RESET** signal is received.

Example: if, for the START-sequence and using 22b words, POWER_UP $(0x02^{\dagger})$ followed by 4 NOPS is to be sent then the first two locations in BRAM could be set to '0x00C08000' and '0x00000000'. 0x00C specifies 3 NOPS ((0x3) << 2 =0xC^*) and 0x080000 is the command (0x02<< 22) with the appropriate padding. The final entry of 0x00000000 is the 4th NOP; its SYNC bit will be automatically set.

7.4.4 Implementation

In order to implement the state-machine two concurrent machines are used: the shift-block serialises and inspects command sequences; and the state-machine that determines when to change state and which command sequences to use. Previously, these responsibilities were split across multiple entities but to simplify state sharing and meet the latency requirements they were merged into a single entity. Figure 7.8 shows a entity diagram of the transmitter, both the shift-block and state-machine are within the state-machine entity.

The shift-entity primarily acts as a smart shift register to serialise the command words being sent to the ASIC. It has two secondary functions that require inspection of the commands being serialised: flagging and setting sync-bits. Some of the other components of the FEM require knowledge of when certain command words are sent for their own functionality (e.g. the slow command needs to know when SYNC_RESET is sent). To do this when a word is loaded from the BRAM (and only from the BRAM) it is compared to the 4 pre-set words given in table 7.12, if it matches one then the appropriate flag is set. These words are tested sequentially, in the order given in the table, so if a word matches multiple triggers only the

[†]See table A.1.

^{*}Where '<<' is a left-shift, the << 2 is account for the 2 MSB of the command.



Figure 7.8: Block diagram of the transmitter.

flag corresponding to the first match will be set. The flag words are set via generics and are not expected to change, flags do not have to be valid ASIC words although the ASIC's response to unknown commands is undefined and should be confirmed prior to use. At the same time as checking for flags the first two bits of each word are set appropriately: for SYNC_RESET they are set to 0b01 and all other words 0b10 as specified in [36]. These syncbits are not included in tests for flags. The shift-entity's functionality for the different states are given in table 7.13, the different states are described below.

Flag	Default	Notes
rsync	0x00069694	Used to keep the slow command line synced to the ASIC.
readout	0x00044400	Used to alert the readout receiver entity of input.
down-scaler start	0x00044000	Moving to the slower clock (either internal or external).
down-scaler stop	0x00055000	Stop using the slow clock (external only).

Table 7.12: Flag information, command words that the shift-entity looks for and will flag along. Default values are for 22b words, see section 7.4.3 for more details. When looking for flags the sync bits (the first two bits) are ignored so do not need to be set.

State	Flags enabled	Sync-bit set	Command source
IDLE	×	×	n/a
NOPS	×	1	State machine, command = $0x0$.
VETO	×	1	State machine, either veto or no-veto.
START	1	1	BRAM: START command sequence.
STOP	1	1	BRAM: STOP command sequence.
RESET	1	1	BRAM: RESET command sequence.

 $^1\,$ i.e. 0x200000 is the full 22b command, including sync-bit.

Table 7.13: Description of the shift-entity's behaviour depending on state.

The second process, the state-machine, is mainly concerned with maintaining position within each command sequence, changing state and controlling the BRAM. Figure 7.9 shows the relationship between the different states and the causes for state changes. Meanwhile figure 7.10 shows how access to the BRAM is determined. During state changes the source of the next word is set according to the sources in table 7.13, obviously this change occurs before the actual state changes. It's important to note that, as discussed in section 7.4.3 the final word of any sequence should sent from the NOP state as control returns to the origin state, e.g. 'START \rightarrow NOPS \rightarrow VETO' is not permitted, if the final word is required to be a NOP then this should be set as a unique BRAM entry.



Figure 7.9: Schematic of the control flow in the transmitter entity.



Figure 7.10: General control logic flow for START, STOP and RESET states.

Chapter 8

Testing

This chapter is concerned with the testing regime used on the design. It is split into two sections, first the methodology is discussed and then the results. The aim of testing is to ensure that firstly the design meets the requirements (see chapter 6) and secondly that the limitations of the system are understood.

Due to the both the nature of the equipment and time constraints no hardware based testing could be carried out. Hardware testing has since been carried out at the both the PETRA and LCLS test beams [40]. The rest of this chapter is concerned with simulated logic testing of the firmware rather than on chip or radiation testing.

