


# Effective Network Complexity



Tolga Uzuner

Kings College

University of Cambridge

A thesis submitted for the degree of

*Doctor of Philosophy*

Yet to be decided

---

## **DECLARATIONS**

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text.

This thesis does not exceed the word limit for the Computer Lab Degree Committee.

To my wife and parents,  
in friendship, gratitude, and admiration

## Acknowledgements

I would like to acknowledge my two advisors Prof. Jon Crowcroft and Prof. Derek McAuley. A thesis is the virtual act of connecting dots. They provided the dots.

Jon vast experience with model and systems, and the small nudges and nods in the right direction averted certain disaster. It is difficult to overestimate the value of his insights and gut intuition in this thesis. He also participated in (i.e. put up with) some of the more experimental episodes of my PhD experience at Cambridge. He is also the original instigator of this thesis. Jon is an adventurer.

Derek has a tremendous experience with physical systems and the research process. He laid out the methodology for this thesis, and spent hours on the blackboard sketching out the backbone. The initial set of questions he set out for me to answer were the compass for this ship.

My only regret is that I was not able to spend more time with them during these four years.

I also acknowledge helpful discussions and collaboration with Richard Gibbens, Ian Pratt, Stephen Allott and Andrew Warfield at the Computer Lab, Cambridge, Derek Stagg at Marconi, and Burkhard Stiller and others at the Dagstuhl 2006, Internet Economics workshop. Yao Song Ng at Credit Suisse was extremely helpful in providing the necessary network data. The guidance of my undergraduate advisors Patrick Winston and Olivier Blanchard, friendships with Gian Carlo Rota and Irvin Schick, and work with Richard Marcus at the Laboratory for Information and Decision Systems at MIT were part of the

inspiration for this thesis. I would also like to acknowledge helpful guidance from my thesis committee members, Prof. Chris Cooper and Emanuele Giovanetti.

---

*In the sound of the bell of the Gion Temple echoes the impermanence of all things. The pale hue of the flowers of the teak-tree show the truth that those who prosper must fall. The proud ones do not last long, but vanish like a spring night's dream. And the mighty ones too will perish in the end, like dust before the wind.*

From the opening lines of the “Tale of the Heike”

## Abstract

Architectural trends in the last decade are leading to increasingly complex control logic that can boost the number of services offered and make use of the Internet Protocols ubiquity. The result is a wider architecture and a larger feature set in the router. Managing the architecture of these networks necessitates a deeper understanding and quantification of this complexity. This thesis transgresses the boundaries of telecommunications into manufacturing and finance to address these issues.

An argument against the naive use of financial pricing and neoclassical market metaphors common in the resource allocation literature is followed by a proposal to ReReshape the Research Agenda around the networks' *Entropy* as an objective. A complexity-centric mechanism for networks based on *Entropic Routing* is proposed, and the computational role of institutions is investigated.

*Effective Network Complexity* is analysed using two models. An information theoretic *Architectural Model* gives insights into typical architectural decisions faced by a network engineer. A practical, real-time *Measurement Model* using techniques from contingent claims analysis adds additional insight, and is able to separately measure idiosyncratic and systematic risks, a key contribution. It can be deployed on any network using easily available *Reference Measurements*. Empirical analysis follows.

Insights into the complexity of networks form the backbone of this thesis. Implications are analysed and the vernacular required to reason about and make decisions in the *Complexity Plane* of a network is developed.

# Contents

<b>Nomenclature</b>	<b>1</b>
<b>1 Introduction</b>	<b>3</b>
1.1 Motivation . . . . .	3
1.2 Formulating the Problem . . . . .	5
1.2.1 Intractability in Networks . . . . .	7
1.2.2 Not So Difficult Problems in Networking . . . . .	11
1.3 The Context . . . . .	12
1.4 Contributions . . . . .	13
1.4.1 Technical . . . . .	13
1.4.1.1 Complexity Measures . . . . .	14
1.4.1.2 Risk: Introducing Slice . . . . .	15
1.4.1.3 Novel Concepts . . . . .	18
1.4.2 Non-Technical . . . . .	18
1.5 The Approach . . . . .	20
1.5.1 Looking over the Fence . . . . .	20
1.5.2 Information Theory . . . . .	22
1.5.3 Market Metaphors . . . . .	24
1.5.4 Modelling . . . . .	25
1.5.4.1 Effective Complexity . . . . .	26
1.6 What Will Not Work ? . . . . .	27
1.6.1 Naive Market Solutions . . . . .	28
1.6.2 Naive Graph Theoretic Approaches . . . . .	29
1.6.3 Negotiation and Fixation on the Agent . . . . .	31
1.6.3.1 The Complexity of Negotiation . . . . .	32



1.6.3.2	The Complexity of Planning . . . . .	32
1.7	Focus Areas . . . . .	33
1.7.1	Traffic Engineering . . . . .	34
1.7.2	Routing . . . . .	36
1.7.2.1	Domain-based Routing . . . . .	38
1.7.3	The Internet Protocol . . . . .	40
1.7.4	Areas Not Addressed . . . . .	40
1.8	Alternative Approaches . . . . .	41
1.8.1	Scheduling Complexity as a Metric . . . . .	42
1.8.2	Novel Graph Theoretic Techniques . . . . .	43
1.8.3	Alternative Entropy Formulations . . . . .	44
1.9	Related Research . . . . .	44
1.9.1	Event Correlation . . . . .	44
1.9.2	Networking Games . . . . .	45
1.9.3	Market-Based Metaphors . . . . .	46
1.10	Relevant Trends . . . . .	48
1.10.1	Routing . . . . .	48
1.10.2	MPLS Developments . . . . .	49
1.10.3	Hardware . . . . .	50
1.10.4	Summary . . . . .	53
1.11	Roadmap . . . . .	53
1.12	ReReshaping the Agenda . . . . .	55
<b>2</b>	<b>Modelling Networks</b>	<b>57</b>
2.1	Imperfect Markets . . . . .	58
2.1.1	The Tâtonnement Problem . . . . .	59
2.1.2	Market Clearing . . . . .	60
2.1.3	Equilibrium and Uniqueness . . . . .	60
2.1.4	Incomplete Markets . . . . .	61
2.1.5	Limits to Computation . . . . .	62
2.1.6	Limits to Information . . . . .	65
2.1.7	Deliberation Costs . . . . .	66
2.1.8	The Implication for Networking Research . . . . .	67

2.2	The Entropic Formulation . . . . .	68
2.2.1	The Minimax Approach . . . . .	70
2.2.1.1	Minimax Applied to Networks . . . . .	73
2.2.2	The Statistical Mechanics Approach . . . . .	74
2.2.2.1	Relationship with Neoclassical Economics . . . . .	76
2.2.2.2	Statistical Mechanics Applied to Networks . . . . .	77
2.2.3	The Informational Efficiency Approach . . . . .	80
2.2.3.1	Market Efficiency and Entropy . . . . .	82
2.3	Dissecting the Entropic Formulation . . . . .	83
2.3.1	The Relationship Between Entropy and Utility . . . . .	86
2.3.1.1	The Role of Arbitrage . . . . .	88
2.3.1.2	Proportional Fairness . . . . .	89
2.3.2	Log Utility is Not Arbitrary . . . . .	92
2.3.3	Entropy Flooding . . . . .	94
2.4	The Analysis of Non-Equilibrium States . . . . .	94
2.4.1	The Role of Institutions . . . . .	96
2.4.2	The Role of Diversity . . . . .	99
2.4.3	Complex Adaptive systems . . . . .	100
<b>3</b>	<b>Dissecting Complexity</b>	<b>103</b>
3.1	Overview of Complexity Measures . . . . .	104
3.1.1	General Complexity Measures . . . . .	105
3.1.2	Higher Order Methods . . . . .	107
3.1.3	System-Specific Complexity Metrics . . . . .	107
3.1.4	Discussion . . . . .	109
3.2	The Indignity of Numerical Simulations . . . . .	110
3.3	Understanding Complexity . . . . .	113
3.3.1	Why Another Measure of Complexity ? . . . . .	113
3.3.2	Why An Information Theoretic Formulation ? . . . . .	114
3.3.2.1	An Axiomatic Definition of a Measure of Informa- tion . . . . .	115
3.3.2.2	Interpretation . . . . .	116
3.3.3	The Architectural Model . . . . .	117

3.3.3.1	Simple Version . . . . .	118
3.3.3.2	Complicated Version . . . . .	119
3.3.4	Breaking Down Architectural Complexity . . . . .	121
3.3.5	Concluding the Architectural Model . . . . .	123
3.4	Measuring Complexity . . . . .	123
3.4.1	The Measurement Model . . . . .	126
3.4.2	Measuring Complexity . . . . .	129
3.4.3	Measuring Risk . . . . .	132
3.4.3.1	A Slice of Risk . . . . .	132
3.4.3.2	Sensitivity Analysis . . . . .	134
3.4.4	Extensions . . . . .	134
3.4.4.1	Heterogeneity of Elements . . . . .	135
3.4.4.2	Heterogeneity of Correlations . . . . .	136
3.4.4.3	Heavy-tailed Distributions . . . . .	137
3.4.4.4	Modelling sub-systems . . . . .	138
3.4.4.5	Dynamic Complexity . . . . .	139
3.4.4.6	Multiple Factors . . . . .	140
3.4.4.7	Model Complexity . . . . .	140
3.4.5	Discussion . . . . .	141
3.4.5.1	The Subjective Nature of Complexity . . . . .	143
3.4.5.2	On Reference Measurements . . . . .	143
3.4.5.3	Assuming Away the Idiosyncracies . . . . .	144
3.4.5.4	On Factor Models . . . . .	145
3.5	Implying Complexity . . . . .	146
3.5.1	Zero Complexity . . . . .	147
3.5.1.1	Implementation . . . . .	149
3.5.2	Implied Multivariate Distribution . . . . .	150
3.5.3	The Complexity Plane . . . . .	152
3.5.4	Implied Parameters . . . . .	152
3.6	Models Discussion . . . . .	155
3.6.0.1	Limitations and Issues . . . . .	156
3.6.1	Effective Complexity: The Full Picture . . . . .	157

<b>4</b>	<b>Insights</b>	<b>160</b>
4.1	The Impact of Complexity . . . . .	160
4.1.1	Simple Version . . . . .	160
4.1.2	Complicated Version . . . . .	163
4.1.2.1	Impact of the Network Precedence Matrices . . .	164
4.1.2.2	Size Matters . . . . .	164
4.1.2.3	Complexity and Architecture . . . . .	165
4.1.2.4	Traffic Mix and Optimal Control Policy . . . . .	165
4.1.2.5	Scheduling Complexity Insights . . . . .	166
4.1.2.6	Architectural Complexity Conclusions . . . . .	168
4.2	Measurement Model Insights . . . . .	168
4.2.1	Heterogeneity of Trigger Probabilities: Impact on <i>Slice</i> . .	169
4.2.2	Coupling Heterogeneity . . . . .	169
4.2.2.1	The Causes of <i>Network Complexity</i> ? . . . . .	170
4.3	Completing the Circle . . . . .	174
4.3.1	Heterogeneity, Complexity and Risk . . . . .	174
4.3.2	Practical Issues . . . . .	176
4.3.2.1	Making Architectural Decisions . . . . .	176
4.3.2.2	Operating an Actual Network . . . . .	177
<b>5</b>	<b>Model Evaluation</b>	<b>179</b>
5.1	Evaluating the <i>Architectural Model</i> . . . . .	179
5.1.1	Connection-oriented vs. Connection-less Paradigms of Net- working: OSPF and MPLS for Traffic Engineering . . . . .	181
5.1.2	Multi Path Routing . . . . .	183
5.1.3	Rolling Features into the Address Space . . . . .	184
5.1.4	ATM vs. MPLS: Flexibility vs. Complexity . . . . .	185
5.1.5	Concluding the Evaluation of the <i>Architectural Model</i> . . .	188
5.2	Evaluating the Measurement Model . . . . .	188
5.2.1	Source Data . . . . .	188
5.2.2	The Creation and Selection of Reference Measurements . .	189
5.2.2.1	Using the Model . . . . .	191
5.2.2.2	Calculating the <i>Slice Curve</i> . . . . .	197

5.2.3	Concluding the Evaluation of the <i>Measurement Model</i> . . .	199
<b>6</b>	<b>Implications</b>	<b>201</b>
6.1	Networking Implications . . . . .	201
6.1.1	Entropic Routing . . . . .	202
6.1.2	The Red Queen Principle and Maximum Entropy Production	204
6.1.2.1	Separating Routing from Routers . . . . .	205
6.1.2.2	End of the End-to-End Argument . . . . .	205
6.1.3	Network Planning . . . . .	206
6.1.4	Element Design . . . . .	206
6.1.5	Infrastructure . . . . .	207
6.2	Financial Implications . . . . .	208
6.2.1	Insurance . . . . .	208
6.2.2	Impact on Strategy . . . . .	209
6.3	Murphy's Law of Networks . . . . .	210
<b>7</b>	<b>Conclusions</b>	<b>211</b>
7.1	Future Work . . . . .	215
7.1.1	Empirical . . . . .	215
7.1.2	Routing . . . . .	216
7.1.3	Modelling . . . . .	217
7.1.4	Applying the Entropic Formulation . . . . .	218
7.1.5	Looking Back over the Fence . . . . .	218
<b>A</b>	<b>Rabbits and Traps</b>	<b>220</b>
<b>B</b>	<b>Entropy</b>	<b>224</b>
B.1	Maximum Entropy and Relative Entropy . . . . .	225
B.2	Axiomatic Formulations of Entropy . . . . .	226
B.3	Principle of Maximum Entropy . . . . .	226
B.4	Jensen's Inequality . . . . .	227
B.5	Properties of Shannon's Entropy . . . . .	227
B.5.1	Kraft Inequality . . . . .	228
B.6	Renyi Entropy . . . . .	229

B.7	Connections to Other Areas . . . . .	230
<b>C</b>	<b>Brief Overview of Statistical Mechanics</b>	<b>231</b>
<b>D</b>	<b>Perfect Markets</b>	<b>234</b>
D.0.1	Neoclassical Economics . . . . .	235
D.0.1.1	Walrasian Auctions . . . . .	236
D.0.1.2	Arrow-Debreu . . . . .	236
D.0.2	Pareto Optimality . . . . .	238
D.0.2.1	First Fundamental Theorem of Welfare Economics	238
D.0.2.2	Second Fundamental Theorem of Welfare Economics	238
D.0.3	Lack of Arbitrage and Market Completeness . . . . .	239
D.0.4	The Premises of Neoclassical Economics . . . . .	240
<b>E</b>	<b>Further Results</b>	<b>242</b>
	<b>Glossary</b>	<b>249</b>
	<b>References</b>	<b>255</b>
	<b>Index</b>	<b>271</b>

# List of Figures

1.1	A Slice of a Distribution . . . . .	17
1.2	Traffic Engineering . . . . .	35
1.3	Routing Protocols . . . . .	37
3.1	Empirical vs. Model CDF . . . . .	130
3.2	Empirical vs. Model Density . . . . .	131
3.3	Effect of Increasing $K_2$ . . . . .	133
3.4	Empirical and Model Entropy, Average Latency 50ms . . . . .	142
3.5	Slice Curve slope impact on the Probability Distribution . . . . .	153
3.6	Changing the Level of the Slice Curve . . . . .	154
3.7	Relationship between Complexity Measures . . . . .	159
4.1	Complexity against Throughput . . . . .	162
4.2	Entropy ag Trigger Probabilities . . . . .	170
4.3	Entropy ag Beta . . . . .	171
4.4	Entropy ag Beta and Trigger Probability . . . . .	172
4.5	Entropy ag Beta and Trigger Probability, Small . . . . .	173
4.6	Relationship between Heterogeneity and Complexity . . . . .	175
5.1	Empirical CDF, Average Latency 50ms . . . . .	192
5.2	Empirical PDF, Average Latency 50ms . . . . .	193
5.3	Sliding Entropy, Average Latency 50ms . . . . .	194
5.4	Network Complexity, Average Latency 50ms . . . . .	195
5.5	Empirical and Model Entropy, Average Latency 50ms . . . . .	196
5.6	Empirical and Model Density Average Latency 50ms . . . . .	197
5.7	Implied Coupling Average Latency 50ms . . . . .	198

## LIST OF FIGURES

---

5.8	Implied Coupling across Reference Measurements . . . . .	198
5.9	Slice Curve Average Latency 50ms . . . . .	199
5.10	Slice Curve, First Derivative Average Latency 50ms . . . . .	200
E.1	Latency Standard Deviation, 25-day Moving Windows, 50-ms Triggers . . . . .	243
E.2	Error Rates, 25-day Moving Windows . . . . .	244
E.3	Discarded Packets, 100-day Windows . . . . .	245
E.4	Firewall Utilisation, 25-day Moving Windows . . . . .	246
E.5	Firewall Average Latency, 25-Day Moving Windows . . . . .	247
E.6	Peak Inbound Traffic, 25-day Moving Windows, 35-ms Triggers . .	248



# List of Tables

2.1	Key Characteristics of Networks . . . . .	<a href="#">81</a>
5.1	Typical Thresholds for Network Measurement Data . . . . .	<a href="#">190</a>

# Nomenclature

## Roman Symbols

$\epsilon_i$	A Normally distributed noise term associated with an <i>Element</i>
$\Lambda$	Network Precedence Matrix
$\lambda$	Elements of the Network Precedence Matrix
$\nu$	Traffic Vector
$\Phi$	Cumulative Distribution function for the Standard Normal Distribution
$\phi$	Density of the Standard Normal Distribution
$\pi$	Vector of operational constraints on each <i>Elements</i>
$\psi$	The overall unit for measuring <i>Architectural Complexity</i> , summarising, for each <i>Traffic Flow</i> , the resource sharing required in the network, the scheduling requirements, and it's overall ratio to other flows.
$\tilde{\pi}$	$\pi$ normalised accross all <i>Elements</i>
$\tilde{\psi}$	$\psi$ normalised accross all <i>Traffic Flows</i> , <i>Elements</i> , and <i>Operations</i> . Can be interpreted as a probability measure.
$\Upsilon$	Operational Matrix
$A_i$	Utilisation of <i>Element i</i> , due to systematic ( $M$ ) and idiosyncratic ( $X_i$ ). When $A_i$ exceeds $C_i$ , a <i>Reference Measurement</i> of (1) is triggered. The probability of this measurement being triggered is $p_i$ .

## LIST OF TABLES

---

$a_i$	A measure of the coupling between <i>Element i</i> and the common factor $M$ , also called <i>Implied Coupling</i>
$b$	number of <i>Traffic Flows</i>
$C_i$	The capacity of <i>Element i</i>
$K_1$ and $K_2$	Lower and upper thresholds for the <i>Slice</i> metric respectively
$l$	index of <i>Traffic Flows</i>
$M$	A factor common to all <i>Elements</i> which drives capacity utilisation due to systematic factors
$m$	number of <i>Elements</i>
$p_i$	Trigger probability for <i>Element i</i>
$q$	number of <i>Operations</i>
$u, v$	index of <i>Operations</i>
$w_l$	Measure of the quantity of a traffic flow $l$
$X_i$	A factor specific to <i>Element i</i> which drives usage in due to idiosyncratic factors not common to other <i>Elements</i>
$z$	index of <i>Elements</i>

# Chapter 1

## Introduction

### 1.1 Motivation

For the past decade, the general approach for improving the performance of networks has been to build architectures with increasingly complex management logic that can boost the number of services offered and make use of the Internet Protocols ubiquity. The increase in complexity results from a wider architecture, a bigger feature set in the management suite, and more complex services.

However, throughput requirements from the data plane conflict with relentless feature creep in the router. Increasingly complex management logic requires longer dependencies and deeper critical paths in the router architecture.

Traffic and technology trends also suggest that the number of invariants characterising the network will decrease over time. Furthermore, direct and immediate control over the high capacity trunks underlying the networks infrastructure are likely to flood existing access architectures not designed to deal with sudden, unpredictable increases in traffic. This suggests that increasing the feature set of existing control units may not deliver higher performance in the future.

In summary, there is a trade-off between network management logic complexity, feature set size, and network performance (speed and robustness, jointly) that needs to be carefully examined while designing improved architectures. This thesis examines ways to measure and understand this trade-off.

The above discussion underscores the need for investigating network architectures that judiciously use complexity for exploiting significant levels of in-

frastructure sharing while permitting a faster and more robust network. Such architectures are risk-return optimised: they attempt to maximise the ratio of risks to rewards rather than push the envelope for each term separately.

While designing for total risk and complexity effectiveness are desirable, the question remains: how do we quantify the architectural structure of a network ?

It is commonplace to measure the performance of a new architecture, typically by using simulation. Such simulations measure network throughput and performance in a direct manner. However, they are limited in scope not only because of sampling biases and the assumption related to architecture and traffic patterns, they usually implicitly assume that network characteristics remain constant over time. What is needed is a framework for quantifying complexity and risk under the assumption of constant change, not in the absence of it.

This thesis takes a step in the direction of casting architectural decisions optimised for complexity and risk in the context of network evolution. Simple analytical models are developed that quantify the tradeoffs in terms of architectural parameters like traffic mix, element-choice, operational flexibility and resource sharing.

An important component of process improvement is the ability to measure. While the need for such metrics is particularly acute during the design stages, for example during green-field build-out for an operational network, or for design exploration, it remains topical during the operational life-cycle as well. Furthermore, it not just network characterisations, but sensitivities to marginal changes, and deviations from optimal control policies that are required for any optimisation to take place around those parameters. Previous research in related fields fails to address the overall issue of architecture: most metrics are developed specific to un-generalisable problems in specific domains, and are too labour-intensive to collect.

This addresses these needs through the development and implementation of a new suite of metrics for quantifying the impact of network architecture in a precise and novel-way. The dual concepts of *complexity-neutral* architectures and *Entropic Routing* are introduced, to make clear the intuition behind a network optimised around a risk-return ratio in the context of complexity.

## 1.2 Formulating the Problem

This thesis attempts to solve the following problem: how does one control and understand the complexity of large-scale networks.

But first, what is meant by the term *complexity* ? In the context of the work in this thesis, the most succinct description of complexity in a network is *Interrogation Costs*: the difficult of determining the state of a network at any point in time. One easy and theoretically appealing measure of this is the expected number of binary questions a hypothetical agent must ask in the process of this determination.

Intensive resource sharing and element diversity is the key property of networks in the two decades. This makes realistic simulation of today's networks practically impossible (see (BT89)) . More importantly, such networks can never provide guarantees: this applies not only to a service crossing interconnected networks, but even to networks internally.<sup>1</sup>

This is not to say that the industry does not try to impose some semblance of discipline: modelling approaches such as Object Modelling Technique (OMT), Booch technique or Unified Modelling Language (UML) offer an abstract, highly structured and understandable view of the functionality of a telecommunication system. Established software and object-oriented design metrics are used to assess

---

<sup>1</sup>From the Associated Press: "New York, NY, Nov. 11, 1996 (AP) – In the latest outage to hit the Internet, a computer glitch at AT&T's on-line service prevented more than 200,000 customers from receiving e-mail for more than a day, AT&T Corp. said Friday.

The trouble, which started at 2:35 p.m. EST Thursday, was traced to a computer that handles electronic mail for AT&T's WorldNet, the nation's second-largest provider of access to the Internet. The problem had not been resolved by Friday night.

The problem affected about half of WorldNet's 425,000 subscribers. It did not hurt their ability to send e-mail, except if they tried to send mail to another WorldNet customer affected by the brownout.

The Internet brownout was the latest in a series of challenges for AT&T, which faces fresh competition in a de-regulated telecommunications industry and criticism from Wall Street for recently hiring a relative unknown as president and likely heir to the chief executive's post.

AT&T said it was the biggest problem to hit its online service since it launched WorldNet last March." Control mutation originating in the switching system propagated through the signalling network causing degradation of the operation and ending in a total shutdown, the financial loss amounting to 1 billion dollars.

complexity and coupling. As a matter of fact, research has demonstrated positive correlation between the complexity of individual elements, and their coupling with risks to operational control. Such metrics have been used to identify risky areas for redesign consideration.

There are two specific issues this thesis attempts to deal with:

- When controlling the complexity of large-scale networks, one of the real difficulties is in building reliable traffic models (see (WKZL96)). For example, the model must be a worstcase characterisation of the source to provide an absolute upper bound on a source's packet arrivals. Second, the model must be parameterised so that a source can efficiently specify its traffic characterisation to the network. Next, the model should characterise the traffic as accurately as possible so that the admission control algorithms do not overestimate the resources required by the connection. Finally, the model must be police-able so that the network can enforce a source's traffic characterisation.<sup>1</sup> In today's real-world networks, it is practically impossible to build traffic models like this. Another characterisation of the traffic is needed.
- Another issue is understanding what large-scale networks actually look like. However, Internet topology at the router level typically generates graphs which are in the megabytes. Even medium-size AS graphs can be 2-300K. It is graphically impossible to understand these graphs. Given the size and nature of these graphs, new analysis methods have to be employed (see (PSF<sup>+</sup>01))<sup>2</sup>.

As a final consideration, high-level analysis of an existing large-scale network composed of tens of thousands of existing elements already in place can be a practical problem faced by many engineers who inherit a legacy network. What if you could only look through a keyhole? The machinery of this thesis implies the performance and complexity characteristics of a network from empirical data,

---

<sup>1</sup>This suggests that policability is another candidate metric.

<sup>2</sup>The Internet and most large-scale data networks are surprisingly resilient to random link and router failures: robustness and resilience are not considered at length.

and hence makes decision-making possible without requiring detailed information on every element in the network.

### 1.2.1 Intractability in Networks

There are three areas of difficulty in data networking:

1. Traffic predictability
2. Network characterisation
3. Optimal Control

**Traffic Predictability** Voice networks are governed by a number of pretty well-defined mathematical “universal laws”:

- call arrivals are mutually independent and follow a Poisson process
- call inter-arrival times are exponentially distributed

These conjoined with the following assumptions about voice telephony make the use of queueing mathematics highly useful and reliable:

- growth rates are highly predictable
- network controls and operations are fully centralisable
- services are strictly regulated and monitored

The resulting traffic models are mathematically tractable and sometimes obviate the need for maintaining measurement infrastructure in voice networks. Crucially, in large voice networks, the law of large numbers can be used to make the aggregate properties of a system over a long period of time much more predictable than the behaviour of individual parts of the system.

Queueing theory, originally developed for circuit-switched networks which serviced voice traffic, is partially applicable to packet-switched networks which carry data.



## 1.2 Formulating the Problem

---

However, a notable difference between voice networks and data networks is self-similarity, a consequence of connections being carried between computers, and not people. This means that packet networks, also increasingly used for voice telephony as well, behave very differently. First of all, data traffic has statistical characteristics which are dramatically different: usually much longer tails to the distribution and significantly higher levels of variation. Furthermore, the use of packet technology causes frequent and bursty congestion. The protocols used for transmitting data all use some form of congestion control which is shaped by the conditions connections have faced in the past. This further introduces significant correlations across time as well as complex interactions. For references, see (FGHW99), (PF95), (WP98) and (CB96).

**Network Characterisations** The third fundamental problem with current networks is measurement.

This difficulty has two subcomponents:

1. Change: Most closed networks face changes due to shifting requirements from applications. Interconnection with new media, such as wireless, are another source of constant change.
2. Evasiveness: Almost all useful metrics for the purposes of understanding something about the graph theoretic characteristics of even small networks require that each node be visited (see (GY02)).

**Optimal Control** Optimising around constraints in a network is a difficult problem. This is the starting point for most networks today, not only at the inter-AS level, but arguably also at the intra-AS level. The sheer size and complexity of most networks today, coupled with hardware trends which emphasise speed over coordination, imply that certain goals are not practically attainable.

The current architecture of most networks, including the Internet, makes certain things not only practically, but also theoretically impossible.

The list below attempts to cover all the major issues:

**Optimal Efficiency** Under the assumption of selfish nodes and the control of central administration only at the switch level, the classic result from (She95) is that no service discipline can guarantee optimal efficiency <sup>1</sup>.

**Optimal IP Routing** Perhaps the most important result in data networking is that associated with Optimal IP routing. Optimal IP routing is defined over a network  $G(V, E)$ , a demand matrix  $D$  and capacities  $c_e$  using source invariant routing as the problem of minimising network congestion. The proof is based on a bijection between this problem and the *subset sum* (see (GJ79)), defined as the problem of finding a subset of element  $i = 1, \dots, n$  with size  $s(i) \in \mathbb{Z}^+$  summing to  $B$ .

**Optimal Flow Routing Problem** The same bijection also works for the Optimal Flow Routing problem, where the goal is the find a routing assignment with maximum throughput.

**Optimal Link Weights** Finding optimal link weights that lead to a routing minimising the total routing cost for a given traffic demand is a NP-Hard problem (see (FT00)), even for the case of a linear objective function. It is NP-hard to find a weight setting within a factor 1.77 from optimality for arbitrary graphs. The worst case may be up to 5000 times the cost of optimal routing.

**Optimal Splitting** minimum average delays for multi-path routing can be guaranteed using distributed computation (see (Gal77)). However, this algorithm is practically impossible to implement because of its infinitely divisible and stationary traffic assumptions.

**Robustness** Equivalent to the task of determining the maximum number of disjoint sets out of some given set, known to be NP-hard.

**Completeness and Consistency** And other issues regarding the convergence of routing algorithms (see Section 1.7):

- Reachability (from one node to another) is NP-complete

---

<sup>1</sup>However, there are service disciplines which do guarantee fairness.

## 1.2 Formulating the Problem

---

- Asymmetry (of a path from one node to another, and back) is NP-complete
- Solvability (the existence of a final state) is NP-complete, Un-solvability, Solvability, and Un-solvability are NP-hard
- Uniqueness is NP-hard

Control-based approaches also face significant practical difficulties. Take for example, the issue of guarantees at the flow level. No matter how the problem is specified, the core will need to perform some kind of regulation and scheduling based on available resources and deadlines. These are expensive calculations: on the data path, flow classification and scheduling is difficult at today's data rates. On the control path, maintaining consistent and dynamic state in today's distributed networks spanning literally tens of thousands of network elements is also very hard.

Control algorithms based on routing, exclusion, and scheduling are broadly divided into two categories:

**Access Control** Use mechanisms at the entry point of the network to enforce pre-agreed constraints.

**Flow Control** Use feedback from the network to maintain a tradeoff between network utilisation and adherence to a pre-agreed constraint

The main difficulty with access control is knowing the state of the network: under incomplete information, access control can lead to inefficient usage of a network, frequently leading to worst-case performance.

Very high speeds are the primary problem with flow control, leading to propagation delays in the feedback loop and algorithmic complexity. In (KL92), the authors claim that algorithmic complexity can potentially become a bottleneck in the future. They claim that simple controls with good enough performance will bring a tradeoff between efficiency and complexity.

These are well-known, and frequently identified problems in networking. Less appreciated, but perhaps more important is the issue of calculating the probability of cascading events in a network. This problem is NP-Complete: one needs to

calculate the probability that no more than a certain number of elements in a network trigger an event which is a Bernoulli-(0,1):  $Prob(X \leq K)$ . This is the same as asking if there exists a subset of all elements whose trigger measurements (a Bernoulli-(0,1)) add to  $K$ . This decision is equivalent to the SUBSET SUM problem, SP13 in (GJ79). In Section 3.4.3, a parametric measure of network risk is developed using conditional independence mechanics.

### 1.2.2 Not So Difficult Problems in Networking

That having been said, not all problems in networking are impossible to solve. When the demand matrix or traffic conditions can be predicted, certain problems become easy.

At the top of these are voice networks governed by the Universal laws of Section 1.2.1. The fact that very few parameters govern their traffic models makes them enormously tractable.

Certain issues one is faced with during green-field build-outs are relatively easy to resolve. For example, constant-approximations for a number of layered network design problems exist:

- Hierarchical caching, where caches are placed in layers and each layer satisfies a fixed percentage of the demand (bounded miss rates)
- Multi-level facility location problems, where constant combinatorial approximations exist.
- Load Balancing, where each open facility must serve at least a certain amount of demand

These problems are typical in the domain of the Access Network Design problem. See (CW99).

Even though most versions of the Steiner tree problem are NP-Complete, a number of special cases with important applications are solvable:

- single-sink buy-at-bulk with variable pipe types between different sets of nodes

- facility location with buy-at-bulk type costs on edges
- constructing single source multi-cast trees with good cost and delay properties

There are a number of other approximation solutions to important problems in networking, some of which are based on the idea of randomisation, and others on limiting the space of possible configurations. (AL95) is a good compendium of these results. For example, the problems above typically involve finding a Steiner tree which optimises the sum of edge costs along one metric and the sum of source-sink distances along an unrelated second metric. This problem has an  $O(\log k)$  randomised approximation scheme.

## 1.3 The Context

John Doyle has argued that there are two great abstractions in the 20th century:<sup>1</sup>

- Separate systems engineering into control, communications, and computing
- Separate systems from physical substrate

This has facilitated explosive growth in both mathematical theory and technology, but has also created a “new Tower of Babel where even the experts do not read papers or understand systems outside their subspeciality.”

Like most engineering fields, this is also true in the field of networking, where specialisation has become so extreme that arguably, engineers focused on one layer of the network are unable to understand what is happening in any other layer. Wireless connectivity only adds to this complexity. This is coupled with increased variation in usage characteristics.

The business environment is also highly competitive, subject to *Red Queen* mechanics: constant innovation is needed just to survive and maintain fitness relative to co-evolving systems (Section 6.1.2), leading to increased structural complexity. This is a constant challenge.

