

**PARTIAL ORDER AND PARTIAL RELIABILITY TRANSPORT SERVICE  
INNOVATIONS IN A MULTIMEDIA APPLICATION CONTEXT**

In Two Volumes

Volume I  
(Pages 1-175)

by

Phillip T. Conrad

A dissertation submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Computer and Information Sciences.

Fall 2000

Copyright 2000 Phillip T. Conrad  
All Rights Reserved

**PARTIAL ORDER AND PARTIAL RELIABILITY TRANSPORT SERVICE  
INNOVATIONS IN A MULTIMEDIA APPLICATION CONTEXT**

by

Phillip T. Conrad

Approved: \_\_\_\_\_  
M. Sandra Carberry, Ph.D.  
Chair of the Department of Computer and Information Sciences

Approved: \_\_\_\_\_  
Conrado M. Gempesaw II, Ph.D.  
Vice Provost for Academic Programs and Planning

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

---

Paul Amer, Ph.D.  
Professor in charge of dissertation

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

---

Charles Boncelet, Ph.D.  
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

---

Errol Lloyd, Ph.D.  
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed:

---

Adarshpal Sethi, Ph.D.  
Member of dissertation committee

## ACKNOWLEDGMENTS

This work was supported in part by grants from the US Army Research Laboratory under the Federated Laboratory Program, Cooperative Agreement number DAAL01-96-2-0002, and by the National Science Foundation (NCR-9314056). I am grateful for this support.

I would like to express my deepest gratitude to the many people who have given their time, energy, support and prayers in support of my efforts on this dissertation. First and foremost, I would like to thank my advisor, Paul Amer, who has been an excellent mentor, guide, teacher, colleague and friend. I thank him for his professional, personal, and moral support, for his confidence in me, and his incredible patience and leadership.

Next, I would like to thank the other members of my committee: Errol Lloyd, Adarsh Sethi, and Charlie Boncelet; these folks have gone beyond the call of duty for me on many occasions, and I am deeply grateful. I would also like to thank my colleagues in the Protocol Engineering Lab at the University of Delaware, especially: Greg Burch, Tom Connolly, Armando Caro, Maruisz Fecko, Ed Golden, Sami Iren, Rahmi Marasli, Gosia Steinder and Mason Taube, for their many valuable suggestions. In particular, I would like to thank Armando, Ed, Sami and Mason for their software development assistance with UTL, ReMDoR, the Lossy Router, and the experiment scripts.

I thank my colleagues in the CIS Departments at the University of Delaware and at Temple University, for their support and confidence during the last three years of this project, when I was teaching full-time in addition to working on the

research. In particular, I would like to extend my thanks to Sandee Carberry, Kathy McCoy, Bob Caviness, Frank Friedman and Giorgio Ingargiola for their guidance and mentoring during my first three years as a full-time faculty member. I would also like to thank the EECIS lab staff group at the University of Delaware, particularly Mike Davis and Ben Miller, for running the finest system administration group I have ever encountered. I am most grateful for their assistance with the many special requests our lab has made in connection with the network experimentation component of this work.

There are so many friends and family members who have offered their support and encouragement that it would be impossible to name them all in the small space allotted for these acknowledgments. However, I would like to particularly thank my mother, Greer Conrad, my father, Roddy Conrad. I would also like to thank the members of West Presbyterian Church, particularly Betty Crocker, for their prayers and gentle reminders to keep on working towards the goal. I am also very grateful for the assistance of Tom Ledbetter and Lynne Hagelin. I would also like to thank a few folks who were directly involved with helping to produce the final product. I particularly thank Sean Crist, who provided an incredible amount of help with proofreading, photocopying, collating, and generally keeping me somewhat sane through the final hours before my dissertation defense. Chris Coons, Dave Hunt, and Jeff Krehbiel also provided invaluable proofreading help at various stages.

Finally, my deepest debt of gratitude is to my partner, Bob Nieder, who has long suffered the pains of being a Ph.D. candidate spouse. His loyalty and patience are far more than I deserve, and I am so grateful for his support, love, encouragement, and faithfulness throughout the last twelve years we have been together, and especially the last eight years of dissertating. I thank him from the bottom of my heart, and dedicate this work to him.

## TABLE OF CONTENTS

LIST OF FIGURES .....	xviii
ABSTRACT .....	xxiv
Chapter	
1 INTRODUCTION .....	1
1.1 Problem statement .....	1
1.2 Key results of this dissertation.....	3
1.2 Structure of the dissertation.....	5
1.3 Background.....	7
1.4 Partial order and partial reliability transport service .....	8
1.4.1 Problems with using TCP or UDP .....	10
1.4.2 Partially-ordered/partially-reliable transport service as an alternative .....	12
1.5 PO/PR Transport service and multimedia document retrieval .....	13
1.5.1 Graceful degradation of multimedia documents .....	14
1.5.2 Traditional transport protocols are not satisfactory.....	15
1.5.3 Traditional authoring systems are not satisfactory. ....	16
1.6 Overview of systems developed for this dissertation .....	17
1.6.1 Universal Transport Library (UTL).....	17
1.6.2 Overview of the <u>Remote Multimedia Document Retrieval</u> system (ReMDoR).....	18
1.6.3 Publications based on these tools .....	20
1.7 Overview of performance experiments .....	22
1.8 Chapter summary.....	23
2 INNOVATIONS IN TRANSPORT SERVICE ORDER AND RELIABILITY .....	26
2.1 Introduction .....	26
2.2 Background: Partially-ordered/partially-reliable (PO/PR) transport service.....	27
2.2.1 PO/PR transport service illustrated: the “Screen refresh” example .....	28
2.2.2 Notation and terminology related to partial orders.....	30
2.2.3 Previous work on designs for a PO/PR transport service.....	32
2.2.4 Introduction to Partial Order Connection version 2 (POCv2).....	33
2.2.5 Relationship of POCv2 to UTL.....	33

2.2.6	Overview of POCv2 transport service, and comparison with POC .....	34
2.2.7	Service primitives of POCv2: Read(), Write(), etc. ....	35
2.2.8	POCv2 sequence numbers: epoch, period, objNum, cellNum .....	37
2.3	The POCv2 stream object abstraction .....	40
2.3.1	Motivation for stream objects.....	40
2.3.2	Advantages of representing the PO at a higher level than packets .....	41
2.3.2	POCv2 stream objects, cells, the streamEnd flag, objNums, and cellNums .....	45
2.4	Service profile management in POCv2 .....	46
2.4.1	The need for service profile negotiation and multiple partial orders.....	46
2.4.2	POCv2 service profiles: notation, formal definition, and representation .....	47
2.4.3	Representation of the service profile.....	48
2.4.4	Periodic partial orders in POC.....	50
2.4.5	Multiple periods and epochs in POCv2.....	51
2.5	Multimedia synchronization: background material .....	53
2.5.1	Coarse-grained vs. fine-grained synchronization .....	54
2.5.2	Temporal scenarios for coarse-grained multimedia synchronization.....	55
2.5.3	Object Composition Petri Nets (OCPN) .....	57
2.5.4	Extended OCPN (XOCPN).....	60
2.6	Transport layer support for multimedia synchronization in POCv2.....	62
2.6.1	Motivation .....	62
2.6.2	Coarse-grained synchronization via explicit release .....	64
2.6.3	Formal definition of explicit release synchronization .....	65
	Basic assumptions: .....	66
	Formal Definition of Explicit Release Synchronization:.....	67
2.6.4	Comparison of explicit release synchronization with OCPN.....	68
	Future work: two-color Petri net delivery semantics. ....	71
2.6.5	Objections to explicit release (and, motivation of data preview).....	73
2.7	Data preview (buffer access) for support of integrated layer processing.....	74
2.8	The relationship between partial order and partial reliability.....	78
2.8.1	Previous work.....	78
2.8.2	The PR reliability class.....	79

	Definition of PR reliability in POCv2 .....	82
2.8.3	PR reliability + explicit release synchronization = graceful degradation .....	83
2.8.4	Unresolved issues and future work for the PR reliability class .....	85
2.9	Providing control over reliability via ADN-cancel.....	86
2.10	Current status of POCv2 implementation and areas for future work .....	89
2.11	Chapter summary.....	90
3	THE UNIVERSAL TRANSPORT LIBRARY (UTL).....	92
3.1	Introduction .....	92
3.1.1	Organization of this chapter .....	93
3.2	Motivation .....	94
3.3	Overview of UTL .....	95
3.3.1	Central principles of UTL .....	96
3.3.2	Common service model: connection oriented, PO/PR message service .....	97
3.3.3	Connection-oriented implies three phase operation—nothing more.....	97
3.3.2	Selecting transport QoS via UTL mechanisms.....	98
3.3.3	Modifying transport QoS via UTL protocol parameters .....	100
3.3.4	UTL is a library providing flexible transport QoS, not a QoS Architecture.....	104
3.3.5	Mechanisms are composed of layers .....	105
3.3.6	Rules for composing mechanisms from layers.....	107
3.3.7	Bottom layers: TXL, KXP, KX2, KX3 .....	108
3.3.8	Upper layers: TOL, POL, NUL, SRL, and layer stacking rules .....	109
	Total Ordering Layer (TOL).....	110
	Partial Ordering Layer (POL).....	110
	Null Layer (NUL) .....	113
	Segmentation Reassembly Layer (SRL).....	114
3.3.9	User level implementation with cooperative multitasking.....	114
3.4	Formal specification of rules for composing UTL layers.....	115
3.4.1	Definition: UTL protocol parameters vector .....	115
3.4.2	Definition: UTL QoS specification .....	116
3.4.3	Definition: Well-formed UTL stack.....	116
3.4.4	Definition and algorithm for function <i>stackQoS</i> .....	117
3.5	Design issues .....	117
3.5.1	User-level vs. kernel-level development .....	118
3.5.2	Service model .....	121
3.5.3	Why all UTL services are connection-oriented.....	122
3.5.4	Why all UTL services are message-oriented .....	123

