

---

# Monographs in Computer Science

---

*Editors*

David Gries  
Fred B. Schneider

Annabelle McIver  
Carroll Morgan

# **Abstraction, Refinement and Proof for Probabilistic Systems**

With 63 Figures

 Springer

Annabelle McIver  
Department of Computing  
Macquarie University  
Sydney, NSW 2109  
Australia  
anabel@ics.mq.edu.au

Carroll Morgan  
School of Computer Science  
and Engineering  
University of New South Wales  
Sydney, NSW 2052  
Australia  
carrollm@cse.unsw.edu.au

*Series Editors:*

David Gries  
Department of Computer Science  
Cornell University  
Upson Hall  
Ithaca, NY 14853-7501

Fred B. Schneider  
Department of Computer Science  
Cornell University  
Upson Hall  
Ithaca, NY 14853-7501

*Cover illustration:* Herman's-graph variant. See Figure 2.8.3 on p. 59.

ISBN 0-387-40115-6

Printed on acid-free paper.

© 2005 Springer Science+Business Media, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America. (MVY)

9 8 7 6 5 4 3 2 1

SPIN 10930342

springeronline.com

# Preface

Probabilistic techniques in computer programs and systems are becoming more and more widely used, for increased efficiency (as in random algorithms), for symmetry breaking (distributed systems) or as an unavoidable artefact of applications (modelling fault-tolerance). Because interest in them has been growing so strongly, stimulated by their many potential uses, there has been a corresponding increase in the study of their correctness — for the more widespread they become, the more we will depend on understanding their behaviour, and their limits, exactly.

In this volume we address that last concern, of understanding: we present a method for *rigorous reasoning about probabilistic programs and systems*. It provides an operational model — “how they work” — and an associated program logic — “how we should reason about them” — that are designed to fit together. The technique is simple in principle, and we hope that with it we will be able to increase dramatically the effectiveness of our analysis and use of probabilistic techniques in practice.

Our contribution is a probabilistic calculus that operates at the level of the program text, and it is *light-weight* in the sense that the amount of reasoning is similar in size and style to what standard assertional techniques require. In the fragment at right, for example, each potential loop entry occurs with probability  $1/2$ ; the resulting iteration establishes  $x \geq 1/2$  with probability exactly  $p$  for any  $0 \leq p \leq 1$ . It is thus an implementation of the general operation *choose with probability  $p$* , but it uses only simple tests of unbiased random bits (to implement the loop guard). It should take only a little quantitative logic to confirm that claim, and indeed we will show that just four lines of reasoning suffice.

```
x := p;  
while 1/2 do  
  x := 2x;  
  if x ≥ 1  
    then x := x - 1  
  fi  
od
```

Economy and precision of reasoning are what we have come to expect for standard programs; there is no reason we should accept less when they are probabilistic.

---

*The cover illustration comes from page 59.*

*The program fragment is adapted from Fig. 7.7.10 on page 210.*

### *Scope and applicability*

Methods for the analysis of probabilistic systems include automata, labelled transition systems, model checking and logic (*e.g.* dynamic or temporal). Our work falls into the last category: we overlay the Hoare-logic paradigm with probabilistic features imported from Markov processes, taking from each the essential characteristics required for a sound mathematical theory of refinement and proof. The aim is to accommodate modelling and analysis of both sequential and distributed probabilistic systems, and to allow — even encourage — movement between different levels of abstraction.

Our decision to focus on *logic* — and a proof system for it — was motivated by our experience with logical techniques more generally: they impose a discipline and order which promotes clarity in specifications and design; the resulting proofs can often be carried out, and checked, with astonishing conciseness and accuracy; and the calculation rules of the logic lead to an algebra that captures useful equalities and inequalities at the level of the programs themselves.

Although we rely ultimately on an operational model, we use it principally to validate the logic (and that, in turn, justifies the algebra) — direct reliance on the model’s details for individual programs is avoided if possible. (However we do not hesitate to use such details to support our intuition.) We feel that operational reasoning is more suited to the algorithmic methods of verification used by model checkers and simulation tools which can, for specific programs, answer questions that are impractical for the general approach that a logic provides.