---

<sup>1</sup>From “Theory of Complex Systems” by John Doyle

## 1.4 Contributions

Contributions are both technical and qualitative.

### 1.4.1 Technical

This is the only work of its kind which takes two methodologies designed to deal with large, complicated systems and pushes them all the way to their logical conclusions in the context of networking:

**Entropy** This thesis explores the Entropic approach exhaustively in the context of network control and architecture. There are very few papers which have considered the use of Entropy as the basis for network control. The thesis also draws parallels and highlights the links with some existing models to clarify how the *Entropic Formulation* fits in with existing research.

**Conditional Independence** Conditional independence models are commonly used in finance to turn hard problems associated with large systems into easy calculations. This is the first time they are systematically applied to large-scale network control.

The exhaustive exploration of all the issues relating to the use of these techniques to control networks is the most important technical contribution of this thesis.

The application of these techniques yield a number of meaningful *Insights*: they are meant to act as a handbook for network engineers. Faced with the practical issue of interconnected networks, changing the infrastructure in an architectural way, and choosing amongst technologies, they are an objective guide to the practical impact of complexity on a network. Complexity is a high level invariant for networks, so pinning it down and being able to reason about it is helpful. Most of the existing technical exposition is either too technical and usually focused on very specific domains.

The insights offered in this thesis minimise the risk integral to operations early in the system architecture life-cycle. Too often, outsiders are unable to understand why experienced network engineers make certain decisions. Sometimes,

even experience does not help: the permutations of technology and paradigms is large. The question of how all this can be turned into a cost function is addressed. Main insights of the analysis presented in this thesis are in Chapter 4.

Finally, this is not the first time the use of methods based on the maximisation of an individuals' utility function, broadly termed the neoclassical paradigm in microeconomics, has been criticised as a useful and valid starting point for the analysis of problems in networking. However, this thesis points out a fundamental information theoretic objection to the use of neoclassical techniques which goes above and beyond the practical problems associated with this approach. This aspect of the problem with the neoclassical paradigm in networking does not appear in any previous research on the issue.

### 1.4.1.1 Complexity Measures

The *Deviations from Simplicity* principle, defined below, is a key contribution to the thinking around network architecture issues. It is meant to deal explicitly with the failings of complexity measures like Kolmogorov Complexity where the difference between chaos and complexity are not clear. It explores the notion that Complexity is subjective<sup>1</sup> and promotes the exploration of complexity in the context of a specific network model. In “Kolmogorovian” fashion, the simplest model rich enough is chosen. . As long as the super-modular structure of the underlying parameters relative to the complexity of the network is preserved, this principle is exportable across domains.

Alongside this, six measures of complexity are introduced:

**Operational Complexity** The complexity associated with variation in operations required to carry data

**Element Complexity** Complexity due to element variation and interaction

---

<sup>1</sup>“As measures of something like complexity for an entity in the real world, all such quantities are to some extent context-dependent or even subjective. They depend on the coarse graining (level of detail) of the description of the entity, on the previous knowledge and understanding of the world that is assumed, on the language employed, on the coding method used for conversion from that language into a string of bits, and on the particular ideal computer chosen as a standard.” Murray Gell-Mann, *What is Complexity ?*, (GM95).

**Scheduling Complexity** Complexity due to flexibility in, and the number of scheduling requirements in the network.

**Architectural Complexity** An overall complexity metric which subsumes *Operational* and *Element Complexity*, crossed with the added complexity due to scheduling requirements

**Implied Complexity** This is the total complexity of the network implied empirically, in the absence of any model.

**Network Complexity** The complexity associated with what cannot be explained easily, which is the difference between *Implied Complexity* and *Architectural Complexity*.

Murray Gell-Man argues that the complexity of any system will contain a strongly subjective element, in relation to the information set of the observer: complexity exists in the presence of and in relation to an observers model of a system. Hence, in order to measure the *Effective Complexity* of a network<sup>1</sup> a *Structural Model* for a network is proposed. The underlying hypothesis of this model has strong implications, which is considered to be a positive, because it makes it easy to separate the impact of coupling and resource sharing on the network from other factors. Finally, ways of extending the model, and hence the benefits of investing in modelling, are explained and analysed.

### 1.4.1.2 Risk: Introducing Slice

Complexity is an over-researched topic but few have noted the importance of *Risk* in networks. A key point of this thesis is that architectural decisions are a tradeoff between complexity and risk (see Chapter 4).

To make this possible, this thesis introduces the notion of a *Slice* of risk. This is defined here, as the qualitative notion is used frequently, before it is formalised in the context of an actual mathematical model.

---

<sup>1</sup>Formally defined in 3.1.1, and informally in Section 1.5.4.1.



*Slice*<sup>1</sup> is a risk metric which measures risks operators care about: cascading failures<sup>2</sup>, large scale operational failures to deliver, and persistent congestion. Using the infrastructure built around the models of Chapter 3, it is also possible to attribute risks to individual elements.

**Definition 1 (Slice)** *A risk metric which measures the probability that a certain Reference Measurement of importance associated with a network will take on a value between two thresholds,  $K_1$  and  $K_2$ .*

Graphically this is illustrated in Figure 1.1.

A *Slice* value can be calculated either empirically over a time-series of *Reference Measurement* realisations over a network, or can be calculated off a model which attempts to explain the behaviour of a network in terms of some measurement of concern.

In financial risk management, *idiosyncratic risk* refers to price changes due to circumstances specific to a certain asset, not related to the overall market. Such risks can be eliminated for free from a portfolio of assets through diversification, meaning that the market pays no premium to investors who take such risks. Likewise, *systematic risks* refers to risks emanating from market-wide circumstances. Such risks cannot be hedged, and hence pay investors a premium.

The two threshold,  $K_1$  and  $K_2$  allow one to specify whether it is idiosyncratic risks or systematic risks which one wishes to measure. One of the key characteristics of the models in Chapter 3 is that they make it very easy to calculate this metric.

Some of the the technical contributions resulting from the *Slice* parameterisation are:

**Separating Risk into Components** Breaking down risk into systematic and idiosyncratic.

---

<sup>1</sup>Tranche is French for slice. The inspiration for *Slice* comes from the financial markets which use the term *Tranche* to refer to classes of bonds collateralised by pools of other bonds.

<sup>2</sup>On 10 August 1996, a fault in two power lines in Oregon led, through a cascading series of failures, to blackouts in 11 US states and two Canadian provinces, leaving about 7 million customers without power for up to 16 hours.

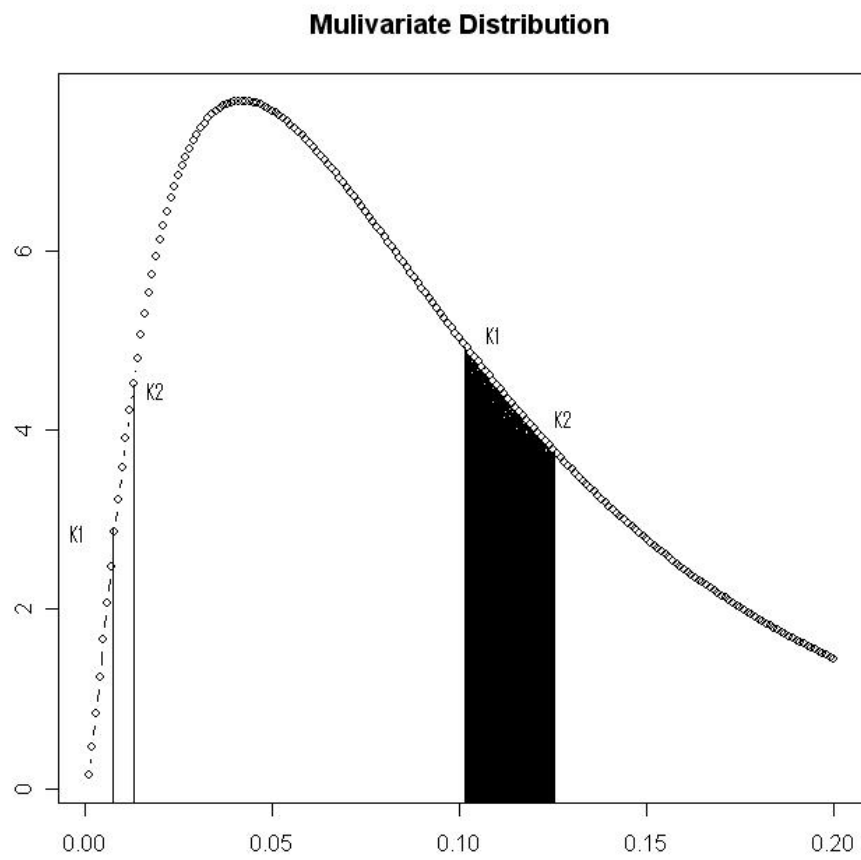


Figure 1.1: A Slice of a Distribution

**Risk Attribution** Attributing the contribution of an individual elements and other decisions (say routing decisions) to the systematic and idiosyncratic profile of a network

### 1.4.1.3 Novel Concepts

The judicious use of a model-based approach leads to three novel concepts:

**Entropic Routing** The use of the *entropy* of the network as a parameter in feedback flow control

**Implied Coupling** Many network models have been built which use the average degree distribution of the graph as an input. However, these models assume that one has perfect information about the network and few models actually mention how degree distributions could be calculated in the absence of detailed data about the networks topology. This thesis proposes a novel way that this could be achieved using the concept of *Implied Coupling*.

**Network Complexity** That part of observed entropy in a network which cannot be explained by a model.

## 1.4.2 Non-Technical

It is also possible to consider the architectural future of networking in a qualitative manner using notions of increasing entropy, and a systems innovative reactions to it. This creates a time-line for a system. Some of these implications are driven by the axiomatic implications of information theory, and should stand the test of time. At the very least, they are the result of pushing certain axiomatic ideas to their logical conclusions and should provoke inquiry. These can be found in Chapter 6.

De-regulation and increased competition increase the relevance of strategic considerations. Strategic deliberation involves an assessment of risk. Twenty-five years ago, Harry Markowitz showed that risk and return are two sides of the same coin. The spot price of services in telecommunications, and the present value of their future evolution are not constants. In a world with fluctuating

prices, no operator can exist without risk management. This is the first work of its kind which attempts to quantify the differences between systematic and idiosyncratic risks in a network, the key ingredients in the analysis of the balance of risk and return in any market. The provision of telecommunication services is an increasingly complex one. The complexity of a network can enable or disable what services can be offered: hence, there is a financial angle to this work<sup>1</sup>

It has been the experience of the author that most venture capitalists do not understand why certain systems technologies succeed where others fail. A very strong technology coupled with a seasoned management team and sound finances can fail for entirely inexplicable reasons. This thesis argues that this usually happens because the true complexity impact and cost of a proposed technology is entirely misunderstood. For example, sometimes, carriers prefer more complexity in their network because they do not know how the future will evolve or they know their networks operate under sub-optimal control policies.

Past attempts at introducing physical trading into telecommunications have failed: they will fail again, in networking and other technical areas, if there is not a better of understanding of why markets exist in the first place and what exactly is the asset that should be traded. One thing is clear, it is not point-to-point capacity. An alternative is proposed. *Slice* has the potential to become the basis of a financial contract for networks.

One of the great innovations of the last 30 years in finance is the notion of “implied volatility”. This is the value the volatility parameter needs to take in the standard Black-Scholes model to arrive at a traded price in the market.

---

<sup>1</sup>At the height of spending in 2000, ILECs spent nearly 30% of revenues on CapEx, and the IXC's, incredibly, were over that mark. Emerging service providers spent, in aggregate, over 100% of their revenues on CapEx. Spending anything like 30% of revenues on new capital equipment is clearly un-sustainable. In the case of WorldCom, capital expenditure was increased based on the common belief that demand would continue at the 8 times annual growth factor the industry was experiencing along with expenses for line capacity, and SONET Rings via leases. These were capitalised based on their reading of accounting regulations. 2002 brought successive quarterly revenue declines for Worldcom (and many other carriers). In June 2002, management concluded that billions of dollars of CapEx could not be amortised because there were no matching revenues. The resulting shift of \$7B from amortised costs to direct expenses wiped out the profitability of this leading carrier and led to the largest bankruptcy in U.S. history.

Simple, yet it has become the vernacular by which an entire industry thinks and reasons about market conditions. Similar mechanics are introduced here, termed “implied complexity” in Section 3.5.

Finally, the study of strategy in telecommunications has long been a business school topic, and has never been tackled from the perspective of the network engineer. The impact of architecture on strategy is quantified.

## 1.5 The Approach

The key statement of this thesis is that, in modelling the macroscopic features of networks with the dual purpose of both understanding how networks behave, and to make decisions about networking matters in the control plane, the modeller must assume that the network behaves as if to maximise the entropy of its configurations, and not the “utility” of its agents.

The overall approach is based on:

- The use of analytical techniques from adjoining fields, such as manufacturing and finance.
- Insights from the field of Communication Theory as they relate to measuring the complexity of a system.
- The use of market metaphors for understanding networks, albeit in a novel way
- The specific use of the technique of conditional independence to simplify the analysis of large scale networks by tractably attributing large scale behaviour to commonly shared factors.

### 1.5.1 Looking over the Fence

In 2001, the Computer Science and Telecommunications Board took a first exploratory examination of what the field of networking research might become in [Looking over the Fence at Networks: A Neighbour’s View of Networking Research](#) (see (SB01)). The main point was that *ossification* risked further advances in Internet research:

- Intellectual ossification: The pressure for compatibility with the current Internet stifling innovative intellectual thinking
- Infrastructure ossification: That the ability of researchers to affect what is deployed in the core infrastructure (which is operated mainly by businesses) is limited
- System ossification: That limitations in the current architecture have led to shoe-horn solutions that increase the fragility of the system.

It is arguable whether this workshop ever succeeded in bringing over anything useful from adjoining fields. The main conclusion seems to have been that different groups have different problems as they pertain to the Internet:

- *Insiders*: interested in the “ilities”: reliability, manageability, configurability, predictability
- *Outsiders*: diversity in the experience, expertise, and desires from the network.
- *Commercial Interests*: diverse roles and complex relationships that cannot be ignored when developing solutions to current and future networking problems

This was followed up two years later by [Network Research: Exploration of Dimensions and Scope \(NREDS\) \(SIG03\)](#), in which the author was a presenter: <sup>1</sup> During that workshop, the author had a chance to take the committee back in time to the goals of the previous workshop, by using metaphors from entrepreneurship in manufacturing as a predictor of the future for networking research. While there

---

<sup>1</sup>Two years ago, the Computer Science and Telecommunications Board took a first exploratory examination of what the field of networking research might become in Looking Over the Fence at Networks: A Neighbour’s View of Networking Research. This workshop is intended to be the next step in that process, beginning to take a more organised look. Not only is it valuable to consider specific directions that research might move, but we also expect to explore “meta-level” issues, such as the nature of our field, how it relates to others and how we evaluate new research.”

were a number of issues raised that are interesting in and of themselves, the main takeaway was that networking is beginning to look like manufacturing.

Existing techniques from the fields of manufacturing and finance form the mathematical backbone of this thesis. They are:

**Using Entropy to Control and Measure Manufacturing Facilities** The use of information theoretic measures to control manufacturing plants is now an established methodology in the field of manufacturing.

**Conditional Independence** This is a technique from the financial markets used to measure, manage, control and price the risk associated with very large pools of assets in finance.

### 1.5.2 Information Theory

The primary contribution of this thesis to the analysis and modelling of networks is the proposal to use *Entropy* as the key macroscopic feature for the purposes of control and prediction.

The concept of entropy in information theory describes how much information there is in a signal or event. An intuitive understanding of information entropy relates to the amount of uncertainty about an event associated with a certain probability distribution.

The overall approach to using entropy as the basis for network control and modelling is called the *Entropic Formulation*:

**Definition 2 (Entropic Formulation)** *The use of the Entropy of a network as the key macroscopic feature for the purposes of control and behaviour prediction.*

The detailed reasoning and justification for this is given in Chapter 2. One high-level justification is suggested by *Ashby's Law of Requisite Variety*:

**Principle 1 (Ashby's Law of Requisite Variety)** *The larger the variety of actions available to a control system, the larger the variety of perturbations it is able to compensate. This implies that the more variety there is in a system the more information has to extract from it in order to understand and control it.*

This suggests that the amount of information and the difficulty with which that information can be extracted from a system is a key macroscopic feature. Information theory is the machinery used for measuring the informational content of a system and entropy is a lower bound for the interrogation costs of determining state in a system <sup>1</sup> (this is related to the interpretation of binary polls as codes; see (CT91)).

Entropy is also a suitable measure for dealing with variety in networks, such as coupled systems, exhibiting diversity and innovations. Indeed, whenever one deals axiomatically with variety under certain reasonable requirements (such as continuity, certain boundary conditions), one ends up with the entropy formulation as an axiomatic consequence. Shannon originally derived entropy from first principles based on such axioms.

There are also some other practical reasons:

- There is a bijection between the entropic hypothesis and a commonly used specification of existing agent based metaphors: log utility maximisation, as in (GK98). This not only suggests that existing results can be reused, but that entropy is not necessarily entirely at odds with some of the existing body of work based on market metaphors.
- Entropy is additive, meaning that it can deal with sub-systems and can be used for comparative decisions<sup>2</sup>.
- Entropy is fractal, and hence invariant to time scales.
- Since entropy is a strictly concave function, when the constraints are linear, the maximisation function always has a unique interior solution, which is practically very convenient and important.
- Shadow prices on the constraint Lagrangians can be interpreted in insightful ways. They are also useful for calibration purposes.

---

<sup>1</sup>The efficiency of the agents polling requirement to infer enough information to act is explicitly dealt with in Section 2.2.1.

<sup>2</sup> The term *additive* refers directly to the outer  $\sum$  in Section 3.3.3.1, Section 3.3.3.2, and



For these and other reasons expanded upon in Chapter 2, the argument for re-basing the use of market metaphors in Section 2.2, the *Architectural Model* of Section 3.3 and the analysis of Chapter 6 are based on Information Theory.

### 1.5.3 Market Metaphors

Market metaphors have been an important tool for resolving resource allocation issues in networking research for over two decades. Networks and markets have a lot in common:

- Agent interactions are the building block of most markets.
- Most decisions are made locally.
- The market exists to resolve constraints on resources in an economy.
- Interactions are non-linear

In these ways, networks should probably be considered as markets to begin with (and in certain ways, markets are networks, see Section 7.1.5), where different agents share infrastructure to achieve their own narrow goals.

Markets face strong information and transaction constraints: in real circumstances, agents do not participate in an infinitude of markets at zero transaction and information costs. The same could be said of networks, where most routing decisions are made locally as it is impossible to collect information about the entire network at any point in time. And just like markets, networks resolve conflicts over resources such as capacity and processing power: if everything was in abundant supply, there would be no need for markets to begin with.

The primary source of non-linearity in markets usually has something to do with optionality, a key characteristic of any derivative contract in the financial markets. In the simplest derivative contracts, one agent has an option, but not the obligation, to purchase or sell an asset at a pre-determined price. The fact that there is no certainty about the transaction creates non-linearities.

It is quite clear that these issues were not clear to most telecommunications operators in the late 90's. Telecom revenues are among the most stable income

streams in any business globally. Yet, the financial valuation of telecom companies, even those operating in the most stable businesses have been tremendously volatile. This volatility has been driven by the perceptions about the optionality in the industry associated with new technologies, such as the Internet and the World Wide Web (WWW). There is non-linearity in the telecommunications industry, and this indicates the existence of optionality in the underlying infrastructure.

Hence, the techniques in this thesis are also based on market metaphors, however they are significantly different from the existing body of work in this field. The existing body of work analyses network effects as the side-effects of utility<sup>1</sup> maximising agents. This thesis instead focuses directly on the macroscopic features of networks, and justifies these using statistical mechanics techniques, while still remaining in the market framework.

### 1.5.4 Modelling

Two models are developed:

**Architectural Model** This model is based on an information theoretic approach to understanding complexity issues in a network. It provides insights into the impact of architectural decisions on the large scale complexity of the network. We develop this model with a view to understanding how networks work. However, the model is only useful as a tool to think and reason with: it is not easily adapted to physical networks composed of large numbers of elements. This infrastructure is based on taking the simplest possible model for a network, and measuring complexity as the distance between the predictions of this model and actual observations. In that sense, it is tied strongly to complexity arguments in (GM94).

---

<sup>1</sup>An economic term referring to the total satisfaction received from consuming a good or service. Utilities in economics are represented using “utility functions”, which represent how economic agents rank the choices over different baskets of good given to them. Rationality is defined precisely in terms of imputed utility maximising behaviour under economic constraints.

**Measurement Model** This model is meant to overcome the tractability issues of an architectural model by using novel techniques from the actuarial sciences. It can be used with standard networking tools to control and measure the actual complexity of a network. It is easy to build, scalable, and computationally tractable.

The models are complementary and occupy opposite ends of a spectrum of tradeoffs between tractability and detail. The *Architectural Model* makes use of not only the graph underlying the entire network, but also the coupling between shared resources and the dependency between operations on data flows. All this comes at the expense of computational and informational tractability: it is not realistic to actually populate all the parameters of this model for any realistic network.

The *Measurement Model* is simple, at the expense of making strong assumptions:

- All elements are homogeneous
- The coupling between elements is homogeneous
- There is one factor driving the systematic component of the networks behaviours

Complexity is then measured as deviations from the simple predictions of this model. The impact of relaxing the assumptions is studied one-by-one.

The primary mathematical tools of the *Architectural Model* are linear algebra, Jackson Network analysis, and Information Theory. The Measurement Model makes use of the law of large numbers, conditional independence, numerical integration and Gram-Charlier expansions to the normal distribution.

### 1.5.4.1 Effective Complexity

*Effective Complexity* is formally defined in Section 3.1.1. Murray Gell-Mann in (GM94) argues strongly that practical measures of complexity must contain a subjective element (see (GM95), and also (Edm99a)). His main point is that,

when analysing a system, any observer must evaluate the complexity in the presence of a model. If a model exists which perfectly describes a systems behaviour, the system does not really have any *Effective Complexity*. The philosophy behind the measure of complexity adapted in this thesis uses entropy in the context of a specific model to measure the *Effective Complexity* of networks. This is an important enough point that it is worth elevating to a *Principle*.

**Principle 2 (Deviations from Simplicity)** *Build the simplest model possible of the system which is just rich enough to explain all observable static characteristics, but no richer. Measure complexity as the deviations of observed dynamics from the predictions of this simple model using Entropy.*<sup>1</sup>

## 1.6 What Will Not Work ?

Previous misguided attempts at solving the problem of scalable network control are reviewed.

The specific problems associated with network control have been investigated along three research areas that all share a distributed control/computation paradigm:

1. Naive Market solutions
2. Naive Graph Theoretic Approaches
3. Negotiation-based systems:
  - Agent based solutions
  - Active Networking

This research has generated many interesting results. However, this thesis argues that it has been misguided where it has been based on neoclassical

---

<sup>1</sup>This is similar to the idea of looking for the shortest model as the basis of complexity, as with Kolmogorov. However, Kolmogorov's measure has a serious shortcoming: it attains a maximum with random sequences, and is in that sense, is really a measure of randomness and not really complexity in the practical sense.

paradigms, as the assumptions are unrealistic and ignore computational and informational intractability issues, primarily associated with the the *tâtonnement* process.

### 1.6.1 Naive Market Solutions

At the turn of the century, a market in traded bandwidth materialised between a number of energy firms, investment banks, and telecom carriers. A derivative market also appeared, in the form of traded swaps and options on fibre capacity.

The primary model for traded capacity was driven by supposed arbitrage conditions around an underlying asset modelled as point-to-point connections. A good example of such a model can be found in (Kep02). This market no longer exists, for a variety of important reasons:

- There were too many tradable assets
- Information and data on all assets was difficult to centralise
- The market was severely incomplete: there was a significant mismatch between the true risks to agents in the market and the assets being traded

The real world is full of imperfect and incomplete markets<sup>1</sup>, and there are ways to deal with them, discussed in Section 2.2. Markets can tolerate a certain level of incompleteness. Sometimes, the very existence of liquidity in transactions completes an otherwise incomplete market, by allowing risk transfer from risk-averse to risk-loving agents.

However, there are limits to how much incompleteness a market can handle. First of all, for the market to be able to step-in and play the role of the liquidity provider of last resort, the agents participating must come from variable backgrounds with different levels of risk aversion. This is what creates liquidity for a transaction<sup>2</sup>.

---

<sup>1</sup>A market is said to be spanned by primary assets if contingent contracts in every state of the economy can be expressed in terms of the primary assets. In the derivatives markets, the underlying asset and a risk-less bond span the option price under the assumptions of Black-Scholes. Option spanning and market completeness are equivalent.

<sup>2</sup>(Arr53) showed that, under some conditions, the ability to buy and sell securities can effectively make up for missing securities and complete the market. (see Section 2.4.2)

Furthermore, the problem with the bandwidth markets was not that there weren't enough assets: arguably, there were too many. Making point-to-point connections the primary asset was probably a mistake to begin with. The problem was that the traded assets did not properly cover the risks of participating agents. The risks faced by most market participants in the telecommunications is not the price of fibre capacity. This comprises a small portion of the total cost of delivery of telecommunication services. The primary non-business risks faced by carriers is congestion in their network, unforeseen cascading failures failure to deliver services for operational reasons. It is these risks which the bandwidth markets did not cover.

This is a mistake which need not be repeated. Simplistic application of market mechanisms to standard telecommunication and computational units, such as bandwidth, fibre-capacity, storage, and processor cycles are not likely to succeed for this reason in the future. The price or physical cost variation of these assets do not span future states against which participants seek protection. Also, speculators do not consider the future pricing of these units as an opportunity.

The same criticism can be levelled at simplistic applications of the neoclassical paradigm to congestion pricing and resource allocation, such as in (FORR98). The neoclassical paradigm is deeply flawed and does not even model real-world markets well: not only are these assumptions usually unrealistic, the resulting calculations required of agents are for most cases entirely intractable. This issue is discussed at length in Section 2.1.

### 1.6.2 Naive Graph Theoretic Approaches

Networks are inherently difficult to understand using simple graph-theoretic techniques:

1. The network diagram can be an intricate tangle.
2. The network diagram could be subject to change over time, as links are set-up and torn down.
3. The links between nodes could be diverse, in capacity, direction, and type.

4. The nodes could be diverse, also in type and size.
5. The nodes could have state which effects network behaviour, and this state could be non-linear.
6. The networks layout might be dependent on historical path-dependent factors, or the connections between various parts of the network might be influenced by the traffic going through the network at any point in time.

There are also analytical problems. In (RV76), the authors prove that every nontrivial monotone graph property is evasive<sup>1</sup>, a conjecture originally due to Karp (see (Ros73)). This poses a significant tractability problem for most graph-theoretic measures of any significance. Due to their size and dynamics, it is not possible to obtain complete graphs of even medium-sized networks, never mind the Internet. Hence, most graph-theoretic metrics are computed under incomplete information<sup>2</sup>

Furthermore, graph-theoretic measures do not take into account complexity in a network arising from operations. All graph-theoretic measures see are edges and nodes, without taking into account the fact that delivery of a certain flow through a network requires the execution of certain operations in exact precedence (such as DNS lookup, routing table queries, and setting up virtual paths).

Graph-theoretic measures are also blind to traversal dynamics and cannot take into account issues associated with aggregation and disaggregation. For example, an application shared amongst users would trigger significant complexity in a network: a graph theoretic measure would be invariant between a network which was capable of delivering such a service and one which wasn't, if they had the same topological layout.

The need for large scale intrusive measurements must be obviated, and operational complexity cannot be ignored. This is achieved in Chapter 3.

---

<sup>1</sup>A function  $f$  is evasive if, in the worst case, an algorithm needs complete information about  $x$  to compute  $f(x)$ . In the context of calculating the value certain *evasive* property of a graph, this is equivalent to a requirement that all nodes be visited.

<sup>2</sup>For example, the authors show that trace-route measurements create a biased view of the Internet, in (LBCX02), casting doubt on the accepted wisdom that degree distributions of nodes follow a power-law tail.

### 1.6.3 Negotiation and Fixation on the Agent

Agents are software designed to be autonomous in pursuit of goals, capable of collaborative and competitive interaction. They typically target complex, distributed processes.

**Definition 3 (Active Networks)** *Active networks allow individual user, or groups of users, to inject customised programs into the nodes of the network. "Active" architectures enable a massive increase in the complexity and customisation of the computation that is performed within the network, e.g., that is interposed between the communicating end points*<sup>1</sup>.

Agent based solutions can be seen as the result of a reductionist philosophy in science which treats differences in market structures where the agents congregate as second-order complications. Most telecom research has modelled "The Market" as a relatively homogeneous and undifferentiated entity, with interactions between agents as the primary determinant of market structure. Focus has been on the technical microeconomics of the agent, not the detailed infrastructure of the market.

In IP networks, making changes to the control plane require changes to forwarding algorithms, making it very difficult to experiment with new control techniques. Active networking can be seen as a response to this difficulty. In some ways, some of the techniques from Active Networking have influenced the development of the Multi-Protocol Label Switching (MPLS) protocol. One reason for the development of MPLS was to allow the control plane of the IP layer to direct Asynchronous Transfer Mode (ATM) switches (see Section 1.10.2), by integrating Layer 2 switching with Layer 3 routing. Another achievement of MPLS was its ability to operate on a wide variety of link-level technologies, including Frame Relay, Packet over SONET, and also LANs such as Ethernet and token ring.

The key separation between agents and active networking is that in active networking, the packets traversing the network make modifications to the network<sup>2</sup>. Active networking departs from traditional architectures which separate

---

<sup>1</sup>Active Networks web-site, MIT

<sup>2</sup>Enabled through nodes in the operating systems which support execution environments, and in certain cases, dedicated *Active* hardware.



network components from data for stability reasons.

However, neither active networks nor agent based solutions have lived up to expectations. Their similarity based on the notion of negotiation and dynamic interaction has sealed their combined fate. Agent-based techniques, involve automatic negotiation by rational agents who optimise an individualised, implied “utility” function ensconced in the assumptions of neoclassical economics. The dynamic interaction of active networks with the network they share implies interaction with each other. The issue is computational intractability.

### 1.6.3.1 The Complexity of Negotiation

Even benign agents pose computational difficulties. Take for example the case where the agents cooperate to achieve a common objective, such as fairness, or throughput<sup>1</sup>. It has been shown that computing the optimal strategies for even cooperating agents is intractable. Furthermore, even if the constraint is relaxed to solutions  $\epsilon$ -distant from the optimal value, i.e. approximate solutions, the problem remains as hard as the optimal version of the problem. See (RGR02).

Automatic contract negotiation is posed as an attempt by agents sharing resources to construct a mutually beneficial optimal reallocation through trading. If contracts are restricted to those in which a limited number of resources can be transferred from one agent to another and are required to be rational (in the sense of strictly improving overall worth of an allocation), then not only is it the case that a suitable contract-path to an optimal allocation may fail to exist, but even deciding if a path from a given allocation to a specified more beneficial allocation is possible is intractable. However, if the cost matrix is restricted to trees, the problem does become tractable. See (Woo00).

### 1.6.3.2 The Complexity of Planning

Automated planning is the problem of determining how an agent achieves a goal given a repertoire of actions. In the early days of agent-based research, strong

---

<sup>1</sup>A realistic example would be fault localisation. In large networks, errors can cross domains, leading fault management systems to trigger incorrectly. Proposed distributed event management nodes governing subsets would not be able to localise these faults collectively.

assumptions were made about the agent's knowledge and control over the world, namely that its information is complete and correct, and that the results of its actions are deterministic and known. When, these assumptions are relaxed, the results are pretty pessimistic. For example, the classical “*probabilistic planning*” problem has been shown to be formally undecidable. These results can be applied to a broad class of stochastic optimisation problems where the agent:

1. operates over an infinite or indefinite time horizon, and
2. has available only probabilistic information about the system's state.

The same set of un-decidability results also apply to corresponding approximation problems with un-discounted objective functions. See (MHC03).

## 1.7 Focus Areas

(Mor02) gives a good overview of various traffic engineering techniques, especially as they relate to the Internet. Network control is divided into two broad areas. These are Network Operations and Network Architecture.

Network operations are comprised of:

1. *Planning*, concerned with long-term physical network planning. Its primary input is traffic growth.
2. *Network Engineering*, dealing with physical link configurations to establish capacity around existing links and nodes.
3. *Traffic Engineering* for optimising an existing topology to meet current demands.

The Measurement Model in Section 3.4 has two outputs useful in all three areas. The first is *Slice*, a direct measure of the current complexity of a network that can be fine-tuned to specific specific flows. The second is the first derivative of *Slice* with respect to a specific element in the network.

The purpose of this thesis is also to aid *Architectural* decisions which involves all three Network Operational functions at different time scales. Network Architecture is involved with not the implementation, but the actual selection of which of many different ways these operations can be performed through:

- *Design principles*: network topology, level of resource sharing
- *Technology choices*: layered vs. integrated elements
- *Methodologies*: connection-less vs. connection-oriented routing

The *Architectural Model* in Section 3.3 is a qualitative model used in this context. The emergence of new technologies, green-field build-outs, and network upgrades are typically the junctions at which architectural decisions are made. Engineers typically turn to experience as a guide. This model does not replace, but complements experience by making it easy to compare decisions.