3.5.5	Minimizing data copies for faster throughput .....	125
3.5.6	QoS negotiation .....	128
3.5.7	CPU scheduling in UTL via cooperative multitasking.....	131
	Multiple processes .....	133
	Multiple threads.....	134
	Signal Handlers.....	134
	Cooperative multitasking.....	135
3.5.8	I/O multiplexing, and the need for a RAW mechanism .....	139
3.5.9	Application-Transport flow control.....	140
	<i>Sender and receiver</i> application-transport flow	
	control.....	141
	<i>Sender-only</i> application-transport flow control.....	142
	Application-transport flow control is mandatory	
	in TCP, advisory in UTL .....	142
3.6	Selected service and protocol details for the KXP, KX2 and KX3	
	layers.....	143
3.6.1	Unordered, <i>k</i> -xmit reliable service .....	143
3.6.2	Flow control .....	144
3.6.3	Packet types .....	144
3.6.4	Acknowledgments .....	144
3.6.5	Sequence numbers .....	145
3.6.6	Congestion control .....	145
3.6.7	RTO calculation .....	146
3.7	UTL development, testing and debugging.....	149
3.7.1	UTL development.....	149
3.7.2	UTL testing.....	150
3.7.3	Debugging macros.....	152
3.7.4	Memory debugging macros .....	153
3.8	Related work.....	156
3.8.1	User-level (user-space) protocol implementations .....	156
3.8.3	Flexible protocol architectures .....	157
3.9	Chapter summary, and future work related to UTL.....	158
4	<b>THE REMOTE MULTIMEDIA DOCUMENT RETRIEVAL SYSTEM</b>	
	(ReMDoR).....	161
4.1	Introduction .....	161
4.2	Overview of ReMDoR functionality .....	162
4.3	Overview of ReMDoR system components .....	163
4.4	Syntax and Semantics of PMSL .....	165
4.5	The PMTP protocol .....	168
4.6	Functions provided by the ReMDoR document compiler.....	169
4.7	Related work.....	171
4.7.1	MEDIADOC/MEDIABASE (Rody and Karamouch, 1995).....	172

4.7.2	Fiets (Rutledge et al., 1998), and HyTime, DSSSL and SMIL.....	172
4.8	Project history.....	174
4.9	Chapter summary and suggestions for future work.....	175
5	RESULTS OF PERFORMANCE EXPERIMENTS.....	176
5.1	Introduction .....	176
5.1.1	Goals and limitations of our investigation .....	176
5.1.2	Organization of this chapter and overview of performance experiments. ....	178
5.1.3	Experimental setup.....	179
5.1.4	The lossy router .....	181
	Validation of the lossy router .....	183
5.1.5	The packet reflector.....	185
	Modified leaky bucket scheme used in the packet reflector.....	186
	Validation of the packet reflector .....	188
5.2	Experiment N1: showing the upper bound for performance gain (U/R vs. PO/R vs. O/R Service using NETCICATS).....	189
	Organization of Section 5.2.....	190
5.2.1	Unordered vs. ordered service: related work.....	190
	Unordered vs. ordered service (Iren, 1999b) .....	191
	Unordered vs. ordered service (Diot and Gagnon, 1999).....	192
5.2.2	Experiment N1: parameters.....	195
5.2.3	Format of performance graphs .....	196
5.2.4	Experiment N1: observations .....	197
5.2.5	Exp. N1 analysis: delay in delivery of pixels vs. bytes .....	198
	Exp. N1 analysis: differences in performance at 0% loss .....	199
5.2.7	Exp. N1 analysis: differences in performance at 10% and 20% loss .....	201
5.2.8	Caveats regarding interpretation of example ReMDoR screen dumps .....	202
	The $p(t)$ function: pixels on the screen as a function of time for ordered service .....	203
	For unordered or partially-ordered service, the $p(t)$ function is problematic .....	204
	Avoiding bias in the choice of screen dumps .....	205
5.2.9	Exp. N1 analysis: interpretation of sample screen dumps.....	206
5.2.10	Exp. N1: conclusions and summary .....	206
5.3	Experiment R1: O/R vs. PO/R for eight parallel GIF images at 9.6kbps.....	212
	Organization of Section 5.3 .....	213

5.3.1	Related Work: other ways of providing parallel flows.....	214
	The parallel TCP connections approach has serious drawbacks.....	214
	The Multi-Stream Protocol.....	214
5.3.2	Experiment R1: parameters .....	215
5.3.3	Why the R2E and T2E mechanisms are used for Experiments R1–R5 .....	216
5.3.4	Experiment R1.1: observations and conclusions.....	217
5.3.5	Experiment R1.2: observations and conclusions.....	219
5.3.6	Experiment R1.3: observations and remarks.....	226
5.3.7	Experiment R1.3: conclusions.....	228
	Why larger windows hurt performance more for T2E, and less for R2E.....	228
	Why finding the optimal window size is a hard problem.....	230
	With partially-ordered service, choosing an oversized window is less detrimental .....	231
	Future work: evaluating gain for dynamic windows, correlating gain with density ....	231
5.3.8	Experiment R1.4: observations and conclusions.....	237
5.3.9	Experiment R1: summary.....	237
	Future Work.....	238
5.4	Experiment R2: O/R vs. PO/R for eight parallel GIF images at 128kbps.....	240
	Organization of Section 5.4 .....	241
5.4.1	Experiment R2: parameters .....	241
5.4.2	Experiments R2.1 and R2.2: observations and conclusions.....	242
5.4.3	Experiment R2.3: observations and conclusions.....	245
5.4.4	Experiment R2.4: observations and conclusions.....	248
	The optimal window size for Experiment R2.4 lies between 16 and 32.....	248
	PO/R outperforms O/R over a range of window sizes at 10% loss.....	250
5.4.5	Experiment R2: summary.....	251
5.5	Experiment R3: O/R vs. PO/R, eight parallel GIF images, various bit rates .....	255
5.5.1	Experiment R3: motivation .....	255
5.5.2	Experiment R3: parameters .....	256
5.5.3	Experiment R3: observations and summary.....	256
	Observations for bitrate 2.4kbps.....	257
	Observations for bitrate 9.6kbps.....	258
	Observations for bitrate 33.6kbps.....	259

	Observations for bitrate 128kbps.....	260
	Experiment R3: summary.....	260
5.6	Experiment R4: images in parallel with audio (O/R vs. PO/R for images and audio from <code>paris.pmsl</code> ).....	266
5.6.1	Background: Three proposed metrics for audio performance.....	266
	1) INT (Absolute number of interruptions) .....	269
	2) FRACPLAY (Fraction Playing).....	270
	3) FRACPLAY <sup>INT</sup> .....	271
5.6.2	Experiment R4: parameters and hypotheses.....	271
5.6.3	Experiment R4.1: observations and conclusions.....	273
	Observations and conclusions for pixels and bytes ....	273
	Interpreting the graphs for the audio metrics.....	275
	Observations and conclusions for audio metrics .....	276
5.6.4	Experiment R4.2: observations and conclusions.....	282
5.6.5	Experiment R4.3: observations and conclusions.....	291
5.6.6	Experiment R4.4: observations and conclusions.....	301
5.7	Experiment R5: a complete multimedia document (PO/R vs. O/R for <code>parisScene1.pmsl</code> ).....	310
5.7.1	Experiment R5: description of document <code>parisScene1.pmsl</code> .....	310
5.7.2	Experiment R5: flow control, and the use of R3 and T3 mechanisms .....	311
5.7.3	Experiment R5: parameters and hypotheses.....	314
5.7.4	Experiment R5: observations and conclusions for pixels and bytes .....	315
	Overall conclusions related to bytes/pixels .....	319
	Experiment R5: observations and conclusions for audio metrics .....	319
	Overall conclusions related to audio .....	321
5.8	Problems in performance measurement .....	330
5.8.1	Tannenbaum's Pitfall #1: Insufficient sample size .....	331
5.8.2	Tannenbaum's Pitfall #2: Non-representative samples.....	332
5.8.3	Tannenbaum's Pitfall #3: Inaccurate time measurements.....	335
5.8.4	Tanenbaum's Pitfall #4: Unexpected interference .....	337
5.8.5	Tanenbaum's Pitfall #5: Artifacts of Caching.....	337
5.8.6	Tanenbaum's Pitfall #6: Misunderstanding what is being measured.....	338
5.8.7	Tanenbaum's Pitfall #7: Unwarranted extrapolation .....	339
5.9	Overall conclusions from our experiments.....	339
6	ALGORITHMS FOR PO/PR PROTOCOLS.....	343
6.1	Introduction .....	343
6.1.1	Organization of this chapter .....	344