8.1 Methodology

The methodology used for testing was to iterate rapidly, working from smaller to larger entities and limiting changes to the outermost entity being tested, changes made in subentities restarted the processes from that entity to confirm functionality at each level.

In order to test the designs 'test benches' were written. A test bench is a entity that wraps another, the Unit Under Test (UUT), and generates test inputs to the UUT so its behaviour can be checked. These simulated inputs can be single signals, words or clocks. The UUT and test bench are then simulated and run using a program called, in this case called 'iSim', which gives complete control over the system, this means that the response of sub-units can be tested as well as top-level entities. iSim generates an output that is similar to an oscilloscope view with green traces representing boolean values ('1' or '0') and other colours for the other std_logic values (e.g. orange corresponds to 'U').

8.2 Results

While each individual entity was tested, for brevity only the top level tests are included here as they describe the full functionality of the firmware. There are six tests that were run: a simple dynamic run (START \rightarrow NO-VETO \rightarrow STOP); veto/no-veto state changes; response to reset

command; stop with down-scaler (a requirement for the v 1.0 ASIC); RDMA interface test; and a LocalLink interface test.

For all of these tests there were a couple of pass requirements:

- Correct command to response matching.
- Synchronicity with the bunch clock.
- Fixed latency between commands and responses.
- Correct state changes.

In the following diagrams the green portions correspond to the iSim traces, the red and blue lines have been added afterwards and indicate chains of causation and serialised command words respectively, the larger yellow dashed lines were also added and indicate bunch trains synchronised to the clk_to_asic. On the left hand side are the signal names as used in the design, note that these may not always match the names given in the previous implementation sections.

In all of these tests the clock speed was set to 100 MHz.

8.2.1 Start, no-veto, stop

Figure 8.1 shows a very simple dynamic sequence with a single no-veto word being sent prior to stopping. The top six lines show the clk, rst signals; the transmitter's state (current_state); the simulated input from the CCC (cmd and veto); and the output to the ASIC (cmd_to_asic).

After the START signal is received from the CCC the appropriate flag (start_from_rx) is set, followed by the delayed flag (start_delayed) to the transmitter, which starts sending the appropriate command set.* The delayed flag then forms the log_start and veto_start signals that indicate when to log the header information and start logging vetos respectively. Meanwhile the NO-VETO signal is received and combined with the veto_start to form flag that the transmitter should send a TRIGGER_FLAG_SET. Finally the STOP signal is received, triggering the stop sequence and ultimately a return to IDLE (not shown).

Based on the general requirements it's clear that the correct responses are being sent (and hence the state-changes are occurring as required), the commands are synchronised with the bunch clock and the latency is as required: 6 clocks for commands and 7 for vetos (the extra length for vetos is due to the three rather than four bit command).

^{*}The first word to the ASIC is RSYNC so begins with a '0' rather than '1'.





8.2.2 Veto/No-veto

Figure 8.2 shows the dynamic veto mode in operation, 3 vetos followed by 6 no-vetos. The main lines to note are between veto_from_ccc and veto_to_tx. The general operation is: first the signal is received and flagged (on either of the two lines, veto_from_rx and nveto_from_rx), the decision is made by the veto-filter (veto_to_tx) and the response sent to the ASIC. The different responses can be seen in the single NOP for VETOs whilst the TRIGGER_FLAG_SET (0x10) is sent for NO-VETO.





8.2.3 Reset Command

Figure 8.3 shows the reset sequence being sent in response to the RESET command from the CCC, first the word arrives, then we see the delayed signal from the receiver and then the three words sent to the ASIC. The first word of the RESET command sequence is a RSYNC so the rsync_sent_flag is asserted. Also it contains a NOP command so the appropriate state change occurs and a single NOP is sent. Figure 8.4 shows the same process being triggered via the control register which is set via RDMA, i.e. the 'reset-mode'.





