

Reversibility for Efficient Computing

(Manuscript, Based on Ph.D. Thesis)

Michael P. Frank

CISE Department
University of Florida
<http://www.cise.ufl.edu/~mpf>

Working version of
Dec. 20, 1999

This document is available online through
<http://www.cise.ufl.edu/~mpf/manuscript>

Abstract

Today's computers are based on *irreversible* logic devices, which have been known to be fundamentally energy-inefficient for several decades. Recently, alternative *reversible* logic technologies have improved rapidly, and are now becoming practical.

In traditional models of computation, pure reversibility seems to decrease overall computational efficiency; I provide a proof to this effect. However, traditional models ignore important physical constraints on information processing.

This thesis gives the first analysis demonstrating that in a realistic model of computation that accounts for thermodynamic issues, as well as other physical constraints, the judicious use of reversible computing can strictly *increase* asymptotic computational efficiency, as machine sizes increase. I project real benefits for supercomputing at a large (but achievable) scale in the fairly near term. And with proposed future computing technologies, I show that reversibility will benefit computing at all scales.

Next, the thesis demonstrates that reversible computing techniques do not make computer design much more difficult. I describe how to design asymptotically efficient processors using an “adiabatic” reversible electronic logic technology that can be built with today's microprocessor fabrication processes. I describe a simple universal reversible parallel processor chip that our group recently fabricated, and a reversible instruction set for a more traditional RISC-style uniprocessor.

Finally, I describe techniques for programming reversible computers. I present a high-level language and a compiler suitable for coding efficient reversible algorithms, and I describe a variety of example algorithms, including efficient reversible sorting, searching, arithmetic, matrix, and graph algorithms. As an example application, I present a linear-time, constant-space reversible program for simulating the Schrödinger wave equation of quantum mechanics.

Acknowledgments

First and foremost, I am enduringly grateful to my advisor Tom Knight and also to Norm Margolus, for taking me into their fold, for their extensive and wise guidance in all areas of this research, for the many things I learned from them, and for providing an exceptionally supportive and stimulating work environment. Thanks very much also to the other readers Gill Pratt and Gerry Sussman, for kindly agreeing to serve on my thesis committee and lend their expertise to the evaluation of this work.

I would also like to thank the Pendulum group's head graduate student, Carlin Vieri, for inventing the Pendulum architecture, which provided the context for much of this work; for taking it upon himself to do most of the project's administrative chores; and for the fun we had working together on the nuts and bolts of various aspects of the project such as the Pendulum ISA and the FLATTOP chip. (Also, thanks for all the Nolios!)

I should also thank a number of other MIT students: Matt DeBergalis created many helpful software tools, such as the Pendulum assembler and emulator. Josie Ammer assisted with chip design, and made important contributions to the theory and algorithms work. Scott Rixner and Nicole Love were my primary design partners on the Tick and FLATTOP processors, respectively. Matt Becker often popped into my office for a friendly greeting and an interesting conversation.

Prof. Michael Sipser deserves thanks for his informal but very helpful guidance during the development of the proof in §3.4. I am also grateful to Alain Tapp and Pierre McKenzie of the University of Montreal, for providing much helpful feedback on that result, and for inviting me to visit their lab. Thanks also to Prof. Tom Leighton for communicating the reversible shortest-path algorithm, and to both him and Prof. Charles Leiserson for feedback on the scaling results of chapter 6.

Finally, I would like to thank Carl R. Witty and Warren D. Smith for helpful editorial comments on drafts of this document, my family and my wife Lori for their love, and all my friends throughout my years at MIT for their companionship and encouragement.

This research was supported by DARPA (the Defense Advanced Research Projects Agency of the United States of America) as part of the Scalable Computing Systems research program [43], under contract number DABT63-95-C-0130.