Thus the impact of our approach is most compelling when applied to programs which are intricate either in their implementation or their design, or have generic features such as undetermined size or other parameters. They might appear as probabilistic source-level portions of large sequential programs, or as abstractions from the probabilistic modules of a comprehensive system-level design; we provide specific examples of both situations. In the latter case the ability to abstract modules’ properties has a significant effect on the overall verification enterprise.

### *Technical features*

Because we generalise the well-established assertional techniques of specifications, pre- and postconditions, there is a natural continuity of reasoning style evident in the simultaneous use of the new and the familiar approaches: the probabilistic analysis can be deployed more, or less, as the situation warrants.

A major feature is that we place probabilistic choice and abstraction together, in the same framework, without having to factor either of them out for separate treatment unless we wish to (as in fact we do in Chap. 11). This justifies the *abstraction and refinement* of our title, and is what gives

us access to the stepwise-development paradigm of standard programming where systems are “refined” from high levels of abstraction towards the low levels that include implementation detail.

As a side-effect of including abstraction, we retain its operational counterpart *demonic choice* as an explicit operator  $\sqcap$  in the cut-down probabilistic programming language *pGCL* which we use to describe our algorithms — that is, the new probabilistic choice operator  $_p\oplus$  refines demonic choice rather than replacing it. In Chap. 8 we consider angelic choice  $\sqcup$  as well, which is thus a further refinement.

Probabilistic and demonic choice together allow an elementary treatment of the hybrid that selects “with probability *at least*  $p$ ” (or similarly “*at most*  $p$ ”), an abstraction which accurately models our unavoidable ignorance of exact probabilities in real applications. Thus in our mathematical model we are able to side-step the issue of “approximate refinement.”

That is, rather than saying “this coin refines a fair coin with probability 95%,” we would say “this coin refines one which is within 5% of being fair.” This continues the simple view that either an implementation refines a specification or it does not, which simplicity is possible because we have retained the original treatment in terms of sets of behaviours: abstraction is inclusion; refinement is reverse inclusion; and demonic choice is union. In that way we maintain the important relationship between the three concepts. (Section 6.5 on pp. 169ff illustrates this geometrically.)

### *Organisation and intended readership*

The material is divided into three major parts of increasing specialisation, each of which can to a large extent be studied on its own; a fourth part contains appendices. We include a comprehensive index and extensive cross-referencing.

Definitions of notation and explanations of standard mathematical techniques are carefully given, rather than simply assumed; they appear as footnotes at their first point of use and are made visually conspicuous by using SMALL CAPITALS for the defined terms (where grammar allows). Thus in many cases a glance should be sufficient to determine whether any footnote contains a definition. In any case all definitions, whether or not in footnotes, may be retrieved by name through the index; and those with numbers are listed in order at page xvii.

Because much of the background material is separated from the main text, the need for more advanced readers to break out of the narrative should be reduced. We suggest that on first reading it is better to consult the footnotes only when there is a term that appears to require definition — otherwise the many cross-references they contain may prove distracting, as they are designed for “non-linear” browsing once the main ideas have already been assimilated.

Part I, *Probabilistic guarded commands*, gives enough introduction to the probabilistic logic to prove properties of small programs such as the one earlier, for example at the level of an undergraduate course for Formal-Methods-inclined students that explains “what to do” but not necessarily “why it is correct to do that.” These would be people who need to understand how to reason about programs (and why), but would see the techniques as intellectual tools rather than as objects of study in their own right.

We have included many small examples to serve as models for the approach (they are indexed under *Programs*), and there are several larger case studies (for example in Chap. 3).

Part II, *Semantic structures*, develops in detail the mathematics on which the probabilistic logic is built and with which it is justified. That is, whereas the earlier sections present and illustrate the new reasoning techniques, this part shows where they have come from, why they have the form they do and — crucially — why they are correct.