6.1.2	Processing overheads in O/R and PO/PR transport protocols .....	345
6.1.3	Packet resequencing at the transport receiver in ordered protocols .....	346
6.1.4	The goal for resequencing in PO protocols: $O(1)$ per operation .....	347
6.2	Topological sort of a directed acyclic graph by incremental delete (DAGITS) .....	349
6.2.1	The TS-DAG and TS-DAG-V problems .....	351
6.2.2	Topological sorting and PO/PR transport protocols .....	351
	Sending order validation in PO transport service (PO-SND-VALID) .....	351
	The PO/R receiver TSDU resequencing problem (PO/R-RCV-RESEQ) .....	352
6.2.2	Incremental TS-DAG (INCR-TS-DAG) .....	353
6.2.3	The usual DFS-approach to TS-DAG is problematic for INCR-TS-DAG .....	354
6.2.4	Topological sorting by incremental delete and the DAGITS algorithm .....	355
6.2.5	The DAGITS algorithm: proof of correctness and running time .....	356
6.2.6	Concluding Remarks .....	359
6.3	Choosing a linear extension in ReMDoR .....	359
	Summary of Section 6.3 .....	362
6.4	Verifying the sending order .....	366
6.4.1	The ISOLE rule .....	366
6.4.2	Enforcement of the ISOLE rule in UTL PO/PR services .....	366
6.4.3	Using the DAGITS algorithm to enforce the ISOLE rule .....	367
	Problem ISOLE: Initial Sending Order must be a Linear Extension .....	367
6.5	Resequencing out-of-sequence PDUs for delivery using partial order .....	368
	The matrix approach to resequencing PDUs for PO service .....	368
	The matrix approach is not acceptable for PO service .....	370
	Incorporating explicit release synchronization into PO resequencing .....	371
6.6	Implementing the PR semantics of POCv2 .....	374
6.7	PO/PR-DEL-BASIC: Basic POCv2 delivery (no stream, no explicit release) .....	380
	Overview of Section 6.7 .....	381

6.7.1	Extending the PO/R-RCV-RESEQ algorithm to incorporate PR.....	382
6.7.2	Proofs of correctness, running time for PO/PR-DEL-BASIC pseudocode .....	391
	Proofs related to maintaining the <i>numCOURPs</i> values.....	391
	Proofs related to computing the $L(y)$ set via a DFS of the $TRPO^T$ .....	397
6.7.3	Comparison/Contrast of PO/PR-DEL-BASIC with PO/R-RCV-RESEQ.....	401
6.7.4	PO/PR-DEL-BASIC: Section Summary .....	402
6.8	PO/PR-DEL-FULL: Adding streams and explicit release to POCv2....	403
	Overview of Section 6.8 .....	404
6.8.1	Integrating POCv2's PR class with explicit release and stream objects .....	404
	Interaction of Reliability Classes with Stream Objects .....	404
	Integration of stream objects with explicit release .....	405
	Streaming vs. Stalled objects, and the underflow notification.....	406
	What running time can be achieved when incorporating all three features? .....	407
6.8.2	The main difficulty: explicit release of streaming objects .....	408
6.8.3	An efficient algorithm that sacrifices two desirable properties .....	409
	Sketch of algorithm PO/PR-DEL-OPTION1 .....	409
	The first property we sacrifice: no premature loss declarations .....	413
	The second property we sacrifice: no deliverable data waiting for unreliable data .....	414
6.8.4	A brute-force algorithm that implements the full POCv2 PR semantics .....	415
	Sketch of algorithm PO/PR-DEL-OPTION2 .....	416
	Running time of algorithm PO/PR-DEL-OPTION2.....	419
6.8.5	Future work related to the POCv2 PR reliability class.....	421
	Approaches to improving the running time of PO/PR-DEL-OPTION2 .....	422
6.9	Representation of partial orders (encodings, data structures).....	423
6.10	Chapter summary and suggestions for future work .....	429
	Future work: linear extension selection.....	430

7	CONCLUSIONS AND SUGGESTIONS FOR FUTURE WORK.....	432
7.1	Summary of main results .....	432
7.1.1	PO/R service can provide benefits in practice.....	432
7.1.2	Other results: transport service.....	435
7.1.3	Other results: multimedia systems.....	435
7.2	Applicability of our results to other applications .....	436
7.2.1	Application of results in this dissertation to work on SCTP .....	437
7.2.2	Applying partial order to the results of SQL queries.....	438
7.3	Future Work.....	439
7.3.1	Experimentation with PO/PR service.....	439
7.3.2	Further experimentation with PO/R service .....	440
7.3.2	Other future work .....	442
	Future Directions in Multimedia Systems .....	442
	Future Work on ReMDoR .....	442
	Future Directions in Transport Protocol Design.....	443
	Future work related to UTL implementation.....	444
	Appendix: EXAMPLE MULTIMEDIA DOCUMENTS .....	445
A.1	The paris.pmsl document: a travelogue of Paris.....	445
A.2	military.pmsl: a briefing on Gulf War weapons systems.....	449
A.3	Listing of example PMSL source: paris.pmsl.....	451
A.4	Excerpts from sample file paris.pmff .....	464
A.5	Description of PMTP output .....	465
	BIBLIOGRAPHY.....	466

## LIST OF TABLES

Table 1.1	Publications related to UTL and ReMDoR implementation .....	20
Table 1.2	Performance Experiment Groups .....	24
Table 2.1	A partial listing of POCv2 service primitives .....	36
Table 3.1	UTL mechanisms.....	101
Table 3.2	UTL QoS Parameters (legend for Table 3.1).....	102
Table 3.3	UTL protocol parameters.....	103
Table 3.4	Bottom layers in UTL (may also serve as top layers).....	111
Table 3.5	Upper layers in UTL (may be middle or top layers).....	112
Table 3.6	Services provided by various transport layers .....	126
Table 3.7	Tick sizes for RTT/RTO calculation in TCP and UTL .....	148
Table 3.7	Production releases of UTL (partial list) .....	154
Table 5.1	Parameters for Experiment N1 .....	196
Table 5.2	UTL Mechanisms used in Experiment N1 .....	196
Table 5.3	Parameters for Experiment R1 .....	215
Table 5.4	UTL Mechanisms used in Experiments R1 (also Exps. R2–R4) .....	216
Table 5.5	Parameters for Experiment R2 .....	241
Table 5.6	First pixel delivered (median) for R2.1 .....	243
Table 5.7	R2.3: Last pixel delivered (median) .....	246
Table 5.8	Parameters for Experiment R3 .....	256
Table 5.9	Parameters for Experiment R4 .....	271

Table 5.10	Parameters for Experiment R5 .....	314
Table 5.11	Number of repetitions for each experiment.....	333
Table 6.1	Audio Parameters Used for Chapter 5 performance experiments .....	362
Table 6.2	A partial listing of POCv2 service primitives .....	425
Table 6.3	Algorithms to convert between PO representations .....	428
Table 6.4	Time to compute transitive closure and transitive reduction in the ReMDoR parser for various values of $n$ on Sun Ultra 10. ....	429

## LIST OF FIGURES

Figure 1.1	Structure of the dissertation.....	7
Figure 1.2	Experimental setup for performance experiments.....	25
Figure 2.1	Screen refresh example from (Amer et al., 1994) .....	30
Figure 2.2	Transitively reduced DAG representation of PO .....	31
Figure 2.3	Excerpt illustrating POCv2 sequence numbers .....	39
Figure 2.4	Example of epoch, period, objNum, cellNums .....	39
Figure 2.5	Storyboard to motivate POCv2 stream objects .....	44
Figure 2.6	Partial order with stream objects .....	44
Figure 2.7	Partial order without stream objects.....	45
Figure 2.8	Example definition of a POCv2 service profile .....	48
Figure 2.9	Periodic partial orders in POC, POCv2.....	51
Figure 2.10	Hamblin’s temporal relations (left); corresponding Petri nets from OCPN (right) (Little and Ghafoor, 1990).....	61
Figure 2.11	Example OCPN for slide presentation adapted from (Little and Ghafoor, 1990) 61	
Figure 2.12	Pseudocode for a “read and display” loop.....	62
Figure 2.13	Example: two element partial order .....	62
Figure 2.14	Example transformation from Lemma 2.1 .....	72
Figure 2.15	Indefinite intervals as a result of network delays .....	73
Figure 2.16	Example PO for explanation of PR reliability class.....	83
Figure 3.1	Selecting from among different transport protocols by a menu (left) or command line option (“-m” for “mechanism”, right). .....	100
Figure 3.2	UTL mechanisms composed of layers.....	106