The model in Section 3.4 is quantitative, making it possible to calculate the impact of:

- sub-system interactions
- network heterogeneity
- traffic variation, and
- the marginal cost of making changes to one element.

### 1.7.1 Traffic Engineering

Traffic engineering primarily uses statistical models of traffic for prediction. Using these predictions, the underlying network is shaped to carry this traffic in the most efficient way possible. During that process, the agents making decisions (be they routers or people) perform an optimisation over user-visible properties of the network in the presence of incomplete information. Traffic engineering has to abide by constraints, such as:

- Bandwidth

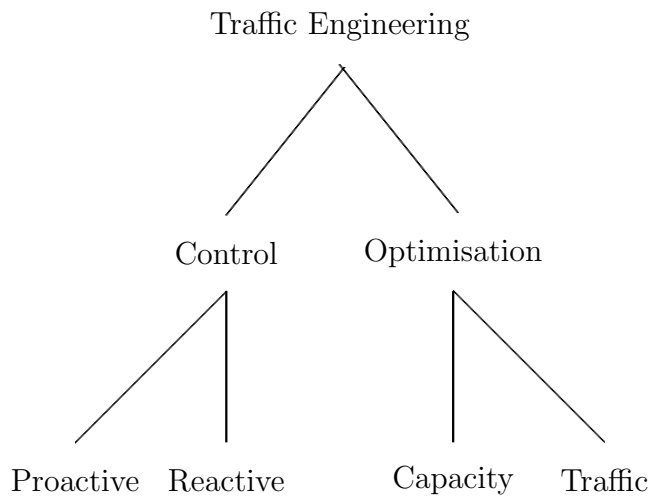


Figure 1.2: Traffic Engineering

- Hop Count
- Setup Priority
- Hold Priority
- Adaptability
- Resilience
- Shared Risk Link group
- Include/Exclude Resources
- Protection

Traffic engineering performs two high-level functions:

- Control:
  - Proactive, meaning prevention
  - Reactive, meaning correction
- Optimisation

- Capacity management
- Traffic management:
  - \* Nodal
  - \* Constraint based

See Figure 1.7.1. Control can be exercised either at nodes through conditioning and scheduling, or through the arbitration of access to network resources. Optimisation is achieved over parameters which determine capacity between nodes and the direction of flow into that capacity.

*Slice*, can play two roles here. First of all, *Slice* is an indicator of the complexity of a network, a desirable feature when it is known that the control policy of a network will be suboptimal (Section 4.1.2.5). On the flipside, complexity can be highly undesirable when the network has to give hard guarantees. Hence, *Slice* can act as a measurement-feedback metric either for adaptive control, or as a constraint requirement. This is termed *Entropic Routing*, covered in Section 6.1.1.

*Slice*, is highly tractable and can be measured in real-time, and has the potential to relay information about the complexity of the network to any element in the network at any point in time.

## 1.7.2 Routing

Routing is the glue holding data networks together. The complexity of routing can be decomposed into two primary operations:

1. *Data Plane*: address lookup and forwarding
2. *Control Plane*: routing table management and signalling

Data plane operations are packet-by-packet operations whereas control plane operations only impact the nodes of a network. There are two paradigms in routing which determine control plane operations:

- *Connection-less*: Routing decisions take place at nodes. This requires heavy coupling between the data plane and the control plane as the address space has to be shared.

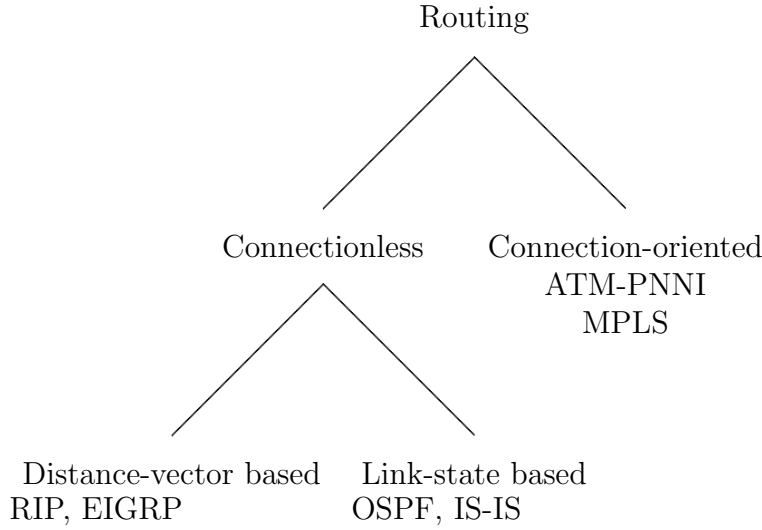


Figure 1.3: Routing Protocols

- *Connection-oriented*: The end points establish an end-to-end connection using a preliminary protocol at initiation, before any data is transmitted. This means that the control plane can decouple from the data plane.

Connection-less routing protocols can maintain state information about the network either using information about nodes, called *distance vector* protocols, or using information about links, called *link-state* protocols (see Section 1.7.2).<sup>1</sup> The choice of protocol has significant implications for the complexity of a network (Section 5.1.1).

**The General Routing Problem** (Awd99) and (FT00) provide a succinct description of the *General Routing Problem*. The network is defined as a directed graph  $G = (N, A)$  whose nodes and arcs represent routers and links. Each arc  $a$  has capacity  $c(a)$  and each network has a demand matrix  $D$  for each pair of nodes,  $(s, t)$ . The problem is the distribution of flow over paths from  $s$  to  $t$ .

The load of an arc  $a$  is  $l(a)$  and utilisation is  $u(a) = \frac{l(a)}{c(a)}$ . The objective is to keep the load within capacity limits.

<sup>1</sup>*Link state* protocols contain complete information about a network, while *distance vector* protocols do not.

Routing must give two guarantees:

- *Consistency*: meaning that every node is reachable from every other node using local information
- *completeness*: meaning that every node is able calculate a path to every other node other path

More complete information leads to faster convergence times, as the chances of inconsistency decrease. This has an automatic computation cost, as more distributed state must be processed. This also limits scalability. As both *Completeness* and *Consistency* limit scalability, large networks have to be structured hierarchically <sup>1</sup>.

The most controversial feature of this specification is the existence of a demand matrix. The mathematical framework of Section 3, which forms the backbone of this thesis, models the interaction between  $l(a)$  and  $c(a)$ .  $l(a)$  is modelled as being driven by an *idiosyncratic* component, ignored due to the Law of Large numbers, and a *systematic* component related to the complexity of the network, which cannot be ignored. The complexity of the entire network is measured indirectly through the incidence of correlation between Bernoulli-(0,1) measurements triggered when  $u(a) > 1$ .

### 1.7.2.1 Domain-based Routing

The primary source of scalability in the Internet is routing hierarchies based on domains:

- *Inter-domain routing*: takes place between autonomous routing domains. Currently, this is implemented with the Border Gateway Protocol (BGP), which exchanges reachability information within and between Autonomous Systems (ASes). These are connection-less protocols due to the lack of a central entity in the Internet.

---

<sup>1</sup>As *Areas* in OSPF and Autonomous Systems (AS)S in BGP

- *Intra-domain routing*: takes place within an AS, meaning that it is possible for the controlling service provider to engineer its performance and use one of a number of different protocols depending on requirements. Typically, intra-domain routing is usually (but not always) achieved using connection-oriented protocols, which give the service provider higher-levels of control over the network.

AS boundaries are stitched together using the Border Gateway Protocol (BGP). BGP was not a protocol designed for traffic engineering, even though it has become the de facto standard for inter-AS Internet-working. The main issues in BGP addressable in our framework are:

- Constraints: Passing network performance requirements into BGP
- Traffic engineering: Passing congestion information into BGP

Some of these issues are dealt with in Section 5.

Intra-domain routing can be achieved via a number of alternatives, of which two are prevalent: MPLS, a connection-oriented protocol, and Open Shortest Path First (OSPF), a connection-less, older alternative. Relevant issues are analysed in Section 5.1.1. Issues common to both domains are:

- Multi-path routing
- Survivability
- Fault Detection and Notification
- Embedding user-interfaces/control indirectly into routing decisions
- Mapping policies to protocols
- Operating under sub-optimal control policies
- Rolling features into the addressing structure in a network, rather than as options (see Section 5.1.3)



### 1.7.3 The Internet Protocol

The structure of IP as a datagram model is simple. However, determination of state in an IP network is very difficult, incurring significant control complexity<sup>1</sup>. IP also has an important strategic impact on service providers: it makes it easier for customers to unbundle services, forcing providers and end-users to participate in combinatorial auctions<sup>2</sup>.

In (MFMZ02), the authors argue that IP has technical shortcomings at the core of the network, where they expect it to be replaced by optical circuit switching. This is also an implication of the entropic hypothesis of this thesis, in Section 6.1.2, albeit for different reasons. This also ties in with the claim of (LOR<sup>+</sup>01), where the authors show that the source invariance of traditional forwarding and routing protocols like such as RIP and OSPF are a significant source of performance degradation, over more recent alternative schemes like MPLS, where the forwarding is based on both the source *and* the destination, in addition to many other factors. These issues are discussed in Section 5.1.1.

### 1.7.4 Areas Not Addressed

**Element Design** Individual element design is also fruitful ground for our analysis. For example, it has been observed that the cost of a router port depends on the amount and kind of memory it uses, its processing power, and the complexity of the inter port-processor communication protocol. Furthermore, there are trade-offs between performance, complexity, and cost. Specifically, router ports built with general-purpose processors and complex communication protocols tend to be more expensive than those built using ASICs and simple communication protocols. SRAMs offer faster access times, but are more expensive than DRAMs, while buffer memory is another parameter that is difficult to size. Careful engineering of the control protocol is necessary to reduce the cost of the port control

---

<sup>1</sup>The model in Section 3.3 is based on entropy, implying that we equate complexity to the number of questions (binary interrogations) needed to capture the relevant state of a network.

<sup>2</sup>In Combinatorial Auctions, bidding takes place on packages. Combinatorial auctions are *hard*: there is no polynomial bound to the computational complexity of a winning bid, the deliberation and strategic costs are cognitively complex, and strategic bidding can lead to inefficient and counterintuitive outcomes.

circuitry and also the loss of command packets, which will certainly need retransmission, as in some designs, a centralised controller sends commands to each port through the switch fabric and the ports internal buffers.

However, straightforward design complexity may bring interface simplicity. In the case of equipment selection, very complicated internal design may significantly simplify interfaces to necessary protocol features.

It is possible to adapt the framework for this purpose, but this is left for further research. However, implications are mentioned in Chapter 6.

**Facility Planning** Similarly, this thesis does not deal with facility planning and long-term traffic planning issues. In Chapter 6, it is argued however that the traffic engineering and the long-term traffic planning functions will merge in the future.

Furthermore, there is a vast literature dedicated to facility planning and the broad based result of this body of work is that facility planning is not a fundamentally difficult problem. Even in cases where there are computationally hurdles, good approximation algorithms exist.

## 1.8 Alternative Approaches

There are two alternative practical approaches to understanding and controlling the complexity of a network, and one theoretical variation.

**Schedulability** Given a set of tasks, their arrival patterns, and a schedule, schedulability checks whether any task will miss its deadline. This is usually done through the analysis of an automaton, checking whether a reachable state exists in which the deadline is missed. It is possible to cast the task of network control as one of guaranteeing schedulability. This would equate the complexity of network control with the complexity of schedulability.

**Novel Graph-Theoretic Techniques** Most existing graph-theoretic measures of complexity do not consider how difficult it is to measure the various properties they advertise. However, some recent work on the evasive properties of graphs could remedy some of these shortcomings.

**Subjective Measures of Entropy** The *Entropic Formulation* in this thesis is based on the classic definition given by Shannon. Renyi has generalised this definition, leading to a wider set of measures for complexity.

### 1.8.1 Scheduling Complexity as a Metric

The easiest way to reduce complexity in a routed network is to reduce the complexity of scheduling algorithms. There is a natural tradeoff between the complexity and the flexibility of the scheduler. The common objective is the *time complexity*. Some proposals for this can be found in (ZF93) and (SZ99).

Hence, one could pose scheduling complexity as a key metric, The fundamental complexity vs. flexibility tradeoff means that one can replace per-flow management with coordination, which points to a tradeoff between connection oriented protocols and coordination protocols (discussed further in Section 5.1.1).

Work in the area of schedulability has some interesting pointers in this direction. In (KP95), the authors show that this fundamental tradeoff admits certain hard conditions. They consider several well known non-clairvoyant scheduling problems, including the problem of minimising the average response time, and besteffort, firm realtime scheduling. Even though there are no deterministic on-line algorithms for these problems with bounded competitive ratios<sup>1</sup>, moderately increasing the speed of the processor used by the nonclairvoyant scheduler effectively gives the scheduler the power of clairvoyance. In this context, the word *clairvoyance* refers to a priori knowledge of the execution and release times, and dependencies of a specific job. The increase in processor speed allows a non-clairvoyant scheduler to match the performance of a clairvoyant scheduler.

The alternative to speed is additional resources. (Edm99b) consider non-clairvoyant multiprocessor scheduling of jobs with arbitrary arrival times and changing execution characteristics. They assume that the “*scheduler is in the dark*”, i.e. has no knowledge about the jobs except for knowing when a job arrives and knowing when it completes, and prove that there exists a scheduler

---

<sup>1</sup>The worst case of the ratio between the cost incurred by an on-line algorithm and the best-case cost.

algorithm “Equipartition” which performs within a constant factor as well as the optimal scheduler as long as it is given at least twice as many processors.

### 1.8.2 Novel Graph Theoretic Techniques

While graph-theoretic techniques for measuring and managing complexity fail to capture operational complexities and do not say much about architectural decisions, some of their shortcoming in the area of tractability could be overcome.

For example, in (GY02), the authors recognise the evasive nature of non-trivial graph-theoretic measures, and propose a meta-metric called  $(\gamma, \sigma)$ -evasiveness which indicates if a metric can be estimated with  $1 - \sigma$  accuracy by using only  $\gamma$  percentage of the data.

They then proceed to discuss:

- how  $\gamma$  and  $\sigma$  vary for various oft-used graph theoretic measures
- sampling issues

Their meta-metric has the making of a complexity measure, as a more complicated network would presumably be proportional to  $\frac{\gamma}{1-\sigma}$ .

Recent developments in the use of statistical methods for the analysis of large networks attempt to glean macroscopic features of networks that can be divorced from the intricacies of how the network is actually wired up. They attempt to answer questions about the network without knowing every detail about the network. This is also an aim of this thesis.

This body of work has three aims:

1. Find statistical properties such as path lengths and degree distributions that characterise the complexity of networked systems.
2. Create models which are explanatory.
3. Predict network behaviour as a function of measured structural properties and local rules governing vertices.

These methods are related to the *Entropic Formulation* proposed in this thesis, as they are statistical in nature, and hence subject to the tools and mechanics of statistical mechanics. Some of these results can be interpreted in an entropic sense. This issue is covered in Section 2.2.2.2.

### 1.8.3 Alternative Entropy Formulations

To deal with the subjectivity of Complexity, the “Deviation from Simplicity” principle juxtaposes a model of the network, and proposes true *Network Complexity* as arising from that part of an empirical distribution which cannot be explained by the model. Shannon’s measure of entropy requires the presence of a model because it is inherently objective. Syntactically, this arises from its definition over probability distributions, and semantically, from the additivity constraint.

A relaxation of this requirement leads to an alternative measure titled Renyi Entropy:

$$H = \frac{1}{1 - \alpha} \log_2 \sum_i p_i^\alpha \quad (1.1)$$

Renyi entropies have been successfully used to measure complexity in dynamic systems with subjective observers. Renyi Entropies also lead to various uncertainty measures. More details are in Appendix B.

## 1.9 Related Research

Three strands of directly related research are reviewed.

### 1.9.1 Event Correlation

There is a large body of research dedicated to event correlations in networks. A good exposition can be found in (SS01b). The vocabulary associated with this literature consists of:

- *events*: exceptional occurrences in hardware or software
- *faults*: events that can be handled directly

- *errors*: discrepancies between computed and observed values, and
- *symptoms*: external manifestations of faults

The primary goal of this literature is fault localisation in the presence of ambiguity, inconsistency and incompleteness. The main issues relate to the occurrence of unrelated faults, multiple hypothesis, and the computational infeasibility of processing and managing a knowledge base of events. Computational infeasibility of managing knowledge bases is covered in (BCFK95), and (KYY<sup>+</sup>95).

Common approaches are covered in (KP97). There are many approaches, but the four leading alternatives are:

- *Model-based reasoning* incorporating a deep knowledge in the form of a model of the underlying system, covered in (BBM<sup>+</sup>93) and (KS95)
- *Fault propagation*, in (HCF95)
- *Case-based reasoning*, in (Lew93), and
- *Model traversal* in (Gru98), (JP93), and (KP97)

The bulk of the research between 1990 and 2000 has focused on low-level faults related to resources. Current challenges are multi-layer faults, real-time performance problem diagnosis and dealing with uncertainty.

### 1.9.2 Networking Games

(ABEA<sup>+</sup>06) provide a very good overview of issues related to networking games in telecommunications. This is a rich field and has resulted in interesting applications both to networking and in transport planning.

The primary tool in the use of games in networking is the concept of a Wardrop Equilibrium. While related to Nash equilibria, the concept is generalised to cover games with infinite players. Wardrop was an English civil engineer interested in road traffic issues. Hence, the primary use of networking games is in the prediction of congestion.

Wardrop defined two notions:

1. Equilibria: No driver can unilaterally change routes to improve travel times
2. System optimality: The average journey time is at a minimum at equilibrium

In this context, economic arguments generally lead to the conclusion that marginal cost road pricing leads to System Optimality.

One of the most interesting results in this field is known as Braess paradox, which states that in a non-cooperative framework, adding capacity to bottleneck links increases delays to all users. This happens when the equilibrium is not optimal. This is frequently the case in non-cooperative networks, such as in (She95) where it is shown that no network switch service discipline exists which can guarantee optimal efficiency. Hence, the primary focus of networking games has been the non-cooperative setting with infinite users.

The key linkage in this field is that between congestion control and pricing, highlighted in (KM99), (CO98), and in a series of papers by Frank Kelly et al starting with (KMT98) and (GK98). One of the most interesting aspects of proportional fairness is the use and implication of log utility in the agents optimisation function. Their model is based on optimising a log utility function associated with network flows. This is related to the entropy based model in Section 3.3, and is discussed in Section 2.3.1 and elsewhere throughout this thesis. The equivalence between the two models and the axiomatic basis of the *Entropic Formulation* means that the choice of a log utility is not arbitrary.

### 1.9.3 Market-Based Metaphors

Market-based solutions received substantial attention from 1995 onwards, as the greater use of computing technology in finance triggered an interest in the use of market mechanisms to solve distributed computation problems. Telecom networks look like markets, large networks of interconnected elements which exchange information, share resources, optimise around constraints and try to stay in equilibrium.

Markets have been used as a basis for survivability, as in the case of the MARX project at the University of Michigan (see (EJK<sup>+</sup>01)). Fault-tolerance, and other

desirable characteristics of a network are and continue to be researched in the mathematical community, as in the case of the RAIN project at CalTECH (see (BFL<sup>+</sup>00)). Previous literature has focused primarily on the pricing of services, and the use of markets and prices to resolve general dynamic equilibrium problems in resource allocation. There has also been some focus on inter-carrier competition based on pricing strategies.

Major research projects which are most relevant are:

- George Fankhauser in (Fan00) presents a network architecture based on Service Level Agreement (SLA) traders which integrate service allocation, routing and pricing functions
- Neil Stratford and Richard Mortier in (SM99): use dynamic pricing as a congestion feedback mechanism based on frequently renegotiated timed resource contracts
- MARX Project (EJK<sup>+</sup>01) is a dynamic computational market designed to enable adaptive allocation of resources in large-scale distributed information systems
- Market Managed Multi-service Internet (M3I, (Man01)): the goal is to design, implement and trial a next generation system which will enable Internet resource management through market forces, specifically by enabling differential charging for multiple levels of service

There is some overlap. For example, Stratford and Mortier talk about the separation of local and global optimisation problems, a primary insights for the use of market mechanisms. The MARX project has survivability as a core focus, enforcing fault-tolerance through market-mechanisms. Both the M3I and the Fankhauser work are squarely focused on Service Provider strategies. The congestion pricing work of Mackie-Mason (MMV94) is related to the model in Section 3.3.

There are also differences.

- Both Stratford/Mortier and MARX project are focused on an actual design for a system based on agents; furthermore, they are focused primarily on resource allocation for computing and capacity



- The M3I and Fankhauser work are primarily focused on SLA trading between counter-parties and on financial contracting between service providers. There are doubts about the viability of these strategies, mentioned in Section 1.6.1. However, if the traded SLA is not idiosyncratic in nature, they may work (Chapter 6).

Also, some of the emphasis in this body of work is on demonstrating capabilities. For example, the MARX project attempts to demonstrate the use of markets for survivability using agents and auctions. The focus here is on mathematical tools for reasoning and thinking about networks.

## 1.10 Relevant Trends

Relevant trends and developments in routing and element design are highlighted, for context.

### 1.10.1 Routing

The primary trend in routing has to do with two developments that are driven by the cost structure of most service providers:

- the removal of lower level protection mechanisms on the basis of costs, such as SONET and SDH protection<sup>1</sup>.
- financial and SLA pressures forcing ISPs to use the shortest distance between two POPs

SDH protection is offered through a Protection Circuit, which backs up the Active Circuit in case there is a break. In SDH, this is called Multiplexed Switching Protection and in SONET, it is called Automatic Protection Switching. As

---

<sup>1</sup>SONET is short for “synchronous optical network”, a standard for communicating digital information using lasers or light emitting diodes (LEDs) over optical fibre. SDH, short for “synchronous digital hierarchy”, is a more recent standard that is popular outside Canada and the USA. SONET and SDH encapsulate earlier digital transmission standards, such as the PDH standard, and can be used directly to support either ATM or Packet over SONET networking. This means that SONET/SDH are more than just a protocol, but an all-purpose transport container for moving both voice and data.

the Protection Circuits do not generate any revenue, alternative solutions which are more efficient are increasingly used.

Another common trend is the de-layering of the network to reduce the costs associated with maintaining personnel and infrastructure for multiple layers. There are three ways to implement IP over fibre optic networks which use Dense Wave Division Multiplexing (DWDM):

1. IP/ATM/SONET/WDM
2. IP/MPLS/SONET/WDM
3. IP/MPLS/WDM

In an IP/MPLS network running directly over DWDM without any ATM nor SDH intermediate layers, it is necessary to achieve the restoration of data paths at the MPLS sub-layer in case of link or node failures.

In MPLS networks, increasingly common in carrier backbones, the current alternative consists of a centralised server which computes backup tunnels, a technique called MPLS Traffic Engineering Fast Reroute. This occurs in the IP layer. While a reduction in the number of layers is welcome, the central server now becomes a potential bottleneck and a new source of risk. As a matter of fact, the network now becomes less resilient against a planned attack as the restoration and protection options have been reduced. This is a common finding in small-world/power law networks, characterised by the presence of a number of large hubs: while they are very resilient against random removals of nodes, they become less resilient against a planned attack which targets the hubs. The de-layering of the network may be opening up networks to such an attack or source of risk.

### 1.10.2 MPLS Developments

Before MPLS, ATM was the preferred methodology for connection oriented routing for IP. IP itself is a hop-by-hop protocol, relying on lower levels for frame forwarding. However, ATM networks failed for a number of reasons, which we investigate in more detail in Section 5.1.4. Specifically, the connection-less model

for IP dictates that the control plane and data plane share the same address space. Because ATM was designed as an overlay model, this led to significant complexity over alternatives<sup>1</sup>.

MPLS was designed as a merged model, taking the routing protocols for IP and merging them with ATM's forwarding plane. In other words, ATM's high-speed switches came under the control of the IP control plane, one of the original goals of MPLS. Even though, the development of dedicated ASICs obviated this need, nowadays, MPLS has been successfully repositioned primarily around traffic engineering, and has become the de facto standard for meeting constraint based routing requirement within domains.

### 1.10.3 Hardware

There are two conflicting trends impacting router complexity (see ([Awe99](#))):

- demand for faster lookups and scheduling, and
- feature creep

**High-speed Lookups and Scheduling** Internet topology is not a tree. This means that routers in the act of routing need to make choices between routes, known as the longest prefix match (LPM). LPM is an expensive operation when performed at a rate of millions of packets per second<sup>2</sup>. The primary hardware has been the use of SRAM over DRAM, with a latency advantage of 5 *ns* versus 50 *ns*, and the overriding trend is the use of virtual output queueing where each input has  $N$ , one for each output to enable high speed crossbar-based switching. These techniques are leading to increasing router complexity as a result.

---

<sup>1</sup>The non-broadcast nature of ATM lead to inefficient multi-hop communications and there was significant overhead associated with servicing the basic requirements for maintaining logical adjacencies. However, an even more important contributor to the failure of ATM was the availability of cheap 100Mb/s “fast” ethernet in the mid 1990's

<sup>2</sup>A router has to deal with packets of minimum size 40 bytes, at link speeds of between 10 Gbits/s to 40 Gbit/s. This means the router has 32 *ns* to decide what to do at 10 Gbits/s, and 8 *ns* at 40 Gbits/s.

**Feature Creep** The trends associated with speed requirements conflict with the routers position in IP networks. In the last decade, ATM was the only concerted effort by both the research and industrial community to introduce features through an overlay model. However, the strong coupling between the control and data planes of the IP protocol made this impossible (again, see Section 5.1.4). Since then, significant feature deployment has shifted to routers<sup>1</sup>. This is perhaps the most pervasive hardware trend in networking today.

Hence routers are frequently at the front-line of feature deployment. A rough list of features required in a modern router<sup>2</sup> are:

- Classification
- Measurement counters
- IPSec, VPN<sup>3</sup> and Firewalls support
- QoS<sup>4</sup> support and traffic isolation
- Shaping and policing of traffic
- IPv6
- Denial of service prevention
- Multi-cast
- Active queue management
- QoS Integrated and Differentiated services

---

<sup>1</sup>It can be argued that Active Networking was also driven by this constraint: see Section 1.6.3

<sup>2</sup>with the exception of IPV6, all these features have also been implemented in switches. However, we just maintain that they are predominantly deployed using routers, especially on smaller networks.

<sup>3</sup>Virtual Private Network.

<sup>4</sup>Quality of Service.

However, these trends will surely exact a cost at some point. One possibility is that clock rates will be held back by gate speeds, meaning at most 12-19% growth per year (see (Asson)) as deeper pipelines reach practical limits. Even if these limits are extended, there will be significant increases in the complexity of the logic required to maintain a large number of features on one device. This will lead to large losses in efficiency. In Chapter 4, we maintain that highly complex structures are very difficult to operate optimally, and this analysis almost surely carries over to this domain.

This feature list places greater and greater processing requirements on routers. As a results, the key architectural trend in current hardware is the offloading of tasks to Network Processors and dedicated hardware (See (Fer03),(FBR<sup>+</sup>04), and (TKS<sup>+</sup>02)). This trend reduces the overall complexity of the router by separating routing from the forwarding function<sup>1</sup>. Some dedicated tasks being offloaded include but are not limited to:

- Packet switching, such as MPLS Label Edge Routing
- Layer 3+ switching
- Aggregation
- Traffic and Qos Management
- VPN, firewalls and security management

Currently, there appear to be roughly three segments in the network router domain:

1. *Core routers* with extremely high throughput
2. *Network edge equipment* such as traffic aggregators and VPNs
3. *Access equipment*, such as cable models and wireless 802.11g routers

---

<sup>1</sup>Even though it does not have an impact on the total complexity of the network.

Some features are being implemented in new dedicated hardware which is defining new segments, such as storage and firewalling. Some of this is also due to new usage patterns which deviate from historical norms. Combined with the trends in offloading tasks to network processors, the number of segments will increase over time.

### 1.10.4 Summary

Overall, the key trends are:

- Reduction in the number of layers
- Increase in the number of segments
- Increase in IP layer complexity
- Offloading of tasks to specialised machinery

The strong coupling between the control plane and the data plane of IP is the common thread running through these trends. Commonly, it is assumed that the reduction in the number of layers is driven by cost considerations. This is possibly true, but comes at the expense of increased brittleness of the network, and gives rise to additional restoration and fault-management complexity within the IP layer. Increases in router complexity are also a function of the difficulty of adding new features to the control plane. Overall, the reduction in the number of vertical layers is offset by an increase in the number of horizontal layers at the IP level. As a result, the total complexity of the network is not being reduced, just shifted around. This may not be resulting in a network which is easier or cheaper to operate when all costs are taken into account.

## 1.11 Roadmap

The first chapter starts with a discussion regarding the unsuitability of the neo-classical paradigm for dealing with real-world markets and networks. An alternative paradigm based on an *Entropic Formulation* for networks is proposed. An argument is made in favour of treating the networks macroscopic characteristics

directly as objective functions, rather than as side-effects of interactions amongst hypothetical agents. The *Entropic Formulation* is justified in various ways, and the way in which networking problems could be solved in this framework are investigated.

Having laid out the justification for considering complexity issues using entropic models, the second chapter builds on this foundation by proposing two models: an *Architectural Model* which aims to illustrate complexity in action “bit-by-bit” in the network, and a Measurement Model which is meant to be a more practical tool for actually measuring complexity in a network with easily measurable data. Both of the models create an entire vocabulary with which to reason about networks, akin to the language of the financial markets. Also, a numéraire for networks is introduced.

Chapter 4 then analyses the models in Section 3 further through various sensitivity tests. Both models can give the engineer strong insights about networking issues. There are a number of interesting points which are highlighted and analysed in this section, but three stand out for special mention: there is a tradeoff between complexity and risk, networks with higher levels of complexity deal better with sub-optimal control policies, while networks with low complexity are more sensitive to changes in operational requirements.

Chapter 5 is meant to validate the models of Chapter 3. The *Architectural Model* is tested with a number of typical high-level issues that come up in networking: connection-oriented vs. connection-less protocols, end-to-end arguments, and rolling new features into the addressing space. The *Measurement Model* is used on actual corporate network data to show the evolution of *implied complexity*, *Implied Coupling* and *Network Complexity* over a two year time period. This vernacular is laid out in Chapter 3.

Chapter 6 discusses the long-term, architectural and financial implications of the proposals in Chapters 2 and 3. The key implication is that traffic and architectural complexity will increase over time, due to increased efficiencies in networks. This is not meant to be a high-level statement: it is a direct consequence of the increasing use of statistical inference in making networking decisions as routers are powered by increasingly more powerful processors and have access to more network data to make decisions.

Section 7 concludes and points the way forward.

## 1.12 ReReshaping the Agenda

In (SCEH96), the authors argued for a reshaping of the research agenda around network pricing:

As the Internet makes the transition from research testbed to commercial enterprise, the topic of pricing in computer networks has suddenly attracted great attention. Much of the discussion in the network design community and the popular press centres on the usage-based vs. flat pricing debate. The more academic literature has largely focused on devising optimal pricing policies; achieving optimal welfare requires charging marginal congestion costs for usage. In this paper we critique this optimality paradigm on three grounds: (1) marginal cost prices may not produce sufficient revenue to fully recover costs and so are perhaps of limited relevance, (2) congestion costs are inherently inaccessible to the network and so cannot reliably form the basis for pricing, and (3) there are other, more structural, goals besides optimality, and some of these goals are incompatible with the global uniformity required for optimal pricing schemes. For these reasons, we contend that the research agenda on pricing in computer network should shift away from the optimality paradigm and focus more on structural and architectural issues.

Pricing in Computer Networks: Reshaping the Research Agenda,  
*S. Shenker, D. Clark, D. Estrin, S. Herzog.*

However, over the ten years since (SCEH96), not much has changed. The existing body of literature based on networking economics continues to base its formulation broadly on the neoclassical paradigm of perfect markets and has sidestepped issues associated with architecture and structure.



In the meanwhile, during the decade, it has become an accepted fact that the neoclassical paradigm of markets is not a good model for financial markets and real world economies, even as a simple starting point.

Markets continue to function in the presence of incompleteness, inconsistency and intractability, at odds with the assumptions of the neoclassical paradigm. Similarly, the Internet has succeeded not because of, but in spite of market economics, as the emphasis has been more on engineering and experimentation. In many ways, it can be said that the existing body of literature has failed to recognise that most medium-size networks and the Internet already operate like markets in the real world: they have emergent properties, take shortcuts, make mistakes, institutionalise long-term memory, and find ways of dealing with incompleteness and inconsistency. The research agenda should focus directly on the macroscopic characteristics of these large systems, as opposed to treating them as side effects of interactions between hypothetical agents.

The ReReshaping the Agenda message is to refocus the research agenda back on the network itself. A good place to start would be by appreciating that networks are markets to begin with.

## Chapter 2

# Modelling Networks

The performance and characteristics of any network are a function of its elements and the way in which the elements interact with each other. The problems in elements design which are independent of the way the elements interact are broadly the domain of semi-conductors, fibre-optics, and low-level hardware design. These are not the focus areas for this thesis.

However, the way in which elements interact to form a network are. This area of research has the aim of attributing macroscopic properties to a large system composed of many elements as a function of the characteristics of those elements, the way the elements are allowed to interact, and finally, the larger environment.

The systems research community has approached issues which relate to interactions amongst elements using two separate techniques:

- In the pre-Internet era, Queueing Theory has been the basis for performance analysis in networks.
- Market Metaphors have been the dominant theme in networking research in the packet-oriented, post-Internet era.