That last point is especially important for students intending to do research in logic and semantics, as it provides a detailed and extended worked example of the fundamental issue of proving reasoning techniques *themselves* to be correct (more accurately, “valid”), a higher-order concept than the more familiar theme of the previous part in which we presented the techniques *ex cathedra* and used them to verify particular programs.

This part would thus be suitable for an advanced final-year undergraduate or first-year graduate course, and would fit in well with other material on programming semantics. It defines and illustrates the use of many of the standard tools of the subject: lattices, approximation orders, fixed points, semantic injections and retractions *etc.*

Part III, *Advanced topics*, concentrates on more exotic methods of specification and design, in this case probabilistic temporal/modal logics. Its final chapter, for example, contains material only recently discovered and leads directly into an up-to-date research area. It would be suitable for graduate students as an introduction to this specialised research community.

Part IV includes appendices collecting material that either leads away from the main exposition — *e.g.* alternative approaches and why we have not taken them — or supports the text at a deeper level, such as some of the more detailed proofs.

It also contains a short list of algebraic laws that demonic/probabilistic program fragments satisfy, generated mainly by our needs in the examples and proofs of earlier sections. An interesting research topic would be a more systematic elaboration of that list with a view to incorporating it into probabilistic Kleene- or omega algebras for distributed computations.

Overall, readers seeking an introduction to probabilistic formal methods could follow the material in order from the beginning. Those with more experience might instead sample the first chapter from each part, which would give an indication of the scope and flavour of the approach generally.

### *Original sources*

Much of the material is based on published research, done with our colleagues, in conference proceedings and journal articles; but here it has been substantially updated and rationalised — and we have done our best to bring the almost ten years' worth of developing notation into a uniform state.

For self-contained presentations of the separate topics, and extra background, readers could consult our earlier publications as shown overleaf.

At the end of each chapter we survey the way in which our ideas have been influenced by — and in some cases adopted from — the work of other researchers, and we indicate some up-to-date developments.

### *Acknowledgements*

Our work on probabilistic models and logic was carried out initially at the University of Oxford, together with Jeff Sanders and Karen Seidel and with the support of the UK's *Engineering and Physical Sciences Research Council* (the EPSRC) during two projects led by Sanders and Morgan over the years 1994–2001.

Morgan spent sabbatical semesters in 1995–6 at the University of Utrecht, as the guest of S. Doaitse Swierstra, and at the University of Queensland and the Software Verification and Research Centre (SVRC), as the guest of David Carrington and Ian Hayes. The foundational work the EPSRC projects produced during that period — sometimes across great distances — benefited from the financial support of those institutions but especially from the academic environment provided by the hosts and by the other researchers who were receptive to our initial ideas [MMS96].

Ralph Back at Åbo Akademi hosted our group's visit to Turku for a week in 1996 during which we were able to explore our common interests in refinement and abstraction as it applied to the new domain; that led later to a three-month visit by Elena Troubitsyna from the *Turku Center for Computer Science* (TUCS), to our group in Oxford in 1997, and contributed to what has become Chap. 4 [MMT98].

David Harel was our host for a two-week visit to Israel in 1996, during which we presented our ideas and benefited from the interaction with researchers there.



## Chapters' dependence on original sources

Chapter 1	<i>see</i>	[MM99b, SMM, MMS00]
Chapter 2	<i>see</i>	[Mor96, MMS00]
Chapter 3	<i>see</i>	[MM99b]
Chapter 4	<i>see</i>	[MMT98]
Chapter 5	<i>see</i>	[MMS96]
Chapter 6		<i>is new material</i>
Chapter 7	<i>see</i>	[Mor96, MM01b]
Chapter 8	<i>see</i>	[MM01a]
Chapter 9	<i>see</i>	[MM97]
Chapter 10	<i>see</i>	[MM99a]
Chapter 11	<i>see</i>	[MM02]

The sources listed opposite are in chronological order of writing, thus giving roughly the logical evolution of the ideas.