						0.604000 us			
ame	Value	0.0 us	1 us 0.2 us	0.5 0.4 0.4 0.4 0.5		.6 us 0.7 us	0.8us	0.9 us	, 1.0
15 ck	0			ען <mark>פועעעעעעעענענייייייייייייייייייייייייייי</mark>					
11 rst 11 dk ho acid	0_0								
and and a state	saca						raceh		4
The end from ecc	0								
The veto from ccc	0								
Le cmd_to_asic	0							•••	
15 tx_ctrl_reg_rdma_wr_en	0								
14 tx_ctrl_reg_rdma_rd_en	0		3						
. 🌄 tx_ctrl_reg_rdma_addr[31:0]	0000000					0000000			
· 🃑 tx_ctrl_reg_rdma_data_in[31:0]	0000000		00000			0000000			
• 📑 tx_ctrl_reg_rdma_data_out[31:0	0000000	(XX 0000000)	× 4000000			0000000			
lle reset_from_rx	0								
Leset_delayed	0								
La rsync_sent_flag	0								
1 readout_sent_flag	0								
🐻 downscaler_start_sent_flag	0								
1 downscaler_stop_sent_flag	0							•••	
• 📑 nvetos_sent[8:0]	00000000			000000	000				
18 II_ck	0								
TIA ILI'rst	0								
1 III l_dst_rdy	0								
7n	¢							•	_
		X1: 0.604000 us							



8.2.4 Stop with down-scaler

For this test the operation of the down-scaler for the v 1.0 ASIC checked. Prior to the yellowdashed line in figure 8.5 the clk_to_asic is running at 100 MHz, after it runs at 1 MHz. The order of events is more clearly shown in the close up, figure 8.6, which corresponds to the region between the dashed and solid lines in figure 8.5. Once the STOP command is received the first word, the down-scaler trigger, is sent to the ASIC, once this is sent the next word (in this case READ_OUT_DATA) is sent at the selected speed. Note that the READ_OUT_DATA flag is set and held for one clock at the down-scaled speed.

			677.625110 us	685.970696 us								
ame	Value	670 us	680 us		sn 002	710 us	720 us	230 us	740 us	750 us	760 us	7
11 <u>s</u> dk	0											
1 rst	0											
1kg current_state	main_stop	main_veto	ļ				main_stop					Ē
1 cmd_from_ccc	0											
🕌 veto_from_ccc	0											
1 a cmd_to_asic	1										2000	
T clk_to_asic	0											
Les stop_from_rx	0						$\left\{ \right\}$					
1ks stop_delayed	0											
1 rsync_sent_flag	0											
1 readout_sent_flag	0											
1 downscaler_start_sent_flag	0											
1 downscaler_stop_sent_flag	0											
📷 nvetos_sent[8:0]	00000000	000000000	X				00000000					Å
11. II_clk	0											8
1 III II III III III IIII IIII IIIII IIII	0											
1 🖥 II_dst_rdy	0											
1 III II Src_rdy	0											
1 III II-sof	0											
TLa Il_eof	0											
📷 II_data_out[255:0]	000000000000000000000000000000000000000	000000000000000	ļ		0	000000000000000000000000000000000000000	000000000000000000000000000000000000000	000000000000000000000000000000000000000	000000			J
1 delay_reg_rdma_clk	0											
1 delay_reg_rdma_wr_en	0											
📙 delay reg_rdma_rd_en	0											

Figure 8.5: The clk_to_asic down-scaled by a factor of 100. The section delimitated by the yellow guides is shown in figure 8.6. Note that the second word sent (after the down-scaler trigger) is the READ_OUT_START command, the flag is also down-scaled.





8.2.5 RDMA

Figure 8.7 shows a simple test of the RDMA accessible registers and BRAMs. The first four addresses of each entity is read out. These were checked by hand against the values written at the start.