Figure 3.3	UTL layer encapsulation example: T3 over Ethernet.....	106
Figure 4.1	User interface of ReMDoR browser.....	163
Figure 4.2	Web architecture, for comparison with ReMDoR.....	166
Figure 4.3	ReMDoR system architecture.....	166
Figure 5.1	Detail of experimental environment.....	181
Figure 5.2	Gilbert-Elliot loss model (two state Markov chain).....	183
Figure 5.3	Illustration of <code>ncg8par.pms1</code> .....	195
Figure 5.4	Illustration of maximum performance gap for Exp. N1 at 0% loss.....	207
Figure 5.5	Experiment N1: Performance graphs .....	208
Figure 5.6	Screen dumps: <code>ncg8par.pms1</code> , 10% loss, 38.4kbps PPP link .....	209
Figure 5.7	Screen dumps: <code>ncg8par.pms1</code> , 10% loss, 38.4kbps PPP link .....	210
Figure 5.8	Screen dumps: <code>ncg8par.pms1</code> , 10% loss, 38.4kbps PPP link .....	211
Figure 5.9	Experiment R1.1: Performance graphs .....	221
Figure 5.10	Screen dumps: R1.1, 9.6kbps PPP link at 10% loss.....	222
Figure 5.11	Screen dumps: R1.1, 9.6kbps PPP link at 20%loss.....	223
Figure 5.12	Experiment R1.1 performance graphs (overlay for Fig. 5.13) .....	224
Figure 5.13	Experiment R1.2 performance .....	225
Figure 5.14	Experiment R1.3: Performance Graphs.....	233
Figure 5.15	Exp. R1.3, Ind. Runs, T2E, win 32, bytes .....	234
Figure 5.16	Exp. R1.3, Ind. Runs, T2E, win 8, bytes .....	234
Figure 5.17	Exp. R1.3, Ind. Runs, T2E, win 16, bytes .....	235
Figure 5.18	Exp. R1.3, Ind. Runs, R2E, win 8, bytes.....	235
Figure 5.19	Exp. R1.3, Ind. Runs, R2E, win 16, bytes.....	236

Figure 5.20	Exp. R1.3, Ind. Runs, R2E, win 32, bytes .....	236
Figure 5.21	Experiment R1.3 vs. R1.4: Performance Graphs .....	239
Figure 5.22	Exp. R2.1, R2.2, pixels, 128kbps, 4 loss rates $\times$ 2 delays .....	244
Figure 5.23	Experiment R2.3: Four one-way delays at 0% and 20% loss.....	247
Figure 5.24	Exp. R2.4: Performance Graphs (128kbps at 0% loss) .....	253
Figure 5.25	Exp. R2.4: Performance Graphs (128kbps at 10% loss) .....	254
Figure 5.26	Experiment R3: Average Performance Graph.....	261
Figure 5.27	Experiment R3: Median Performance Graph.....	262
Figure 5.28	Experiment R3: Quartiles Performance Graph .....	263
Figure 5.29	Experiment R3: Quartiles Performance Graph .....	264
Figure 5.30	Experiment R3: Scaled Quartiles Performance Graphs .....	265
Figure 5.31	Exp. R4.1: bytes, pixels performance graphs .....	278
Figure 5.32	Experiment R4.1: audio interruptions .....	279
Figure 5.33	Experiment R4.1: FRACPLAY metric .....	280
Figure 5.34	Experiment R4.1: FRACPLAY <sup>INT</sup> metric .....	281
Figure 5.35	Experiment R4.2: bytes performance graphs .....	283
Figure 5.36	Experiment R4.2: pixel performance graphs.....	284
Figure 5.37	Experiment R4.2: audio interruptions at 0% loss.....	285
Figure 5.38	Experiment R4.2: audio interruptions at 20% loss.....	286
Figure 5.39	Experiment R4.2: FRACPLAY metric at 0% loss .....	287
Figure 5.40	Experiment R4.2: FRACPLAY metric at 20% loss .....	288
Figure 5.41	Exp. R4.2: FRACPLAY <sup>INT</sup> metric at 0% loss .....	289
Figure 5.42	Exp. R4.2: FRACPLAY <sup>INT</sup> metric at 20% loss .....	290

Figure 5.43	Experiment R4.3: bytes performance graphs .....	293
Figure 5.44	Experiment R4.3: pixel performance graphs.....	294
Figure 5.45	Experiment R4.3: audio interruptions at 0% loss.....	295
Figure 5.46	Experiment R4.3: audio interruptions at 20% loss.....	296
Figure 5.47	Experiment R4.3: FRACPLAY metric at 0% loss .....	297
Figure 5.48	Experiment R4.3: FRACPLAY metric at 20% loss .....	298
Figure 5.49	Experiment R4.3: FRACPLAY <sup>INT</sup> metric at 0% loss.....	299
Figure 5.50	Experiment R4.3: FRACPLAY <sup>INT</sup> metric at 20% loss.....	300
Figure 5.51	Exp. R4.4: LR 0% bytes, pixel perf. graphs.....	302
Figure 5.52	Exp. R4.4: LR 10% bytes, pixel perf. graphs.....	303
Figure 5.53	Experiment R4.4: audio interruptions at 0% loss.....	304
Figure 5.54	Experiment R4.4: audio interruptions at 10% loss.....	305
Figure 5.55	Experiment R4.4: FRACPLAY metric at 0% loss .....	306
Figure 5.56	Experiment R4.4: FRACPLAY metric at 10% loss .....	307
Figure 5.57	Exp. R4.4: FRACPLAY <sup>INT</sup> metric at 0% loss .....	308
Figure 5.58	Exp. R4.4: FRACPLAY <sup>INT</sup> metric at 10% loss .....	309
Figure 5.59	Exp. R5.1: bytes, pixels perf. graphs.....	322
Figure 5.60	Exp. R5.2: bytes, pixels perf. graphs.....	323
Figure 5.61	Exp 5.1 and 5.2, plotting the advantage of R3 over T3.....	324
Figure 5.62	Experiment R5.1: audio interruptions .....	325
Figure 5.63	Experiment R5.1: FRACPLAY metric .....	326
Figure 5.64	Experiment R5.1: FRACPLAY <sup>INT</sup> metric .....	327
Figure 5.65	Experiment R5.2: audio interruptions .....	328

Figure 5.66	Experiment R5.2: FRACPLAY metric .....	329
Figure 5.67	Experiment R5.2: FRACPLAY <sup>INT</sup> metric .....	330
Figure 5.68	Pseudocode for Experiment Loop .....	334
Figure 6.1	Topological sort via repeated delete.....	355
Figure 6.2:	DAGITS (DAG incremental topological sort) .....	355
Figure 6.3:	operation <i>next()</i> .....	356
Figure 6.4	Algorithm ReMDoR-LESTAB inputs, output .....	363
Figure 6.5	Algorithm ReMDoR-LESTAB Pseudocode .....	364
Figure 6.6	ReMDoR-LESTAB, implementation of procedure <i>finish()</i> .....	365
Figure 6.7	Problem ISOLE (operations).....	367
Figure 6.8	Algorithm PO-SENDER-ISOLE.....	369
Figure 6.9	Algorithm PO/R-RCV-RESEQ, specification .....	372
Figure 6.10	PO/R-RCV-RESEQ, operation <i>processIncomingTPDU()</i> .....	373
Figure 6.11	PO/R-RCV-RESEQ, operation <i>getNextTPDU()</i> .....	373
Figure 6.12	PO/R-RCV-RESEQ, operation <i>isAnythingDeliverable()</i> .....	373
Figure 6.13	PO/R-RCV-RESEQ, procedure <i>releaseSuccessors()</i> .....	374
Figure 6.14	Algorithm PO/PR-DEL-BASIC, specification.....	385
Figure 6.15	PO/PR-DEL-BASIC, operation <i>isAnythingDeliverable()</i> .....	386
Figure 6.16	PO/PR-DEL-BASIC, operation <i>getNextTSDU()</i> .....	387
Figure 6.17	PO/PR-DEL-BASIC, operation <i>processIncomingTPDU()</i> .....	387
Figure 6.18	PO/PR-DEL-BASIC, operation <i>init()</i> .....	388
Figure 6.19	PO/PR-DEL-BASIC, <i>updateNumCOURPsofSuccessorNodes()</i> .....	388
Figure 6.20	PO/PR-DEL-BASIC, procedure <i>releaseSuccessors()</i> .....	389

Figure 6.21	PO/PR-DEL-BASIC, <i>fillDeliverOrDeclareLostQueueWithSortedLSet(x)</i> .....	390
Figure 6.22	Finite State Automata for stream states.....	411
Figure 6.23	Modified psuedocode for <i>getNextTSDU()</i> , option 1.....	413
Figure 6.24	Procedure <i>findWaitingObjectWithNoWaitingOrStreamingProperPreds()</i> .....	417
Figure 6.25	Modified pseudocode for <i>isAnythingDeliverable()</i> , option 2.....	417
Figure 6.26	Modified pseudocode for <i>getNextTSDU()</i> , option 2.....	418
Figure 6.27	Brute-Force approach to PO/PR-DEL-OPTION2.....	420
Figure 6.28	Alternate pseudocode for PO/PR-DEL-OPTION2.....	423
Figure A.1	paris.pmsl: storyboard (part 1) .....	447
Figure A.2	paris.pmsl: storyboard (part 2) .....	448
Figure A.3	military.pmsl: partial storyboard. ....	450

## ABSTRACT

This dissertation investigates an innovation in computer communications called *Partially Ordered and Partially Reliable (PO/PR) Transport Service*. PO/PR service bridges the gap between two traditional forms of transport service: Ordered/Reliable (O/R) service, such as the Internet's Transmission Control Protocol (TCP), and Unordered/Unreliable (U/U) service, such as the Internet's User Datagram Protocol (UDP). Some applications—in particular, multimedia applications—require services that lie in between these two extremes. U/U service is insufficient for these applications, yet O/R service is too restrictive and may cause the application to pay a performance penalty.

