

Contents

Kurzfassung	i
Abstract	v
Contents	vii
List of figures	xi
List of tables	xv
Abbreviations and symbols	xvii
1 Introduction	1
1.1 The Internet and its future evolution	1
1.2 Problem statement and contributions of this thesis	2
1.3 Thesis structure	4
2 Communication network architectures	7
2.1 Packet networks	7
2.1.1 Fundamentals and terminology	7
2.1.2 Packet network technologies	8
2.1.3 Performance and Quality of Service	9
2.1.4 Network Quality of Service mechanisms	9
2.2 Internet	11
2.2.1 Historical evolution	11
2.2.2 Design principles	11
2.2.3 Network architecture	12
2.2.4 Protocols	13
2.2.5 Services	14
2.2.6 Network Quality of Service mechanisms	14
2.3 IP-based telecommunication networks	17
2.3.1 Historical evolution	17
2.3.2 Design principles	17
2.3.3 Network architecture	17
2.3.4 Protocols	18
2.3.5 Services	19
2.3.6 Network Quality of Service mechanisms	19

2.4	Future evolution	20
2.4.1	Technological progress and trends	20
2.4.2	Architectural challenges for the Future Internet	20
2.4.3	Research directions	22
3	Performance of broadband interactive applications	23
3.1	Internet and Web applications	23
3.1.1	Classification of applications	23
3.1.2	Web application technologies	24
3.1.3	Performance requirements, metrics, and service level agreements	26
3.1.4	Broadband interactive applications	28
3.2	Existing mechanisms for performance improvement and assurance	29
3.2.1	Server mechanisms	29
3.2.2	Distribution mechanisms	31
3.2.3	Client mechanisms	33
3.3	Performance evaluation methodology	34
3.3.1	Role of performance evaluation	34
3.3.2	Methods for performance modeling, prediction, and measurement	35
3.3.3	Traffic and workload modeling	36
3.3.4	Accurate simulation of real network stacks	38
4	Fast startup congestion control mechanisms	43
4.1	Systematic classification of congestion control methods	43
4.1.1	Terminology and fundamentals	43
4.1.2	Congestion control requirements and design space	47
4.1.3	Classification of end-to-end congestion control methods	49
4.1.4	Classification of network-supported congestion control methods	50
4.1.5	Differences to related mechanisms	53
4.2	State-of-the-art and open issues of Internet congestion control	55
4.2.1	Internet standard solution	55
4.2.2	Interaction with network entities	59
4.2.3	Survey of experimental new algorithms and enhancements	60
4.2.4	Impact of network stack implementations	65
4.2.5	Remaining open and unsolved issues	66
4.3	Fast startups: Definition, motivation, and design principles	70
4.3.1	Definition of fast startup congestion control	70
4.3.2	Motivation and challenges of fast startups	71
4.3.3	Design principles of fast startup congestion control	74
4.4	End-to-end fast startup mechanisms	76
4.4.1	Existing Slow-Start enhancements without speedup	76
4.4.2	Existing optimistic fast startup mechanisms	77
4.4.3	Design of enhanced and new optimistic fast startup schemes	79
4.4.4	Other approaches	82
4.5	Network-assisted fast startup mechanisms	82
4.5.1	Overview of the Quick-Start protocol	82
4.5.2	A new admission control concept: Approval control for Quick-Start	84
4.5.3	Design of improved approval control algorithms for Quick-Start	88

4.5.4	Other approaches	92
4.6	Network-controlled congestion control	92
4.6.1	Network control as a clean slate approach	92
4.6.2	Overview of the eXplicit Control Protocol (XCP)	93
4.6.3	Overview of the Rate Control Protocol (RCP)	94
4.6.4	Other approaches	95
4.7	Functional comparison of fast startup schemes	96
4.7.1	Systematic comparison of Quick-Start, XCP, and RCP	96
4.7.2	Other published comparative studies and related work	100
5	Application integration and implementation issues and solutions	103
5.1	Solutions for the interface design	103
5.1.1	Application interfaces	103
5.1.2	Network interfaces	108
5.2	Proposed mechanisms to avoid interactions with flow control	109
5.2.1	State-of-the-art receive window auto-tuning	109
5.2.2	Possible interactions with fast startup congestion control	109
5.2.3	Proposed solution and its implications	111
5.2.4	More disruptive alternative solutions	112
5.3	Realization complexity and feasibility	113
5.3.1	Overview of the implementation work	113
5.3.2	Implementation in endsystems	114
5.3.3	High-speed implementation in network components	117
5.3.4	Lessons learned from the case study	118
6	Performance evaluation	121
6.1	Evaluation methodology and tools	121
6.1.1	Simulation scenarios	121
6.1.2	Measurement setups	124
6.1.3	Experiments in a test network	126
6.2	Simulation tool validation	126
6.2.1	Scenario selection	126
6.2.2	Validation scenario: Response function	127
6.2.3	Validation scenario: Head-of-line blocking	129
6.2.4	Summary of the validation experiments	130
6.3	Study of functional design aspects of fast startups	131
6.3.1	Comparison of the startup behavior	131
6.3.2	Differences in convergence and bandwidth sharing	133
6.3.3	Design and evaluation of new algorithms for Quick-Start	135
6.4	Quantification of the potential performance improvement	141
6.4.1	New analytical model for fast startup performance	141
6.4.2	Comparison of the potential speedup	144
6.4.3	Bandwidth sharing properties	148
6.4.4	Studies with synthetic source-level Web models	152
6.4.5	Studies with trace-based workloads	153
6.4.6	Summary of the performance experiments	156
6.5	Robustness, fairness, and risk	157

6.5.1	Dealing with small buffers	157
6.5.2	Fairness compared to TCP's default congestion control	158
6.5.3	Robustness against imprecise information	160
6.6	Complexity and costs of network support	163
6.6.1	Computational overhead of endsystem functions	163
6.6.2	Computational overhead of router functions	164
7	Applicability case studies	167
7.1	A new Web performance requirement signaling architecture	167
7.1.1	Congestion control with response time deadlines	167
7.1.2	Proof-of-concept realization	169
7.1.3	Exemplary measurement results	170
7.2	Speedup of 3D visualization applications	171
7.2.1	Network challenges of interactive 3D applications	171
7.2.2	Architecture of an interactive 3D visualization application	172
7.2.3	Exemplary measurement results	173
8	Conclusion	175
8.1	Summary	175
8.2	Outlook	179
A	Appendix: Mathematical background	183
A.1	Distribution functions	183
A.1.1	Exponential distribution	183
A.1.2	Pareto distribution and variants	183
A.1.3	Log-normal distribution	184
A.2	Analytical lower bound for TCP resequencing delays	184
A.2.1	Scope of the model	184
A.2.2	Model details	184
B	Appendix: Documentation of parameters	185
B.1	Measurement setups	185
B.2	General Linux kernel configuration	185
B.3	Default configuration of the fast startup schemes	186
B.3.1	Jump-Start	186
B.3.2	Initial-Start	186
B.3.3	Mega-Start	186
B.3.4	Quick-Start	186
	Bibliography	189