		0.013615 us				
Name	Value	0.00 us , , , , , , , , , , , , , , , , , ,	0.05 us 1 1 1 1	0.10 us	0.15 us , , , , ,	0.20 us
📘 delay_reg_rdma_clk	0					
ا 🖬 delay_reg_rdma_wr_en	0					
ا 🖬 delay_reg_rdma_rd_en	T					
🕨 📲 delay_reg_rdma_addr[31:0]	4			1		
🕨 📑 delay_reg_rdma_data_in[31:0]	0			0		
🕨 😽 delay_reg_rdma_data_out[31:0]	0000054	X00000X00000		0000001		
1 tx_cmd_bram_rdma_wr_en	0					
114 tx_cmd_bram_rdma_rd_en	0					
🕨 📑 tx_cmd_bram_rdma_addr[31:0]	0	0	× 4 × 3 × 2 ×	1		
🕨 📲 tx_cmd_bram_rdma_data_in[31:	0			0		
🕨 📢 tx_cmd_bram_rdma_data_out[3	00569694	0) 00569694	00044)(000f0)(00266)	8000	1000	
📘 tx_ctrl_reg_rdma_wr_en	0					
ا 🖥 tx_ctrl_reg_rdma_rd_en	0					
🕨 📑 tx_ctrl_reg_rdma_addr[31:0]	0		4	3 X 2 X	1	
🕨 📢 tx_ctrl_reg_rdma_data_in[31:0]	0		J			
🕨 📲 tx_ctrl_reg_rdma_data_out[31:0	00000000	0000000	X00000X	(00000)	00000000	
ا 🖥 pattern_bram_rdma_wr_en	0					
📘 pattern_bram_rdma_rd_en	0					
🕨 📑 pattern_bram_rdma_addr[31:0]	0		0	(4)	2 (1	
🕨 📲 pattern_bram_rdma_data_in[31:	0)	$\Big)$	
🕨 【 pattern_bram_rdma_data_out[3	a5a5a5a5	0X	a5a5a5	001f00	0 X33333X fofofofo	
ا 🖬 pattern_id_reg_rdma_wr_en	0					
📘 pattern_id_reg_rdma_rd_en	0					
🕨 🕎 pattern_id_reg_rdma_addr[31:0	0		0			-
🕨 📲 pattern_id_reg_rdma_data_in[3:	0			0		
🕨 📲 pattern_id_reg_rdma_data_out(0000000		00000000		(02000)(04000)(03000)	02000
les stop_from_rx	0					
les start_from_r×	0					
🕨 🌄 train_from_rx[31:0]	000000000000000000000000000000000000000		000000000000000000000000000000000000000	000000000000000000000000000000000000000		
🕨 📢 bunch_id_from_rx(7:0]	0000000			0000000		
lks veto_from_rx	0					
		X1+0.013615.uc				
		55 01001010 11U				

Figure 8.7: Read back of the first 4 locations of every entity accessible via RDMA.

8.2.6 Local Link

A sample readout of the veto-log is shown in figure 8.8. The LocalLink specification requires that data transfer starts as soon as both source and destination are read. In this test all bunches were vetoed apart from the first n of each 256 bunches (where n is the n^{th} block of 256 bunches), this resulted in the top bits being filled with 0s; the first word is 0-padded and contains the header information (i.e. the train ID and the bunch pattern ID).

			677.679200 us								
Name	Value	677.66 us	677.68 us	577.70 us	677.72 us	677.74 us	677.76 us	677.78 us	677.80 us	677.82 us	677.84 u
11a ck	1										
11 <u>s</u> rst	0										
1 distribution of the second o	1										
1 current_state	main veto			main_veto			×		main_stop		
1 cmd_from_ccc	0										
11 veto_from_ccc	0										
1 cmd_to_asic	0										
1 reset_from_rx	0										
1 🛃 reset_delayed	0										
1 rsync_sent_flag	0										
1 readout_sent_flag	0										
📙 downscaler_start_sent_flag	0										
ا 🖥 downscaler_stop_sent_flag	0										
🕨 📢 nvetos_sent[8:0]	00000000					000000000					
1 <mark>1</mark> 4 II_cIk	1										
11 ILrst	0										
1 🖪 II_dst_rdy	0										
La ILsrc_rdy	0										
16 IL sof	0		۲ ۲								
La ILeof	0								1		
🕨 📲 II_data_out[255:0]	100000000000000000000000000000000000000	000000000000000000000000000000000000000	000000000000000000000000000000000000000	0 X 7fffff) 3fff	ff X 1ffffff X Offf	Fff X07Fffff X03F	<u> 166 Xolffff Xooff</u>	fff\007fff\003	Fff X 001 fff X 000	ff) 000000000000000	00000
1🛃 delay_reg_rdma_clk	0)								
TR	c										
		V1+ 677 679200 IIE									
		an 2000 2000 2000 100									


Chapter 9

Conclusion

EuXFEL is a huge project, 3.4 km of systems and sub-systems that must be kept in sync at all times in order to work. One of these many sub-systems is the LPD, one of three planned 2D pixel detectors, that will record up to $5,120 \times 1$ Mpixel images every second. Control of this detector is achieved via 16 FEMs each of which uses a command signal and clock signal, distributed to its supermodule of 512 ASICs, to maintain order. To provide a simple interface between the 2D detectors and the rest of EuXFEL the CCC uses three lines to distribute a ~99 MHz clock, a fast control signal, and a veto signal whilst receiving status information.