Previous investigations of PO/PR transport service used analytic and simulation modeling to investigate the performance of an abstract PO/PR transport service called Partial Order Connection. We build on this work by first describing several innovations in PO/PR transport service developed by the dissertation author, including a new approach to coarse-grained multimedia synchronization based on extending the Object-Composition Petri Net (OCPN) of Little and Ghafoor. We then provide empirical evidence that an implementation of Partially Ordered/Reliable (PO/R) transport service called Partial Order Connection version 2 (POCv2) can provide better Quality of Service (QoS) tradeoffs to applications by providing a transport service better matched to an application's needs. In particular, we show that for multimedia document retrieval, PO/R service provides performance benefits over O/R transport service when the network loses packets.

Our experiments compare O/R service to PO/R service for a variety of documents, bit rates, window sizes, and propagation delays, with packet loss rates ranging from 0% to 30%. The results show that between 5% and 20% loss, user-perceivable improvements in progressive display are observed for bit rates between 2.4kbps to 512kbps. These results suggest that transport services providing reliable delivery over independent streams (such the emerging Internet protocol SCTP) can provide important performance benefits over lossy networks (including wireless nets or combat net radios.)

We also describe two systems developed by the dissertation author for experimenting with transport services: (1) the Universal Transport Library (UTL), a framework for implementation and testing of transport services, and (2) a Remote Multimedia Document Retrieval System (ReMDoR) that uses UTL transport services.

# Chapter 1

## INTRODUCTION

### 1.1 Problem statement

This dissertation investigates an innovation in computer communications called *Partially Ordered and Partially Reliable (PO/PR) Transport Service*. This innovative technique includes several features that are designed to provide computer applications with more flexibility in the way they use a packet-switched communications network such as the Internet. In particular, PO/PR transport service bridges the gap between two traditional forms of transport service:

- Ordered/Reliable (O/R) service, such as the Internet's Transmission Control Protocol (TCP), and
- Unordered/Unreliable (U/U) service, such as the Internet's User Datagram Protocol (UDP)

The premise of this work is that some applications—in particular, multimedia applications—require services that lie in between these two extremes.

Unordered/Unreliable service is insufficient for these applications, yet Ordered/Reliable service is too restrictive and may cause the application to pay a performance penalty. The goal is to determine whether better Quality of Service (QoS) tradeoffs and/or performance improvements can be obtained by using a transport service in between these two extremes—one that is better matched to an application's needs.

Previous investigations of PO/PR transport service used analytic and simulation modeling to investigate the performance of an abstract PO/PR transport service called Partial Order Connection (POC) (Marasli, 1997). By contrast, the **problem statement** for this dissertation is:

*To determine through experimentation with real systems the extent to which PO/PR transport service can provide performance benefits for real applications.*

Towards this end, the author designed a PO/PR transport service called Partial Order Connection, version 2 (POCv2). POCv2 is a second version of the abstract PO/PR transport service (POC), first proposed in (Amer et al., 1994), and then investigated by Marasli through analysis and simulation (Marasli et al., 1996, 1997a, 1997b). POC provides a good basis for reasoning about partial order protocols and doing simulation and analysis. However, as we shall show later in this chapter, POC lacks certain features necessary for experimentation and/or deployment with real applications.

To perform performance experiments comparing POCv2 and other PO/PR transport services to traditional transport services, the author designed:

- a framework for implementation and testing of experimental transport services (including POCv2 and others),
- an application called ReMDoR that benefits from a PO/PR service, and
- a framework for repeating performance experiments with ReMDoR under controlled conditions.

The author then supervised and participated in the development of these systems, and then used them to carry out performance experiments comparing Partially

Ordered/Reliable (PO/R) service (a subclass of PO/PR service) to ordered/reliable (O/R) and unordered/reliable (U/R) service.<sup>1</sup> The remainder of this dissertation describes the systems that were developed to conduct these experiments and the results obtained. The next section outlines our key results.

## 1.2 Key results of this dissertation

In keeping with the problem statement above, our *central* result is the analysis of *performance experiments illustrating a range of parameters where PO/R transport service can provide better performance than O/R service for a multimedia document retrieval system*. This result is important for at least two reasons. First, prior to the publication of this data, all claims about the benefits of PO/PR service have been based on either intuitive arguments, as in (Amer et al., 1994), or simulation or analysis of abstract applications, as in (Marasli et al., 1997a, Marasli 1997b). Putting this data in a real application context provides an important grounding for delineating the practical benefits available from PO transport service. Second, the fact that there is benefit from a PO/R service over an O/R service provides important motivation for future work to extend our empirical study to an investigation of the benefits of PO/PR service.

The experiments (described in detail in Chapter 5) compare O/R service to PO/R service (and in one case, also to U/R service) for a variety of documents, bit rates, window sizes, propagation delays. The experiments were carried out at loss rates ranging from 0% to 30%, with the emphasis on loss rates between 5% and 20%. We defer the presentation of specific numerical results until Chapter 5; nevertheless,

---

<sup>1</sup> We defer experimentation with PO/PR service to future work; this is consistent with the approach taken in a previous dissertation in this area (Marasli 1997b),

we can preview for the reader some of the general trends that were observed. The two most important trends are ones that confirm the value of PO/R service:

- (1) At 0% loss, there is little to no benefit or penalty for using a PO/R service rather than an O/R service, regardless of the values of any other parameters.
- (2) At loss rates higher than 0%, PO/R service provides a clear performance advantage over O/R service under some conditions, for some documents.

In particular, we note the following about the conditions under which PO/R service outperforms O/R service, and the nature of the advantages:

- (3) In general, PO/R service provides faster progressive display of information than O/R service when the loss rate is larger than 0%. As the loss rate increases, the advantage of PO/R service steadily increases. Thus PO/R service can make an application more robust to a certain amount of packet loss, up to a certain point. As the loss rate increases, after a certain loss rate is reached, the gain may be moot since both PO/R and O/R services are unacceptably bad.
- (4) The size of the gain due to partial order is sensitive to changes in the flow control and/or congestion control schemes, including the sender and/or receiver window sizes, and whether or not TCP-friendly mechanisms such as slow start and congestion avoidance are used. However, we can show gains *both* when TCP-friendly mechanisms are used, *and* when they are not.
- (5) PO/R service provides benefits for several different kinds of documents, including: small documents with no temporal dimension (similar to web pages with multiple GIF or JPEG images), simple multimedia documents with images and audio in parallel, and larger documents with complex synchronization relationships among multiple data streams.

In addition to this central result, we also present several other key results:

- We describe two *experimental systems* that were designed and developed as part of this dissertation work:

- a Universal Transport Library (UTL) providing a framework for development and testing of experimental transport services, and
- a Remote Multimedia Document Retrieval System (ReMDoR) that can operate over multiple transport services.

While these systems were originally developed for the author's dissertation work, they proved useful for other research as well—particularly research into *network-conscious image compression* (explained further in Section 1.6.3). In addition to this author's four ReMDoR/UTL related publications, five other authors have published a total of eight other journal articles, conference papers, and/or PhD, MS or BS honors theses that either: (1) included experimental results derived using ReMDoR and/or UTL, and/or (2) described design, architecture and implementation issues related to the development of ReMDoR and/or UTL.

- We describe several *innovations in PO/PR transport service* developed by the dissertation author, thus extending the previous work on PO/PR service. Some of the innovations described here were actually implemented as part of this dissertation, as they were necessary for the completion of the performance experiments. Others have not been implemented to date; for these we provide intuitive arguments why they should provide performance benefits. The latter set of innovations provides a number of opportunities for future study (see Chapter 7).

## 1.2 Structure of the dissertation

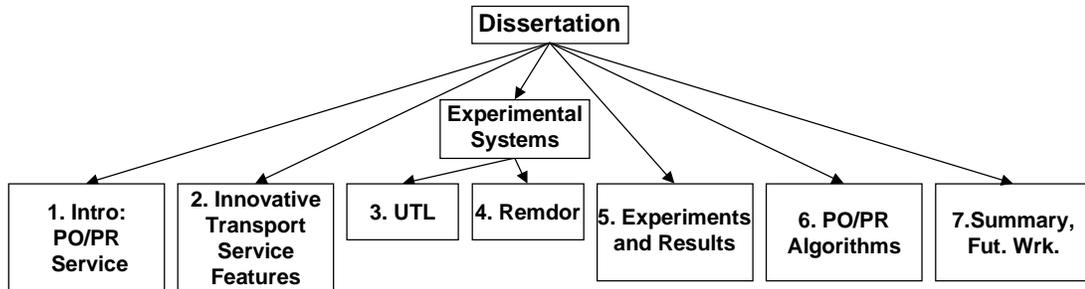
The dissertation is divided into seven chapters, as shown in Figure 1.1. Chapter 1 provides an introduction to the dissertation, including the problem statement, an overview of the motivation for studying partially ordered and partially reliable transport services, and an overview of the motivation for studying the performance of these services in the context of multimedia applications. Chapter 2 provides a more in-depth look at some of the innovations in PO/PR transport service proposed in this dissertation. Chapters 3 and 4, respectively, describe the design and implementation of the Universal Transport Library (UTL), a framework for the

development and testing of experimental transport services, and the Remote Multimedia Document Retrieval System (ReMDoR) that can operate over multiple transport services.