to my father, Patrick Gene Frank

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction and background | 17 |
| 1.1 | What this thesis is about | 17 |
| 1.2 | Motivation | 18 |
| 1.3 | Brief history of reversible computing | 20 |
| 1.3.1 | Early thermodynamics of computation | 20 |
| 1.3.2 | Development of reversible models of computation | 21 |
| 1.3.3 | Development of physically reversible logic devices | 21 |
| 1.3.4 | Previous reversible computing theory | 22 |
| 1.3.5 | Optimal scaling of physical machines | 22 |
| 1.3.6 | Programming reversible machines | 23 |
| 1.4 | Major contributions of this thesis | 23 |
| 1.5 | Overview of thesis chapters | 24 |
| 1.6 | Overall message of thesis | 27 |
| | | |
| I | Foundations of reversible computing | 29 |
| | | |
| 2 | Physical constraints on computation | 31 |
| 2.1 | Propagation speed limits | 31 |
| 2.2 | Information density limits | 32 |
| 2.2.1 | Entropy bounds from black hole physics. | 33 |
| 2.2.2 | Entropy bounds for a photon gas. | 34 |
| 2.2.3 | More reasonable mass densities. | 36 |
| 2.2.4 | More reasonable temperatures. | 36 |
| 2.3 | Information flux rate limits | 38 |
| 2.4 | Computation rate limits | 40 |
| 2.5 | Reversibility of physics | 41 |
| 2.5.1 | Physical reversibility and information erasure | 41 |
| 2.5.2 | Reversibility, entropy, and the second law | 43 |
| 2.5.3 | Entropy and energy | 46 |

| | | |
|----------|---|-----------|
| 2.5.4 | Logical irreversibility and energy dissipation | 46 |
| 2.6 | Quantum computation | 48 |
| 2.7 | Physical constraints—conclusion | 49 |
| 3 | Reversible computing theory | 51 |
| 3.1 | Models of computation | 52 |
| 3.1.1 | Computability | 53 |
| 3.2 | Computational complexity and efficiency | 54 |
| 3.2.1 | Computational efficiency vs. computational complexity | 54 |
| 3.2.2 | Characterizing computational complexity | 54 |
| 3.2.3 | Complexity classes | 59 |
| 3.3 | Review of existing reversible computing theory | 59 |
| 3.3.1 | Reversible models of computation | 59 |
| 3.3.2 | Computability in reversible models | 59 |
| 3.3.3 | Time complexity in reversible models | 61 |
| 3.3.4 | Reversible entropic complexity | 62 |
| 3.3.5 | Reversible space complexity | 63 |
| 3.3.6 | Miscellaneous developments | 67 |
| 3.4 | Reversible vs. irreversible space-time complexity | 67 |
| 3.4.1 | General definitions | 69 |
| 3.4.2 | Oracle results | 72 |
| 3.4.3 | Non-relativized separation | 84 |
| 3.4.4 | Decompression algorithm | 86 |
| 3.4.5 | Can this proof be carried farther? | 87 |
| 3.5 | Summary of reversible complexity results for traditional models | 87 |
| 4 | Quantum computation | 91 |
| 4.1 | Some fundamental quantum concepts | 92 |
| 4.2 | Quantum complexity theory | 95 |
| 4.3 | Outline of Shor's Algorithm | 97 |
| 4.4 | Important open problems | 100 |
| 4.4.1 | Is $\mathbf{BQP} \supseteq \mathbf{NP}$? | 100 |
| 4.4.2 | Is $\mathbf{BQP} \supset \mathbf{BPP}$? | 102 |
| 4.4.3 | Does quantum computing scale? | 103 |
| 4.4.4 | How do we build it, physically? | 103 |
| 4.5 | Summary of quantum computation | 105 |

| | | |
|-----------|---|------------|
| 5 | Ultimate physical models of computation | 107 |
| 5.1 | What is a model of computation? | 108 |
| 5.2 | Existing models of computation | 110 |
| 5.3 | Problems with the existing models | 112 |
| 5.4 | Some candidates for an ultimate model | 115 |
| 5.4.1 | The reversible 3-D mesh (R3M) model | 116 |
| 5.4.2 | The ballistic 3-D mesh (B3M) model | 117 |
| 5.4.3 | The quantum 3-D mesh (Q3M) model | 117 |
| 5.5 | A “tight” Church’s thesis | 118 |
| 5.6 | Ultimate computational complexity | 119 |
| 5.7 | Summary of discussion of ultimate models | 120 |
| 6 | Reversibility and physical scaling laws | 121 |
| 6.1 | Types of architectures studied | 122 |
| 6.1.1 | Shared properties | 122 |
| 6.1.2 | Fully irreversible architecture | 122 |
| 6.1.3 | Time-proportionately reversible architecture | 123 |
| 6.1.4 | Ballistic reversible architecture | 123 |
| 6.2 | Analyses under various physical costs | 124 |
| 6.2.1 | Entropy cost | 125 |
| 6.2.2 | Area-time product | 127 |
| 6.2.3 | Time cost | 131 |
| 6.2.4 | Spacetime cost | 137 |
| 6.2.5 | Mass-time product | 139 |
| 6.2.6 | (Area + mass) \times time | 139 |
| 6.2.7 | Entropy + mass-time | 139 |
| 6.3 | Generalizing the results | 139 |
| 6.3.1 | Speedups for irreversible computations on reversible machines | 140 |
| 6.4 | Summary of scaling results | 141 |
| II | Engineering reversible computational systems | 145 |
| 7 | Adiabatic circuits | 147 |
| 7.1 | Maximizing the efficiency of iCMOS | 148 |
| 7.1.1 | Basic iCMOS review | 148 |
| 7.1.2 | iCMOS entropy generation. | 151 |
| 7.1.3 | The SIA semiconductor roadmap | 156 |
| 7.1.4 | Minimizing permanent energy dissipation in iCMOS | 162 |
| 7.1.5 | Maximizing per-area processing rate for iCMOS | 162 |