The LPD-CCC interface firmware runs on the FEM and provides translation from the three received signals to the two that each ASIC expects. This interface was split into three entities that dealt with receiving, vetoing and transmitting the information from the CCC. Using a rigorous testing regime each entity was developed according to the requirements of both LPD and CCC. The final designs use a combination of state-machines and BRAM to respond rapidly and accurately to each in-bound signal.

Whilst EuXFEL is several years off of operation the interface is already proving its mettle having been used in a recent beam test of the LPD detector at LCLS and in test benches for the CCC board itself.

Appendix A

EuXFEL Appendix

A.1 ASIC command words

The table A.1 provides a full list of all commands implemented by the LPD ASIC, for discussion of what they do see the LPD manual [36].

Word	Name	Word	Name
0x00	NOP	0x0D	RESET_TRIGGER_POINTER
0x01	STAND_BY	0x0E	START_WRITE_POINTER
0x02	POWER_UP	0x0F	START_TRIGGER_POINTER
0x03	ON_CHIP_RESET_DISABLE	0x10	TRIGGER_FLAG_SET
0x04	ON_CHIP_RESET_ENABLE	0x11	READ_OUT_DATA
0x05	RESET_PRE_AMP	0x12	REMOVE_RESET_PRE_AMP
0x06	RESET_GAIN_FRONT	0x13	REMOVE_RESET_GAIN_STAGE1
0x07	RESET_GAIN_BACK	0x14	REMOVE_RESET_GAIN_STAGE2
0x08	Reserved	0x15	CLOCK_DIV_SEL
0x09	TEST_MODE_D	0x16	SELF_TEST_EN
0x0A	TUNE_MODE	0x17	STOP_READ_OUT
0x0B	CLEAR_SKIP_REGISTER	0x18	RESET_STATE_MACHINE
0x0C	RESET_WRITE_POINTER	0x5A5A5	SYNC_RESET

Table A.1: ASIC command words, see [36] for a full description and recommended use of these commands.

A.2 RDMA interface

The RDMA interfaces used throughout the project have the interface seen in table A.2. The appropriate mask for the address is given in appropriate interface notes. For typical BRAMs (of $1,024 \times 32b$ words) 10b of address (9:0) were needed, this is in comparison to the registers which never needed more than 4b of address (3:0).

Name	Direction	Type	Notes
clk		sl	The RDMA clock (this can be separate from e.g. the
	in		CCC clock).
rst		$_{\rm sl}$	Reset the memory to some default state.
rd_en		$_{\rm sl}$	Enable read operations at the address.
wr_en		$_{\rm sl}$	Enable write operations at the address.
addr		slv (31:0)	The address the MSB will be masked.
data_in		slv (31:0)	Used for writing and otherwise ignored.
data_out	out	slv (31:0)	Data out, only guaranteed for read operations.

Table A.2: Standard RDMA interface.

A.3 Local Link interface

The Local Link [41] interface is used only to read out the veto log. The details of the frame composition are given in section 7.3. The interface used is minimal i.e. no optional features are used) and given in table A.3.

Name	Direction	Type	Notes
clk		sl	The LL clock (this can be separate from e.g. the
	in		CCC clock).
\mathbf{rst}		\mathbf{sl}	Abort.
dst_rdy		sl	Destination ready.
src_rdy		sl	Source ready i.e. this entity.
sof	out	$_{\mathrm{sl}}$	Start of frame flag.
eof		\mathbf{sl}	End of frame flag.
data_out		slv (255:0)	Data out.

Table A.3: Minimal local link interface as used by the veto logger.