Some researchers have considered an alternative: the use of statistical mechanics techniques from the the physics of thermodynamics and the related use of Information Theory. These two are closely interrelated through the notion of Entropy and would seem to be a natural candidate for the analysis of a large-scale system meant to carry information efficiently. An early example of such an effort

is the collaboration between Derek Mcauley and Ian Leslie at the University of Cambridge, John Lewis at The Dublin Institute for Advanced Studies, and Nils Bjorkman and Alexander Latour-Henner at Telia in the MEASURE project in 1996 (see (BLHM<sup>+</sup>96)), based on the relationship between thermodynamics and Large Deviation Theory (see (DLO<sup>+</sup>95)). This work has been used primarily for traffic prediction using non-parametric techniques.

Nonetheless, even the MEASURE project and related work have only focused on the use of statistical mechanics techniques for non-parametric traffic prediction. The use of information theoretic techniques to actually implement traffic engineering on the actual graph of the network, and to drive routing in networks is an approach which has rarely been considered. This thesis explores this third alternative.

This chapter has two goals: to critique the existing body of research based on market metaphors using neoclassical economics, and to introduce an alternative means for analysing large-scale networks using techniques from Information Theory and statistical mechanics.

Appendix D contains a brief overview of the main tenets of neoclassical economics. This Section starts out with a critique of neoclassical economics. This is followed by the most important Section of this thesis, which proposes the argument in favour of an alternative formulation of networking problems called the *Entropic Formulation*. This is followed by a smaller section which investigates other alternatives.

## 2.1 Imperfect Markets

In the real world, markets do not obey most of the assumptions on which the Neoclassical Paradigm is based. Markets are strongly imperfect in the neoclassical sense for a number of reasons:

- Markets do not clear
- They are incomplete
- The requisite calculations are intractable

- The requisite information is not easily attainable
- Even with perfect information and ignoring computational costs, agents have deliberation costs

Of these, the key imperfection in real-world markets is market clearing. Even though goods sold and bought in an economy must be equal, this does not mean that markets have cleared, and hence market prices do not have to reflect equilibria. The realisation of the eventuality is most commonly experienced in the labour markets, as Keynes noted in the early 20th Century.

### 2.1.1 The Tâtonnement Problem

Even before Arrow-Debreu, economists wrestled unsuccessfully with the implications of Walrasian Auctions. Walrasian auctions have theoretical and practical failings. The tâtonnement<sup>1</sup> process makes unrealistic assumptions: prices are announced, no transactions and no production take place until an equilibrium is achieved, prices are lowered/raised until there is no excess supply/demand, and then reach equilibrium. The key failing of the Walrasian framework is the assumption that no transactions take place outside of equilibrium, whereas in the real world, most transactions take place out of equilibrium. This remains an unresolved issue in the pure Neoclassical setting to this day.

Furthermore, Walrasian fixed point algorithms are exponentially complex in the number of commodities in an economy. They can be ruled out as part of equilibrium price formation in real world markets.

They are also dismissable due to their implications: for example, they obviate the need for money. Money clearly exists, and plays an important role in all markets. To summarise, Walrasian Auction setting has three failings:

- Information costs
- Transaction costs
- Computation costs

---

<sup>1</sup>Groping, finding the market price.

### 2.1.2 Market Clearing

Another failing of Neoclassical Economics is the assumption that markets clear. Market clearing either refers to the price equilibrium in an economy when demand equals supply, or the process of getting there. In the early 20th Century, the majority of the economics profession had doubts and objections against equilibrium assumptions: it did not come to be a standard assumption perhaps until the work of Walras and later Arrow-Debreu became so prevalent. For 150 years prior, the clearing of markets was also taken for granted, until the Great Depression of the 1930s and Keynes. Keynes key contribution to understanding markets was the notion of “price stickiness”, that rigidities can break down markets because prices are unable to adjust. Keynes was particularly concerned with disequilibrium conditions persisting for long periods, especially in relation to chronic unemployment<sup>1</sup>.

### 2.1.3 Equilibrium and Uniqueness

Since the work of Arrow-Debreu, it has come to be understood that with an even small relaxation of the assumptions underlying the Second Theorem (Section D.0.2.2), multiple equilibria are possible in the Arrow-Debreu economy. This is easily attained in even simple markets with only two goods and two agents with homogeneous utility functions.

As a matter of fact, Keynes argues strongly in the General Theory that markets usually spend most of their time out of equilibrium, adjusting to shocks and innovations. Any realistic model of real-world markets needs to be centred around disequilibrium, price adjustment processes, and the costs and means of gathering information.

Why then is equilibrium so important for neoclassical economics ? Without equilibrium:

1. price determination through supply and demand cannot be realised, and

---

<sup>1</sup>Amongst other economists objecting to the notion, one can count Marshall (who invented partial equilibria and demand/supply curves), Kaldor, Robinson, and Hahn.

2. demand and supply function determination through prices cannot be realised

. The recursive relationship between equilibrium and price determination which clear demand and supply is fundamental, and breaking the link breaks the price mechanism of neoclassical economics.

Yet, in the real world, very few markets are in equilibrium: prices are not stable, even in the absence of new information, and they can exhibit significant stickiness. Some of this stickiness is due to:

- regulation
- imperfect competition
- slow adjustment of prices, for example in the labour markets

When even one of these conditions is present, equilibrium and the assumption of market clearing can be a bad starting point of analysis.

**Demands and Transactions** The existence of equilibrium is deduced via an argument that since total quantity of goods sold must equal total quantity of goods bought, real world prices are equilibrium prices. However, this is not a strong argument. (DS86) point out that while obviously, the amount sold and bought must be equal, this does not mean the market has cleared. (Ben84) separates demands and transactions, pointing out that the latter is what is observed in exchange markets while the former is composed of the signals transmitted to markets by agents before a transaction is realised.

### 2.1.4 Incomplete Markets

The key impact that incomplete markets relates to the Second Theorem. Risk-neutral probabilities are unique if and only if the market is complete. An incomplete market provides many solutions under this framework, and the problem then arises as to which of these is the right one<sup>1</sup>

---

<sup>1</sup>A frequent proposal for selecting the right solution is to choose the one which minimises the Kullback-Leibler (KL) distance between a prior. While this may seem arbitrary, it is related

In reality, many markets are incomplete: the existing set of primary assets do not span the entire space of contingent outcomes for the economy. This is also true of instances when the market metaphor is applied to telecommunication networks. Most markets that have existed or continue to exist for the provision of actual telecommunication services generally do not allow participants to hedge all their risks.

Related to the issue of market completeness is the existence of a fundamental paradox with the financial markets. While a full exposition of derivatives pricing is outside the scope of this thesis, in short, it can be said that derivatives pricing is based on a *replication* argument. A portfolio of securities (Arrow-Debreu Securities) is formed which dynamically replicates the value of an option at all time periods. The value of this option is then equal to the value of the replicating portfolio. The concept of a risk-neutral price arises out of the notion that a short position in the replicating portfolio combined with the derivative will be risk-free.

It is this notion which creates the following paradox:

**Definition 4 (Hakansson's Paradox)** *If markets are complete options are not needed; if they are incomplete then according to financial theory, they cannot be priced.*

No satisfactory solution to this paradox exists within the confines of the neo-classical paradigm.

### 2.1.5 Limits to Computation

Mathematics is unreasonably ineffective in solving the problems associated with neoclassical economics <sup>1</sup>. Until recently, the only appeal to computational com-

---

to the concept of an agent who maximises log-utility, which has some fundamental properties. The *Entropic Formulation* proposed in Section 2.2 below is also related to this proposal.

<sup>1</sup>From (Vel05) “In this paper, I attempt to show that mathematical economics is unreasonably ineffective. Unreasonable, because the mathematical assumptions are economically unwarranted; ineffective because the mathematical formalisations imply non-constructive and uncomputable structures. A reasonable and effective mathematisation of economics entails Diophantine formalisms. These come with natural undecidabilities and uncomputabilities. In the face of this, I conjecture that an economics for the future will be freer to explore experimen-

plexity in economics has been to impose the minimal requirement that any economic problem be *effectively computable*, i.e. that an algorithm terminate in a finite amount of time on a finite state machine<sup>1</sup>

The two early results in this arena are:

1. Rabin ([Rab57](#)) in 1957 who proved that there exist games whose equilibrium strategies are not effectively computable
2. Binmore ([Bin92](#)) established similar results in the area of non-cooperative games

None of this is that profound, as it simply involves adapting the halting problem to a different context. However, effective computability is a very weak requirement and the the narrower and more relevant issue of whether rational behaviour is possible in practical contexts is really related to intractability. The optimisation constraint faced by a rational expectations economy is for almost all practical purposes, intractable.

It has been the application of standard techniques from computability and inapproximability to the field of economics since the mid-80s that has really brought forth the issues association with calculation. On the issue of computability as it relates to game theory, the specific set of results, in rough chronology, are as follows:

1. Lewis, 1985, 1986, 1992: key economic concepts such as individual demand correspondences, Walrasian general equilibria, Stackelberg equilibria, and Hurwiczian resource allocation problems are not effectively computable, even when we restrict attention to the class of problems where the descriptions of primitives of the problem (endowments, preferences, choice sets, etc.) are recursively presentable i.e. all functions describing preferences and production technologies are restricted to be effectively computable and all choice sets, etc. are restricted to be recursively enumerable.

---

*tal methodologies underpinned by alternative mathematical structures. The whole discussion is framed within the context of the celebrated Wignerian theme: The Unreasonable Effectiveness of Mathematics in the Natural Sciences."*

<sup>1</sup>The *Halting Problem* is an example of a problem which is not *effectively computable* ([GJ79](#)).



2. Tsitliklis, Athans, 1985: Team Decision Problems are NP-Hard
3. Nachbar, 1993: the problem of determining a best response to any fixed strategy of an opponent is not effectively computable in games where the discount factor is sufficiently close to 1, in the context of infinitely repeated games.

Many of these problems are studied in the context of a branch of computation termed Information Based Complexity (IBC). (Tra88) provides a good exposition. Most IBC problems suffer the "curse of dimensionality" in the worst case deterministic setting, meaning complexity is exponential in the dimension of the problem.

However, in spite of these results, it is interesting that even though many computer scientists have considered the use of market-based solutions to networking problems, the issue of tractability has been largely ignored. A problem is intractable when:

$$comp(\epsilon, d) = o\left(\left(\frac{1}{\epsilon}\right)^d\right) \quad (2.1)$$

where  $comp$  is the complexity function representing the minimal computation time of producing an  $\epsilon$ -approximation with the problem dimension  $d$ .

A large number of mathematical problems including multivariate integration, non-linear root finding, fixed point finding, solving PDEs, etc. have been shown to be intractable in this sense. This implies that the entire infrastructure associated with neoclassical economics is intractable, including but not limited to:

- utility maximisation
- maximising expected utility
- finding Walrasian equilibrium
- finding a Brouwer fixed point
- calculation of rational expectations equilibria
- computation of option values

- solving infinite horizon dynamic programming problems
- calculation of Nash equilibria

Naturally, these results about the incomputability of the most basic elements of the infrastructure of Neoclassical Economics have important implications for mechanism design:

- social planning (FO95)
- combinatorial auctions (FLBS99)
- dynamic programming (BT00)

are all intractable.

### 2.1.6 Limits to Information

There are also information aspects of computational intractability. When every agent knows the payoffs and strategies of all other agents, they have *complete information*. Rationality implies complete information: without it, perfectly competitive markets cannot exist as actors could not calculate the impact their actions would have on others.

Obviously, this is not a realistic assumption. The problems are not only practical. For example, in *The Team Decision Problem*, teams are not able to exchange any information, be it through a central planner or some other mechanism, and try to solve a problem in unison: given signals  $Y_1$  and  $Y_2$ , actions  $A_1$  and  $A_2$ , compute decision rules  $\alpha_i : Y_i \rightarrow A_i$  which minimise:

$$J(\alpha_1, \alpha_2) = \sum_{y_1 \in Y_1} \sum_{y_2 \in Y_2} c(y_1, y_2, \alpha_1(y_1), \alpha_2(y_2)) p(y_1, y_2) \quad (2.2)$$

under a discrete probability density  $p(y_1, y_2)$ . It has been shown that this problem is NP-Hard, and is an important starting point for the analysis of cooperative schemes in networking in large-scale networks where the cooperating agents do not share the same information space.

When information can be exchanged, the problem becomes solvable in polynomial time. Furthermore, not even perfect information exchange is needed. The

“price vector” in the First Theorem of Welfare is a “sufficient statistic”, so just knowledge of market prices would be enough.

However, this is not as promising an outcome as one would think. Solving a simple cost function is easy, but computing a Walrasian equilibrium or a solution to a social planning problem can be shown to be intractable even if the agents are able to share information and an exact solution to the problem they are trying to solve is not required, i.e. an approximate solution would do.

### 2.1.7 Deliberation Costs

Even in the presence of complete information, preference-formation complexity, dubbed *Deliberation Costs*, can also be substantial. Agents need to have preferences to participate in mechanisms. Determining preferences may have a deliberative overhead. In most markets, agents do not know their preferences and must make decisions not just about prices and goods, but also about their actual demand functions.

Deliberation overhead has important strategic and mechanism implications. It can be shown that no direct market mechanism exists which can simultaneously satisfy the constraints that:

- the market mechanism should not solve the agents deliberation problem
- agents should not deliberate on other agents preferences
- there should be no strategic misrepresentation of preferences
- the outcome should be affected by agents deliberations

and at the same time, allow agents to reveal their preferences truthfully (see (LS04)). Not only is there a practical problem with understanding everyones demand function, there are also strategic implications: agents may deliberate strategically causing others to believe that their true preferences are impossible. The main point is that, even if the resources were there to discover and collect all the information necessary, the information which comes back may not be truthful.

### 2.1.8 The Implication for Networking Research

A growing body of research in the last two decades has aimed to debunk the use of Neoclassical Economics as even a starting point for economic analysis: the failings are so fundamental and the assumptions so unrealistic that one would be better served with an entirely different paradigm.

The issues which plague Neoclassical Economics in financial markets are equally relevant in the realm of Telecommunications Networks:

- Any element in the network which is modelled as an agent in the Neoclassical setting, be it a router or a bandwidth broker, faces the same incalculability issues faced by economic agents in financial settings.
- Similarly, markets for telecom resources are also unlikely to be complete, or to clear. Nowhere has this been more apparent than in the markets for traded bandwidth. However, this issue is also likely to be relevant in the context of market-based resource allocation techniques which are based on shadow prices<sup>1</sup>.
- Can “transactions”, i.e. allocations of bandwidth or capacity to network elements take place outside of equilibrium ? They must, if the network is going to operate in real-time, as there will not be enough time at each epoch for a central auctioneer to calculate optimal allocations in the Walrasian style.

The implications of incalculability and limits to information would imply that according to Neoclassical theory, markets should not be functioning. Yet, in practice, markets do work. Every day, markets in the real world complete millions of transactions, allow risks to be taken, controlled and transferred, and provide liquidity. Furthermore, they resolve difficult problems which require cooperation between selfish agents, like transaction and information retrieval in elegant ways.

---

<sup>1</sup> Shadow prices are usually defined as the change in the objective value of an optimal solution when the constraints are relaxed by a unit. The Lagrangian multipliers usually correspond to shadow prices.

## 2.2 The Entropic Formulation

To manage the complexity of a network, a derivation of the macroscopic characteristics of a network is needed.

The key building block of Neoclassical Economics is the rational agent who maximises utility in the presence of complete, infinitely liquid, markets. Cranking the mathematics of convex optimisation, under severe constraints and notwithstanding intractability, leads to the First and Second Welfare Theorems, which are the key macroscopic features of a neoclassically perfect market (see Section D.0.2.1, and Section D.0.2.2). It is here where one must begin the process of steering the agenda in an alternative direction.

A good example of a market metaphor being applied to networks is the celebrated work by Richard Gibbens and Frank Kelly in (GK98), where the network is composed of agents maximising log utility<sup>1</sup>

The alternative to this approach, proposed here, divorces the macroscopic analysis from the side effects of agents interactions, and instead focus on the impact that informational efficiency has on the orderliness of the system, and the probabilistically likely configurations a system can realise. These are the key building blocks of Information Theory and statistical mechanics, and the macroscopic feature that is the result of this analysis is the Entropy of the configuration. This Section attempts to justify the use of this metaphor in networking.

The argument goes briefly like this:

- In Information Theory, it can be shown that the smallest encoding of the probability distribution of states in a system is limited to its entropy. The length of the encoding is closely related to how much work it takes to fully characterise a system. If a system were maximally efficient and *arbitraged* all information such that there were no *free lunches*, the minimal encoding of the system would be increased as much as possible, making it harder and harder to characterise the system. The system would behave as if to maximise its entropy.

---

<sup>1</sup>It turns out that the fortuitous use of log utility equates this framework to the alternative paradigm proposed in this thesis. See Section 2.3.1.

## 2.2 The Entropic Formulation

---

- Likewise, in statistical mechanics, entropy is related to probability: more probable states (observations) reflect higher entropy, and the most probable state is the one with maximum entropy. If one were to make any assumption about which state would be the prevailing one, the one with maximum entropy would be the right choice.

The rest of the analysis is the result of cranking the handle on the statistics and probability of configurations and constraints specific to networking.

It can be shown that:

- Maximum Entropy is implied from the behaviour of agents utilising statistical inference
- Maximum Entropy is also implied when market making agents are forced to achieve informational efficiency
- Maximum Entropy implies market efficiency
- The Entropic hypothesis can also be based in the context of a market microstructure, and does not conflict with rationality, and can be related to a specific instance of the Agents problem. (see Section 2.3.1).

The freedom to operate outside the Neoclassical framework brings with it significant benefits:

- agents face information retrieval costs
- arbitrage opportunities do exist
- the market is not complete,
- markets may not clear, and
- transactions may, and frequently do, take place outside of equilibrium
- it is possible to make relatively insightful statements about a networks behaviour outside of equilibrium

In sum, the departure from the neoclassical framework means that modellers should assume that the network behaves as if to maximise the *Entropy* of its configurations, and not the *utility* of its agents.

In the following sections, the *Entropic Formulation* for a network model is justified in three ways:

**Minimax Approach** A Minimax Approach which mirrors the rational agent formalism of the Neoclassical Paradigm, using the setting of a zero-sum game.

**Statistical Approach** A Statistical Approach which mirrors the formalisms of thermodynamics.

**Information Efficiency Approach** An Information Efficiency Approach based on the use of Maximum Entropy as a tool for Statistical Inference.

### 2.2.1 The Minimax Approach

In order to bridge the gap with the Neoclassical Paradigm, a good beginning point for justifying the entropy of a networks behaviour as a macroscopic feature is one which also uses the agent paradigm in the context of a market.

Assume that:

- A market exists for some kind of asset<sup>1</sup>
- The market is composed of Market Makers and Investors
- Market Makers broker transactions between, and interact with Investors

In a multi-period setting, the price of the asset at maturity is a random variable, while it's current price is a function of what Investors, as a collective, believe the price at maturity will be. A zero-sum game unfolds between Investors and Market-Makers as the asset is sold and bought from one period to the next.

In Game theory, a 2 Player Strategic Game is defined as a game in which :

- Each player has a strategy

---

<sup>1</sup>This could be capacity, or a network.

## 2.2 The Entropic Formulation

---

- Each player has a payoff or utility

A Zero-Sum Game is a game in which total utility of all players is zero.

In this hypothetical setting, information discovery and transactions follow a cycle:

1. Initially, Market Makers learn Investors beliefs by interacting with them. Investors do not interact with each other. Hence, Market-Makers know more.
2. Conditional on this information, they announce a Price
3. Investors update their beliefs, and transact.
4. The cycle begins anew.

Consider further that there is some kind of constraint on Market Makers which forces them to minimise the number of questions they can ask to infer the probability distribution associated with asset prices. This could be for the following reasons:

- Market Makers compete with each other to announce prices as soon as possible, as the earliest to announce will win the largest number of transactions, and transactions create profits.
- Investors do not wish to share information with Market Makers.
- Information is costly to collect and process.

Hence, this setting admits strategic behaviour by market participants, and also transaction costs, in the form of information acquisition costs. Furthermore, information efficiency is a source of competitive advantage: Market Makers who reveal the most information with the least number of questions are more profitable.

**Definition 5 (Market-Makers problem)** *Find the least number of questions to ask investors to discover their collective belief about future prices.*



## 2.2 The Entropic Formulation

---

In a zero-sum setting, meaning that if Market Makers profit, Investors will lose. This means that one can use the minimax method<sup>1</sup>, where Investors have the dual problem of making it as hard as possible for Market Makers to infer the correct probability distribution for asset prices (for a good reference, see (FT91)).

**Definition 6 (Investors Problem)** *Make it as hard as possible for Market Makers to infer the probability distribution of asset prices.*

The Market Maker and the Investor face a reciprocal problem: the Market Maker tries to minimise the binary poll leading to the statistical distribution of asset prices, and Investors try to maximise the binary poll that Market Makers must carry out. This is the Minimax problem in a zero-cost game.

A classic result from Information Theory, the Market Makers cost is well approximated by the entropy of the assets probability distribution,  $H(g)$  to a constant factor where  $g$  is the distribution of future prices<sup>2</sup>. The proof of this statement is based on the use of binary codes as representations for the encoding of a random variables values, illustrating that the expected length of any code cannot be shorter than the entropy of the random variable, using the Kraft Inequality. See (CT91) for an exposition.

And, leading to the key conclusion of this Section, as the Investor tries to maximise this cost, this leads to asset prices which have the maximum entropy possible under given constraints.

Due to the binary nature of the poll, the cost to the Market Maker can be measured in bits: the minimum number of binary yes/no questions a Market Maker must ask to determine price. This ties in with the notion of *Interrogation Costs* playing a role in the determination of the complexity of a system, stated in Chapter 1.

---

<sup>1</sup>Minimax is a method in decision theory for minimising the maximum possible loss. Alternatively, it can be thought of as maximising the minimum gain (maximin). It started from two player zero-sum game theory, covering both the cases where players take alternate moves and those where they make simultaneous moves. It has also been extended to more complex games and to general decision making in the presence of uncertainty. See (FT91).

<sup>2</sup>In the Arrow-Debreu market, an economy reaches a market equilibrium which maximises individuals utilities. This result hinges on firms being price-takers, and that they do not consider the impact of their actions on prices.

This simple setting already resolves a number of problems associated with the Neoclassical setting:

- Transaction costs do not have to be zero
- Information costs do not have to be zero
- The market does not have to be complete for transactions to take place
- Arbitrage is possible
- Inconsistency is possible
- Multiple equilibria are permissible, depending on the a priori information sets associated with various market makers
- T tonnement problem of Walras (Section [D.0.1.1](#)) is also resolved
- Strategic behaviour is permissible
- The assumptions are realistic: Investors are reluctant to reveal information, so prefer shorter polls
- It is ensconced in the classical setting of an exchange with market makers.

### 2.2.1.1 Minimax Applied to Networks

Here is one hypothetical ways the minimax game version of the Entropic Formulation could be applied practically to a common networking problem: in a multi-path setting, a service provider and a user play a zero-sum game in which the provider allocates resources at links to limit the transmission capacity provided to a user, using some kind of sampling. This is akin to setting probability distributions over the set of paths from the transmitting node to the receiving node. The user then tries to minimise the maximum probability of his packets being hit by transmission constraints. The service provider tries to maximise the minimum probability distribution that the user may choose.

A minimax solution for this zero-sum game exists in the Game Theoretic framework. An optimal transmission constraint/sampling strategy can then be devised for the service provider.

### 2.2.2 The Statistical Mechanics Approach

Another approach which justifies Maximum Entropy as a macroscopic feature of networks comes from statistical mechanics.

It is based on an analogy with the *Second Law of Thermodynamics*, and in essence, posits that the maximum entropy state of the market is the most likely outcome amongst all possible outcomes as the end result of irreversible transformations between agents<sup>1</sup>.

The field of thermodynamics studies the behaviour of energy flow in natural systems. From this study, a number of physical laws have been established:

**Zeroth Law** If system  $A$  is in thermal equilibrium with system  $C$ , and system  $B$  is in thermal equilibrium with system  $C$ , then system  $A$  is in thermal equilibrium with system  $B$ .

**First law** Total energy of a system is conserved.

**Second law** Entropy of an isolated system can only increase.

The Second Law is the key link between reversible and irreversible transformations in thermodynamics. In essence, it allows a system to recover irreversibility from a model based on reversible mechanics and statistics.

---

<sup>1</sup> In thermodynamics, processes that are not reversible are termed irreversible. All natural processes, to some degree, are irreversible. The phenomenon of irreversibility results from the fact that if a thermodynamic system of interacting molecules is brought from one thermodynamic state to another, the configuration or arrangement of the atoms and molecules in the system will resultantly change. A certain amount of “transformation energy”, also referred to as “free energy”, will be used as molecules change from one state to the next. During this transformation, energy will be lost due to intermolecular friction and collisions, energy that will not be recoverable if the process is reversed.

Walrasian auctions, or any other monotonically utility-improving trading algorithm, represent irreversible transformations in the language of thermodynamics. The trading history cannot be reversed because agents will generally refuse to accept trades that reduce their utility, and with suitable, standard assumptions, opposite sides of the same trade can never both be utility-improving for an agent in a given initial state.

Another way thermodynamics admits transformations is to require that they be reversible: those in which a system can be made to pass through a set of states, each arbitrarily close to some equilibrium, in both a forward and its reverse sequence. Likewise, reversible transformations in economics will be those that move agents within their indifference surfaces.

Analogously, it is also possible to set-down a similar set of Fundamental Laws for networks and markets, and crank the handle of statistics to arrive at a number of interesting macroscopic results<sup>1</sup>:

**0th Law/Encapsulation** Agents will not transact if their valuation for all goods are the same.

**1st Law/Constraint** Goods are not created through transactions.

**2nd law/Preference** Agents will only do transactions that increase their utility,

It is possible to conceptualise a system of markets or networks arising from interactions between agents who have mutually advantageous trades to make or data flows to exchange. Just as in thermodynamic system, such a system of markets or networks would also reach an equilibrium which would be irreversible, due the Law of Preference. But what would characterise this equilibrium ?

Instead of focusing on the actual interactions between agents, the statistical mechanics approach focuses instead on the likelihood of configurations achievable by an existing network. Any network, be it an actual data network, or the internal mechanics of a router, is wired up in a specific way. This wiring and interactions between elements in the network places limitations on achievable configurations. If one could attribute probabilities to different *equivalent configurations*, it would be possible to come up with the most likely configuration as a function of the configuration type which can be reached in the largest number of ways. This would be the configuration with the highest Entropy, as a result of a probabilistic analysis of possible permutations<sup>2</sup>.

---

<sup>1</sup>See (Lio04) for a holistic analysis of this issue; interesting as the analogy may be, the parallels were not the reason for considering entropy as the basis for measuring complexity in this thesis.

<sup>2</sup>Using Stirlings Approximation a la Boltzmann. In statistical mechanics, Stirlings Approximation applied to configurations of gas molecules leads to the most likely configuration, characterised macroscopically as that which has the highest *Entropy*. A brief summary is given in Appendix C.

### 2.2.2.1 Relationship with Neoclassical Economics

The difference in the result of this analysis with the neoclassical framework would be that, in the Walrasian equilibrium, the equilibrium configuration would be the Pareto Optimal configuration, which need not be (and most likely isn't) the configuration with the highest Entropy.

The application of this overarching idea to networks involves the specification of how configurations come about and what is meant by *equivalent configurations*. In thermodynamics, two configurations of gas molecules are considered to be *equivalent configurations* if in each, the same number of gas molecules occupy certain energy states. In a network, two configurations could be considered equivalent when they share some similarity in their node or edge degree distributions if the network is able to reconfigure itself dynamically, or if they are similar in the distribution of flows along paths.

Maximum Entropy is a much more realistic and easily justifiable assumption for the statistical characteristics of a graph when there are no constraints. This is the statistical mechanics basis of the argument for the use of Entropy as an objective function in network modelling and control. Where there are constraints on a graph characteristic, one minimises the Kullback-Leibler (KL) distance between the desired distribution and the achievable distribution.

In this formulation, unlike that of Sections 2.2.3.1 and 2.2.1, it is important to see that there are no “economic” forces moving the network to the maximum entropy equilibrium: indeed, the network may arrive at other equilibria from which no further configuration changes would happen.

However, even in the absence of dynamics, this is still an important improvement over the Walrasian Auctioneer. While both Walrasian Equilibria and Entropic Equilibria represent configurations which no element in the network would want to change, and hence are both *Pareto Efficient*, the maximum entropy equilibrium permits bilateral exchanges between elements outside equilibrium. Walrasian competitive equilibrium explicitly rules out trading at disequilibrium prices. Furthermore, the maximum entropy equilibrium does not require elements to end up in the same configuration in the network, whereas in a Walrasian equilibrium, all elements of a given type would end up with the same configuration.

## 2.2 The Entropic Formulation

---

In Walrasian equilibria, all agents of the same type end up consuming the same bundle of goods. This implies that the entropy of a Walrasian equilibrium is zero. This is an important difference between the Walrasian Equilibrium and the equilibrium proposed in this setting: while the equilibrium in the *Entropic Formulation* is the most likely outcome, the Walrasian Equilibrium is a probabilistically highly unlikely outcome, as it has zero entropy. When one considers the information requirements associated with this constraint, one immediately recognises that the reduction in entropy in an economy with heterogeneous agents would be enormous.

The vast amount of information required to sustain this very unlikely allocation is part of the *tátonnement* problem associated with the Walrasian *Auctioneer*, an agent whose only purpose was to calculate equilibrium prices without using any of the resources of the economy, and without any actual trading was meant to overcome this problem. Not only does the Walrasian approach assume informational costs are zero, it also posits an equilibrium which has the highest informational requirement amongst all equilibria as its entropy is zero. This formalises the informational argument against the neoclassical approach.

### 2.2.2.2 Statistical Mechanics Applied to Networks

In some ways, it is relatively easy and straightforward to apply statistical mechanics techniques to networks. The advent of powerful computers and the availability of communication networks, it is now not uncommon to see networks with millions of vertices. This has led to a spate of recent research in statistical techniques for quantifying large networks.

The format for using the tools of statistical mechanics to solve networking problems would be:

1. Enumerate all important characteristics of a network: degree distributions, average/minimum/maximum paths between nodes, clustering coefficients, etc.
2. Determine the constraints amongst these characteristics due to some technical requirements: for example, if one wishes to have a high level of resilience

against the random removal of vertices, a power-law distribution would be desired for node degrees.

3. Relate these characteristics to some function of the networks performance which is an objective.
4. Aim to control the network such that the desired characteristics are reproduced and to make sure the graphs are maximally random in all other respects.

Some important topology characteristics which could be considered include but are not limited to:

- Distance/Geodesic (shortest path length) distribution:
  - Performance parameters of most modern routing algorithms depend solely on distance distribution
  - Prevalence of short distances makes routing hard (one of the fundamental causes of BGP scalability concerns (86% of AS pairs are at distance 3 or 4 AS hops))
- The replication of Small-world effects: for example, to create networks with good search properties
- Transitivity effects, for creating networks with good clustering properties
- Degree distributions, for creating scale-free networks, which are resilient to random node-removals and which also show good search properties
- Degree correlations, to replicate observed properties such as the linkage of back-bone networks with other back-bone networks, the linkage between ISPs and end-users, and the absence (or, perhaps presence depending on the context) of linkages between end-users

A good starting point would be the degree distribution of random graphs, which is given by:

$$p_k = \binom{n}{k} p^k (1-p)^{n-k} \simeq \frac{z^k e^{-z}}{k} \quad (2.3)$$

## 2.2 The Entropic Formulation

---

where  $n$  is the number of vertices in the graph, and  $z = p(n-1)$ , the mean degree of the network. While a random graph reproduces well the small-world effect seen in most communication (and other) networks, it fails to reproduce some other common and desirable effects, such as:

1. Random graphs have low clustering coefficients, i.e. the probability of two vertices being connected is  $p$  regardless of whether they are both connected to another common vertex.
2. They have Poisson degree distributions, unlike most communication networks which have scale-free<sup>1</sup>/power-law properties
3. No correlation between degrees of adjacent vertices.
4. No searching properties using local algorithms.

The application of statistical mechanics probably would start with a random-graph overlaid with these constraints to come up with likely configurations and preferred algorithm for attachment and detachment to and within existing networks. Possibly:

**The Entropy of Links** In analysing a network using the *Entropic Formulation*, one possible candidate is to consider the entropy of the links:

**Constraint** Constrain the total number of links between nodes, to capture the notion of a network limited by bandwidth.

**Partition function** Describes the configuration of nodes in various link-buckets: the number of nodes with  $l_0$  links,  $l_1$  links, and so on.

**The Entropy of Nodes** Similarly, one could constraint the number of nodes and consider the probabilistic configuration of the network in terms of the entropy of the links.

**Constraint** Constrain the total number of links between nodes, to capture the notion of a network limited by the number of elements.

---

<sup>1</sup>Any functional form  $f(x)$  that remains unchanged to within a multiplicative factor under a rescaling of the independent variable  $x$ .



**Partition function** Describes the configuration of links in various node-buckets: the number of links with  $nd_0$  nodes,  $nd_1$  nodes, and so on.

The constraint could be derived from the assumption that nodes prefer to link to other nodes with higher degrees than their own. To provide for transit/peering constraints, one could cast this in the unidirectional graph framework.