Chapter 5 presents the core of this dissertation: performance results from experiments designed to evaluate PO/PR service. This chapter includes an overview of the experimental methods used, and results from experiments designed to validate the experimental framework. This is followed by a complete discussion of experiments comparing ordered/reliable (O/R) service, partially ordered/reliable (PO/R) and unordered/reliable (U/R) service. We present results obtained by measuring a specific application: remote multimedia document retrieval via ReMDoR.

In Chapter 6 we turn to a different question: the communications and processing overhead of PO/PR service. Two algorithmic questions are considered: how to encode a partial order for effective transmission and processing, and how to design efficient algorithms for the extra processing that sending and receiving transport entities must perform to provide PO/PR service. Finally, Chapter 7 provides a summary of the key results from this dissertation, and suggestions for future work.

The appendix describes two sample multimedia documents (named “`paris.pmsl`” and “`military.pmsl`”) that are used throughout this dissertation as examples to illustrate ideas, and as the basis of the performance experiments.



**Figure 1.1 Structure of the dissertation**

### 1.3 Background

It is assumed that the reader is familiar with concepts of computer networking, and with the transport layer in particular. For example, the reader is assumed to be familiar with:

- the OSI model reference model; in particular, the transport layer of that model
- terms such as Protocol Data Unit (PDU) and Service Data Unit (SDU)
- the concepts of service, protocol, implementation and peer entities
- basic concepts of the TCP/IP protocol suite, such as the role of TCP, UDP and IP
- the concept of Quality of Service (QoS)
- what an Internet *Request for Comments (RFC)* document is

The dissertation author is a co-author of (Iren et al., 1999a) which provides a tutorial and survey on the transport layer; the tutorial section in particular provides a good summary for readers not familiar with the terminology above.

## 1.4 Partial order and partial reliability transport service

A distinction can be made between *qualitative* QoS parameters and *quantitative* QoS parameters. Examples of qualitative QoS parameters include:

- **order:** Is the sending order preserved; that is, are messages received out-of-order resequenced before delivery?
- **reliability:** Are lost messages recovered through re-transmission or forward error correction?
- **duplication:** It is possible for a message to be delivered more than once?
- **flow-control:** Are mechanisms in place to prevent a fast sender from overwhelming a slow receiver, resulting in either data loss or waste of network resources?

Examples of quantitative QoS parameters include:

- **delay:** How long does it take for messages to travel from sender to receiver?
- **throughput:** How many messages are delivered per unit of time?
- **jitter (or burstiness):** Do messages arrive in a predictable steady flow, or in huge bursts with long quiet periods in between? More formally, what is the variance of packet interarrival times?

One fundamental transport layer design problem is making appropriate tradeoffs between qualitative and quantitative QoS parameters. Choosing an appropriate tradeoff is important, because while transport services are often called upon to provide a QoS that is an enhancement of the underlying network, improving the performance as measured by one QoS parameter usually involves degrading the performance of another QoS parameter. For example, TCP provides a reliable service on top of the unreliable IP network protocol by means of retransmissions, but does so at the expense of introducing additional end-to-end delay. The selection of which

transport mechanisms are appropriate for a given application is often a matter of considerable debate. For example, while the prevailing view is that retransmission is inappropriate for multimedia data because of its real-time nature, (Little and Ghafoor, 1991), some authors describe circumstances in which retransmission is appropriate (Dempsey et al., 1996.)

Commonly, transport protocol selection is a choice between extremes. For example, in the Internet protocol suite we note that the service provided by TCP (RFC793) is ordered, reliable, no-duplicates and flow-controlled, while the service provided by the *User Datagram Protocol* (UDP, RFC768) is unordered, unreliable, may duplicate messages, and is not flow-controlled. The fact that these protocols provide service at the extremes of each of these four QoS axes creates a dilemma for the designer of an application whose needs reside in the middle. When designing an application that will run over Internet protocols, today's implementer typically has three choices: (1) use TCP, (2) use UDP as is, or (3) implement a custom transport protocol on top of UDP (a considerable software development effort.)

As explained below, choosing either TCP or UDP when neither is appropriate has negative consequences. Thus, many applications utilize the third choice: building the needed transport functionalities from scratch, as an application specific transport protocol layered on top of UDP, for example, see (Jacobs and Eleftheriadis, 1997; RealNetworks, 1997.) Implementing a custom protocol allows the application designer to choose exactly what features to implement based on the requirements of the application. However, implementing transport protocol features at the user-level of the operating system is non-trivial. Two issues are particularly difficult: (1) managing the context switching between asynchronous protocol events,

such as timer expiration and packet arrival, and the rest of the application code, and (2) getting flow-control/congestion-avoidance to operate correctly. In the latter case, to truly test the correctness of the design and implementation requires simulation and/or implementation on a wide scale. The complexity of designing and testing the operation of retransmission timers, resequencing of data, buffers, round-trip-delay estimation, and flow-control/congestion avoidance may exceed the complexity of the application itself!

We argue that given a reasonable alternative, application designers would prefer to avoid programming transport layer functionality. Therefore we present an alternative: a standardized transport service providing the flexibility to specify reliability, ordering, flow-control, and duplication at a finer granularity than either TCP or UDP. This flexibility allows application designers to focus their efforts on their application rather than on transport layer details.

#### **1.4.1 Problems with using TCP or UDP**

In a protocol environment where only the extremes of transport service are provided, some applications cannot find a perfect home. Consider the retrieval of objects that are part of a multimedia presentation. For some objects, no loss is permissible (e.g., text, some still images, control information) while for other objects some loss may be permissible (e.g., audio and video streams, images that are merely decorative). Also, the order of presentation of objects may be defined by a partial order rather than a total order, as is the case for document synchronization requirements described by an Object Composition Petri Net (OCPN) (Little and Ghafoor, 1991.) We describe the service required by such an application as *partially-ordered/partially-reliable*. Partially-reliable refers to the notion that some objects

must be delivered reliably, while others may be lost if necessary. Partially-ordered refers to the fact that data sequencing requirements are expressed as a partial order rather than as a linear order.

For such an application, TCP provides more reliability and resequencing than is necessary at the expense of extra delay and reduced throughput. Extra delay may result in annoying discontinuities in the playback of continuous media such as audio or video data. However, TCP has the advantage of providing flow control and congestion control algorithms that have been tested for nearly two decades, and scale well to a global internet.

UDP, on the other hand, provides a “best effort” service with no guarantees whatsoever. On a lightly loaded LAN where the underlying network is inherently reliable (at least, for most practical purposes, i.e., packet loss probabilities of  $10^{-9}$  or less), a best-effort service may be perfectly acceptable. Over longer Internet distances, where packet loss probabilities routinely range anywhere from negligible to over 50%, UDP may be fine one day and completely unacceptable the next. Our anecdotal experience with the MBONE tools for Internet video-conferencing (i.e., nv, vat) suggests that sustained packet drop rates between 7-15% are common, and that drop rates as high as 50% do occasionally occur. There are several studies that support our anecdotal experience. (Diot and Gagnon, 1999) cite similar experiences with packet loss in the Internet. (Bolot, 1993) observed packet loss rates around 10% on connections between INRIA, Sophia-Antipolis, France, and the Univ. of Maryland, College Park MD, USA. While (Paxson, 1996) does not directly measure packet losses, the high frequency of routing anomalies he cites lends support to the notion that

the Internet provides a highly variable quality of service with a significant rate of failure.

Some argue that since loss in the Internet is commonly due to buffer overflows, bandwidth reservation and improved congestion control methods are needed, and if implemented, will eliminate packet loss as a significant problem. RSVP (Zhang et al.,1993) and YESSIR (Pan and Schulzrinne, 1998) are two examples of such reservations schemes. Our sense is that in spite of excellent research efforts related to guaranteed QoS through reservations, there will always be environments where reservation mechanisms are infeasible to implement, or fail to provide the necessary QoS (e.g., a disaster situation or battlefield scenario involving intermittent jamming). In these cases the loss rate of the underlying network may be higher than an application's tolerance for loss. Furthermore, because UDP is not flow-controlled, unless the application implements its own flow-control mechanism, an application using UDP may flood the network and/or the receiver with packets at a rate faster than either can handle, thus creating another source of packet loss.

#### **1.4.2 Partially-ordered/partially-reliable transport service as an alternative**

What is desirable is a standardized transport service, or a library of functions, modules, or objects, that applications can utilize to gain flexible control over the ordering and reliability of individual objects. Such a service or library would allow applications to achieve an appropriate balance among various QoS parameters without having to “reinvent the transport-layer wheel” with every application. Such an approach is consistent with Application Level Framing (ALF) as proposed in (Clark and Tennenhouse, 1990). This dissertation presents two technologies to address the need for additional flexibility at the transport layer. The first is the POCv2 protocol,

which provides a *partially-ordered/partially-reliable* (PO/PR) transport service. Chapter 2 includes an overview of partially-order/partially reliable transport services, including a summary of previous and related work. The second is a more general mechanism: the *Universal Transport Library* (UTL), which provides a framework for the development of transport services that provide flexible QoS tradeoffs. UTL is described in more detail in Section 1.7.1, and Chapter 3.