Subsequently we have continued to work with Sanders and with Ken Robinson, Thai Son Hoang and Zhendong Jin, supported by the *Australian Research Council* (ARC) over the (coming) years 2001–8 in their *Large Grant* and *Discovery* programmes, at the Universities of Macquarie and of New South Wales.

Joe Hurd from the Computer Laboratory at Cambridge University visited us in 2002, with financial assistance from Macquarie University; and Orieta Celiku was supported by TUCS when she visited in 2003. Both worked under McIver's direction on the formalisation of *pGCL*, and its logic, in the mechanised logic *HOL*.

Hoang, Jin and especially Eric Martin have helped us considerably with their detailed comments on the typescript; also Ralph Back, Ian Hayes, Michael Huth, Quentin Miller and Wayne Wheeler have given us good advice. Section B.1 on the algebraic laws satisfied by probabilistic programs has been stimulated by the work (and the critical eyes) of Steve Schneider and his colleagues at Royal Holloway College in the U.K.

We thank the members of IFIP Working Groups 2.1 and 2.3 for their many comments and suggestions.

*LRI* Paris  
May 2004

Annabelle McIver  
Carroll Morgan

## List of sources in order of writing

- [MMS96] C.C. Morgan, A.K. McIver, and K. Seidel. Probabilistic predicate transformers. *ACM Transactions on Programming Languages and Systems*, 18(3):325–53, May 1996.
- [Mor96] C.C. Morgan. Proof rules for probabilistic loops. In He Jifeng, John Cooke, and Peter Wallis, editors, *Proceedings of the BCS-FACS 7th Refinement Workshop*, Workshops in Computing. Springer-Verlag, July 1996.
- [MM01b] A.K. McIver and C.C. Morgan. Partial correctness for probabilistic programs. *Theoretical Computer Science*, 266(1–2):513–41, 2001.
- [SMM] K. Seidel, C.C. Morgan, and A.K. McIver. Probabilistic imperative programming: a rigorous approach. Extended abstract appears in Groves and Reeves [GR97], pages 1–2.
- [MMT98] A.K. McIver, C.C. Morgan, and E. Troubitsyna. The probabilistic steam boiler: a case study in probabilistic data refinement. In J. Grundy, M. Schwenke, and T. Vickers, editors, *Proc. International Refinement Workshop, ANU, Canberra*, Discrete Mathematics and Computer Science, pages 250–65. Springer-Verlag, 1998.
- [MM01a] A.K. McIver and C.C. Morgan. Demonic, angelic and unbounded probabilistic choices in sequential programs. *Acta Informatica*, 37:329–54, 2001.
- [MM99b] C.C. Morgan and A.K. McIver. *pGCL*: Formal reasoning for random algorithms. *South African Computer Journal*, 22, March 1999.
- [MM97] C.C. Morgan and A.K. McIver. A probabilistic temporal calculus based on expectations. In Groves and Reeves [GR97], pages 4–22.
- [MM99a] C.C. Morgan and A.K. McIver. An expectation-based model for probabilistic temporal logic. *Logic Journal of the IGPL*, 7(6):779–804, 1999.
- [MMS00] A.K. McIver, C.C. Morgan, and J.W. Sanders. Probably Hoare? Hoare probably! In J.W. Davies, A.W. Roscoe, and J.C.P. Woodcock, editors, *Millennial Perspectives in Computer Science*, Cornerstones of Computing, pages 271–82. Palgrave, 2000.
- [MM02] A.K. McIver and C.C. Morgan. Games, probability and the quantitative  $\mu$ -calculus qMu. In *Proc. LPAR*, volume 2514 of *LNAI*, pages 292–310. Springer-Verlag, 2002.