**The Entropy of Buffers** While the previous two are obvious, perhaps a more interesting candidate is to consider “free buffer” space in the network. Free buffer space in routers will be a function of the total bandwidth of the network less the amount of data travelling through it, and the buffer space of individual elements.

**Constraint** Constrain the total amount of free buffer space available in the routers of the network.

**Partition function** Describes the configuration of buffers with available bandwidth.

This line of analysis would be most appropriate for the analysis of Quality of Service (QoS) in networks.

Some further possibilities are listed in the Table [2.2.2.2](#).

### 2.2.3 The Informational Efficiency Approach

A key characteristic of systems in which agents face informational costs and hence have to employ statistical inference is that the system is driven to maximum entropy.

The argument for this is based on the realisation that the act of asking of a question is an admission of uncertainty: hence, questions are a probability distribution. This then begs the question of which question to ask so as to obtain an unbiased result. The answer to this is the question which minimises entropy the most, also known as the Maximum Entropy Method in statistics, due to Jaynes (see ([Jay57](#)))<sup>1</sup>. The Maximum Entropy distribution for a constraint is the

---

<sup>1</sup>In his famous 1957 paper, Ed. T. Jaynes wrote:

## 2.2 The Entropic Formulation

---

Network Characteristic	Mathematical Representation	Comment
The probability of a vertex having degree $k$	$p_k = \binom{n}{k} p^k (1-p)^{n-k} \simeq \frac{z^k e^{-z}}{k!}$	Poisson random graph, $z = \sum_k k p_k$ is the mean degree in the network
The probability of a vertex having degree $k$ in a small-world network	$p_j = \binom{L}{j-2k} \left[ \frac{2kp}{L} \right]^{j-2k} \left[ 1 - \frac{2kp}{L} \right]^{L-j+2k}$	$L$ is the number of vertices, $k$ is the largest neighbour distance for small-world attachments
The probability of a vertex having degree $k$ in a power-law network	$p_k = \frac{2, (m+1)}{(k+2)(k+1)k}$	In the limit of large $k$ , $p_k \sim k^{-3}$ . $2m$ is the mean degree of the network
The probability distribution of the degree of a vertex $i$ with age $a$ measured as the number of vertices added after vertex $i$	$p_k(a) = \sqrt{1 - \frac{a}{n}} \left( 1 - \sqrt{1 - \frac{a}{n}} \right)^k$	For the case $m = 1$ .

Table 2.1: Key Characteristics of Networks

least committal, most *random*, most inherently uncertain distribution, making the smallest number of additional assumptions beyond what is known.

Under the assumption that agents use statistical inference, it can then be shown that *entropic dynamics* based on statistical inference move any closed system continuously and irreversibly along the entropy gradient. For a proof of this statement, see (Cat01) and (Cat04).

In summary:

1. Agents use statistical inference to make decisions
2. The most unbiased form of statistical inference is the Maximum Entropy method, re Jaynes
3. This leads the dynamics of the state along the entropy gradient: the act of asking entropy reducing questions and acting accordingly leads to the evolution of state to that which maximises entropy (see (Cat01)).
4. Hence, the maximum entropy objective.

### 2.2.3.1 Market Efficiency and Entropy

Another related way to justify the use of the entropic hypothesis is to consider it in the context of market efficiency. A highly competitive marketplace is associated with a lack of predictability, i.e. where there are no *easy* opportunities for agents to partake in profitable speculative behaviour. This does not mean there are no opportunities: however, over time, as agents using statistical inference techniques decrease the number of opportunities, prices become increasingly harder to predict and the systems behaviour can be characterised as leading to maximum entropy. Low entropy implies predictability, and exploitable opportunities through simple statistical inference.

---

Information theory provides a constructive criterion for setting up probability distributions on the basis of partial knowledge, and leads to a type of statistical inference which is called the maximum entropy estimate. It is least biased estimate possible on the given information; i.e., it is maximally noncommittal with regard to missing information.

---

## 2.3 Dissecting the Entropic Formulation

In the financial literature, this is known as the *Efficient Market Hypothesis* (EMH). The bedrock of EMH is the characteristic of asset prices being a Brownian motion fractal. Unsurprisingly, a Brownian motion fractal has infinite entropy.

The driving force behind price changes for any commodity is new information. In efficient markets, price adjustments take place instantaneously: this implies they are 100% efficient Brownian motion fractals. Markets form institutions to make this possible: they maximise the impact of new information by creating announcement and centralisation rules. This is related to the law of inference in the previous section, which relates inference based on entropy to the evolution of a state towards maximum entropy.

Price adjustments do not have to be instantaneous: if there were some time delay, entropy would be lower and there would be some predictability. Such predictability could be exploited through faster computation, or access to more information. There would be a simple tradeoff between the cost of this information and the arbitrage available from exploiting it. This relates to the market makers problem.

## 2.3 Dissecting the Entropic Formulation

The underlying mathematic tools used to maximise entropy are Lagrangian multipliers, as in Equation C.7. Suppose one were to find the discrete probability distribution with maximal information entropy. Then

$$f(p_1, p_2, \dots, p_n) = - \sum_{k=1}^n p_k \log_2 p_k. \quad (2.4)$$

The sum of these probabilities equals 1, meaning  $g(p) = 1$  with

$$g(p_1, p_2, \dots, p_n) = \sum_{k=1}^n p_k. \quad (2.5)$$

At this point, Lagrange multipliers can be used to find the point of maximum entropy (depending on the probabilities). For all  $k$  from 1 to  $n$ :

$$\frac{\partial}{\partial p_k}(f + \lambda g) = 0, \quad (2.6)$$

## 2.3 Dissecting the Entropic Formulation

---

which gives

$$\frac{\partial}{\partial p_k} \left( - \sum_{k=1}^n p_k \log_2 p_k + \lambda \sum_{k=1}^n p_k \right) = 0. \quad (2.7)$$

Differentiating the  $n$  equations:

$$- \left( \frac{1}{\ln 2} + \log_2 p_k \right) + \lambda = 0. \quad (2.8)$$

This shows that all  $p_i$  are equal (because they depend on  $\lambda$  only). By using the constraint  $\sum p_k = 1$ :

$$p_k = \frac{1}{n}. \quad (2.9)$$

As an aside, this implies, in the absence of any other constraints, that the Uniform Distribution is the distribution with the greatest entropy.

The addition of any constraints would be carried out by adding additional Lagrangian multipliers into Equation 2.7. Each new constraint would come with a new multiplier. These multipliers can be interpreted as shadow prices.

The key constraints that one would posit in the context of a network might be, from Section 2.2.2.2:

- Number of Links
- Number of Nodes
- Total availability of Buffer Space

The Lagrangian shadow prices associated with these constraints would represent the entropy prices of the links, nodes, or buffers, i.e. the increase in entropy of the network as a result of increasing the number of links, nodes or buffer space by one unit.

The likelihood of any configuration change for type (where type should be interpreted as a bucketing or some other characterisation of node degree, link capacity, or buffer space) is governed by the same entropy prices across the whole network (or perhaps, sub-system). This reflects the fact that all elements are indeed linked in the same network.

There are a number of interesting points to be made about the interpretation of these entropy prices as applied to networks:

## 2.3 Dissecting the Entropic Formulation

---

**Insight 1 (Likelihood of Configuration Changes)** *The likelihood of a configuration change is inversely proportional to its entropy value (or cost) at the entropy prices. Specifically, configurations which involve a large amount of elements are relatively unlikely. Entropy maximisation makes high entropy cost changes in the configuration of a network unlikely because they reduce the degrees of freedom in the assignment of other elements resources in their offer sets.*

**Insight 2 (Sensitivity to Entropy Price)** *When Entropy Prices are very high, the network reaches an equilibrium which becomes very difficult to change.*

In the context of a market metaphor: the decreasing likelihood of high entropy cost transactions is a reflection of the fact that every transaction must have an actual counterpart: larger transactions must satisfy more constraints in terms of finding actual counterparts, and hence are less likely.

The last point with regards to the magnitude of entropy prices is linked to a point which will be made in Chapter 4: networks with higher complexity are less sensitive to changes in their architecture.

In all this, it is important to point out that:

1. The number of agents of each type must be very large, so that statistical considerations dominate.
2. The analysis here pertains to a single epoch in the life of the network, when the surrounding environment does not change.

**The Role of the Market** In the *Entropic Formulation*, the key role of the network is connection between all elements through a *market clearing* condition, akin to market clearing in the setting of the *Neoclassical* setting. In the *Entropic Formulation*, the network exposes all elements to the same set of probabilities in equilibrium. In this way, the network itself plays the role of a market.

When two elements interact in a network, the combinatorial possibilities for other elements are reduced. However, the exact configuration realised as a result of this exchange between two elements is statistically impossible to determine.

## 2.3 Dissecting the Entropic Formulation

---

**The Role of the Market Maker** The statistical mechanics interpretation also posits an important role for *market makers* in a network. Sometimes, the realisation of a specific Pareto Efficient outcome may leave open the possibility that a higher entropy state is not realised due to network configuration constraints. Some of these constraints and their economic impact are surveyed in (DG04).

In the thermodynamic literature, entropy is a measure of the energy available to do work in a closed system. As entropy increases, this free energy is released, through irreversible transformations. When all transformation become reversible, the entropy level in a system reaches it's maximum equilibrium. By creating opportunities for transformations previously not possible, a market maker allows a system to release more free energy. This free energy is the profit available to a market maker who expands resources to bring about a higher level of entropy in the system.

**The Role of Interconnections** This analysis also allows for an interpretation of what can happen to a network which becomes interconnected with a larger network (say, the Internet). The conservation of the resources in the network in exchanges with the outside world shows that the entropy change in the network must be offset by a corresponding entropy change in the larger network, just as an entropy gain or loss of a thermodynamic system must be balanced by a corresponding loss or gain of the reservoir. In this sense, the open network can be interpreted as being attached to a heat reservoir at fixed temperature.

**Insight 3 (The Complexity of Interconnections)** *If a network opens up to the rest of the world the entropy of the world (i.e the economy and its environment) will increase. However, the result for the smaller network is actually indeterminate.*

### 2.3.1 The Relationship Between Entropy and Utility

The departure from the utility based framework is a significant deviation from accepted practice. However, the *Entropic Formulation* is not a construct based on an arbitrary assumption about a market or a magical assumption about a networks tendencies (towards maximum entropy). Neither is it a desire to stretch the

## 2.3 Dissecting the Entropic Formulation

---

*Second Law of Thermodynamics* beyond its intended purposes. Like the neoclassical utilitarians, it is also based on rational agents who employ statistical inference in their decision making processes. The difference with the neoclassical utilitarian set-up is that the actions of these agents are distributed probabilistically to account for real-world constraints, implying that prices are also distributed probabilistically.

Nonetheless, there are important correspondences between utility and entropy. The problem of finding conditions which imply the existence of a utility function that represents a given preference relation is called a *Utility Representation Problem*. In two fundamental papers ((Deb54) and (Deb64)), Debreu laid out the conditions for the existence of a continuous real-valued utility representation of a preference relation on a topological space, using the concept of a *gap*<sup>1</sup>:

**Theorem 1 (Debreu's Theorem)** *Let  $\preceq$  be a preference relation on a topological space  $(X, \tau)$  that is second countable. If  $\preceq$  is  $\tau$ -continuous then there exists a real-valued continuous utility function on  $(X, \tau, \preceq)$ .*

Similarly, in (Coo67), Cooper studies the foundations of thermodynamics by developing an appropriate mathematical order-theoretic structure. Given an accessibility relation on a state space  $S$ , the *Entropy Representation Problem* aims to find general conditions on the state space  $S$  and the accessibility relation which imply that there is a continuous order-preserving function  $f$  on  $S$  with the property that  $x \preceq y \Leftrightarrow f(x) \leq f(y)$ .

In (MMIC01), the authors show that the mathematical structure of the *Entropy Representation Problem* is exactly the same as that of the *Utility Representation Problem*. It is similar to the method of proof used by Debreu (Deb54) and consists in first constructing a strictly isotone function on  $S$ , which may not be continuous, and then modifying it by removing the untoward “gaps” to get a continuous entropy function on  $S$ .

---

<sup>1</sup>Let  $\overline{\mathbb{R}}$  denote the extended real line. A degenerate set in  $\overline{\mathbb{R}}$  is a set with at most one point. A **lacuna** of a subset  $S$  of  $\overline{\mathbb{R}}$  is a non-degenerate interval that is disjoint from  $S$  and has a lower bound and upper bound in  $S$ . A *gap* is a **maximal lacuna**.



### 2.3.1.1 The Role of Arbitrage

There is also an important correspondence between utility and entropy in the context of incomplete markets. In a market without any arbitrage opportunities, a unique valuation measure exists for the pricing of all contracts and the market is complete<sup>1</sup>. In such complete markets, there is no need for full-equilibrium models.

However, in an incomplete market, the lack of a unique valuation operator means that one must either work in the presence of a full-equilibrium model (which means making ad hoc choices about the utility functions of representative agents) to arrive at a unique valuation measure, or one must make ad hoc choices about the probability distributions asset prices (such as maximum entropy) are likely to follow:

- Choose a representative agent, and assume the market sets prices to maximise the utility of that agent. This requires an assumption about the utility function for that agent, and further that there is one representative agent in the entire economy.
- Choose a valuation measure which is not directly related to the utility preferences of individual agents, but is related to the statistical properties of asset prices. This is the method this thesis defends, and of many different characteristics, proposes the maximum entropy of asset prices as the right choice.

In (MP97), the authors illustrate a relationship between these two methods by showing that behind the selection criterion for the distributions of asset prices is its “shadow” problem, which can be shown to be connected to the utility maximisation problem of the profits of an arbitrageur subjected to a zero budget constraint. This is slightly different from the usual utility maximisation problems

---

<sup>1</sup>It can be shown that the absence of strictly acceptable opportunities among liquid assets, i.e. the lack of arbitrage, is equivalent to the existence of a representative state price density: this is equivalent to market completeness. An opportunity is strictly acceptable if its *expected* payoff is positive in at least one set of prices Arrow-Debreu prices (a valuation measure). This is a weaker form of arbitrage. See (Ros76).

## 2.3 Dissecting the Entropic Formulation

---

of agents who face budget constraints. In this case, the budget is zero, and only self-financing portfolios are allowed. When the arbitrageur is assumed to be maximising the utility of uncertain wealth and the utility function is exponential, the equivalent criterion minimises the relative entropy function.

The key to this result is duality theory: Lagrange multipliers of a convex problem are the optimal solutions of the dual problem. This permits the interpretation of risk-neutral probabilities (the probability distribution to be used to value financial contracts) to be interpreted as the normalised shadow prices of the arbitrage profit maximisation problem for the arbitrageur with zero-budget constraints.

The symmetry between the primal and dual problems in duality theory makes it possible to establish a one-to-one correspondence between each selection criteria and a particular utility function: the primal arbitrageurs problem of utility maximisation finds a dual problem in which a valuation operator is chosen based on a selection criterion that is the convex conjugate of the negative of the arbitrageurs utility function.

Nonetheless, it must be noted that, while, the entropy function of a network in equilibrium resembles a neoclassical utility function, one must be careful not to associate this with a representative agents preference ordering of alternative resource bundles. The entropy function is just the maximum of the statistical entropy over the set of probability distributions compatible with constraints on the characteristics of the network.

### 2.3.1.2 Proportional Fairness

A good model to consider in light of the highlighted correspondences between utility and entropy would be the log utility maximising networks of Frank Kelly et al in (KMT98). In this paper, the authors use logarithmic utility, a common choice due to its concavity (implying diminishing marginal utility) and monotonicity.

In the Kelly model, a utility function associated with network flows is optimised. Applying Lagrangian techniques, Kelly was also able to interpret certain variables as shadow prices and usage charges. Of all the different utility functions, a logarithmic (or diminishing returns) utility function for each flow favors

## 2.3 Dissecting the Entropic Formulation

---

short flows more than long flows, which Kelly termed *Proportional Fairness*. This resembles the shadow pricing based on Entropy in Section 2.3.

Define:

$U$  vector of utility functions of users

$A$  The adjacency matrix of the network

$c$  The capacity constraints on edges

$r$  An index for users

$x_r$  Amount of flow user  $r$  has been assigned

In this setting the *Systems* problem is to maximise Utility:

$$\begin{aligned} \text{maximize} \quad & \sum_r U_r(x_r) \\ & Ax \leq c \\ & x \geq 0 \end{aligned} \tag{2.10}$$

The *User* problem is:

$$\begin{aligned} \text{maximize} \quad & U_r(x_r) - \lambda_r x_r \\ & x_r \geq 0 \end{aligned} \tag{2.11}$$

where  $\lambda_r = \sum_j \mu_j$  in the Lagrangian of Equation 2.11:

$$L_1 = \sum_r U_r(x_r) - \sum_j \left( \sum_r x_r - c_j \right) \tag{2.12}$$

So, whereas in the entropic setting, the Investor is trying to learn as much as possible about future prices while revealing as little as possible, in the log-utility setting, the User is trying to maximise his own utility by pushing through as much flow as possible while trying to minimise (shadow) prices which are impacted by flows.

## 2.3 Dissecting the Entropic Formulation

---

The key insight in their work is that a first order condition for the maximisation of the sum total of all agents utilities in the face of congestion costs depends only on their utility and the marginal cost of congestion, also identified as *shadow prices*, and not on their preferences. The complementarity condition then leads to the problem of minimising the congestion in the network.

This leads to the problem for the Network:

$$\begin{aligned} \text{maximize} \quad & \sum w_r \log x_r \\ & Ax \leq c \\ & x_r \geq 0, \forall r \end{aligned} \tag{2.13}$$

Kelly et al use a log utility function as the basis for their argument.

There is a linkage between this implication and Kelly's Criterion from the world of gambling. The Kelly Criterion (see (Kel56)) maximises the long-term growth rate of a gambler <sup>1</sup> engaged in repeated plays and determines that the percentage of his wealth  $f^*$  that should be invested in the gamble is:

$$f^* = \frac{(bp - (1 - p))}{b} \tag{2.14}$$

where:

- $b$  is the odds of the bet
- $p$  is the probability of winning

In a perfect market,  $p=50\%$  and  $b = 1$  and the gambler bets nothing.

The goal is to maximise the expected value of growth of capital:

$$\frac{B(N)}{NB(N)} \tag{2.15}$$

as  $N$  grows without bounds.

The bankroll  $B(N)$  after  $N$  hands is equal to the product

$$B(0)_{i=1}^N \prod (1 + b_i \times R_i) \tag{2.16}$$

---

<sup>1</sup>Furthermore, the Kelly gambler has no risk of ruin. Keynes, the treasurer of Kings College, Cambridge, and arguably one of the first hedge fund managers of the century, is known to be a Kelly investor.

## 2.3 Dissecting the Entropic Formulation

---

where  $b_i$  is the fraction of bankroll wagered on hand  $i$ , and  $R_i$  is the per-unit result of hand  $i$ . To maximise that product, it is sufficient to maximise its logarithm:

$$\sum_i^N \log(1 + b_i \times R_i) \quad (2.17)$$

which is maximised when

$$E[\log(1 + b \times R)] \quad (2.18)$$

is maximised for all situations.

When moments of a random quantity  $Q$  ( $1 + b \times R$  above) are known, the expected value  $E[f(Q)]$  can be approximated for a function  $f$  ( $\log$  above) if the powers in the Taylors series for  $f$  are replaced with moments. Using Taylor's Expansion, and replacing the powers in the series with moments, one gets the result above using the first two terms ( $E[R]$  and  $E[R^2]$ ).

Once again John Kelly's setting is that of the log utility agent of Frank Kelly and Richard Gibbens in (GK98) and the entropy maximising market maker of Section 2.2.1.

### 2.3.2 Log Utility is Not Arbitrary

This may be more than a coincidence. In the theory of portfolio investing, it has been shown that the entropy criterion treats the investor's desired long-term portfolio growth rate target as a parameter. Given a value for that, the criterion determines the portfolio with the lowest probability of under-performing that growth rate over suitably long horizons. If the investor's target growth rate is the maximum feasible long-term growth rate, then maximising the criterion is equivalent to maximising expected log utility (see (KS02))<sup>1</sup>. This has been accomplished via both Lagrangian duality techniques and using results from geometric programming. See (KS02),

---

<sup>1</sup>If the investor's target growth rate is lower than the maximum feasible long-term growth rate, then maximising the criterion is equivalent to maximising an expected power utility, with a coefficient of risk aversion greater than (log utility's value of) one, that is also chosen to maximise that expected utility. The result is a more conservative portfolio that has a smaller chance of under-performing that target than does the expected log utility maximising portfolio.

## 2.3 Dissecting the Entropic Formulation

---

There is a particular relationship between the log-family of utilities and the *Entropy Formulation*. In (FS03), the authors consider the problem of learning a probabilistic model from the viewpoint of an expected utility maximising decision maker/investor who would use the model to make decisions. They express each Pareto optimal model as the solution of a strictly convex optimisation problem and its strictly concave (and tractable) dual. Each dual problem is a regularised maximisation of expected utility over a well-defined family of functions. Each Pareto optimal model is robust: maximising worst-case out-performance relative to the benchmark model. They show that the method of selecting the Pareto optimal model with maximum (out-of-sample) expected utility reduces to the minimum relative entropy method if and only if the utility function is a member of a three-parameter logarithmic family.

Similarly, under the Kelly Criterion, the agent wishes to maximise terminal wealth. Applying any monotonically increasing function to the growth of some resource would not change it's maximum. Of these, the log function is particularly convenient. Growth over time is a product. However, it is easier to work with sums than products. Specifically, one is frequently confronted with the decision of how much *utility* one should add at decision epochs, rather than the more difficult decision of what fraction to multiply wealth on average. Furthermore, it is easier to calculate the derivative of sums than of products.

Most error scores are logarithms because one typically wants to optimise posterior likelihoods conditioned on observations. For example, when one performs a "least-squares" regression analysis, it is implicitly assumed that the noise in the output is Gaussian, having the form,

$$Noise = \frac{1}{\sqrt{2\pi}s} \exp(-.5 \frac{x^2}{s^2}) \quad (2.19)$$

(or it's multidimensional counterpart), where  $s$  is the standard deviation of the noise, and  $x$  is the difference between the predicted value and the observed *training* value. Working with that noise distribution to obtain the joint likelihood of the estimates as a product of their individual likelihoods would be quite unwieldy. So instead, one takes the logarithm, discards some additive and multiplicative constants which do not change the maximum of the score, yielding:

$$Score = - \sum x \times x \quad (2.20)$$

---

## 2.4 The Analysis of Non-Equilibrium States

Adding rather than multiplying leads to the familiar *sum of squares* error score. This is the logarithm of assumed Gaussian noise, disregarding constants.

### 2.3.3 Entropy Flooding

The analysis above leads to an even stronger statement about closed systems with entropy maximising agents. Since entropy maximising agents maximise the long-term growth rate of their wealth, in a closed system, they will dominate the market over time with probability one. Other agents are eventually driven out of the market. This includes agents who base their decisions on the Capital Asset Pricing Model (CAPM, (CH05)). Log utility is not an arbitrary choice for an agents optimisation function, and neither is entropy maximisation.

**Insight 4 (Entropy Flooding)** *In closed systems, the behavioural rules of agents maximising entropy will dictate the dynamical behaviour of the system over time.*

## 2.4 The Analysis of Non-Equilibrium States

The maximum entropy state of a network is a statement about an equilibrium condition. In the absence of any interaction with it's environment, a system attains equilibrium at the most probable state, the one with maximum entropy<sup>1</sup>.

When the system is not closed and or not yet at equilibrium, the system is characterised by non equilibrium states. An analysis of non-equilibrium states leads to an understanding of how networks can be expected to interact with other networks, what kind of complex adaptive systems, such as institutions, can be expected to evolve, and what role such mechanisms play in a networks evolution.

One approach to analysing non-equilibrium states is based on the principle of *Maximum Entropy Production*:

**Principle 3 (Principle of Maximum Entropy Production)** *A system tries to maximise entropy as fast as possible.*

---

<sup>1</sup>In the realm of thermodynamics, if a system is closed, it will dissipate all the free energy it has, equivalent to maximising its entropy. Thermal equilibrium is the state of maximum entropy.

## 2.4 The Analysis of Non-Equilibrium States

---

In a closed system, the *Second Law* implies that the free energy of an isolated system is reduced over time, and the system reaches a state of maximum entropy. This transition to higher entropy is a transition to more probable states. Theoretical and applied studies have also suggested that the maximisation of entropy production during non-equilibrium processes (see (MS06)). The term *Maximum Entropy Production* has become a statement about how fast a system achieves maximum entropy, and the various structures it can maintain to bring about maximum entropy in a system as soon as possible.

In open systems, Ilya Prigogine has shown that the *Second Law* becomes a continuity equation where the total change in entropy is the sum of the entropy contributions from the system and its environment (see (PS84)). This essentially means that organised internal states can be maintained at the expense of discarding high entropy fluxes out of the system. This is also sometimes referred to as *Extropy*<sup>1</sup>. Further, Dewar (Dew03) has also derived the Maximum Entropy Production Principle and Self-Organised Criticality from Jaynes MaxEnt Formalism, implying that *Maximum Entropy Production* is the most probable outcome for certain complex systems<sup>2</sup>.

The Principle of Maximum Entropy Production can be related to a number of interesting constructs used in the complexity literature to study complex organisations. Maximum Entropy Production implies that:

**The Choice of a Trajectory** A system will evolve along trajectories which reach maximum entropy in the fastest way possible.

**Order from Chaos** A system may exhibit transient internal order if this makes it possible to attain maximum entropy states sooner<sup>3</sup>

---

<sup>1</sup>Coined by Tom Bell (T. O. Morrow) in 1988, extropy is a metaphor, and not a technical term. It is not simply the opposite of entropy, but broadly refers to the extent to which an organisational system exhibits intelligence, and functional order.

<sup>2</sup>Dewar uses Fluctuation Theory, to show that, under the assumption that non-equilibrium process can be considered as a superposition of changes in macroscopic parameters and microscopic spontaneous processes, non-equilibrium processes, in general, are accompanied by maximum entropy production.

<sup>3</sup>The classic example of this is sand piles in the context of the Self-Organised Criticality literature. See (BTW88).



## 2.4 The Analysis of Non-Equilibrium States

---

**Increasing Dimensions** A system may increase the dimensionality of its variables if this allows it to achieve a higher level of entropy<sup>1</sup>

A good direct example of a model in which an increase in dimension directly impacts network structure can be seen in the work of Giovanetti in (Gio02). Giovanetti studies the impact of network interconnections on retail and access prices. New connections between autonomous networks change the existing architecture of the Internet, adding new dimensions in which providers who were formerly monopolies are forced to compete against their downstream customers. When the downstream industry is poorly differentiated, this leads to lower retail and access prices.

Giovanetti also gives a good example of self-organised criticality in networks in (GA98). Firms located on a mono-dimensional lattice imitate the technology of immediate neighbours, where individual payoffs depend on local substitutability. Adoption of new technologies take place when agents obtain higher productivity. The length of the sequence of adoptions along the lattice is driven by the configuration of nodes, leading to self-organised aggregate dynamics characterised by catastrophic events over short period and inactivity over long periods.

A number of interesting insights can be gleaned from the perspective of the principle, expanded below. These includes:

- The Role of Institutions
- The Role of Diversity
- Complex Adaptive Systems

### 2.4.1 The Role of Institutions

The Principle of Maximum Entropy Production justifies structural complexity as an emergent phenomena in many complex systems which are out of equilibrium. This is a case of transient order being established to smooth the overall systems attainment of maximum entropy. In the case of markets, and networks as markets,

---

<sup>1</sup>Here, the classic example is Benard Convection. See (Swe89).

## 2.4 The Analysis of Non-Equilibrium States

---

these are institutions like Internet Corporation For Assigned Names and Numbers (ICANN), exchanges, legal systems, rules and best-of-practice traditions.

Institutions play very important roles in the smooth functioning of market-like structures. They:

- disseminate data
- route orders
- execute orders
- achieve efficient price discovery
- perform custodial and delivery functions
- and keep records

Perhaps, an appropriate analogy would to an operating system. Institutions solve two fundamental problems:

- Most problems in real markets are incomputable for either theoretical or practical reasons. Practically, if every market participant had to recalculate the fundamental value of every firm every morning from scratch, the stock markets could not function. The record keeping function of markets obviates this requirement. Theoretically, comparable information and records of previous days prices act as good starting points for the non-linear optimisation associated with the incorporation of new information into the valuation of a security.
- The key source of sub-optimality in systems is lack of coordination. In most markets, system wide shared operational schema emerge over time to resolve such sub-optimal outcomes. Such operational schema rarely operate without any infrastructure <sup>1</sup>. Good examples are the existence of clearing-houses for securities, operations for sharing ATM machines amongst banks, and ad hoc cooperative behaviour such as the LTCM bail-out of 1998 with the aid of the Federal Reserve.

---

<sup>1</sup>A good counterexample is the usage of common grazing grounds in certain parts of the United Kingdom, which are not over-grazed even in the absence of any explicit legal framework.

## 2.4 The Analysis of Non-Equilibrium States

---

Furthermore, institutional arrangements decentralise computations and information requirements <sup>1</sup>. In (Axt99), the authors formalise the benefits of decentralising computations across exchanges by showing that parallel distributed agent based models of k-lateral exchange can restore polynomial complexity in calculation of equilibrium prices and allocations.

Legal institutions, rights, government institutions can all be seen as constituting approximate solutions to problems: in that sense, they resemble “good” starting points for a non-linear optimiser, in the sense that the optimal solution is usually close-by and only small iterations are required to find it. Some of these institutions seem to have converged to some kind of stable point, such as market institutions, which simplify the computation burdens facing economic agents. In this sense, they resolve the computational issue faced by the hypothetical agent, by posing only easy and tractable calculations, as opposed to hard and intractable ones. The *raison d’être* for most market institutions is their computational efficiency. Once they are in place, agents need to exercise very little intelligence to make the right decisions.

Another advantage of institutions is that they act as carrier of knowledge capital through time: in that sense, they perform a *memoisation* function. If the problem has been solved before, the rules, regulations, and notion of best practice directs the agent to the previously found optimal solution. Structurally complex and efficient societies develop ways and means for transferring information and experience through time, through their schools of higher education, notions of best-practice, traditions and procedures.

All of this can be seen in the spirit of the Principle of Maximum Entropy Production. Brokers, arbitrageurs, and advertising media function reduce the likelihood of trading at disequilibrium prices by disseminating information faster. However, it should be noted that in real markets, as in the real world of thermodynamics, these devices expend real resources, establishing a practical tradeoff

---

<sup>1</sup> Hayek, 1945: “We cannot expect that this problem will be solved by first communicating all this information to a central board which, after integrating all knowledge, issues its orders . The problem is to show how a solution is produced by interaction of people each of whom has partial knowledge.” (see (Hay45))

## 2.4 The Analysis of Non-Equilibrium States

---

between entropy reduction and real resources available for other uses. This trade-off is explicitly ignored in the Walrasian framework (see Section 2.2.2.1).

The neoclassical paradigms reductionist agenda and agent-based approach makes it difficult to directly recognise the role of institutions as permanent features of markets essential to their continued evolution.

In the case of the Internet, there is tacit agreement that monopolistic practices by incumbents and firms which enjoy market power, and the lack of liberal telecommunication industries in developing countries contributes significantly to both the structure of interconnections and the digital divide between Western and emerging countries. For example, high Internet Access prices in developing countries is a function of not only local access monopolies, but also asymmetric relationships. Large Internet Providers refuse to sign peering agreements with their counterparts in developing countries due to anticompetitive reasons: local providers are in turn forced to bear the total cost of a links formation leading to high access prices and low usage, even though the entire network benefits through an increase in the number of users and content. Hence, there is an important and crucial role to be played by an authority which is able to monitor and penalise anticompetitive practices, leading to lower access prices and increased efficiencies.

### 2.4.2 The Role of Diversity

As pointed out in Section 1.6.1, markets can live with a certain amount of incompleteness: if the market place is composed of investors with varying degrees of risk aversion and not all risks are hedgeable, transactions will take place between risk-averse and risk-loving investors in contracts with undhedgeable contingencies. However, for this to happen, there must be significant diversity amongst agents in terms of their risk preferences and/or endowments.

The Principle of Maximum Entropy Production also suggests a role for diversity, as a means to explore alternative means of achieving higher states of entropy. In essence, diversity in a network increases the number of trajectories by which that system can attain maximum entropy, and may even enable even higher lev-

## 2.4 The Analysis of Non-Equilibrium States

---

els of entropy not achievable due to configuration restrictions in an otherwise un-diversified network ecology.

In the financial markets, the agents who step in when the market is incomplete are called speculators. They are diverse players with agenda's that are not necessarily coincident with the actual *raison d'être* of the market place. Yet, they play an important role. Most importantly, shallow markets with no speculators will not be able to deal efficiently with incompleteness.

Markets value derivatives through a replication argument: the value of the contract is equal to the value of a dynamically replicating portfolio of alternative assets which would end up with the same value at expiry. The value of such a contract has a dynamic relationship with the underlying asset which is related to the shifting probability associated with that contracts likelihood of having value on its very last day. In the financial markets, these shifting probabilities over time are referred to as the *gamma path* of the derivative contract. The replicating portfolio is required to follow the same *gamma path* in all future states of the world. However, the primary assets available for trading do not always span the contingencies of all contracts and perfect replication is not possible. In such cases, the market itself *completes* the transaction by providing a medium for the exchange of risks between asymmetrically risk-tolerant participants.