In addition to allowing an application to request the precise level of reliability and ordering required, POCv2 and the UTL provide an additional benefit that neither TCP nor UDP provides: a mechanism that facilitates coarse-grained synchronization of multimedia objects. This synchronization feature is described in more detail in Chapter 2.

## **1.5 PO/PR Transport service and multimedia document retrieval**

Previous work on partially-ordered service has produced quantitative results for delay and throughput gains for PO/PR transport service. (Marasli et al. 1996, Marasli et al., 1997a, Marasli 1997b). However, these quantitative results were derived for an abstract PO/PR service, and not directly related to any concrete application. A goal of this dissertation is to put such quantitative measures into an application context so that such results can be interpreted in terms of their impact on an end user. The application chosen for this purpose was remote multimedia document retrieval; that is, retrieving a multimedia document from a server on the Internet, and presenting this document as it is retrieved from the server. The remainder of this section explains why this application is particularly suited to a study of PO/PR services. It also introduces two key themes of our work: (1) the benefits of providing graceful degradation of multimedia documents, and (2) the benefits of

integrating the transport order and reliability features with the coarse grained synchronization mechanism of the multimedia application.

### **1.5.1 Graceful degradation of multimedia documents**

Many systems now exist that allow authors to construct pre-orchestrated multimedia documents. One of the most popular commercial systems is Macromedia Director. The proceedings of the IEEE and ACM Multimedia conferences contain examples of research systems; a survey of such systems appears in Chapter 4.

Multimedia documents consist of objects such as still images, text, audio clips and video clips, which are pre-arranged according to a temporal scenario. Various schemes exist for expressing temporal scenarios (Pérez-Luque and Little, 1995). During the playback of such a document, object presentation proceeds according to this temporal scenario until some event occurs which stops or resets it-- for example, a user interaction point is reached, or a user presses a pause button. Typically, a multimedia workstation with sufficient CPU, memory, and I/O capabilities can present a document in compliance with its temporal scenario, provided that the channel delivering the information is ordered and error-free.

However, suppose the document is stored on a remote file server, and the channel delivering this information is the Internet. In this case, network errors and delays may wreak havoc with attempts to present the document correctly.

We propose that in such situations, it is appropriate to provide for *graceful degradation* of the multimedia document presentation. Graceful degradation is helpful in multimedia documents where not all objects have equal importance, or the same quality-of-service requirements; that is, some objects are essential to document content, while others are nice to have, but optional.

Graceful degradation is also helpful when some objects must be presented in a specific order, while other objects can be presented in an order different from their transmission order, with no loss of quality. For example, in a document describing a simple repair to a piece of equipment, “step 1” should be presented before “step 2”. Now, suppose the same document also contains three images that should be presented roughly simultaneously. If two of them show up, and one has to be retransmitted, in many cases it is desirable to go ahead and present the images that arrived while waiting for the retransmission of the missing image.

### **1.5.2 Traditional transport protocols are not satisfactory**

Given that we want to provide for graceful degradation, what transport protocol should be used for multimedia objects? We argue that classic transport services such as TCP and UDP are ill suited to this application, and investigate PO/PR service as an alternative. The ReMDoR system developed for this dissertation provides users with the capability to author multimedia documents and place them on a server for remote retrieval and display via the Internet. The purposes of the ReMDoR system are:

- to show how a PO/PR transport service facilitates coarse-grained synchronization of multimedia objects and graceful degradation during times of network stress,
- to demonstrate the mechanisms needed to implement a PO/PR transport protocol in practice, and
- to demonstrate and quantify performance improvements when a PO/PR transport service is used instead of an ordered/reliable service (e.g., TCP) or an unordered/unreliable service (e.g., UDP).

### **1.5.3 Traditional authoring systems are not satisfactory.**

Note that the manner in which a document should be gracefully degraded is largely dependent on the intentions of the document author; it cannot be deduced solely from structural aspects of the document such as the type of objects (graphics, sound, text, etc.) For example, consider a sound clip. In one document, a sound clip may be only a sound effect to draw attention to a certain visual image on the screen. If this image is also highlighted in other ways (e.g., with color), then the sound effect may be helpful, but non-essential. It would be acceptable, for example, to allow for small gaps in the sound object (say, gaps of up to 20 ms); these might create small cracks or pops in the sound, but the sound would still serve its intended purpose. On the other hand, if the document were intended as a tutorial for peacekeeping troops to teach useful phrases in the local language, a sound clip might be the most important content in the document. In this case, even a small amount of infidelity in the sound could render the document useless for its intended purpose.

Current commercial authoring systems (e.g., Macromedia Director and Authorware) generally assume either

- a perfect, high-bandwidth channel (e.g., from a local CD-ROM or DVD device), or
- transmission over the Internet (i.e., the Web) using a reliable ordered channel (that is, TCP), with a medium-to-high bitrate (i.e., 33.6–56kbps to 1.544 Mbps or higher).

In general, current systems do not provide the capability to author for graceful degradation; that is, authors cannot specify the relative importance of objects in a document, or specify multiple orders for object presentation.

Some systems for Web authoring do provide the capability to make multiple versions of a document available: for example, a graphics intensive version

for those connecting via a high-speed connection to the Internet, and a text-only version for those with a slow connection. There are new features that make this easier in the newest version (v1.1) of the Hypertext Transfer Protocol (HTTP) used for retrieving Web documents. In these systems, the user (or the user's browser) must choose *in advance* to retrieve a lower-quality version of the document content.

It is certainly useful for a user or a browser to be able to convey to the server that the user already *knows* network conditions are less than ideal, and that the user will accept a lower quality version of the document. However, providing multiple documents for users to choose under different network conditions is fundamentally different from the kind of graceful degradation we address. We address the case where network conditions are either unknown a priori, or suddenly become worse in the middle of a transmission. Our vision of graceful degradation is that when performance is (perhaps) unexpectedly bad, the application and the transport protocol should have already marked the most important information and sequence constraints for preservation. The application and transport entities can then immediately take steps to preserve the most important document elements and relationships, even while discarding or disregarding others. It is this kind of graceful degradation that we explore in the ReMDoR system as we consider how PO/PR transport may be of benefit.

## **1.6 Overview of systems developed for this dissertation**

### **1.6.1 Universal Transport Library (UTL)**

The central goal of this dissertation is to investigate whether PO/PR transport service provides a measurable benefit for some application. To show a

measurable benefit, we need an application that can be demonstrated over more than one transport service. The ReMDoR application described in Chapter 4 was developed for exactly this purpose.

However, early in the development of ReMDoR, we recognized that designing an application that can run over multiple transport protocols presents certain difficulties (as we explain in Section 3.1.) To overcome these difficulties, we developed the Universal Transport Library (UTL). UTL is a library of transport layer software that can be linked in with an application, to provide a range of transport services through a single API. The transport services provided in UTL include simple wrappers for TCP and UDP, as well as a range of PO/PR transport services. The transport layer functionality in UTL is implemented at user-level rather than in the kernel, and sits in between the application and the regular UDP and TCP services provided by the operating system.

UTL provides benefits both for developers of new transport layer services, and developers of applications that want to take advantage of various kinds of transport layer service. For developers of transport layer services and protocols, UTL provides a framework for rapid prototyping of transport layer implementations. For application writers, UTL provides a library of various transport services that can be accessed through a single API.

### **1.6.2 Overview of the Remote Multimedia Document Retrieval system (ReMDoR)**

ReMDoR is a multimedia document retrieval system that allows authors to specify synchronization requirements and varying degrees of reliability for multimedia elements (Conrad et al., 1996, Conrad et al., 1998). The key motivation for ReMDoR

was the investigation of partial order and partial reliability transport protocols in the context of multimedia document retrieval. Because it was infeasible to incorporate partial order and partial reliability into existing multimedia document retrieval systems, it was necessary to develop a simple multimedia document retrieval system in order to carry out this research.

The basic model is similar to that of the World Wide Web; documents are available on a server and are retrieved via a browser. The ReMDoR browser has a look and feel that is similar to traditional web browsers, making experimentation convenient, and helping to demonstrate the idea of multimedia document retrieval in a familiar context. However, unlike Web documents, ReMDoR documents are *temporal*—they have a time dimension requiring synchronization of multimedia elements. ReMDoR has capabilities that support experimentation with innovative protocols and data compression techniques, such as:

- the ability to select from a wide range of transport services and transport service features (via UTL),
- the ability to record statistics about performance on an object-by-object basis
- features to support the automation of repeated performance experiments, and
- the ability to easily incorporate new image formats, such as the formats required for research into network-conscious image compression: Network Conscious GIF (Amer et al., 1999), the SPIHT Wavelet format (Said and Pearlman, 1996; Iren, 1999), and the Network Conscious Wavelet format. (Iren et al., 1998; Iren and Amer, 2000).