# Contents

<b>Preface</b>	<b>v</b>
<b>List of definitions <i>etc.</i></b>	<b>xvii</b>
<hr/>	
<b>Part I Probabilistic guarded commands</b>	<b>1</b>
<hr/>	
<b>1 Introduction to <math>pGCL</math></b>	<b>3</b>
1.1 Sequential program logic . . . . .	4
1.2 The programming language $pGCL$ . . . . .	7
1.3 An informal computational model for $pGCL$ . . . . .	11
1.4 Behind the scenes: elementary probability theory . . . . .	16
1.5 Basic syntax and semantics of $pGCL$ . . . . .	18
1.6 Healthiness and algebra for $pGCL$ . . . . .	28
1.7 Healthiness example: modular reasoning . . . . .	32
1.8 Interaction of probabilistic- and demonic choice . . . . .	34
1.9 Summary . . . . .	35
Chapter notes . . . . .	36
<b>2 Probabilistic loops: invariants and variants</b>	<b>37</b>
2.1 Introduction: loops via recursion . . . . .	38
2.2 Probabilistic invariants . . . . .	39
2.3 Probabilistic termination . . . . .	40
2.4 Invariance and termination together: the loop rule . . . . .	42
2.5 Three examples of probabilistic loops . . . . .	44
2.6 The Zero-One Law for termination . . . . .	53
2.7 Probabilistic variant arguments for termination . . . . .	54

2.8	Termination example: self-stabilisation . . . . .	56
2.9	Uncertain termination . . . . .	61
2.10	Proper post-expectations . . . . .	63
2.11	Bounded <i>vs.</i> unbounded expectations . . . . .	68
2.12	Informal proof of the loop rule . . . . .	74
	Chapter notes . . . . .	77
<b>3</b>	<b>Case studies in termination</b>	<b>79</b>
3.1	Rabin’s choice coordination . . . . .	79
3.2	The dining philosophers . . . . .	88
3.3	The general random “jump” . . . . .	99
	Chapter notes . . . . .	105
<b>4</b>	<b>Probabilistic data refinement: the steam boiler</b>	<b>107</b>
4.1	Introduction: refinement of datatypes . . . . .	107
4.2	Data refinement and simulations . . . . .	108
4.3	Probabilistic datatypes: a worked example . . . . .	110
4.4	A safety-critical application: the steam boiler . . . . .	117
4.5	Summary . . . . .	123
	Chapter notes . . . . .	124
<hr/>		
	<b>Part II Semantic structures</b>	<b>127</b>
<hr/>		
<b>5</b>	<b>Theory for the demonic model</b>	<b>129</b>
5.1	Deterministic probabilistic programs . . . . .	130
5.2	The sample space, random variables and expectations . . . . .	133
5.3	Probabilistic deterministic transformers . . . . .	135
5.4	Relational demonic semantics . . . . .	137
5.5	Regular transformers . . . . .	141
5.6	Healthiness conditions for probabilistic programs . . . . .	145
5.7	Characterising regular programs . . . . .	149
5.8	Complementary and consistent semantics . . . . .	154
5.9	Review: semantic structures . . . . .	157
	Chapter notes . . . . .	164
<b>6</b>	<b>The geometry of probabilistic programs</b>	<b>165</b>
6.1	Embedding distributions in Euclidean space . . . . .	166
6.2	Standard deterministic programs . . . . .	166
6.3	Probabilistic deterministic programs . . . . .	167
6.4	Demonic programs . . . . .	168
6.5	Refinement . . . . .	169

6.6	Nontermination and sub-distributions . . . . .	171
6.7	Expectations and touching planes . . . . .	172
6.8	Refinement seen geometrically . . . . .	174
6.9	The geometry of the healthiness conditions . . . . .	175
6.10	Sublinearity corresponds to convexity . . . . .	176
6.11	Truncated subtraction . . . . .	177
6.12	A geometrical proof for recursion . . . . .	177
	Chapter notes . . . . .	180
<b>7</b>	<b>Proved rules for probabilistic loops</b>	<b>181</b>
7.1	Introduction . . . . .	182
7.2	Partial loop correctness . . . . .	184
7.3	Total loop correctness . . . . .	186
7.4	Full proof of the loop rule . . . . .	189
7.5	Probabilistic variant arguments . . . . .	191
7.6	Finitary completeness of variants . . . . .	193
7.7	Do-it-yourself semantics . . . . .	195
7.8	Summary . . . . .	214
	Chapter notes . . . . .	216
<b>8</b>	<b>The transformer hierarchy</b>	<b>217</b>
8.1	Introduction . . . . .	217
8.2	Infinite state spaces . . . . .	219
8.3	Deterministic programs . . . . .	221
8.4	Demonic programs . . . . .	224
8.5	Angelic programs . . . . .	227
8.6	Standard programs . . . . .	231
8.7	Summary . . . . .	238
	Chapter notes . . . . .	242
<hr/>		
<b>Part III</b>	<b>Advanced topics</b>	<b>243</b>
<b>9</b>	<b>Quantitative temporal logic: an introduction</b>	<b>245</b>
9.1	Modal and temporal logics . . . . .	245
9.2	Standard temporal logic: a review . . . . .	247
9.3	Quantitative temporal logic . . . . .	252
9.4	Temporal operators as games . . . . .	254
9.5	Summary . . . . .	262
	Chapter notes . . . . .	263