Under the neoclassical paradigm, Arrow (see (Arr53)) showed that, under some conditions, the ability to buy and sell securities can effectively make up for missing securities and complete the market. The existence of diverse speculators increases the transactional capability of a market. This is termed *Dynamic Completeness*.

### 2.4.3 Complex Adaptive systems

In a recent article in the New York Times<sup>1</sup> scientists demonstrated the step-by-step progression of how evolution creates a new molecular machinery by reusing and modifying existing parts. This showed how a progression of small changes could produce the intricate mechanisms found in living cells.

---

<sup>1</sup>“Study, in a First, Explains Evolution’s Molecular Advance”, The New York Times, April 7, 2006

## 2.4 The Analysis of Non-Equilibrium States

---

The study of complex adaptive systems begins with the realisation that physical dynamical systems are bound by the same limitation of computation as computing devices, notably the halting problem. The halting problem, encompassed in Church's and Gödel First Incompleteness Theorem imply that Turing machines can produce innumerable encodings.

Yet, such encodings exist, and the area of Complex Adaptive Systems studies them. Indeed, the existence of such encodings is held up to be a “hard” theory of evolution as in the introductory example.

The study of Complex Adaptive Systems is a burgeoning field, and is not central to this thesis. However, it justifies the ReReshaping argument in favour of analysing the behaviour of the market as a whole, as opposed to the behaviour of the constituent agents, so it is briefly covered. The central themes of the Complex Adaptive Systems argument are:

- *Innovation*: Innovation plays two roles in the Complex Adaptive Systems paradigm. First, it is nature's panacea to increasing entropy. No market can survive forever against the tyranny of increasing entropy, as without diversity, there would be no transactions and prices would not emerge. Systems develop structural complexity to combat entropy. The second role of Innovation is as an explanation for disequilibrium.
- *Self Organisation and Emergence*: Institutions are a form of self-organisation in a market. So is the emergence of money. The primary tenet of self-organisation is that “*Order is not pressure which is imposed on society from without, but an equilibrium which is set up from within.*”<sup>1</sup>
- *Decentralisation*: Complex Adaptive Systems exist to make calculations for agents easier. This happens through the emergence of decentralised algorithms which turn non-polynomial problems into polynomial ones. Axtell's algorithm in Section 2.4.1 is one such example.
- *Evolution*: The focus of evolutionary economics is long-run , as opposed to instantaneous change in response to marginal adjustments. The key

---

<sup>1</sup>Jos Ortega y Gasset (see (OyG27)).

## 2.4 The Analysis of Non-Equilibrium States

---

focus areas for evolutionary economics is variation and diversity, and the role these play in the health of the market. Disequilibrium can persist for arbitrarily long periods of time. For all these reasons, evolutionary economics are much closer to real-world economies. A fundamental set of results in evolutionary economics relate to complexity. The formation and emergence of complexity at all levels are explained as a result of evolution and the interactions between competing agents through *Innovation*.

As with transient order and other means by which a system increases its ability to achieve higher levels of entropy, Complex Adaptive Systems can also be analysed in the guiding light of the Principle of Maximum Entropy Production. In essence, Complex Adaptive Systems are transient organisations and functions which evolve to allow a system to explore a larger space of trajectories in more efficient ways.

## Chapter 3

# Dissecting Complexity

The main goal of this chapter is to measure *Network Complexity* in the spirit of *Effective Complexity*. This chapter presents two analytic models: they are too simple to capture every detail about networking, but rich enough to tell an interesting and meaningful story. The entropic justification of the previous chapter is used as the basis for selecting information theoretic measures as the basis for measuring complexity.

The *Architectural Model* is not meant to be implemented in practice on a any real network, but provides strong insights and is a useful tool for thinking qualitatively about architectural decisions: comparing OSPF with MPLS, considering the separation of routing from forwarding, and element design are some examples of when the *Architectural Model* can be used to systematically consider networking issues.

The *Measurement Model* is simple and tractable, and is designed to be implemented on a real network. It could feed the flow control in a router. It is forklifted from the world of actuarial sciences and the trading of collateralised debt obligations. This model makes strong assumptions: yet, it has the advantage that even with a parsimony of parameters, it is able to subsume many complications involving heterogeneity within its simple structure. The strength of its structural assumptions may seem like a limitation but in the context of measuring *Effective Complexity*, it turns out to be a preferable feature. The ultimate measure of complexity derived in this chapter, termed *Network Complexity*, measures the



### 3.1 Overview of Complexity Measures

---

complexity in a network in relation to this model. The strength of the hypothesis underlying the model makes this measure more reliable as a result.

One of the key goals of this chapter is to create a rich vernacular for thinking and communicating in the *Complexity Plane* of a network. The expressiveness of a language plays a key role in its success as a medium for communication. This chapter aims to create such a language for communicating the complexity and riskiness of a network.

The *Architectural Model* is qualitative, even though its output is in bits, proportional to the number of binary questions required to determine the state of the network. This is in keeping with the *Interrogation Costs* theme as a measure of complexity. In comparison, the *Measurement Model* is a quantitative model which could operate in a network in real-time.

This chapter only presents the models. The key insights and sensitivity analysis are in Chapter 4. The linkage between these two models and the entropic setting of Section 2.2.1 is discussed later in Section 4.3, once insights and empirical data illustrating applicability have been presented.

Before proceeding, *Implied Complexity*, already introduced as a concept in Chapter 1, is defined precisely. *Network Complexity* is related to this measure and it is a concept used throughout the Chapter.

**Definition 7 (Implied Complexity)** *The entropy implied by the histogram of actual network Reference Measurement data.*

## 3.1 Overview of Complexity Measures

Before proceeding with the introduction of a new measure, an overview of existing complexity measures for networks are presented. There are many good reviews in the literature (for example, see (Wec97), (Edm99a), and (DFT03)), so what follows is a brief overview.

### 3.1.1 General Complexity Measures

**Entropic Estimates** Using the classical definition of entropy  $H = \sum(-p_i \log(p_i))$  where  $\log_2$  implies bit-representation (see (CT91)) as a starting point:

**Approximate Entropy** Given groups of  $N$  points in a series, *Approximate Entropy* is related to the probability that two sequences which are similar for  $N$  points remain similar at the next point.

**Sample Entropy** *Approximate Entropy* corrected for self-matches<sup>1</sup>.

**Fourier Entropy** A power-spectral-density of a time series (or portion thereof) is computed, normalised to produce a probability-like distribution, and the *Sample Entropy* calculated as above.

**Wavelet Entropy** Related to the number of wavelet components needed to fit the signal.

**Renyi Entropy (the use of)** Compute the time-frequency spectrogram of a signal, then count the connected regions above some threshold energy. The idea here is that one peak in time-frequency space is an elementary event, so counting peaks gives an estimate of complexity.

**Chaos-based estimates of complexity** Using the mechanics of Chaos-based measures.

**Lyapunov Exponent** Measures the rate at which the system generates new information.

**Permutation Entropy** Estimate complexity as the entropy of the distribution of permutations of groups of time samples. Take each member of the group and give it a sequence number (in the series) from 1 to  $n$ . Reorder into ascending order, note the new order of the sequence numbers: the new order serves as a bin number into which the total count of all groups with a given sequence is accumulated. The result is a histogram of number occurrences of each sequence order. Normalise to a probability distribution and compute the entropy.

---

<sup>1</sup>In Chapter 5, this measure is used for threshold selection.

### 3.1 Overview of Complexity Measures

---

**Kolmogorov estimates (algorithmic complexity)** The grandfather of complexity measures, This is the minimum length of a Turing machine program needed to generate a pattern.

**Lempel-Ziv** Relate the length of the original sequence to the (generally shorter) length of the description when repeated sequences are correctly noted and enumerated.

**Hidden Markov** Estimate the minimal set of Markov states which can statistically reproduce the behaviour of a data set.

**Logical depth** The computational resources (chiefly time) taken to calculate the results of a program of minimal length, a combination of computational complexity and Kolmogorov complexity. Aimed at the complexity of the process and not the results.

**Thermodynamic Depth** The entropy of the ensemble of possible trajectories leading to the current state

**Effective Complexity** Length of a highly compressed description of its regularities, meant to distinguish between regularities and those features that are random or incidental.

**State Machine Models** Model as finite state-machines the system being measured, which is allowed to evolve. Complexity is measured as the ratio of total number of states to the total number of arrows connecting states. This is sort of an average index of branching in the system.

**Graph-theoretic Measures** These are divided into two: *Local metrics*, using information extracted from individual nodes, and *global metrics* which reflect properties of the entire network. Metrics can be distinguished in terms of the scope of the information they provide. *Local metrics* are extracted from individual nodes, whereas *global metrics* reflect the properties of an entire network. A single global value can be extracted from the local values by using simple statistics.

**Degree** Defined as the number of the connections of a single node for an undirected graph. It is a *local measure*.

---

### 3.1 Overview of Complexity Measures

**Clustering Coefficients** defined as the percentage of connection between neighbours of node  $i$ :

$$C_i = \frac{2E_i}{k(k-1)} \quad (3.1)$$

where  $k$  is the number of neighbours of node  $i$  and  $E_i$  is the existing connections between neighbours. It is a *local measure*. However, the averages for the network can be computed.

**Distance** Defined as the shortest hop distance between nodes. *Diameter* is a *global metric* defined as the maximum of minimum distances between any two nodes and *closeness* is a *local metric* which measures the connectedness of a network, defined as the average of the distances between a node  $i$  and all other nodes. *Closeness* indicates the centrality of a node. *Average path length* is a *global metric* defined as the average of the closeness values for all nodes. The *Hop-plot Exponent*,  $H$  is the number of node pairs within  $h$  hops, and the *Effective Hop Diameter* is defined as  $\left(\frac{N^2}{N+2E}\right)^{\frac{1}{H}}$ .

#### 3.1.2 Higher Order Methods

Any of the entropy methods above can be generalised further in the following way: break the data into parts, whether by times series or through an arbitrary partitioning of the network, compute the entropy of each part using one of the information or graph theoretic measures above, then treat the values as a time series and compute the entropy of that sequence. A good discussion of higher order methods can be found in (GP02).

Another alternative is to adapt the measures above to different time-scales. This can be done by analysing the data at different bandwidths, or by equivalently, averaging together groups of points. For example, (TG00) estimate *instantaneous complexity* using sliding windows of varying width.

#### 3.1.3 System-Specific Complexity Metrics

There are a significant number of proposals which rely on system-specific metrics to measure complexity. In (GM94), Murray Gell-Mann argues that complexity

### 3.1 Overview of Complexity Measures

---

is subjective and defines the notion of *Effective Complexity* to deal with this randomness:

It would take a great many different concepts - or quantities - to capture all of our notions of what is meant by complexity (or its opposite, simplicity.) However, the notion that corresponds most closely to what we mean by complexity in ordinary conversation and in most scientific discourse is “effective complexity.” In nontechnical language, we can define the effective complexity (EC) of an entity as the length of a highly compressed description of its regularities. For a more technical definition, we need a formal approach both to the notion of minimum description length and to the distinction between regularities and those features that are treated as random or incidental.

From “Effective Complexity”, *Murray Gell-Mann and Sety Lloyd*

This approach is used to a certain extent in the *Measurement Model* in Section 3.4. Similar approaches that rely on system-specific complexity metrics are reviewed below:

- In (ZHS99), the authors identify the timing constraint as the key deciding factor amongst architectural alternatives. They introduce the concept of states and propose to use state cycles to measure the probability of an entire system being feasible. They find that the state based probability is a good predictor of system timing behaviour. This has similarities with the proposal for *Entropic Routing* in Section 6.1.1.
- In (TKS<sup>+</sup>02), the authors tackle the problem of evaluating network processor architectures in the face of constantly evolving applications and technologies. The main principle is termed *benchmarking*, which they claim applies not only to network processors, but also to other domains as well. This has similarities to the concept of *Deviations from Simplicity* introduced in Section 1.5.2.

- In (CKT03), the authors present a real-time calculus for analysing system properties.
- In (TCGK02), the authors perform a search over the design space of the architecture of a packet processor, using models specific to packet processing tasks. They are able to measure the performance of network processors under different traffic patterns. While this is an interesting approach in the context of network processors and is in some ways similar to the approach in the *Architectural Model* in Section 3.3, a full blown implementation for an entire network is not realistic.
- Finally, in (AK02), the authors focus on the special conditions pertaining to multiple wireless connections, and look at certain cost functions and decision processes. Their approach is simulation based and considers various controller strategies.

#### 3.1.4 Discussion

The classic benchmark for all algorithmic complexity measures remains that of Kolmogorov's, the shortest program which can generate a code.

Generally, it has been shown that Kolmogorov's measure is not computable. Furthermore, it is now universally recognised that Kolmogorov's measure is really not a measure of complexity, but of randomness, as it is maximised by random codes.

In order to fix this shortcoming, various authors have proposed variations on the Kolmogorov theme. The first variation on Kolmogorov's measure is Bennett's *logical depth* measure, which is the running time of the shortest program and Lloyd and Pagels' *thermodynamic depth*, the entropy of all trajectories by which one can reach the current state from past states.

In reality, there is an overabundance of complexity measures in the complexity literature. Usually, the measures are not related to any specific domain, or are un-relatable in a practical way, and this is their main failing. In (GM94), the author points out that complexity is a subjective measure, otherwise one ends up measuring randomness: any practically useful measure needs to account for this.

### 3.2 The Indignity of Numerical Simulations

---

Murray Gell-Mann's *Effective Complexity* is also related to Kolmogorov Complexity, but is abetted with a model. This is an algorithm for relating data to a probability distribution, and hence can be interpreted subjectively. The approach in this thesis, termed *Deviations from Simplicity* follows the spirit of Gell-Mann's *Effective Complexity* with an entropic model. In that sense, it is a mixture of two different ways of looking at complexity, the algorithmic and entropic.

### 3.2 The Indignity of Numerical Simulations

... of my strongest stylistic prejudices in science is that many of the facts Nature confronts us with are so implausible given the simplicities of nonrelativistic quantum mechanics and statistical mechanics, that the mere demonstration of a reasonable mechanism leaves no doubt of the correct explanation. This is so especially if it also correctly predicts unexpected facts such as the correlation of the existence of moment with low density of states, the quenching of orbital moment for all d-level impurities as just described, and the reversed free-electron exchange polarisation which we shall soon discuss. Very often such, a simplified model throws more light on the real workings of nature than any number of ab initio calculations of individual situations, which even where correct often contain so much detail as to conceal rather than reveal reality. It can be a disadvantage rather than an advantage to be able to compute or to measure too accurately, since often what one measures or computes is irrelevant in terms of mechanism. After all, the perfect computation simply reproduces Nature, does not explain her.

Nobel Lecture, 1977, *Philip Anderson*.

A useful methodology for understanding complexity in a network must be analytic, general enough to be applicable to a broad range of architectural, operational and technical issues, yet granular enough to capture relevant specifics and be practically useful.

### 3.2 The Indignity of Numerical Simulations

---

This is not possible with one model: if it is too general, it cannot easily be applied to a real network if real results are required. At the same time, a model meant for real-time deployment is not necessarily useful for high-level architectural discussions. It is possible with two models and the spirit of universality can be preserved if they share the same heritage. That heritage is the *Entropic Formulation* of Section 2.2.1.

A complexity model must accounts:

- Interrogation/Deliberation costs , i.e. some measure of the number of questions one has to ask to deduce the state of the network
- The Riskiness of a Network.
- Domain Independence

and the model should be able to reflect what the network engineer knows about:

- The number of relevant relationships between the system's component
- Data-flow information, related to aggregation and disaggregation
- Operational complexity, i.e. the precedence and coupling of operations performed on data-flows
- The number of elements

These issues are now discussed at length.

**Interrogation Costs** The metric must be related to the size of the computer output needed to determine network state. For example, in the simplest case, the complexity measure must have some congruence with the number of expected (binary) interrogations to determine a certain fault in the network. The ideal dynamic complexity metric would be proportional to the time required by an agent, be it human or machine, to understand the cause of a change in the state of the network given a set of monitoring devices. However, this would be difficult to perform directly, so indirect methods are needed.



### 3.2 The Indignity of Numerical Simulations

---

**Risk Attribution** A good complexity model must also contribute to an understanding of the riskiness of a network. There are two things which one must understand about risk:

1. Is it systematic or idiosyncratic
2. What is the contribution of a specific element to the idiosyncratic and systematic risks in a network ?

Both of these goals are achieved with the *Measurement Model* and the notion of a *Slice*. A novel idea of this thesis is the attribution of risk to individual elements and the separation of idiosyncratic risks from systematic risk.

To understand why this matters, suppose one has two kinds of elements in the network

- Elements which share operations with many other parts of the network
- Elements which operate almost independently

Elements of the former type would contribute more to systematic risks and the latter to idiosyncratic risk. The *Measurement Model* allows this to be quantified at the level of the individual element.

**Domain Independence** Networks are varied and becoming more diverse. Domain specific metrics cannot stand the test of time. In another domain, if the underlying *Measurement Model* here were not deemed appropriate, one could just consider an alternative but still use the same template. Naturally, one would lose the mathematical machinery in this thesis specific to the *Measurement Model* but the analysis based on *Deviations from Simplicity* and the overall methodology could still be applied.

Nonetheless, the structural nature of the underlying model on which the *Effective Complexity* of the *Measurement Model* is based is relatively powerful yet parsimonious: “as simple as possible, but not simpler”<sup>1</sup>. It is possible that this model could be used with most systems which process data using some set of *Reference Measurements* quite easily.

---

<sup>1</sup>Albert Einstein.

## 3.3 Understanding Complexity

### 3.3.1 Why Another Measure of Complexity ?

There is nothing inherently wrong with the formulations of Section 3.1 for measuring complexity. However, they are not suitable for the overall goals of this thesis: the systematic analysis of architectural issues specific to networks, and practical measurement of complexity in a live network.

Networks generate substantial time-series data, so in theory, this data could be use to generate entropy measures. However, these measures would not have anything to say about the architecture of the network in terms meaningful to a network engineer:

- What is the impact of adding a new *type* of element ?
- What is the impact of combining layers ?
- What are the differences between connection-less and stateful protocols ?

Classical graph-theoretic measures are especially useless in this context: they say nothing about the precedence of operational requirements in the context of processing a specific flow through the graph. In that sense, they are “blind” to whether resources are shared and there is significant co-dependence on shared resources.

Finally, graph-theoretic are also blind to the characteristics of data-flow, especially in the context of multi-homed, multi-path flows. Flow state is an important source of complexity. They do not take into account the added complexity associated with prioritising, routing and shepherding the flow through the complex maze of the network. In short, graph theoretic measures ignore:

- The coupling between operations (not elements)
- Aggregation and disaggregation of data flows
- Sharing of resources

### 3.3.2 Why An Information Theoretic Formulation ?

The Information Theoretic Formulation  $\sum_i p_i \log p_i$  forms the mathematical basis for both defining theoretical architectural complexity, and also for practically inferring the realised complexity of networks in subsequent Sections. There are strong theoretical reasons for this choice, laid out in Chapter 2.

The Information Theoretic framework also satisfies some important and easy to work with conditions for a complexity metric:

1. The metric is continuous in probabilities associated with states.
2. It attains its highest value when everything is equally likely, i.e. equally random.
3. It is additive.

Axiomatically, the only functional representation of these conditions is the form given by Shannons Entropy (see (CT91)).

Finally, from a practical standpoint, the strictly concave nature of the Entropy function has two key benefits:

- The simple mechanics of dual optimisation can be brought to bear on optimisation problems with results guaranteed to be unique. The entropy function is just the maximum of the statistical entropy over the set of probability distributions compatible with the resource constraints on mean actions. Since entropy is a strictly concave function, when the constraints are linear, the maximisation function always has a unique interior solution. This is a very convenient result as it relates to the existence and uniqueness of equilibria.
- Jensen's Inequality can be used to reason about the sub-synergistic effects of variables<sup>1</sup>.

---

<sup>1</sup>*Sub-synergistic Effects*: Variables have sub-synergistic effects when the sum of their separate effects is greater than their joint effect.

### 3.3.2.1 An Axiomatic Definition of a Measure of Information

In order to illustrate the axiomatic nature of Shannon's measure, suppose:

- $I(p)$  is to be a measure of *Information* representing the physical occurrence of an event  $X = x$ , which a priori was assumed to have probability  $p$ .
- Further, posit that  $I(p)$  is always non-negative, as the occurrence of an event at the worst would provide no *Information*.
- Finally, a desired property of  $I(p)$  is that it decreases with increasing  $p$ , as there is no information provided by the occurrence of an event with 100% probability<sup>1</sup>

If  $x$  and  $y$  are independent, then:

$$P(X = x) = p \quad \text{and} \quad P(Y = y) = q \quad (3.2)$$

$$P(X = x; Y = y) = pq \quad (3.3)$$

and hence, when  $X = x$  and  $Y = y$  are independent:

$$I(pq) = I(p) + I(q) \quad (3.4)$$

Differentiating Equation 3.4 by  $p$  and  $q$ , and dividing one by the other, one gets:

$$\frac{q}{p} = \frac{\frac{\delta}{\delta p} I(p)}{\frac{\delta}{\delta q} I(q)} \quad (3.5)$$

and since  $q$  and  $p$  are independent, rearranging so that  $q \frac{\delta}{\delta q} I(q) = p \frac{\delta}{\delta p} I(p)$ , it is implied that:

$$I(q) = -c \ln(q) + A \quad (3.6)$$

Since an event that happens with certainty carries no information  $A = 0$ , and  $c$  must be positive due to the non-negativity and monotonicity of  $I(q)$ . Since  $c$  is just a scaling factor, this can be set to  $c = 1$ , and hence:

$$I(p) = -\ln(p) \quad (3.7)$$

---

<sup>1</sup>The sun not rising tomorrow morning carries much more information than otherwise.

$I(p)$  is sometimes also known as a *Surprisal*, as the smaller an event's probability, the bigger the surprise when it occurs.

Defining *Entropy* as the **Expected Information** from the occurrence of an event in its distribution, one gets:

$$H(X) = \sum_{i=1}^n p_i \log(p_i) \quad (3.8)$$

#### 3.3.2.2 Interpretation

Just because a quantity can be defined or measured on a system, it does not have to mean something, even if it is highly correlated with observable characteristics.

There are also semantic reasons for the *Entropic Formulation* which relate to the way in which its results can be interpreted:

- Entropy is a measure of *Interrogation costs*: log base 2, it can be interpreted as a lower bound on the number of binary questions required to determine state.
- Entropy is a powerful statistical inference rule: maximally informative statistics minimise the Kullback-Leibler distance between posterior distributions. This coincides with an underlying microstructure which assumes the use of statistical inference by agents facing information costs.

In the models in this chapter, Information Theory is utilised in two ways:

- Architecturally, the mathematical formulation of entropy is used as a proxy for *Interrogation Costs*: this contrasts with the primary interpretation of Entropy in Chapter 2 as a measure of the most likely configuration for a network.
- Practically, the notion of *Relative Entropy* (see Section B.1) is used as a complexity measure by metricising the distance between the predictions a network model makes about the behaviour of a network, and actual, empirically observed behaviours.

Applying the ideas in information theory to a network allows the right questions to be asked, so long as there is agreement on state definitions. Furthermore, it enables a quantitative comparison of paradigms and operational practices, so that comparisons can be made.

#### 3.3.3 The Architectural Model

The goal of the *Architectural Model* is to cast decisions about traffic flow variation, element design (what to put in and what to take out of a network), and protocol choice, in a single frame of reference.

The *Architectural Model* cannot be used in a practical network: its data requirements are unrealistic. However, it can act as a sanity check for qualitative thinking. It needs to enable quick, consistent, concise thinking about architectural issues, at all levels of networking analysis: its ability to succinctly capture and analyse such decisions is its litmus test.

The entropy of a network can be calculated by attributing probabilities to states. As an outline, one would, for each element  $j$ :

1. Measure variations across flows
2. Classify data according to similar levels of variation
3. Calculate the probabilities for similar states
4. Obtain a complexity index by calculating the entropy of the probability distribution

This is the starting point for a complexity metric based on Information Theory. In the next section, this model is successively enhanced to cover differences between serial and parallel networks, element and operational variability, and scheduling constraints.

The vernacular used in the following sections is as follows:

- Networks are composed of *Elements*, which perform *Operations*, on *Traffic Flows*.
- There can be multiple types of each, and they can vary over time.

- A *Reference Measurement* is an observation on *Elements* which network engineers care about. They are not continuous, but discrete. In the case of the simplest model below in Section 3.3.3.1, they can have more than two values. In the context of Section 3.4, they are a Bernoulli-(0,1).

#### 3.3.3.1 Simple Version

Formalising this with a simple model to begin with, let the complexity of the network be measured by:

$$H_{\text{Network}} = - \sum_{i=1}^m \sum_{j=1}^Q p_{ij} \log_2 p_{ij} \quad (3.9)$$

where:

- $m$  is the number of elements in the network, and
- each element can take one of a number of states in set  $Q$  with probability  $p_{ij}$ .

States can be defined as a range over levels of utilisation:

- Over-utilised: 75-100% Utilisation
- Normal: 25-75% Utilisation
- Underutilised: 0-25%

Using Jensen's Inequality (see (CT91)), it is easy to show that complexity can be reduced by:

- decreasing the number of elements
- decreasing the the number of states
- increasing the variability of states.

However, this simple metric has significant shortcomings:

- There is no mix of traffic: Only one kind of flow is serviced.

- Dependence between states is not accounted for.
- Linkage between elements is not accounted for.
- Precedence requirements between operations is not accounted for.
- Each operation involves only one element.

The next section complicates this simple model to address these shortcomings.

#### 3.3.3.2 Complicated Version

In the context of a network, infrastructure complexity is a result of the following factors:

1. Traffic mix
2. Variety in the infrastructure of the network, like routers, switches, firewalls, etc.
3. Interaction requirements (e.g. between the control plane, the signalling plane, the data plane, routers, etc.)
4. Sharing of infrastructure, and
5. Operational complexity, meaning precedence requirements between operations.

Any complexity measure would need to capture changes in any of these factors. The conditions required for such a measure are that it:

- Increase with traffic mix
- Increase with operational complexity
- Increase with scheduling flexibility
- Increase with infrastructure sharing



A more complicated version of the model in Section 3.3.3.1 can satisfy these requirements. The framework below is also applicable across any domain which involves transfers: in a network, traffic could be any one of IP, ATM, or Optical traffic. In a router, it could be packets. While *Elements* (see below) in a network consist of routers, switches, add-drop multiplexers, they might be an instruction cache, a buffer, a bus, and so on in a processor.

**Traffic Mixture, Operational and Element Complexity** Allow for a mixture of *Traffic Flows* represented in a *Traffic Vector*  $v$ :

$$\nu_l = \frac{w_l}{\sum_{j=1}^b w_j} \quad (3.10)$$

where  $w_j$  is a measure of the quantity of flow  $j$ .

*Operations* (DNS Lookup, Routing Lookup, etc) are performed by *Elements* (routes, switches, fibre, etc.) on *Traffic Flows* (either by destination, source or type) where:

- $b$  types of *Traffic*, indexed by  $l$ , on which
- $m$  *Elements*, indexed by  $z$ , perform
- any of  $q$  *Operations*, indexed by  $u$ , and  $v$ .

The operational requirements for each *Traffic* type  $l$  are summarised in an *Operational Matrix*  $\Upsilon$ :

$$\Upsilon_l = \begin{bmatrix} v_{11} & \cdots & v_{1m} \\ \vdots & \ddots & \vdots \\ v_{q1} & \cdots & v_{qm} \end{bmatrix} \quad (3.11)$$

where  $v_{ij}^l$  is 1 if *Element*  $z$  performs *Operation*  $u$  or  $v$  on *Traffic* flow  $l$ . This matrix has *Operations* going top to bottom and *Elements* left to right.

**Scheduling Constraints** *Operations* do not happen randomly, but in certain order. A DNS lookup is followed by a routing table prefix match followed by a forwarding operation. The information in the *Operational Matrix* is used to define the *Network Precedence Matrix*  $\Lambda$ :

$$\Lambda_l = \lambda_{uvzl}, \forall u, v, z, l \quad (3.12)$$

such that, if *Operation*  $u$  precedes  $v$  for *Traffic*  $l$  on *Element*  $z$ , then  $\lambda_{uvzl} = 1$ , else  $\lambda_{uvzl} = 0$ . If there is no precedence relationship, then  $\lambda_{uvzl} = \lambda_{vuzl} = 1$ .

In order to cross this information with the *Traffic Vector*, for each type of *Traffic*  $l$ , the *Traffic Precedence Matrix* is defined:

$$\hat{\lambda}_{uvz} = \sum_{l=1}^b \lambda_{uvzl} \quad (3.13)$$

Define:

$$\psi_{uvzl} = \nu_l \hat{\lambda}_{uvz} \quad (3.14)$$

Normalise:

$$\tilde{\psi}_{uvzl} = \frac{\psi_{uvzl}}{\sum_{u=1}^q \sum_{v=1}^q \sum_{z=1}^m \sum_{l=1}^b \psi_{uvzl}} \quad (3.15)$$

where  $\sum_{u=1}^q \sum_{v=1}^q \sum_{z=1}^m \sum_{l=1}^b \tilde{\psi}_{uvzl} = 1$ .

$\tilde{\psi}$  can be interpreted as a probability in the frequentist sense. It represents the likelihood of observing the *Operational* sequence  $u \rightarrow v$  on *Element*  $z$  being performed on *Traffic Flow*  $l$ . Likewise,  $\hat{\lambda}$  represents the number of times a particular *Operational* sequence on an *Element* is required over all *Flows*: it measures the similarity between sequences of *Operations* on *Flows*, and the sharing of resources to carry out those *Operations*.  $\Upsilon$  represents the complexity arising not from the sequence of *Operations*, but the way in which they are spread out amongst *Elements*.

*Architectural Complexity* of the model is defined using these tools:

**Definition 8 (Architectural Complexity)**

$$H = - \sum_{u=1}^q \sum_{v=1}^q \sum_{z=1}^m \sum_{l=1}^b \tilde{\psi}_{uvzl} \log \tilde{\psi}_{uvzl} \quad (3.16)$$

#### 3.3.4 Breaking Down Architectural Complexity

The intermediate steps leading to the *Architectural Complexity* metric admit meaningful interpretations.  $\Upsilon$  captures both the *Operational* and the *Element Complexity* of the Network:

**Definition 9 (Operational Complexity)** *Construct a column vector by summing across each row. Calculate Operational Complexity by interpreting the sum as a frequency of operations by elements, and normalise to a probability distribution by dividing each element of the column vector by the sum of the vector. Calculate entropy on the resulting probability distribution.*

**Definition 10 (Element Complexity)** *Same as in Operational Complexity, but sum across columns.*

$\hat{\lambda}$  captures the complexity associated with scheduling *Operations* on individual *Elements*. Using it, one can define:

**Definition 11 (Scheduling Complexity)** *Construct a vector of operational constraints on each element:*

$$\pi_z = \sum_{uv} \hat{\lambda}_{uvz} \quad (3.17)$$

*Normalise this vector (in a frequentist sense):*

$$\tilde{\pi}_z = \frac{\pi_z}{\sum_z \pi_z} \quad (3.18)$$

*Calculate its entropy:*

$$H(\tilde{\pi}) = - \sum \tilde{\pi} \log(\tilde{\pi}) \quad (3.19)$$

Finally, it is relatively straightforward to define *Traffic Complexity* as the entropy of  $\nu$ :

**Definition 12 (Traffic Complexity)** *Define Traffic Complexity as:*

$$H(\nu) = - \sum \nu \log(\nu) \quad (3.20)$$

An increase in any of these metrics would lead to an overall increase in *Architectural Complexity*.

### 3.3.5 Concluding the Architectural Model

It is relatively easy to reason about the *Architectural Model*, an important design goal for any complexity metric. When confronted with any architectural issue, one would need to consider the following:

- The impact on the number of columns in the *Operational Matrix*  $\Upsilon$ , which is related to the *Element Complexity* of a network.
- The impact on the number of rows in  $\Upsilon$ , which is related to the *Operational Complexity* of a network.
- The impact on the size and variability of the *Traffic Vector*  $\nu$  which is related directly to overall *Architectural Complexity*.
- The impact on the number of zeros in the *Network Precedence Matrix*  $\Lambda$  which is related to the *Scheduling Complexity* of a network.

An increase in any of these sub-components will lead to an increase in the overall *Architectural Complexity* of the network. This makes it easy to pin down how and why the complexity of a network will increase due to an architectural decision.

Furthermore, the impact of this additional complexity on the performance and riskiness of a network can also be considered, once one knows what happens to these parameters when the complexity of a network increase. This analysis is in Chapter 4. An empirical evaluation is in Chapter 5.

## 3.4 Measuring Complexity

The *Architectural Model* is not a true measure of *Network Complexity*. While it is a description of its static complexity arising out of the way the elements are wired together, it says nothing of the emergent properties of the network and how it's complexity relates to an engineers understanding of how a network should behave.

A simple Sample Entropy analysis of the time series behaviour is also not a measure of a networks complexity. Both approaches have an important short-coming: they assume implicitly that the network engineer knows nothing about networks. Specifically, they assume the engineer does not have a model of network behaviour<sup>1</sup>, such as the *Architectural Model* here.