### 1.6.3 Publications based on these tools

UTL and ReMDoR have proven useful in research beyond this dissertation, chiefly, in the NETCICATS project (Iren et al., 1998a,1998b,1998c,1999b; Amer et al. 1999). In addition, the implementation work has provided opportunities for one MS Thesis, and one undergraduate Honors thesis. Table 1.1 provides a summary of the various publications, theses, and dissertations based on research that uses UTL and ReMDoR. The remainder of this section summarizes highlights of this related work.

**Table 1.1 Publications related to UTL and ReMDoR implementation**

1st Author	Publications	Citations
Caro	Undergraduate Honors Thesis	CIS Dept., U. Del. (Caro, 1998)
Conrad	Two Conference Papers Two Workshop Papers	MMCN'96 (Conrad et al., 1996) MILCOM'98 (Conrad et al., 1998) IWQoS'97 (Conrad et al., 1997) Sync'95 (Conrad et al. 1995)
Golden	MS Thesis	CIS Dept., U. Del (Golden, 1997)

The *Network Conscious Image Compression and Transmission System* (NETCICATS) project, which is the Ph.D. dissertation work of Sami Iren (Iren, 1999b) uses UTL and ReMDoR as the central tools for conducting performance experiments. *Network-conscious image compression* is an approach to image compression that seeks not solely to maximize compression, but rather to optimize *overall* performance when compressed images are transmitted over a lossy packet-switched network such as a battlefield network. Using an Application Level Framing philosophy, an image is compressed into path-MTU sized Application Data Units (ADUs) at the application layer.

Performance experiments in NETCICATS typically involve comparison of a traditional image compression scheme requiring a more restrictive transport protocol (e.g., Ordered/Reliable service) against an image compression scheme that may require more bits per pixel, but which allows a less restrictive transport protocol. The key idea is to find the loss rates at which the best trade-off is made; that is, the loss rates above which the *penalty* of more bits per pixel is outweighed by:

- the *benefit* of being able to deliver information out-of-order (thus providing better progressive display), *and/or*
- the *benefit* of being able to tolerate a certain degree of loss while still maintaining acceptable image quality.

Thus, for evaluations of NETCICATS, three things are essential:

- (1) The ability to easily put a single application on top of various transport services (an ability provided by UTL)
- (2) The ability to retrieve and display different image formats over various transport services, and add new image formats easily (an ability provided by ReMDoR.)
- (3) Facilities for the automation of repeated performance experiments, and the gathering of performance statistics (also provided by ReMDoR).

Thus, in addition to providing a platform for the evaluation of PO/R transport service vs. O/R transport service, ReMDoR and UTL have been critical to the network-conscious image research of Iren.

UTL and ReMDoR have also provided an opportunity for significant participation by MS and undergraduate students in research, and holds the promise of continuing to do so for future students of the dissertation author. For his MS thesis, (Golden 1997), described the first implementation of the KXP mechanism of UTL (see Sections 3.3.7, 3.6), and introduced a new transport protocol, the Timed Reliability

Unordered Message Protocol (TRUMP). In this protocol, messages can be given timestamps after which their reliability expires. TRUMP is particularly appropriate for military applications such as the Fact Exchange Protocol, where the value of a particular message becomes less valuable with time (Chamberlain, 1994). UTL provides a suitable environment for experimenting with this technique; implementation of TRUMP is among the projects suggested for future work. As an undergraduate honors thesis, (Caro, 1998) did most of the programming work for the second major version of ReMDoR, including the addition of several key features we describe in Chapter 4; he is currently working on a Java interface for UTL. The dissertation author plans to involve students in several UTL and ReMDoR projects following the completion of this dissertation; including porting UTL and ReMDoR to Linux, and building a Java version of ReMDoR over Caro's Java interface to UTL.

## **1.7 Overview of performance experiments**

The performance experiments presented in this dissertation are divided into two major groups, labeled N for NETCICATS, and R for ReMDoR. The NETCICATS group of experiments compares O/R protocols with PO/R and U/R protocols. NETCICATS is built on top of ReMDoR, but concentrates on fixed images with no temporal dimension, and thus uses only a subset of ReMDoR's full functionality. The ReMDoR group of experiments focuses on the difference in performance between O/R and PO/R service. The ReMDoR group incorporates experiments with temporal documents that cover a range of complexity in terms of multimedia content and synchronization requirements. Chapter 5 provides details concerning these experiments. All use the basic setup illustrated in Figure 1.2. This architecture consists of a server and client, with an unreliable network in between.

Two experimental tools developed for this dissertation provide control over the properties of this network: a lossy router for simulating packet loss, and a packet reflector for simulating various bitrates and propagation delays. Section 5.3 provides more details concerning these tools.

## **1.8 Chapter summary**

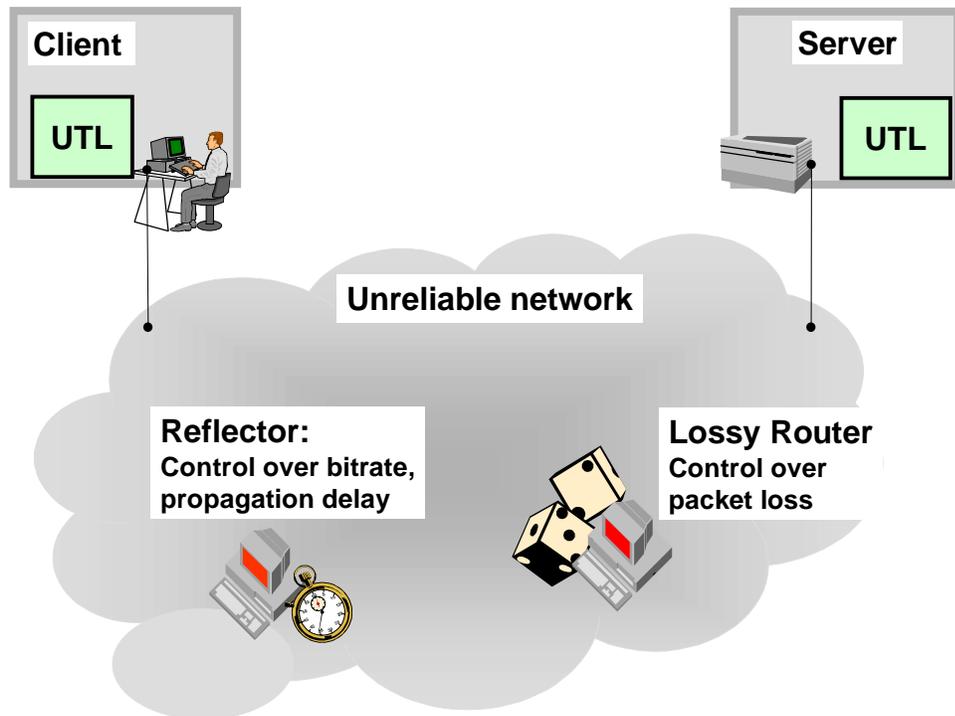
We have presented the problem statement for this dissertation, which is to determine through experimentation with real systems the extent to which partially-ordered/partially-reliable (PO/PR) transport service can provide performance benefits for real applications. The motivation for PO/PR service is to provide applications with more flexibility in making tradeoffs between QoS parameters such as order and reliability, and quantitative QoS parameters such as throughput, delay, and buffer utilization. Previous work in this area used analytic modeling and simulation to predict the performance benefits of PO/PR service, while this work focuses on putting these benefits into a concrete application context so that the impact on the end user can be assessed. Multimedia document retrieval is used as an example application that can benefit from PO/PR service.

We have described two experimental systems designed to conduct experiments comparing the performance of a remote multimedia document retrieval system over PO/PR transport service: (1) UTL and (2) ReMDoR. We provided an overview of the design and results of performance experiments conducted for this dissertation using these tools. We also noted that while the author designed UTL and ReMDoR for his own purposes, they have proven useful beyond the scope of the author's work.

The next chapter (Chapter 2) focuses on *transport protocol ideas*. This is followed by two chapters (Chapters 3 and 4) focusing on *experimental systems*, covering UTL and ReMDoR, respectively. This is followed by a chapter on *performance results* (Chapter 5), a chapter on *algorithms* for PO/PR protocols (Chapter 6), and a *summary and description* of future work (Chapter 7).

**Table 1.2 Performance Experiment Groups**

<b>Grp</b>	<b>Explanation</b>	<b>Evaluation of</b>	<b>number of experiments</b>	<b>described in sections</b>
<b>N1</b>	NETCICATS as motivation for unordered service, limiting case for gains from partial order	U/R vs. PO/R vs. O/R	1	5.2
<b>R1</b>	ReMDoR application, 8 parallel images, slow PPP link, combat net radio speeds.	PO/R vs. O/R	4	5.3
<b>R2</b>	ReMDoR application, 8 parallel images, Narrowband ISDN speed (128kbps)	PO/R vs. O/R	4	5.4
<b>R3</b>	ReMDoR application, 8 parallel images, Narrowband ISDN speed (128kbps), effect of various bitrates	PO/R vs. O/R	1	5.5
<b>R4</b>	ReMDoR application, audio in parallel with images.	PO/R vs. O/R	4	5.6
<b>R5</b>	Excerpt from a complete document with audio, parallel images streams, complex synchronization relationships	PO/R vs. O/R	2	5.7
<b>Total</b>			16	



**Figure 1.2** Experimental setup for performance experiments