<b>10 The quantitative algebra of <math>qTL</math></b>	<b>265</b>
10.1 The role of algebra . . . . .	265
10.2 Quantitative temporal expectations . . . . .	268
10.3 Quantitative temporal algebra . . . . .	277
10.4 Examples: demonic random walkers and stumblers . . . . .	283
10.5 Summary . . . . .	289
Chapter notes . . . . .	291
<b>11 The quantitative modal <math>\mu</math>-calculus <math>qM\mu</math>, and games</b>	<b>293</b>
11.1 Introduction to the $\mu$ -calculus . . . . .	293
11.2 <i>Quantitative</i> $\mu$ -calculus for probability . . . . .	295
11.3 Logical formulae and transition systems . . . . .	295
11.4 Two interpretations of $qM\mu$ . . . . .	298
11.5 Example . . . . .	301
11.6 Proof of equivalence of interpretations . . . . .	304
11.7 Summary . . . . .	308
Chapter notes . . . . .	310
<hr/>	
<b>Part IV Appendices, bibliography and indexes</b>	<b>311</b>
<hr/>	
<b>A Alternative approaches</b>	<b>313</b>
A.1 Probabilistic Hoare-triples . . . . .	313
A.2 A programming logic of distributions . . . . .	316
<b>B Supplementary material</b>	<b>321</b>
B.1 Some algebraic laws of probabilistic programs . . . . .	321
B.2 Loop rule for <i>demonic</i> iterations . . . . .	328
B.3 Further facts about probabilistic <i>wp</i> and <i>wlp</i> . . . . .	331
B.4 Infinite state spaces . . . . .	332
B.5 Linear-programming lemmas . . . . .	341
B.6 Further lemmas for <i>eventually</i> . . . . .	342
<b>Bibliography</b>	<b>345</b>
<b>Index of citations</b>	<b>357</b>
<b>General index</b>	<b>361</b>

# List of definitions *etc.*

Definitions and notations for standard mathematical concepts are given in footnotes, rather than in the main text, so that they do not interrupt the flow. They can be found directly, via the index, where they are indicated by bold-face page references. Thus for example at “fixed point, least” we find that the *least fixed-point* and its associated “ $\mu$  notation” are defined in Footnote 33 on p. 21, and a second form of the  $\mu$  notation is defined in Footnote 32 on p. 102.

---

Definition 1.2.1	9	Lemma 2.10.2	66
Figure 1.3.1	14	Figure 2.11.1	69
Figure 1.3.2	15	Figure 2.11.2	71
Figure 1.5.1	22	Figure 2.11.3	73
Definition 1.5.2	24	Figure 2.11.4	75
Figure 1.5.3	26	Figure 3.1.1	83
Definition 1.6.1	28	Figure 3.2.1	90
Definition 1.6.2	29	Figure 3.2.2	91
Figure 1.6.3	33	Figure 3.2.3	92
Lemma 1.7.1	33	Lemma 3.2.4	94
Definition 2.2.1	39	Figure 3.2.5	98
Lemma 2.4.1	42	Figure 3.3.1	101
Figure 2.5.1	45	Definition 4.2.1	109
Figure 2.5.2	47	Definition 4.2.2	110
Figure 2.5.3	52	Figure 4.3.1	111
Lemma 2.6.1	54	Figure 4.3.2	111
Lemma 2.7.1	55	Figure 4.3.3	115
Figure 2.8.1	57	Figure 4.3.4	116
Figure 2.8.2	58	Lemma 4.3.5	117
Figure 2.8.3	59	Figure 4.4.1	118
Figure 2.10.1	65	Figure 4.4.2	119