| | | |
|----------|---|------------|
| 7.1.6 | Maximizing iCMOS cost-efficiency | 165 |
| 7.2 | Historical development of adiabatic circuits | 166 |
| 7.3 | A comment on terminology | 170 |
| 7.4 | Basic principles of adiabatic circuits | 171 |
| 7.5 | The SCRL technique | 174 |
| 7.5.1 | Basic SCRL components | 174 |
| 7.5.2 | SCRL pipelines | 179 |
| 7.5.3 | Timing disciplines | 180 |
| 7.6 | SCRL circuit analyses | 180 |
| 7.6.1 | A simple SCRL model for analysis | 183 |
| 7.6.2 | Switching losses as a function of technology parameters | 184 |
| 7.6.3 | Minimizing the sum of switching and leakage energy | 190 |
| 7.6.4 | Fixing a problematic case for plain SCRL | 197 |
| 7.7 | Experimental SCRL Circuits | 199 |
| 7.7.1 | The Billiard Ball Model | 201 |
| 7.7.2 | The Billiard Ball Model Cellular Automaton | 203 |
| 7.7.3 | Logic minimization | 204 |
| 7.7.4 | FlatTop array design | 205 |
| 7.7.5 | Minimum energy estimation | 207 |
| 7.8 | Resonant power supply techniques | 208 |
| 7.9 | Scaling SCRL to future technology generations | 209 |
| 7.10 | Mostly reversible computation | 211 |
| 7.11 | Adiabatic Circuits—Conclusion | 212 |
| 8 | Future reversible device technologies | 215 |
| 8.1 | Cooling technologies | 215 |
| 8.2 | Irreversible device technologies | 217 |
| 8.3 | Reversible technologies | 218 |
| 8.4 | Future device technologies—Conclusion | 222 |
| 9 | Design and programming of reversible processors | 225 |
| 9.1 | Context of this work | 226 |
| 9.1.1 | Previous reversible architectures | 226 |
| 9.1.2 | Pendulum architecture | 226 |
| 9.2 | Reversible instruction set architectures | 227 |
| 9.2.1 | Asymptotic efficiency | 227 |
| 9.2.2 | Use of paired branches | 227 |
| 9.2.3 | Reversible logic/arithmetic operations | 230 |
| 9.2.4 | Data transfer operations | 231 |