Hence, these are not *Effective* measures of *Network Complexity*. If a model exists for a system which can explain a significant source of the variability in its behaviour with a few parameters, that part of the variability which is explained should not be deemed to be complex: by carrying out sensitivity analysis on the model inputs, and comparing them to observed behaviour, the engineer can determine state without incurring any interrogation costs. Complexity is that behaviour of the network which cannot be explained by a simple, structural model, i.e. which incurs *Interrogation Costs*. This is the *Deviations from Simplicity* principle in action, and the essence of Murray Gell-Manns' notion of *Effective Complexity*.

The *Architectural Model* is also not a practical model. It is unrealistic to assume that any network traffic engineer could build the adjacency matrices in Section 3.3.3.2 for any network of significant size. Also, most practitioners deal with networks that have a significant legacy component to them: the data needed to populate the *Network Precedence* and *Operational Matrices* is typically not available. Even if SDL-like<sup>2</sup> semantics were to be considered, mapping shared resources into the matrix would be daunting.

What are easily available are measurements on elements, termed *Reference Measurements*:

**Definition 13 (Reference Measurements)** *Measurements which proxy the complexity of the network and are common to all elements. These measurements*

---

<sup>1</sup>Arguably, the same criticism can be levelled at the bulk of the “Power-Law” literature which tries to explain the incidence of high kurtosis and skewness in observed network traffic data solely through the existence of queues.

<sup>2</sup>Specification and Description Language (SDL) is a specification language targeted at the unambiguous specification and description of the behaviour of reactive and distributed systems. It is defined by the ITU-T (Recommendation Z.100.) Originally focused on telecommunication systems, its current areas of application include process control and real-time applications in general.

*should be easy to calculate and not specific to individual elements or technologies. Reference Measurements are Bernoulli-(0,1) and triggered when Capacity, determined as a threshold over some characteristic of the Element, is exceeded in usage.*

It would be much more practical and realistic if a model could be built which made use of these empirical measurements. This is the approach taken by the *Measurement Model*.

Some candidate *Reference Measurements* are:

- Error rates, such as dropped packets
- Average Inbound/Outbound Traffic
- Throughput
- Utilisation

Using just this data, and a simple, tractable model of the network, it is possible to extract true *Network Complexity* and parametric Risk Metrics.

There is another benefit to focusing the analysis of a networks complexity and risks to common *Reference Measurements*: it makes it much easier to author financial contracts based on the values they realise. This is a topic explored further in Chapter 6.

The goals of a *Measurement Model* are twofold:

1. Explain network behaviour due to static complexity arising from *Operational, Element, and Scheduling Complexity* in a simple, easy to measure way.
2. Characterise risk.

Even though the *Measurement Model* aims to be simple, it is easily extended in many directions while remaining tractable. The key mathematical features are:

- A central systematic factor is modelled which is meant to act as a proxy for the inherent coupling of the network, essentially acting as a proxy for the *Network Precedence Matrix* of Section 3.3

- Very simplistic assumptions are made about the structure of *Operations* and *Elements* as a starting point.

In section 3.3, the primary tools were from the field of manufacturing. In this section, the primary tools are from finance and the actuarial sciences.

### 3.4.1 The Measurement Model

**Definition 14 (Measurement Model)** *A parsimonious model which uses Reference Measurements, describes as much of the networks behaviour as possible, and acts as an easy to use proxy for the Architectural Model.*

The right choice of model is arbitrary, but Occam’s Razor<sup>1</sup> narrows the choices. This metric is built on the simplest possible model of a network:

- Networks are composed of homogeneous elements
- The level of coupling between elements is homogeneous
- The dynamics of each element has just two components:
  - An idiosyncratic component
  - A single systematic factor

Once again, reusing some of the same variable definitions as in Section 3.3.3.2:

- assume the network is composed of  $m$  *elements*. Elements could be nodes or edges. *Capacity*<sup>2</sup>, as it relates to the variable being measured, is fixed for each element.
- let  $M$  be a factor common to all *Elements* which drives usage in an element due to systematic factors.
- let  $X_i$  be a factor specific to *Element*  $i$  which drives usage due to idiosyncratic factors not common to other *Elements*

---

<sup>1</sup>In Latin, “entia non sunt multiplicanda praeter necessitatem,” which translates to “entities should not be multiplied beyond necessity.”

<sup>2</sup>Capacity could mean bandwidth for a link, CPU speed for a firewall, or SRAM memory. The model is applicable in many domains.

- let  $C_i$  be the capacity of *Element*  $i$
- let  $a_i$  be a measure of the coupling between *Element*  $i$  and the common factor  $M$ .

The goal of the model is to characterise the multivariate distribution of the network. Each element is modelled by a 2-state Bernoulli-(0,1), which takes the value 1 if some utilisation threshold is cleared: this is called a *trigger*, and does not have to be 100% <sup>1</sup>

Utilisation of the the element, which would be Link Utilisation in the case of edges and CPU utilisation or buffer space in the case of nodes, are composed of two parts, one part idiosyncratic and one part systematic:

$$A_i = a_i M + \sqrt{1 - a_i^2} X_i \quad (3.21)$$

where  $M, X_i, i = 1, \dots, m$  are independent standard normally distributed variables.

A measurement is triggered when utilisation of an element  $i$  exceeds a certain trigger level,  $C_i$ , implied by the element's trigger probability  $p_i$ :

$$p_i = P[A_i \leq C_i] = \Phi(C_i) \quad (3.22)$$

where  $\Phi$  is the Cumulative Distribution function for the Standard Normal Distribution. Hence,

$$C_i = \Phi^{-1}(p_i) \quad (3.23)$$

This means that the probability that the  $i^{th}$  element triggers conditional on the factor  $M$  is:

$$p_i(M) = \Phi\left(\frac{C_i - a_i M}{\sqrt{1 - a_i^2}}\right) \quad (3.24)$$

In the actuarial literature, this is known as “conditional independence”.

$a_i$  is the sensitivity of element utilisation to some common factor, common to all elements in the network. In the simplest specification of the mode,  $a_i = a$  for all elements. This assumption will be relaxed in Section 3.4.4.2.

---

<sup>1</sup>Network engineers typically set alarms at significantly lower thresholds, usually around 10%.



All elements are homogeneous, i.e. they have the same probability of triggering a *Reference Measurement* which also means they have the same capacity (this assumption will be relaxed in Section 3.4.4.1). Under this assumption, the probability of a trigger is:

$$p(M) = \Phi\left(\frac{C - aM}{\sqrt{1 - a^2}}\right) \quad (3.25)$$

Under the homogeneity assumption, the multivariate trigger probability  $L_k = \frac{k}{m}N$  is equal to the probability that exactly  $k$  out of  $m$  elements trigger:

$$P[L = L_k | M] = \binom{m}{k} \Phi\left(\frac{C - aM}{\sqrt{1 - a^2}}\right)^k \left(1 - \Phi\left(\frac{C - aM}{\sqrt{1 - a^2}}\right)\right)^{m-k} \quad (3.26)$$

This is the conditional distribution, conditional on the systematic factor  $M$  being realised at a certain value,  $u$ . To get the unconditional distribution, integrate over all possible values of  $M$ :

$$P[L = L_k] = \binom{m}{k} \int_{-\infty}^{\infty} \Phi\left(\frac{C - au}{\sqrt{1 - a^2}}\right)^k \left(1 - \Phi\left(\frac{C - au}{\sqrt{1 - a^2}}\right)\right)^{m-k} d\Phi(u) \quad (3.27)$$

This is a costly computation for large  $m$ . A common portfolio limit approximation from the insurance industry due to Vasicek significantly reduces the computational burden: see (RO04). Consider the cumulative probability of a percentage of the network triggering a measurement not exceeding  $\theta$ :

$$F_m(\theta) = \sum_{k=1}^{[m\theta]} P[L = l_k].$$

Substituting  $s = \Phi\left(\frac{C - au}{\sqrt{1 - a^2}}\right)$ , one gets:

$$F_m(\theta) = \sum_{k=1}^{[m\theta]} \binom{m}{k} \int_0^1 s^k (1 - s)^{m-k} d\Phi\left(\frac{\sqrt{1 - a^2}\Phi^{-1}(s) - C}{a}\right) \quad (3.28)$$

Using:

$$\lim_{m \rightarrow \infty} \sum_{k=1}^{[m\theta]} \binom{m}{k} s^k (1 - s)^{m-k} = \begin{cases} 0, & \text{if } \theta < s \\ 1, & \text{if } \theta > s \end{cases}$$

the Cumulative Distribution Function (CDF) of trigger probabilities for a large network becomes

$$F_{\infty}(\theta) = \Phi\left(\frac{\sqrt{1 - a^2}\Phi^{-1}(\theta) - C}{a}\right) \quad (3.29)$$

The advantage of this specification is parsimony: there are only two parameters.

- Average trigger probability
- Coupling

These two parameters and the model give a complete multivariate distribution for trigger probabilities in the network.

### 3.4.2 Measuring Complexity

With Equation 3.29, it is possible to characterise certain aspects of the multivariate distribution of a *Reference Measurement* in relation to a network. There is no *Effective Complexity* in that characterisation (especially in the context of this simple model). What is complex is what is not explained by this model.

*Network Complexity* is meant to measure what is not explained by the model. This is the essence of using Murray Gell-Mann's *Effective Complexity* notion in the domain of networks. Hence, it is defined as that part of a network's *Implied Complexity* which deviates from the predictions of the model the engineer has built of the network. Ideally, the engineer would build the *Architectural Model* of Section 3.3, however, that is not practical. Hence, Equation 3.29 is used as a proxy for the engineers understanding of the network's architecture, and *Network Complexity* is quantified as the KL distance between that and the empirically observed complexity of the network called *Implied Complexity*.

Quantitatively:

**Definition 15 (Network Complexity)** *The entropic distance (in practice, the Kullback-Leibler distance) between the empirical behaviour of the network (proxied through Reference Measurements) and the behaviour implied by a simple model of the network.*

In practice, the steps to calculate it are:

1. Calculate the empirical distribution of the *Reference Measurement* across the entire network.
2. Calibrate the average trigger probability as the empirically observed average trigger probability of all elements

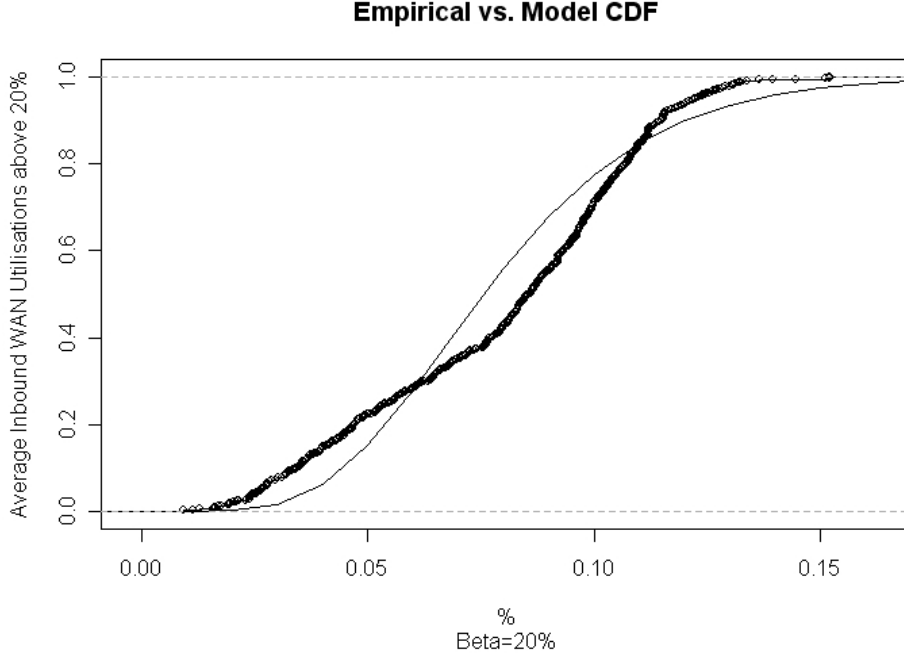


Figure 3.1: Empirical vs. Model CDF

3. Calibrate the Coupling parameter such that the Kullback-Leibler distance between the actual multivariate distribution and the predicted multivariate distribution is at a minimum
4. *Network Complexity* of the network is the remaining entropy distance between the model and the empirical distribution.

In Figure 3.1 the empirical CDF of Average Inbound WAN statistics are graphed against the CDF implied by Equation 3.29: due to simplicity of the model specification and the strength of the underlying hypothesis, the architect will have significant confidence in aspects of the network explained by the model.

In theory, there is no reason why this methodology cannot be applied to any system which generates multivariate distributions over *Reference Measurements* and where the structural model above can be reused as a meaningful approximation. some candidates are processors, routers, road traffic, a manufacturing facility, etc.

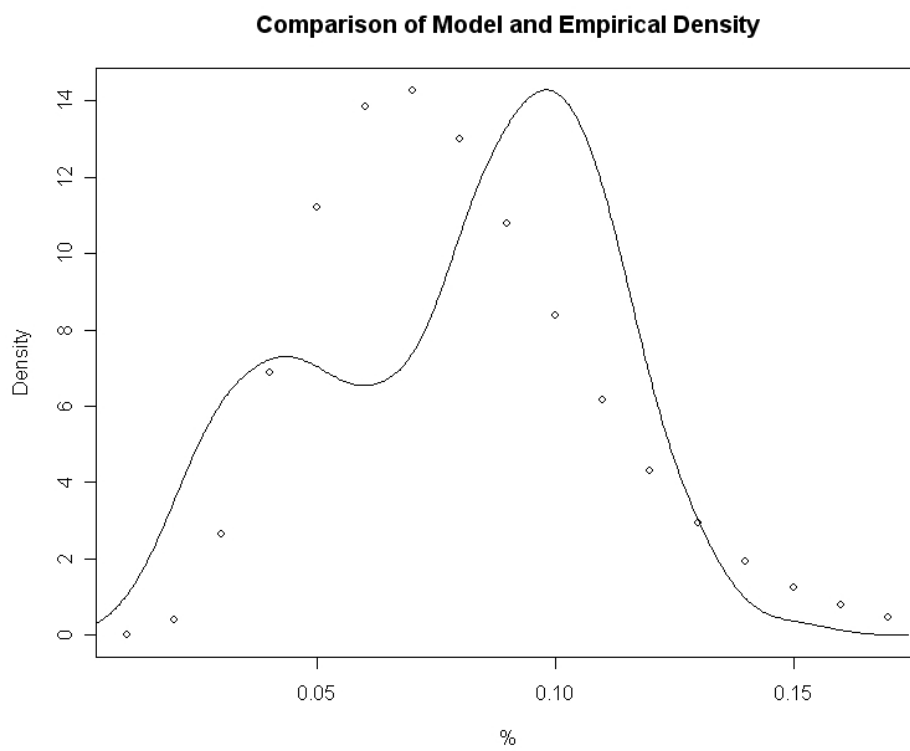


Figure 3.2: Empirical vs. Model Density

### 3.4.3 Measuring Risk

So far, the *Measurement Model* only acts as a proxy for the *Architectural Model*. However, the *Measurement Model* is also useful for analysing risk in the context of a network.

The *Measurement Model* makes it easy to calculate two different kinds of risks:

**Definition 16 (Idiosyncratic Risk)** *The risks due to the unique characteristics of individual elements, as measured via  $X_i$  in Equation 3.21<sup>1</sup>.*

**Definition 17 (Systematic Risk)** *Risks inherent to the entire market, due to  $M$  in Equation 3.21<sup>2</sup>.*

#### 3.4.3.1 A Slice of Risk

Risk is measured through *Slice*, defined in Chapter 1. By varying the thresholds  $K_1$ , and  $K_2$ , *Slice* is able to characterise both systematic and idiosyncratic risks. Furthermore, these risks can be graphically illustrated by graphing *Slice* with its parameters on the sub-axis.

Mathematically:

$$\text{Slice}_{(K_1, K_2)} = \frac{1}{K_2 - K_1} \left( \int_{K_1}^1 (x - K_1) dF(x) - \int_{K_2}^1 (x - K_2) dF(x) \right) \quad (3.30)$$

In the context of the homogeneous approximation above, this integral evaluates quite simply to:

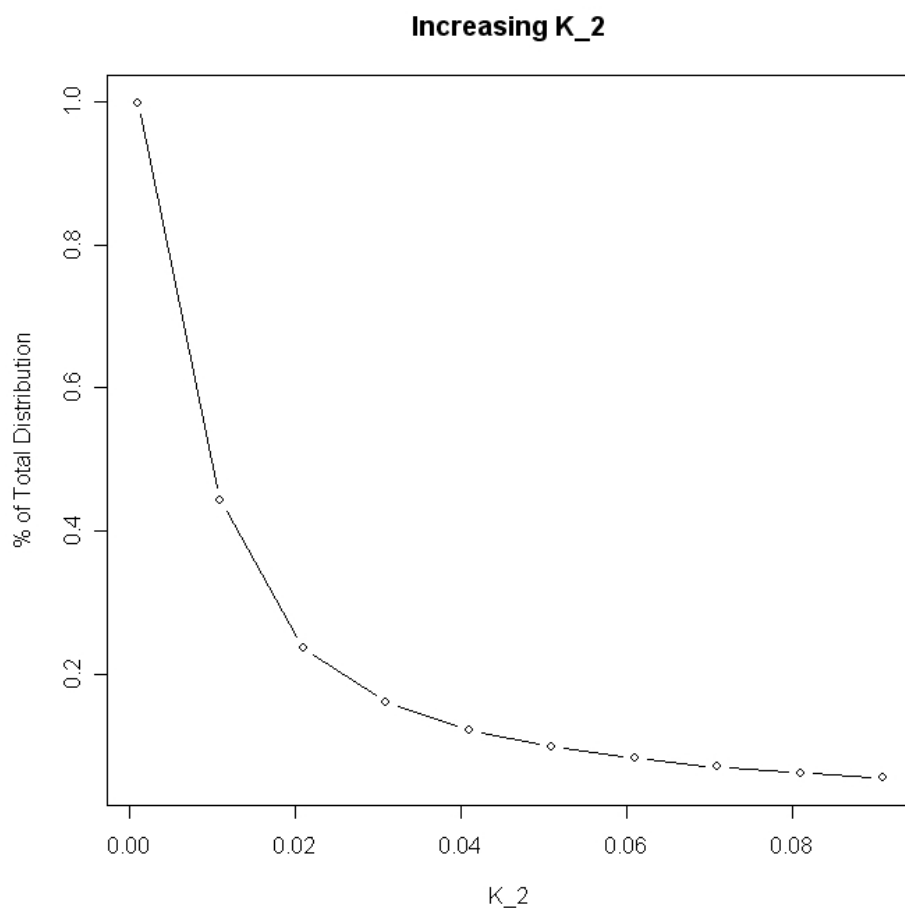
$$\text{Slice}_{(K_1, K_2)} = \frac{\Phi_2(-\Phi^{-1}(K_1), C, -\sqrt{1-a^2}) - \Phi_2(-\Phi^{-1}(K_2), C, -\sqrt{1-a^2})}{K_2 - K_1} \quad (3.31)$$

where  $\Phi_2$  is the bivariate normal distribution.

---

<sup>1</sup>In the financial markets, this is also known as the the risk of price change due to the unique circumstances of a specific security, as opposed to the overall market. This risk can be virtually eliminated from a portfolio through diversification. also called unsystematic risk. Hence, the financial markets do not pay an investor a premium for taking this risk, as it can be hedged for free using diversification. This is a key tenet of the Capital Asset Pricing Model.

<sup>2</sup>In the financial markets: "The risk inherent to the entire market or entire market segment. Also known as "un-diversifiable risk" or "market risk."

Figure 3.3: Effect of Increasing  $K_2$ 

The parametric nature of *Slice* makes it possible to define and measure systematic and idiosyncratic risks in a network. Also, *Slice* permits a fine grain picture of the complexity of the network, makes it easy to talk about the impact of decisions on the network complexity, and is the key adjective in a vernacular meant to make thinking about complexity, discussed in Section 3.5.3.

The relationship between *Slice* and idiosyncratic and systematic risks is:

- When  $K_1$  is low, *Slice* measures idiosyncratic risk
- When  $K_1$  is high, *Slice* measures systematic risk

### 3.4.3.2 Sensitivity Analysis

The riskiness of a metric measured by *slice* depend on a number of factors:

**Lower barrier,  $K_1$**  The higher this is, the more faults have to take place before one enters the window, lowering *Slice*.

**Higher barrier,  $K_2$**  The wider the width of the window, the more unlikely faults it will cover, hence *Slice* will fall.

**Marginal Probabilities of Triggering** Monotonically increasing.

**Time Horizon** Depends on the time dependence of marginal trigger probabilities. Assuming the likelihood of a trigger increases over time, this means that a longer time horizon for this metric will mean a higher value of *Slice*.

**Trigger Correlations** If trigger correlations are high, elements tend to fault together and this increases *Slice* measurements with higher attachment points. However, elements tend to survive together decreasing the values of *Slice* metrics with lower attachment points.

### 3.4.4 Extensions

*Effective Complexity*, the bedrock of *Network Complexity*, is based on the difference between the actual behaviour and the predicted behaviour of the network: it is that part of the networks behaviour which cannot be explained. Later, the *Implied Complexity* of a network will be a term used to represent the actual behaviour of the network, which leads to the following qualitative relationship:

$$\text{Effective Network Complexity} = \text{Implied Complexity} - \text{Architectural Complexity} \quad (3.32)$$

This relationship is quantified through the KL distance between the two measures.

This means that there are two contributors to *Network Complexity*:

**Architectural Complexity** The architecture of the network and traffic variability causes complexity by making it more difficult to determine absolute state.

**Model Simplicity** A simple model can not explain the total behaviour of a model.

This means that, in addition to architectural considerations, the engineer also needs to consider modelling. By increasing the complexity of his model, he can reduce *Effective Network Complexity*.

In this spirit, extensions to the assumptions of the underlying model are investigated. The main assumptions in the *Measurement model* are good candidates for extensions:

- All elements are homogeneous: they have the same trigger probability.
- The coupling between elements is homogeneous.
- The systematic factor follows a Normal Distribution.

Extensions to the model also introduce the notion of *Model Complexity*, defined in Section 3.4.4.7.

#### 3.4.4.1 Heterogeneity of Elements

Extending the model to cover heterogeneous elements is relatively straightforward using a recursive model, even though it could be an unrealistic requirement in real life: the trigger probability for each element would need to be known. One can imagine a number of ways this could be achieved either though none are really satisfactory. For example, trigger probabilities could be related to the number of links for each node.

Nonetheless, this is a useful exercise because it reveals some interesting insights (Chapter 4). Consider a network composed of  $N$  elements. Build the multivariate distribution of outcomes for all elements using recursion. The distribution for just one asset is:

$$\begin{aligned}p^1(0) &= 1 - p_1 \\ p^1(1) &= p_1\end{aligned}$$



Adding another element gives:

$$\begin{aligned} p^2(0) &= (1 - p_1)(1 - p_2) \\ p^2(1) &= p_1(1 - p_2) + p_2(1 - p_1) = p^1(0)p_2 + p^1(1 - p_2) \\ p^2(2) &= p_1p_2 = p^1p_2 \end{aligned}$$

Generalising:

$$p^{i+1} = \prod_{j=1}^{i+1} (1 - p_j) \quad (3.33)$$

$$p^{i+1}(k) = p^{i1}(k - 1)p_{i+1} + p^i(k)(1 - p_{i+1}), \text{ for } 0 < k < i + 1 \quad (3.34)$$

This gives the unconditional multivariate distribution of the network. To evaluate *Slice*, one would need the conditional distribution: substitute  $f_j(p_j, x)$  for  $p_j$  above, the trigger probability for an element given  $x$  for the systematic factor from Section 3.24. To calculate the *Slice* metric, one would have to numerically integrate the conditional multivariate distribution using Equation 3.30.

To calculate *Network Complexity*, one would find the nearest (in entropy terms) coupling parameter for the unconditional distribution, and calculate the remaining distance.

***Slice Sensitivity*** A benefit of this recursion formula is that it allows for the calculation of the sensitivity of both *Network Complexity* and *Slice* to a single element. This allows for risk attribution to individual elements in a network. This is done simply by iterating the recursion formula backwards for that one element in Equation 3.33, and recalculating the values for *Network Complexity* and *Slice*.

#### 3.4.4.2 Heterogeneity of Correlations

Another assumption of the *Measurement Model* is that the coupling between all elements is the same. This is equivalent to assuming a flat correlation matrix when evaluating the multivariate distribution of a vector of Gaussian variables. It is relatively simple to assume the correlation matrix of coupling between elements is heterogeneous.

There are a number of ways this could be done:

- Build the correlation matrix from pair-wise correlations, using some kind of statistical averaging, or
- Build the correlation matrix based on *factors*:
  - each element is grouped along *factors*; say sub-systems, or groups of elements which are tightly coupled
  - an assumption about correlation between each element and the group is made
  - a further assumption about correlations between groups is made
  - a *factor* correlation matrix is generated

Once one has a correlation matrix, one now has to resort to the *Indignity of Numerical Simulations*:

1. Using Cholesky Decomposition of the correlation matrix, correlated random variables is generated for each element
2. A further random variable, signifying a realisation of the systematic factor would be generated.
3. Using Section 3.23 and Equation 3.21, one would then generate a set of Bernoulli-(0,1) samples

These steps can then be repeated as many times as necessary to generate a Monte-Carlo simulation of the networks behaviour.

#### 3.4.4.3 Heavy-tailed Distributions

An obvious extension is to challenge the distributions of the idiosyncratic and the systematic component of Equation 3.21, both assumed to be  $N(0,1)$ . This results in:

$$F_{\infty}(\theta) = T \left( \frac{\sqrt{1 - a^2} T^{-1}(\theta) - C}{a} \right) \quad (3.35)$$

#### 3.4.4.4 Modelling sub-systems

Networks are composed of sub-systems in a hierarchy. An interesting extension of these models is to allow for the existence of sub-systems with elements that overlap multiple systems. The key output of such an analysis would be measurement for *Slice* parameters.

For this, an alternative way of specifying the  $\text{Slice}_{K_1, K_1}$  metric is used

$$\text{Slice}_{K_1, K_1} = \min(\max(L - K_1, 0)K_2 - K_1) \quad (3.36)$$

An accurate means of estimating the distribution of  $L$  determines *Slice*. Assume that the multivariate distribution is approximated well by a normal distribution, and solve for the conditional mean and variance:

$$\mu^j|M = E[L_j|M] = \sum_i N_i^j (1 - S_i^M) \quad (3.37)$$

$$\text{Cov}(L_j, L_{\dot{j}}|M) = \sum_i N_i^j N_i^{\dot{j}} S_i^M (1 - S_i^M) \quad (3.38)$$

where  $S_i^M$  is the probability that the measure is not triggered on element  $i$  conditional on  $M$ .

As before, take a *Slice* of the multivariate distribution of the inter-provider network as specified above in Equation 3.36 as the risk metric for the network.

A closed form solution is no longer possible for the conditional value of  $L$ , conditional on  $M$ . This once more means resorting to the *Indignity of Numerical Simulations*. Thankfully, a low-dimension Monte-Carlo can generate the multivariate distribution from Equations 3.37 and 3.38. Once  $E[L|M]$  is calculated, a second step numerically integrates over  $M$  to give the unconditional distribution. Similarly, to get  $\text{Slice}_{K_1, K_2}$ , evaluate Equation 3.27 numerically.

Note that neither the Monte-Carlo integration, nor the numerical integration steps will be too costly. The Monte-Carlo integration is over the number of sub-systems, not over all elements in all sub-systems: a substantial reduction in the dimensionality of the problem. The numerical integration which forms an outer loop to the Monte-Carl integration can be performed with something as easy as the trapezoidal rule.

Step-by-step:

1. Set-up an outer loop numerically integrating over the common factor  $M$ , say from -10 to +10 with .1 intervals.
2. In the inner loop, set-up a Monte-Carlo integration to calculate the conditional value of a multi-variate normal from Equations 3.37 and 3.38, at each value of  $M$ .

To get  $\text{Slice}_{K_1, K_2}$ , use Equation 3.27.

As an aside, this set-up gives an alternative means of calculating the Slice metric for single networks. Equation 3.36 is a call spread<sup>1</sup> on a normally distributed variable which also has a closed form formula:

$$E[L|M] = (\mu_M - K_1)\Phi\left(\frac{\mu_M - K_1}{\sigma_M}\right) + \sigma_M\phi\left(\frac{\mu_M - K_1}{\sigma_M}\right) - (\mu_M - K_2)\Phi\left(\frac{\mu_M - K_2}{\sigma_M}\right) - \sigma_M\phi\left(\frac{\mu_M - K_2}{\sigma_M}\right)$$

where  $\phi$  is the density of the Standard Normal Distribution.

#### 3.4.4.5 Dynamic Complexity

Another extension of this model challenges the assertion that expected trigger probabilities are constant over time. This is clearly not a realistic assumption.

The assumption is implicit in the way the systematic factor is specified: as can be seen from the steps required to Monte Carlo the model in Section 3.4.4.2, the systematic factor is assumed to realise a value only once, at time=0.

A relatively straightforward way to introduce dynamics into the model would be to allow the systematic factor to follow a Brownian motion path. Rather than to draw a single value for the systematic factor, one would draw a sample path, and a set of realisations of the idiosyncratic factor for each element at each time step along the way, correlated either through a flat or a heterogeneous correlation matrix.

Finally, the correlation matrix used along the path of the factor could also be functional or stochastic:

---

<sup>1</sup>Derivatives parlance for a combination consisting of going long one call option and short another with similar maturities but differing strikes.

- A single correlation value could be related to the systematic factor. A reasonable assumption would be for negative realisations of the factor to be correlated with positive realisations of correlation.
- Correlation itself could be stochastic, meaning the model would become a 2-factor model.

### 3.4.4.6 Multiple Factors

Finally, it is also possible to extend the model to multiple systematic factors by rewriting Equation 3.21 as:

$$A_i = \sum_{j=1}^J a_i^j M_j + \epsilon_i \quad (3.39)$$

where  $\epsilon_i$  is the noise of *Element i*, and:

- $M \sim \Phi(0, \Omega_M)$
- $\epsilon_i \sim N(0, \omega_i^2)$

This simple extension could be quite useful for modelling multiple Autonomous Systems in BGP simultaneously. A factor common to all the elements within one AS could be used to measure risk in that domain. A separate factor could be used for a different AS. An Internet-wide factor could still be shared amongst all elements. This idea can be layered as many times as desired:

- On the basis of shared connections in an AS, and
- Within geographies and times zones

See (LKSS00).

### 3.4.4.7 Model Complexity

There is always a tradeoff associated with increased modelling complexity, encoded in the notion of *Occam's Razor*. *Model Complexity* is explicitly defined so that the engineer can keep track of the costs of the complexity introduced into a model.

The various extensions to the model in Section 3.4.4 lead naturally to a definition of *Model Complexity*:

**Definition 18 (Model Complexity)** *Model Complexity is an information theoretic distance between the simplest possible model of a system and extensions which aim to increase its explanatory power. In this context, it can be measured as the Kullback-Leibler distance between the CDF implied in Equation 3.29 and various extensions in Section 3.4.4.*

The advantage of defining *Model Complexity* formally is that it allow investments in model development to be measured quantitatively. For example, using the recursion in Equation 3.33, it may be possible to model each and every individual element in an entire network. This may not be a valuable investment if the same effects can be more easily captured simply through the distributional changes in Section 3.4.4.3.

Model Complexity is meant to quantify the complexity associated with extending a base model for a network to account for known details in a networks configuration which deviate from the base case assumptions. Qualitatively:

$$\text{Model Complexity} = \text{Entropy of Base Model} - \text{Entropy of Extended Model} \quad (3.40)$$

This can easily be quantified as the KL distance between the two models.

### 3.4.5 Discussion

The distribution of the model in Equation 3.29 is Gaussian, as both the systematic factor  $M$  and the idiosyncratic factor  $X_i$  in Equation 3.21 are Normally distributed.

This is by design: Gaussian Distributions maximise entropy among all distributions with a known mean and variance, making them the natural choice of underlying distribution for data summarised in terms of sample mean and variance.

This implies that if we were to graph the entropy of a system and the entropy of a Gaussian model which best fits it, the Gaussian model would have higher entropy. This point is illustrated below in Figure 3.4 and again in Section 5.2.2.1.