Figure 4.4.3	120	Figure 6.7.1	173
Figure 4.4.4	121	Figure 6.10.1	176
Figure 4.4.5	122	Lemma 6.12.1	178
Lemma 4.4.6	122	Definition 7.2.1	184
Lemma 4.4.7	122	Lemma 7.2.2	185
Definition 5.1.1	130	Lemma 7.3.1	186
Definition 5.1.2	131	Lemma 7.3.2	187
Definition 5.1.3	131	Theorem 7.3.3	188
Definition 5.1.4	132	Lemma 7.5.1	191
Definition 5.1.5	132	Lemma 7.6.1	193
Definition 5.2.1	134	Theorem 7.6.2	194
Definition 5.2.2	134	Definition 7.7.1	196
Definition 5.2.3	135	Definition 7.7.2	198
Definition 5.4.1	138	Lemma 7.7.3	200
Definition 5.4.2	138	Lemma 7.7.4	202
Definition 5.4.3	139	Theorem 7.7.5	202
Definition 5.4.4	139	Lemma 7.7.6	203
Definition 5.4.5	140	Lemma 7.7.7	204
Definition 5.4.6	140	Lemma 7.7.8	205
Definition 5.4.7	140	Figure 7.7.9	208
Lemma 5.4.8	141	Figure 7.7.10	210
Figure 5.5.1	143	Figure 7.7.11	211
Definition 5.5.2	144	Definition 8.3.1	222
Definition 5.5.3	145	Definition 8.3.2	222
Definition 5.6.1	146	Lemma 8.3.3	222
Lemma 5.6.2	146	Lemma 8.3.4	223
Lemma 5.6.3	147	Theorem 8.3.5	223
Lemma 5.6.4	147	Figure 8.3.6	224
Lemma 5.6.5	147	Figure 8.4.2	226
Lemma 5.6.6	147	Figure 8.4.3	226
Figure 5.6.7	148	Theorem 8.4.1	226
Lemma 5.6.8	148	Definition 8.5.1	227
Definition 5.7.1	149	Definition 8.5.2	228
Lemma 5.7.2	150	Lemma 8.5.3	228
Lemma 5.7.3	151	Definition 8.5.4	229
Lemma 5.7.4	151	Lemma 8.5.5	229
Lemma 5.7.5	153	Lemma 8.5.6	229
Lemma 5.7.6	153	Lemma 8.5.7	229
Theorem 5.7.7	153	Theorem 8.5.8	230
Lemma 5.8.1	154	Figure 8.5.9	231
Lemma 5.8.2	156	Figure 8.5.10	232
Figure 5.9.1	158	Definition 8.6.1	232
Figure 5.9.2	160	Definition 8.6.2	232
Figure 6.2.1	167	Lemma 8.6.3	233
Figure 6.3.1	167	Definition 8.6.4	234
Figure 6.4.1	168	Lemma 8.6.5	234
Figure 6.4.2	169	Definition 8.6.6	234
Figure 6.5.1	169	Definition 8.6.7	235
Figure 6.6.1	171	Lemma 8.6.8	235