| | | |
|------------|---|------------|
| 9.2.5 | Hardware-guaranteed reversibility | 231 |
| 9.3 | Simple example PISA program: Multiplication algorithm | 232 |
| 9.3.1 | Discussion | 235 |
| 9.4 | Reversible programming languages | 236 |
| 9.4.1 | General issues | 236 |
| 9.4.2 | “R,” a reversible language | 238 |
| 9.4.3 | The R compiler | 239 |
| 9.5 | Reversible algorithms | 239 |
| 9.5.1 | Sorting | 240 |
| 9.5.2 | Arithmetic | 240 |
| 9.5.3 | Matrices | 240 |
| 9.5.4 | Searches | 241 |
| 9.5.5 | Graph problems | 241 |
| 9.5.6 | Physical simulations | 241 |
| 9.6 | Operating system issues | 244 |
| 9.7 | Parallelism | 246 |
| 10 | Alternative applications for reversibility | 249 |
| 10.1 | Auditable/verifiable/trustable computation | 249 |
| 10.1.1 | Detecting transient errors | 250 |
| 10.1.2 | Logging or limiting effects of unwelcome intrusions | 251 |
| 10.2 | Program debugging | 252 |
| 10.3 | Transaction processing and database rollback | 253 |
| 10.4 | Speculative execution in multiprocessors | 253 |
| 10.5 | Numerical stability in physics simulations | 254 |
| 10.6 | Alternative applications: Conclusion | 254 |
| 11 | Conclusion and Future Work | 257 |
| 11.1 | Summary of Contributions | 257 |
| 11.2 | Major areas for future research | 259 |
| 11.3 | Final words | 263 |
| III | Appendices | 265 |
| A | FlatTop processor schematics and layouts | 267 |
| A.1 | High-level blocks | 267 |
| A.2 | Detailed gate schematics | 269 |
| A.3 | Cell layout | 269 |

| | | |
|----------|---|------------|
| B | The Pendulum instruction set architecture (PISA) | 275 |
| B.1 | Overall organization | 275 |
| B.2 | List of Instructions | 279 |
| B.3 | Arithmetic/logical ops | 280 |
| B.4 | Ordinary branches | 288 |
| B.5 | Special instructions | 291 |
| C | The R reversible programming language | 295 |
| C.1 | Introduction | 295 |
| C.2 | What type of language is R? | 296 |
| C.3 | Overview of R Syntax | 296 |
| C.4 | User-level Constructs | 296 |
| C.4.1 | Program Structure | 297 |
| C.4.2 | Control Structure | 298 |
| C.4.3 | Variables | 301 |
| C.4.4 | Data Modification | 302 |
| C.4.5 | Expressions | 304 |
| C.4.6 | Static Data | 307 |
| C.4.7 | Input/Output | 308 |
| C.5 | Example Programs | 309 |
| C.6 | Compiler Internals | 309 |
| C.7 | Conclusions | 309 |
| D | The R language compiler | 311 |
| D.1 | R Compiler User's Guide | 311 |
| D.2 | Compilation technique | 312 |
| D.3 | Internal compiler constructs | 313 |
| D.3.1 | Intermediate-level internal constructs | 313 |
| D.3.2 | Low-level constructs | 321 |
| D.4 | Compiler LISP source code | 333 |
| D.4.1 | loader.lisp | 335 |
| D.4.2 | util.lisp | 335 |
| D.4.3 | infrastructure.lisp | 336 |
| D.4.4 | location.lisp | 341 |
| D.4.5 | environment.lisp | 342 |
| D.4.6 | regstack.lisp | 346 |
| D.4.7 | variables.lisp | 347 |
| D.4.8 | branches.lisp | 350 |
| D.4.9 | expression.lisp | 354 |
| D.4.10 | clike.lisp | 359 |

| | | |
|----------|---|------------|
| D.4.11 | print.lisp | 362 |
| D.4.12 | controlflow.lisp | 362 |
| D.4.13 | subroutines.lisp | 364 |
| D.4.14 | staticdata.lisp | 367 |
| D.4.15 | program.lisp | 367 |
| D.4.16 | library.lisp | 368 |
| D.4.17 | files.lisp | 369 |
| D.4.18 | test.lisp | 370 |
| E | Reversible Schrödinger wave simulation | 377 |
| E.1 | Derivation of discrete update rule | 377 |
| E.2 | Reversible C implementation | 384 |
| E.3 | Source code in R language | 396 |
| E.4 | Compiled PISA code | 397 |
| F | Units, Constants, and Notations | 407 |