Bibliography

- [1] Y. Kuno, M. Yoshida and A. Sato, The MuSIC project under the center of excellence of sub atomic physics, 2010.
- [2] J. Beringer et al., Phys. Rev. D 86 (2012) 010001.
- [3] T. Suzuki, D. Measday and J. Roalsvig, Phys. Rev. C 35 (1987) 2212.
- [4] B. Rossi and D.B. Hall, Phys. Rev. 59 (1941) 223.
- [5] R. Pohl et al., Nature 466 (2010) 213.
- [6] J. Coulombe et al., T2K ND280 Conceptual Design Report (The T2K Collaboration, 2005).
- [7] T. Prokscha et al., Nucl. Instr. and Meth. A 595 (2008) 317.
- [8] T. Matsuzaki et al., Nucl. Instr. and Meth. A 465 (2001) 365.
- [9] G. Marshall, Muon beams and facilities at TRIUMF, The Future of Muon Physics, pp. 226–231, Springer, 1992.
- [10] T. Behnke et al., ILC TDR and DBD, 2013.
- [11] S. Cook et al., JPCS Vol. 408, p. 012079, IOP Publishing, 2013.
- [12] K. Hatanaka et al., J. of the Korean Phys. Soc. 54 (2009) 1937.
- [13] Y. Cui et al., Conceptual Design Report for Experimental Search for Lepton Flavour Violating μe Conversion at Sensitivity of 10^{-16} with a Slow-Extracted Bunched Proton Beam (COMET) (The COMET Collaboration, 2009).
- [14] A. Ilakovac, A. Pilaftsis and L. Popov, Phys. Rev. D 87 (2013) 053014.
- [15] U. Bellgardt et al., Nucl. Phys. B 299 (1988) 1.
- [16] L. Torrisi, A. Desiderio and G. Foti, Nucl. Instr. and Meth. B 166 (2000) 664.
- [17] D. Groom, N. Mokhov and S. Striganov, At. Data. Nucl. Data Tables 78 (2001) 183.

- [18] D. Measday, Phys. Rept. 354 (2001) 243.
- [19] R.W. Huff, Ann. Phys. 16 (1961) 288.
- [20] B. Goulard and H. Primakoff, Phys. Rev C 10 (1974) 2034.
- [21] K. Ford and J. Wills, Nucl. Phys. A 35 (1962) 295.
- [22] A. de Gouvea, Nucl. Phys. B 188 (2009) 303.
- [23] W. Bertl et al., Eur. Phys. J. C 47 (2006) 337.
- [24] MEG Collaboration, J. Adam et al., Phys. Rev. Lett. 107 (2011) 171801.
- [25] S. Agostinelli et al., Nucl. Instr. and Meth. A 506 (2003) 250.
- [26] T. Roberts and D. Kaplan, PAC. IEEE, pp. 3468–3470, IEEE, 2007.
- [27] Eljen Technology, 1300 W Broadway Sweetwater TX 79556 USA, EJ-212 Plastic Scintillator, 2007.
- [28] Saint-Gobain Crystals, 104 Route de Larchant, BP 521, 77794 Nemours Cedex, France, Scintillation Products: scintillating optical fibers, 2011.
- [29] Hamamatsu, 1126-1, Ichino-cho, Higashi-ku, Hamamatsu City, 435-8558, Japan, MPPC Technical Information, 2009.
- [30] S. Ritt, P. Amaudruz and K. Olchanski, Maximum Integration Data Acquisition System (MIDAS)), 2012.
- [31] E. Motuk et al., J. Instrum. 7 (2012) C01062.
- [32] A. Blue et al., Nucl. Instr. and Meth. A 607 (2009) 55.
- [33] M. Altarelli et al., The European X-ray Free-Electron Laser (EuXFEL): Technical design report (European XFEL GmbH, 2007).
- [34] K. Tiedtke et al., New J. Phys. 11 (2009) 023029.
- [35] H. Graafsma, J. Instrum. 4 (2009) P12011.
- [36] M. Hart, D. Braga and M. Key-Charriere, LPD project ASIC manual aka XFEL ASIC.
- [37] B. Henrich et al., Nucl. Instr. and Meth. A 633 (2011) S11.
- [38] G. Lutz et al., Nucl. Instr. and Meth. A 624 (2010) 528.
- [39] J. Coughlan et al., J. Instrum. 6 (2011) C11029.
- [40] A. Koch et al., J. Instrum. 8 (2013) C11001.

[41] Xilinx, LocalLink Interface Specification, , 2.0 ed., 2005.