Lemma 8.6.9	235	Lemma 11.4.1	299
Lemma 8.6.10	235	Lemma 11.4.2	300
Lemma 8.6.11	236	Figure 11.5.1	302
Lemma 8.6.12	236	Definition 11.6.1	305
Corollary 8.6.13	237	Lemma 11.6.2	306
Theorem 8.6.14	237	Lemma 11.6.3	307
Lemma 8.6.15	237	Lemma 11.6.4	307
Lemma 8.6.16	237	Theorem 11.6.5	308
Figure 8.7.1	240	Lemma 11.6.6	308
Figure 8.7.2	241	Figure A.1.1	315
Definition 9.2.1	251	Figure A.1.2	315
Definition 9.2.2	251	Lemma B.2.1	329
Definition 9.2.3	252	Theorem B.2.2	329
Definition 9.2.4	252	Fact B.3.1	331
Definition 9.3.1	252	Fact B.3.2	331
Definition 9.3.2	253	Fact B.3.3	331
Definition 9.3.3	253	Fact B.3.4	331
Definition 9.3.4	253	Fact B.3.5	331
Lemma 9.4.1	256	Fact B.3.6	331
Lemma 9.4.2	257	Definition B.4.1	332
Lemma 9.4.3	257	Lemma B.4.2	333
Lemma 9.4.4	257	Lemma B.4.3	333
Lemma 9.4.5	258	Lemma B.4.4	334
Lemma 9.4.6	258	Lemma B.4.5	335
Theorem 9.4.7	260	Lemma B.4.6	335
Figure 10.1.1	266	Lemma B.4.7	336
Lemma 10.1.2	266	Lemma B.4.8	337
Figure 10.2.1	270	Lemma B.4.9	337
Figure 10.2.2	272	Lemma B.5.1	341
Figure 10.2.3	274	Lemma B.5.2	341
Figure 10.3.1	278	Lemma B.5.3	342
Definition 10.3.2	278	Lemma B.6.1	342
Definition 10.3.3	279	Lemma B.6.2	342
Lemma 10.3.4	279	Lemma B.6.3	343
Figure 10.3.5	281	Lemma B.6.4	343
Figure 10.3.6	282	Lemma B.6.5	343

# Part I

## Probabilistic guarded commands and their refinement logic

<b>1</b>	<b>Introduction to <math>pGCL</math></b>	<b>3</b>
1.1	Sequential program logic . . . . .	4
1.2	The programming language $pGCL$ . . . . .	7
1.3	An informal computational model for $pGCL$ . . . . .	11
1.4	Behind the scenes: elementary probability theory . . . . .	16
1.5	Basic syntax and semantics of $pGCL$ . . . . .	18
1.6	Healthiness and algebra for $pGCL$ . . . . .	28
1.7	Healthiness example: modular reasoning . . . . .	32
1.8	Interaction of probabilistic- and demonic choice . . . . .	34
1.9	Summary . . . . .	35
	Chapter notes . . . . .	36

<b>2</b>	<b>Probabilistic loops: invariants and variants</b>	<b>37</b>
2.1	Introduction: loops via recursion . . . . .	38
2.2	Probabilistic invariants . . . . .	39
2.3	Probabilistic termination . . . . .	40
2.4	Invariance and termination together: the loop rule . . . . .	42
2.5	Three examples of probabilistic loops . . . . .	44
2.6	The Zero-One Law for termination . . . . .	53
2.7	Probabilistic variant arguments for termination . . . . .	54
2.8	Termination example: self-stabilisation . . . . .	56
2.9	Uncertain termination . . . . .	61
2.10	Proper post-expectations . . . . .	63
2.11	Bounded <i>vs.</i> unbounded expectations . . . . .	68
2.12	Informal proof of the loop rule . . . . .	74
	Chapter notes . . . . .	77
<b>3</b>	<b>Case studies in termination</b>	<b>79</b>
3.1	Rabin’s choice coordination . . . . .	79
3.2	The dining philosophers . . . . .	88
3.3	The general random “jump” . . . . .	99
	Chapter notes . . . . .	105
<b>4</b>	<b>Probabilistic data refinement: the steam boiler</b>	<b>107</b>
4.1	Introduction: refinement of datatypes . . . . .	107
4.2	Data refinement and simulations . . . . .	108
4.3	Probabilistic datatypes: a worked example . . . . .	110
4.4	A safety-critical application: the steam boiler . . . . .	117
4.5	Summary . . . . .	123
	Chapter notes . . . . .	124