List of Figures

| | | |
|------|---|-----|
| 2.1 | Forward and reverse determinism | 42 |
| 2.2 | Information “erasure” under reversible physics | 44 |
| 2.3 | Venn diagram of entropy and information | 47 |
| 3.1 | Configuration graphs in reversible and irreversible models of computation | 60 |
| 3.2 | Bennett’s 1989 reversible simulation algorithm | 64 |
| 3.3 | Euler tour of an irreversible machine’s configuration tree | 66 |
| 3.4 | Structure of permutation oracles | 71 |
| 3.5 | Encoding outdegree-1 directed graphs in self-reversible oracles | 74 |
| 3.6 | Problem graph defined by oracle | 76 |
| 3.7 | Optimal reversible pebble game strategy | 79 |
| 3.8 | Triangle representation of oracle queries | 80 |
| 3.9 | Visualizing the definition of the set of pebbled nodes | 81 |
| 3.10 | Decompression algorithm | 88 |
| 4.1 | Example of Shor’s algorithm: initial superpositions | 99 |
| 4.2 | Superposition after Fourier transform | 101 |
| 6.1 | Speed limit for reversible machines of given dimensions | 129 |
| 6.2 | An optimal irreversible machine for 3-D CA simulations | 134 |
| 6.3 | A faster reversible machine for 3-D CA simulations | 135 |
| 6.4 | “Folding” a column of processors to minimize volume | 138 |
| 7.1 | Ordinary CMOS inverter | 149 |
| 7.2 | Energy dissipation in conventional switching | 150 |
| 7.3 | Charging with constant current | 172 |
| 7.4 | Adiabatic charging in CMOS | 173 |
| 7.5 | Voltage curves for slow and fast charging | 173 |
| 7.6 | SCRL inverter | 176 |
| 7.7 | SCRL generalized inverter | 177 |
| 7.8 | SCRL bidirectional latch | 178 |
| 7.9 | SCRL pipeline | 181 |

| | | |
|------|--|-----|
| 7.10 | Full SCRL timing diagram | 182 |
| 7.11 | Simplified SCRL circuit model | 183 |
| 7.12 | Scaling of energy/op in SCRL with speed, given nonzero leakage . . . | 192 |
| 7.13 | How minimum energy scales with threshold voltage | 196 |
| 7.14 | A problematic case for SCRL | 198 |
| 7.15 | Fixing the problem case | 200 |
| 7.16 | Two logic gates in the physical billiard ball model | 202 |
| 7.17 | The billiard-ball model cellular automaton | 203 |
| 7.18 | Boolean logic form of BBMCA update rule | 204 |
| 7.19 | Grid of FLATTOP processing elements | 206 |
| 9.1 | Reversible control-flow structures | 229 |
| 9.2 | Reversible assembly-language multiplication routine. | 233 |
| 9.3 | Multiplication routine in R language | 238 |
| 9.4 | Initial state in Schrödinger simulation | 242 |
| 9.5 | State after 1000 simulation steps | 243 |
| A.1 | Block diagram of PE cell | 268 |
| A.2 | Icon for a single FLATTOP cell | 268 |
| A.3 | One corner of an array of FLATTOP PEs | 270 |
| A.4 | The full 20×20 array of PEs | 271 |
| A.5 | Stage 1 logic gate | 272 |
| A.6 | Logic for stages 2 and 3 | 273 |
| A.7 | Complete layout of a single FLATTOP PE | 273 |
| B.1 | “Non-expanding” arithmetic/logical operations | 276 |
| B.2 | “Expanding” arithmetic/logical operations | 277 |
| B.3 | Branch and I/O operations | 278 |

List of Tables

| | | |
|-----|---|-----|
| 2.1 | Limits on entropy density from various analyses | 38 |
| 2.2 | Physical constraints on computation | 50 |
| 3.1 | Some existing theoretical models of computation | 53 |
| 3.2 | Some measures of cost or complexity | 55 |
| 6.1 | Three classes of architectures | 124 |
| 6.2 | Summary of asymptotic scaling results | 142 |
| 7.1 | SIA semiconductor roadmap | 157 |
| 7.2 | Minimum entropy generation for irreversible CMOS | 158 |
| 7.3 | Dielectric constants | 159 |
| 7.4 | Outer-area times time, for irreversible CMOS | 164 |
| 7.5 | Device parameters for the HP14 fabrication process | 207 |
| 8.1 | Maximum entropy flux estimates | 216 |
| 8.2 | Maximum per-area rate for various irreversible technologies | 218 |
| 8.3 | Entropy coefficients for various reversible technologies | 220 |
| 8.4 | Reversible device density to beat irreversible technologies | 221 |
| 8.5 | Thicknesses above which reversible machines win | 221 |
| 9.1 | Registers used in MULT routine | 234 |
| F.1 | Unit magnitude prefixes | 407 |
| F.2 | Fundamental units used in this document | 408 |
| F.3 | Fundamental physical constants used in this document | 408 |
| F.4 | Asymptotic order-of-growth notation | 409 |