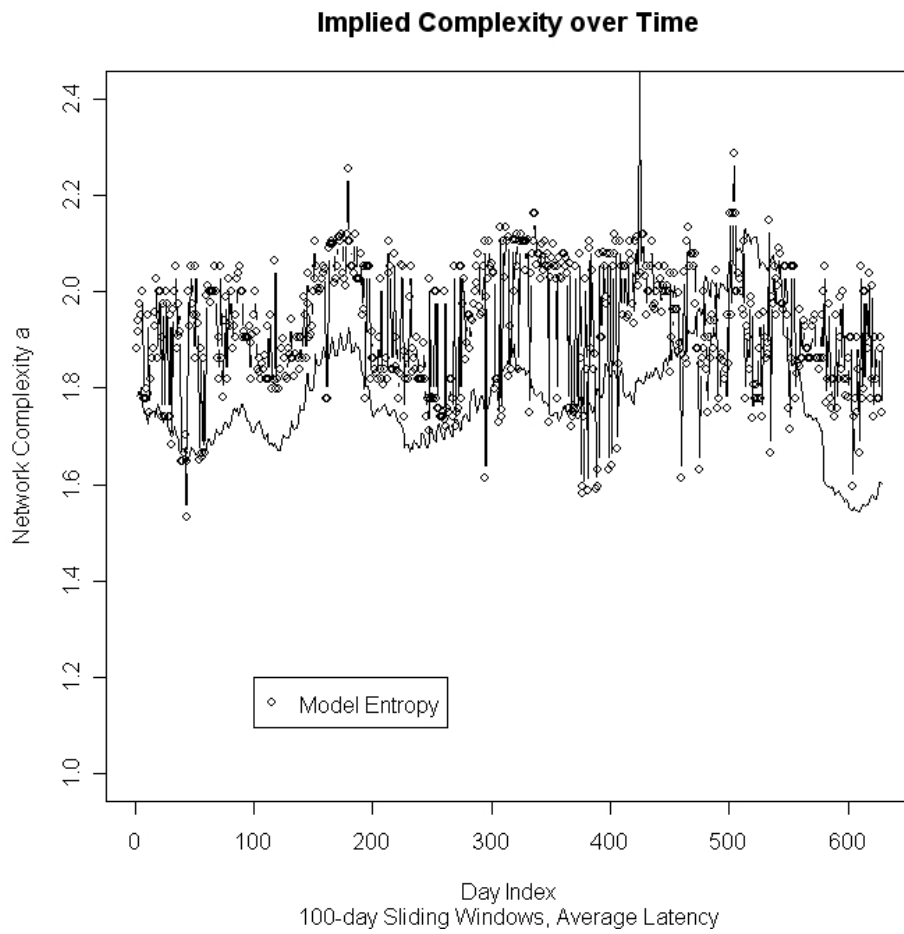


Figure 3.4: Empirical and Model Entropy, Average Latency 50ms

In the real-world, real networks are non-random in important ways. For example, they exhibit power-laws in their degree distributions. A totally random graph is Poisson in its degree distribution, and hence has maximum degree distribution entropy. Deviations from this statistical distribution happen for good reasons. For example, networks whose degrees are distributed with power-laws are extremely resilient to random removals of vertices. Networks which are subject to such idiosyncratic faults adapt their structure to minimise such disruptions.

### 3.4.5.1 The Subjective Nature of Complexity

One of the key insights of Murray Gell-Mann([GM94](#)) is that complexity is highly subjective and model dependent. There is no complexity in a system if the observer finds a simple model explains a significant part of the variability in a system.

Hence, complexity cannot be defined as a universal constant independent of the observer. Instead complexity is a relationship between a model of a system and an observer. A system is only complex when the observers model does not capture the complex, emergent properties observed in a network.

This is why the complexity metric is built on deviations from the predictions of a simple, base-case model. This model is open to criticism and could be specified in alternative ways. However, the goal with this model was to find the simplest model with the richest parameterisation, so any competing model would either have to be simpler and still retain the explanatory power of this specification, or be no more complex yet be more descriptive.

### 3.4.5.2 On Reference Measurements

In today's networks, almost all networks are forced to reduce their network monitoring to data which is common across all elements. such data is typically limited to:

- Utilisation rate
- Error rate
- Latency



- Input/output measurements

In most large-scale networks, legacy and element heterogeneity would make it absolutely impossible to build a complexity measurement metric based on anything other than the most generic network measurements. Further, network behaviour which is not measurable, or just too difficult to measure, deserves to be called complex anyway.

Nonetheless, which *Reference Measurement* to use remains an unanswered question in this thesis. Perhaps, it is preferable to measure all of them, as the calculations are extremely easy, and use the measurements in the right context. To make this concrete, assume that one considered a routing protocol as an extension of the work in this thesis which used *Slice* to make certain decision about how to route traffic amongst sub-systems. In the context of the *Entropic Routing* discussions of Section 6.1.1, traffic sensitive to latency could be routed on the basis of latency complexity and risk.

#### 3.4.5.3 Assuming Away the Idiosyncracies

One of the key lessons to be learnt from financial theory in the last 50 years is the Capital Asset Pricing Model (CAPM, (Sha64)). While CAPM also makes some unrealistic assumptions which are subject to the criticisms of Section 2.1, its two primary messages are grounded in common sense:

**Risk equals return** More risky instruments should earn a higher rate of return.

**Diversification reduces risk** The market pays no premium for taking idiosyncratic risk.

Most students of finance understand the first point, but do not really understand the second. The implications of the second point are pretty strong: no market participant should ever expect to be paid for taking risks which are entirely idiosyncratic, as under the assumptions of CAPM, these risks can be covered through diversification for free. A very sophisticated model is not needed to make sense of these two statements. This is the driver behind the idiosyncratic/systematic separation of risks.

### 3.4.5.4 On Factor Models

Models based on conditional independence<sup>1</sup> are indispensable for a number of reasons.

The specification of full joint trigger probabilities is a hard problem: While there are only four joint events for two elements (none trigger,  $A$  triggers,  $B$  triggers and both trigger), there are  $2^N$  joint default events for  $N$  elements. For a realistic number of elements it is impossible to enumerate these probabilities. Put another way, the calculation of the multivariate fault distribution from the network, when each element is modelled by a 2-state Bernoulli-(0,1) which takes the value 1 if some threshold is cleared is NP-Hard (see Section 1.2.1). This is the main reason why conditional independence models which predict measurement correlations from more fundamental variables must be used.

Furthermore, the *conditional independence* framework admits the following requirements.:

**Credible Modelling of Effects** There are two main interaction effects, common factors that influence all elements, such as intra-day cycles, and idiosyncratic factors effecting individual elements. The model must cater to each separately.

**Efficient Computation** Many models that are convenient for a single element or a few of them may be computationally expensive to use for large networks. The use of classical queueing theory based models with copula dependence structure are intractable when the number of elements increase. Simulation must show little sensitivity in computational terms to the number of elements.

**Ease of Calibration** Together with the computational issue, a desirable property of a model is easy calibration, essentially meaning a small number of orthogonal factors which span well the state space of possible outcomes.

Any factor model, such as Equation 3.21 involves a systematic term(s), and an idiosyncratic term. Conditional Independence only solves half the problem.

---

<sup>1</sup>The terms *conditional independence* and *factor models* are used interchangeably.

Large-sample approximations in order to achieve efficient computational algorithms have to be used where possible, to retain tractability

## 3.5 Implied Complexity

The concept of *implied volatility* is one of the most important developments in finance in the last thirty years. *Implied volatility* is the value which, when used in the Black-Scholes equation for an option, returns the current market price for traded options.

Under the stylised assumptions of the Black-Scholes model, this parameter is supposed to take the same value for all options. This rarely happens: usually every financial market values options with different characteristics at widely different volatilities.

To a first order, this does not invalidate the core of the Black-Scholes model although it does suggest extensions: using a local or a stochastic volatility model, it is possible to re-price options in the market using one parameter for volatility, at the expense of additional parameters (stochastic volatility, correlation, and additional structure associated with the underlying distribution, such as skewness and kurtosis).

There are in infinitude of reasons for why volatility is not constant, meaning that the underlying simple Black-Scholes model is wrong. In any case, the financial markets still use implied volatilities from the Black-Scholes model to exchange pricing information about options. This vernacular is not just model misspecification: it is part of a systems (the market) response to hard computations. By creating a formalism shared by all participants, the system reduces the cost of transmitting information and reduces computational burdens. This point is argued in Section 2.4.1. Such vernacular can also be seen as a case of *emerging complexity* common to Complex Adaptive Systems, as in Section 2.4.3.

The same construction is possible in networking. The machinery in Section 3.4 permits a parameter space which can play the role of a vernacular for thinking about and relaying information about complexity issues in a network in a succinct way. Most importantly, due to the information theoretic foundations of this

system of complexity measures, the parameters are comparable between networks, and even systems,

Hence, it is possible to forklift the same inversion from price to parameter into the domain of networking. This enables a language which allows network engineers to communicate and reason about:

1. The *Architectural Complexity* of the network
2. Idiosyncratic and Systematic Risks in their network designs through *Slice Curves*<sup>1</sup>
3. The *Operational, Element, and Scheduling Complexity* of their networks through their *Implied Coupling*<sup>2</sup>

Combined, these mathematical relationships define the *Complexity plane* of the network. A number of different ways of characterising and communicating the complexity of a network are analysed in the rest of this Section.

### 3.5.1 Zero Complexity

In the *Measurement Model*, the systematic factor has a normal distribution. This is chosen because from all distributions with a given variance, the normal distribution has the highest entropy. It also leads to tractable, simple mathematics.

In theory, Section 3.31 could be used on *Slices* of the empirically observed multivariate distribution of a network to imply a level for  $a$ . However, one would see (see Section 5) that a different level of  $a$  would result for different instantiations of  $K_1$  and  $K_2$ . This happens because it is rare that the *Measurement Model* captures the entire *Implied Complexity* of the network and is perfectly Gaussian.

It would be interesting to see what exact shape the factor would have to take so that one correlation parameter, could explain all *Slice* measurements with one  $a$  parameter.

An oft employed technique from the financial markets (and others), a Gram-Charlier Type A expansion of the normal distribution <sup>3</sup> can be used for this

---

<sup>1</sup>Defined in Section 3.5.2.

<sup>2</sup> Defined in Section 3.5.4.

<sup>3</sup>Alternatives are Type C expansions, and an Edgeworth expansion

calculation. See (JR99). A low order version of this expansion is easy to adapt, and the mathematics remains tractable.

Assume that the systematic factor driving the performance of individual elements can be written as:

$$A_i = a_i M + \sqrt{1 - a_i^2} X_i \quad (3.41)$$

as before. Again, let  $x_i \tilde{N}(0, 1)$  but, departing from the previous specification, now the distribution of  $M$  is given by a Gram-Charlier Type A series expansion, such that:

$$\begin{aligned} f(x) &= \sum_{j=0}^{\infty} c_j He_j(x) \phi(x) \\ c_r &= \frac{1}{r} \int_{-\infty}^{\infty} f(x) He_r(x) dx \\ \phi(x) &= \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \end{aligned}$$

where  $He_i(x)\phi(x) = (-D)^i \phi(x)$ ,  $D^i$  denotes the  $i^{th}$  derivative, and  $He$  is the Hermite polynomial.

As before, element homogeneity is assumed, and as the network size grows, the idiosyncratic noise  $x_i$  are diversified away, each  $a_i$  are identically distributed and in the limit, the fraction of elements triggering a measurement becomes:

$$F \approx \Phi \left( \frac{C - au}{\sqrt{1 - a^2}} \right) \quad (3.42)$$

exactly as before. However, now, the CDF of trigger probabilities for a large network becomes:

$$F_{\infty}(\theta) = \Upsilon_{GCparameters} \left( \frac{\sqrt{1 - a^2} \Phi^{-1}(\theta) - C}{a} \right) \quad (3.43)$$

where  $\Phi$  has been replaced by  $\Upsilon_{GCparameters}$ , the normal distribution expanded by Type A Gram-Charlier parameters.

So far, except for the specification of the systematic factor and the implied CDF for the network, this is a repeat of previous results. The key result needed

in order to implement the factor assumption is an explicit relationship between the trigger probabilities and individual element capacity, which is as follows:

$$p[x_i \leq C_i] = \Phi(C_i) - \sum_{j=1}^{\infty} a_i^j c_j \phi(C_i) He_{j-1}(C_i) \quad (3.44)$$

For more details, see (Sch04). To extract the implied distribution of the factor, assume that the correlation structure is homogeneous across all elements and truncate the Gram-Charlier Expansion after the fourth moment. Practically, this means that the factor distribution will have non-zero skew and kurtosis. As before, *Network Complexity* could then be measured as the Kullback-Leibler distance between this distribution and a  $N(0, 1)$ .

#### 3.5.1.1 Implementation

In (JR99), the authors point out that even though Gram-Charlier expansions allow for a certain flexibility over skewness and kurtosis, they have the unfortunate drawback of sometimes yielding negative densities. They then proceed to derive various positivity constraints which can be numerically implemented. In terms of Skewness and Kurtosis, this amounts to constraining the optimisation as follows:

$$\begin{aligned} \text{Skewness} &= [-1.0493, 1.0493] \\ \text{Kurtosis} &= [3, 7] \end{aligned}$$

To find the exact parameters, flat correlation is assumed ( $a_i \equiv a$  for all  $i$ ), and Powell's method (PTVF92) is used to find the  $a$ , skewness and kurtosis which best fits (squared relative error terms) the empirical distribution of the networks measurements.

A Type A Gram-Charlier expansion is particularly easy to implement, hence its common usage in the financial markets. The CDF of a Type A expansion is simply:

$$\Phi(c) - \frac{Skew}{Fact(3)} \frac{d^3 \Phi}{dx^3} + \frac{Kurtosis}{Fact(4)} \frac{d^4 \Phi}{dx^4} \quad (3.45)$$

The role of Equation 3.44 is to allow the calibration of  $C_i \equiv C$  to the skew and kurtosis of the systematic factor so that conditional and unconditional probabilities across the entire network match a certain a-priori expectation.

### 3.5.2 Implied Multivariate Distribution

Another technique from the financial markets reverses the previous question: what is the implied distribution of the systematic factor in a model with zero complexity, i.e. a model which perfectly explains network *Reference Measurements* data ?

In the large homogeneous model of the network used as a base case, the Law of Large Numbers was invoked to produce a simple measurement of the network, *Slice*. *Slice* has two parameters,  $K_1$  and  $K_2$  which are the lower and upper bound of the cutoff of the distribution being measured.

For the percentage of the total network trigger probabilities to exceed a certain threshold  $K$ , the systematic factor  $M$  has to be less than  $A(K)$  given by:

$$A(K) = \frac{1}{a} \left( C - \sqrt{1 - a^2} \Phi^{-1} K \right) \quad (3.46)$$

The call option<sup>1</sup> payoffs on this distribution, known as stop-loss transforms in the actuarial literature, are easy to calculate:

$$F(K) = E \left[ [L - K]^+ \right] \quad (3.47)$$

It turns out that one can use *Slice* to convey all the information about a networks characteristics (in relation to a *Reference Measurement*) using *Slice*. Consider only  $Slice_{0,K_2}$  measurements, i.e. thicker and thicker *Slices* all indexed to start at the origin of the distribution<sup>2</sup>.

A classical method from the options markets by (BL78) can then be used to infer the multivariate distribution, given a set of  $Slice_{0,K_2}$  measurements. Breeden and Litzenbergers' key insight is that  $Slice_{0,K_2}$  can be written in a way which resembles the payoff of a call option contract:

$$F(K) = \int_0^1 [x - K]^+ f(x) dx \quad (3.48)$$

---

<sup>1</sup>A *Call Option* is a derivative instrument which is the right, but not the obligation, to buy a certain security at a certain pre-specified price on a certain date in the future.

<sup>2</sup>In the markets for collateralised debt obligations, this measurement is known as “base correlations”. The machinery here is taken directly from those markets.

Taking derivatives:

$$F'(K) = \int_0^1 H(x - K)f(x)dx$$

$$F''(K) = \int_0^1 \delta(x - K)f(x)dx = f(K)$$

where  $H$  and  $\delta$  are the Heaviside function and Dirac distribution respectively. Using Equation 3.31 and the chain rule:

$$f(x) = \alpha_0 + \left( \alpha_{10} + \alpha_{11} \frac{da}{dK} \Big|_{K=x} + \alpha_2 \frac{d^2a}{dK^2} \Big|_{K=x} \right) \quad (3.49)$$

where:

$$\alpha_0 = \frac{\phi(A)\sqrt{1-a^2}}{\phi(B)a}$$

$$\alpha_{10} = -2\phi(A) \frac{Ca - A}{a * \sqrt{1-a^2}}$$

$$\alpha_{11} = \frac{\phi(A)\phi(a)}{\sqrt{1-a^2}}$$

$$\alpha_2 = \frac{\phi(A)\phi(B)}{\sqrt{1-a^2}}$$

and

$$A = \frac{1}{a} \left( C - \sqrt{1-a^2} \Phi^{-1}(K) \right)$$

$$B = \Phi^{-1}(K)$$

This means that  $a = a(K_2)$ : the coupling parameter changes value for each measurement  $Slice_{0,K_2}$ . This defines the *Slice Curve*:

**Definition 19 (Slice Curve)** *The coupling curve parameterised by  $a = a(K_2)$  which fits the Measurement Model in Section 3.4.1 to a series of chained Slices of the the empirical distribution perfectly.*



### 3.5.3 The Complexity Plane

It can be shown that in the context of Gaussian coupling in the specification for the *Measurement Model*, there is a one-to-one relationship between the coupling parameter and a certain specification of a series of *Slice* measurements (see (BG05)). Note that the coupling parameter in the case of the *Measurement Model* is the parameter  $a$  which measures the sensitivity of element utilisation to the systematic factor.

This makes it possible for the engineer to think about the *Slope* and *Level* of the *Slice* curve as fundamentals:

**Slope** Architectural and routing decision which increase the slope of the *Slice curve* (Section 3.5.2) simultaneously increase idiosyncratic and systematic risks in the network, while decreasing the likelihood of intermediate states which are *complexity-neutral*. This can be seen in Figure 3.5.

**Level** increasing the level increases systematic risks reduces idiosyncratic risks (Section 3.5.3), but can be *Complexity Neutral* states unchanged. This can be seen in 3.5.3.

This gives rise to the *Complexity Plane* of the network:

**Definition 20 (Complexity Plane)** *A succinct and effective description of a networks complexity which acts as a tool for directly relating architectural decisions to the risk characteristics of a network.*

### 3.5.4 Implied Parameters

The final parameter which can be implied from the *Measurement Model* is it's *Implied Coupling*.

*Implied Coupling* is meant to measure how tightly the elements in a network are coupled together. In the context of the model, this is parameter  $a$  in Equation 3.21.

**Definition 21 (Implied Coupling)** *The value of the coupling parameter ( $a$  in Equation 3.21) in a Measurement Model which minimises the Kullback-Leibler distance between the model and empirically observed data.*

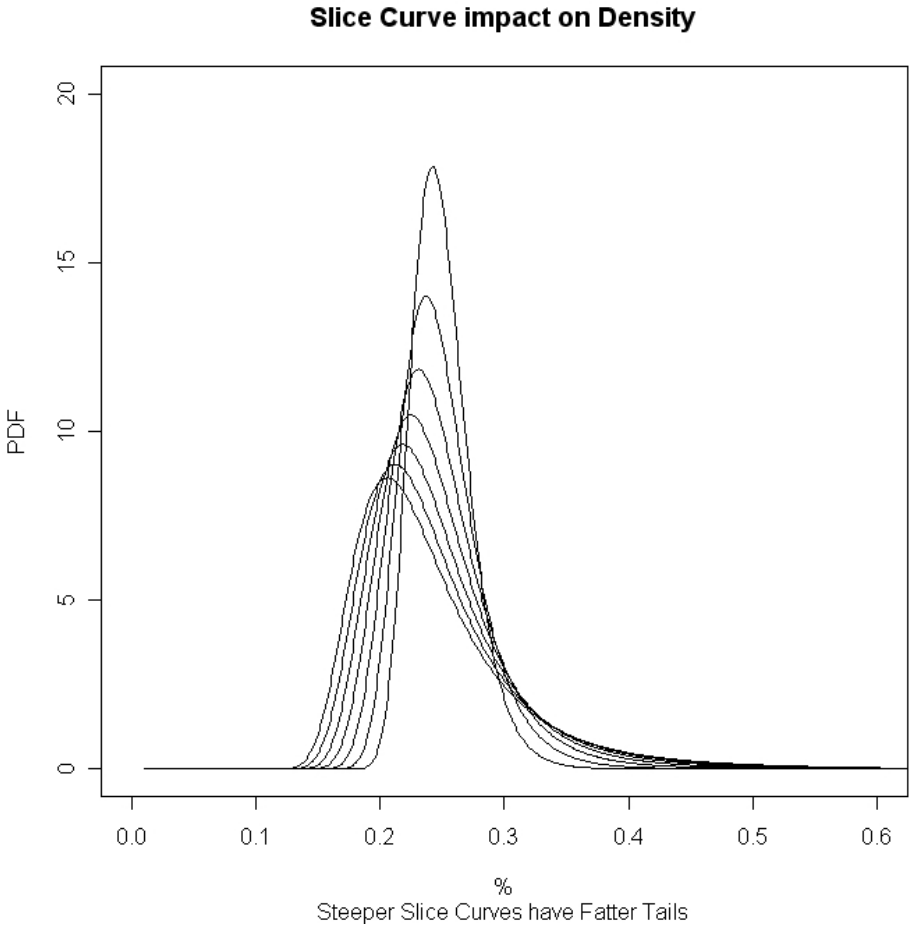


Figure 3.5: Slice Curve slope impact on the Probability Distribution

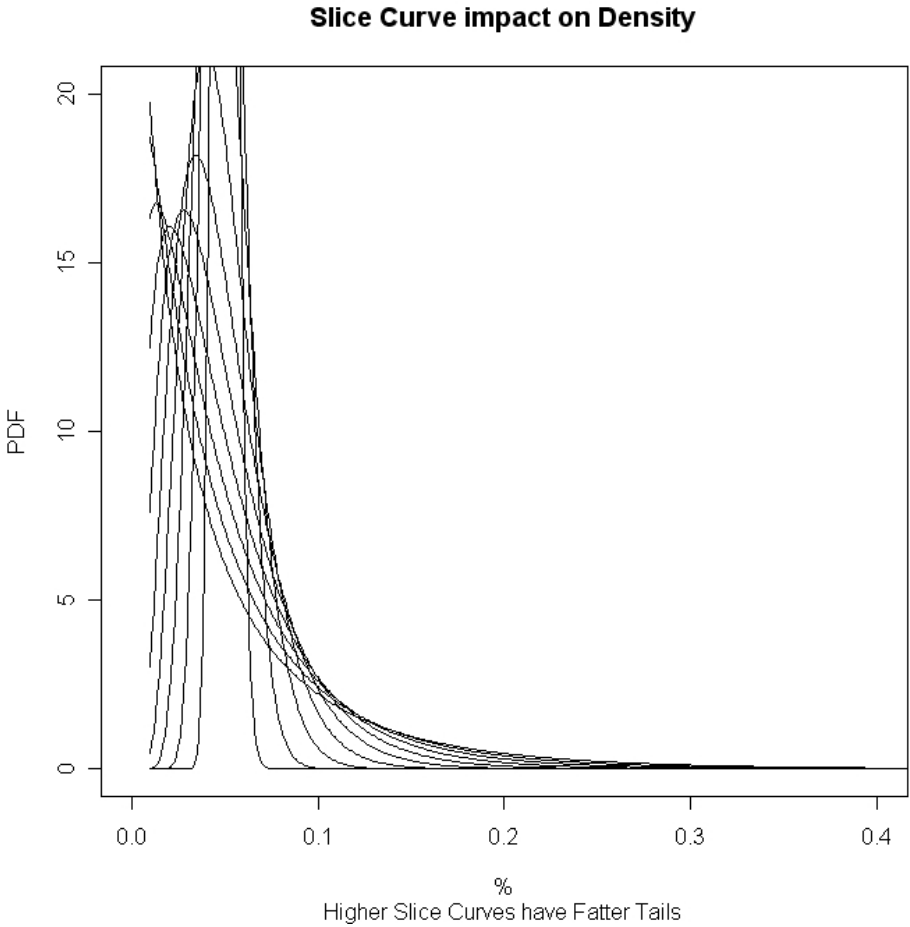


Figure 3.6: Changing the Level of the Slice Curve

This is one of the most novel contributions of this thesis. Many network models which use graph-theoretic models use the average degree distribution of the graph as an input. However, there is scant mention of how this could be measured on a network, or a sub-system, if one does not actually have access to the graph. *Implied Coupling* of a network could easily be measured using the *Measurement Model* and *Reference Measurements*.

Finally, to conclude this section, the concept of *Complexity Neutrality* is defined. *Complexity Neutrality* is meant to encompass all architectural and control decisions about a network which are not expected to impact the *Implied Complexity* of the network.

**Definition 22 (Complexity Neutral)** *Control decisions which do not impact the Implied Complexity of the network.*

This can happen either because they do not impact trigger probabilities and the Implied Coupling of the network, or because the impacts offset.

## 3.6 Models Discussion

The universal laws of Section 1.2.1 have remained valid for the last fifty years. These models have been tremendously successful for one simple reason: their dynamics can be described by very few parameters. This is a highly desirable property of any network model (see (WP98) for a discussion). The primary benefits of the model in Section 3.4 are parsimony in parameters and tractability.

An advantage of the *Measurement Model* is that the base case model is easily extended to cover for its shortcomings. These extensions cover the heterogeneity of elements, couplings, the existence of sub-systems and potentially, stochastic complexity as well.

Another practical benefit of the specification for the *Measurement Model* is that it can play a role in creating a financial market for the transfer, pricing and insuring of network risks. In Section 1.6.1, it was pointed out that one of the primary failings of previous attempts to introduce market metaphors into telecommunications was the mismatch between the true risks faced by market participants and the actual asset the markets tried to trade. The *Measurement*

*Model* makes it possible to separate out the risks which the market is not concerned about, i.e. idiosyncratic risks, from those for which the market will pay a premium, systematic risks. Finally, the *Measurement Model* admits a rich vernacular with which to reason about the market, a key requirement for the smooth functioning of any financial market.

In Section 2.3.1, it is pointed out that there is an equivalence between models based on optimising log utility and those maximising entropy. This posits the question: why not build on the work of (KMT98) and focus instead on the log utility framework in the context of the *Architectural Model* in Section 3.3 ?

This has been partially answered in Chapter 2. There is a further practical issue which should be highlighted, which impacts model design. In reality, the entropy function of an equilibrium system looks like a neoclassical utility function. However, even if there were only one agents in a system, the associated entropy function would not represent the agents preference ordering of alternative resource bundles. Hence, when considering networking issues, this makes it possible to abstract oneself from the need to consider the preferences of individual agents. For any model to be practicable, this is an important requirement.

Furthermore, in Chapters 6 and 4, architectural implications of increasing entropy in networks are discussed and various insights into the impacts of optimal control are analysed. Many of these insights are a practical function of the concave shape of the entropy function.

Finally, the entropy function is just the maximum of the statistical entropy over the set of probability distributions compatible with the resource constraints on mean actions. Since entropy is a strictly concave function, when the constraints are linear, the maximisation function always has a unique interior solution. This is a very convenient result in terms of tractability.

### 3.6.0.1 Limitations and Issues

A key limitation of the *Measurement Model* is its reliance on *Reference Measurements*. This model is limited to describing and controlling the complexity of a network that is a function of those *Reference Measurements*. For example, the complexity of a network as viewed from the perspective of latency measurements

may be different from that of throughput measurements. This leaves open the question of which *Reference Measurement* to use.

To make progress, this is a requirement: different fields suppress certain complications while highlighting others. Some of these issues were mentioned in Section 1.6.2: the diversity in links, nodes. The relationships between the complexity of a network vis-a-vis different *Reference Measurements* is a question left unanswered in this thesis, and is a rich area for future research.

### 3.6.1 Effective Complexity: The Full Picture

The practical goal of the Engineer in the field is to understand:

- what causes complexity
- the impact of complexity on network performance, and
- how architectural decisions will impact the network due to increases and decreases in complexity.

What is needed is some measure of the networks *Architectural Complexity*. However, a full calculation of a networks *Architectural Complexity* may not be possible due to the intractability of the data collection requirements. A number of complexity measures have been introduced in this Chapter to enable these calculations:

- The *Architectural Model* leads to a notion of *Architectural Complexity*, which reflects the total complexity of the network due to four components:

**Operational Complexity** Complexity arising from the number and heterogeneity of *Operations*.

**Element Complexity** Complexity arising from the number and heterogeneity of *Elements*.

**Scheduling Complexity** Complexity arising from the scheduling of *Operations* on *Elements* in order to abide by precedence requirements.

**Traffic Complexity** Arising from the number and heterogeneity of traffic flows.

- The *Measurement Model* gives a set of tools with which to physically measure complexity from actual network data, leading to:

**Implied Complexity** The entropy of the measured behaviour of the network, which is further broken down and analysed as:

- *Implied Coupling*: A measure of how tightly the network is coupled together and resources are shared across *Operations* and *Elements*.
- *Slice Curve*: A practical tool for visualising and communicating the idiosyncratic and systematic risk inherent in the coupling of the networks resources.

**Network Complexity** Describing that part of the behaviour of the network which cannot be explained by a model of the network.

**Model Complexity** Describing the complexity of model extensions.

Hence, as an initial point of reference, *Implied Complexity* can be taken as a proxy for *Architectural Complexity*, when it is not possible to actually build the *Precedence* and *Operational Matrices* required for the calculations in Section 3.3. The *Implied Coupling* of the network is a proxy for that part of the networks *Architectural Complexity* due to the interrelationships between *Operations* and *Elements*. The *Slice Curve* is a further tool which allows the engineer to measure and visualise the networks idiosyncratic and systematic risks. The relationships outlined between the idiosyncratic and systematic risks in Section 4.3 make practical the suggestion that reasoning about complexity can impact real traffic requirements.

The role of *Network Complexity* is to bring in modelling as a factor in network management: when there is a simple model which is able to explain network data, the networks *Effective Complexity* is reduced as far as the Network Engineer is concerned. In the presence of a model, *Network Complexity* becomes the true proxy for the *Effective Architectural Complexity* of a network. These relationships are graphically illustrated in Figure 3.6.1.

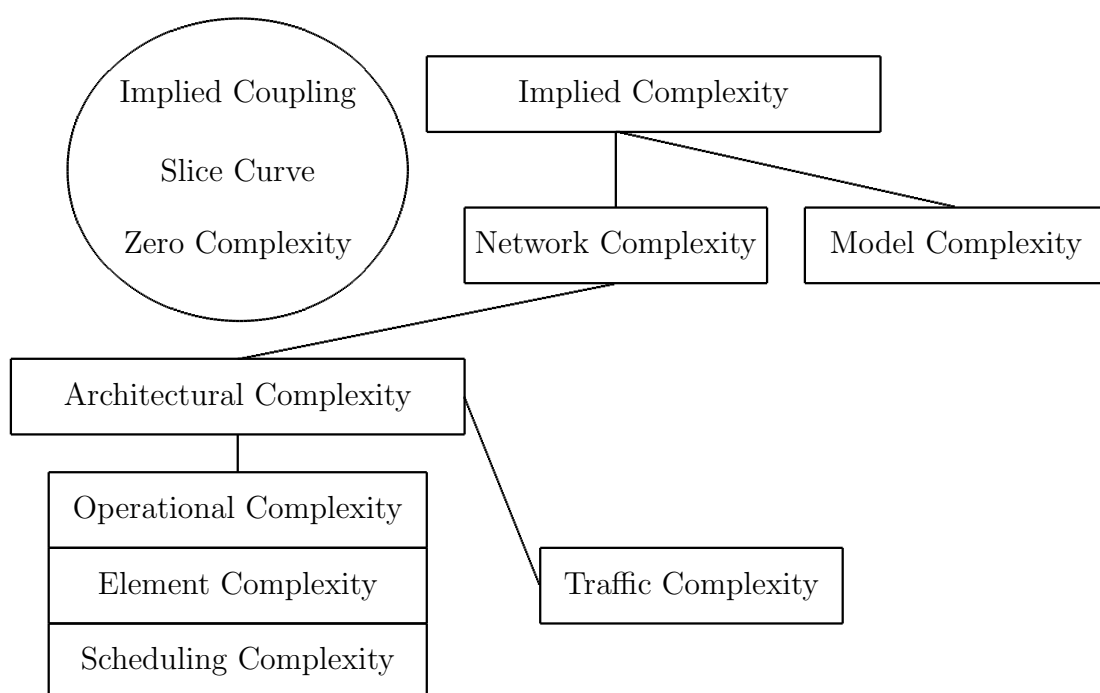


Figure 3.7: Relationship between Complexity Measures



# Chapter 4

## Insights

This Chapter performs an in depth analysis of various aspects of the *Architectural* and *Measurement Model*. The models provide a variety of valuable insights for the purposes of network analysis and design. The Chapter ends with an attempt to bring together all the various components into a complete picture for the network engineer.

A large part of information theory consists in finding bounds on certain performance measures. Most of the key insights in this Chapter substitute simpler bounds for what would otherwise be very complicated functions, had it not been for the use of Entropy and the application of Jensen's Inequality (see Section B.4). As mention in the Introduction, while this does not justify the *Entropic Formulation*, it is one of its great benefits.

### 4.1 The Impact of Complexity

#### 4.1.1 Simple Version

Even the simple model of Section 3.3.3.1 provides some meaningful insights. Consider a simple two-state system ( $A$  and  $B$  with probabilities  $P_A$ , and  $P_B$ )) with just one element. Its complexity will be:

$$H = -P_A \times \log(P_A) - P_B \times \log(P_B) \quad (4.1)$$

## 4.1 The Impact of Complexity

---

Expected Flow though the system will be:

$$F = A \times P_A + B \times P_B \quad (4.2)$$

Finally,  $P_A + P_B = 1$ . Manipulating, one can solve for Complexity as a function of Flow:

$$H = - \left( \frac{F - A}{B - A} \right) \log \left( \frac{F - A}{B - A} \right) - \left( 1 - \left( \frac{F - A}{B - A} \right) \right) \log \left( 1 - \left( \frac{F - A}{B - A} \right) \right) \quad (4.3)$$

Figure 4.1 illustrates the interacting impact of increasing flow and variability between states ( $A = 1 - B$ ) highlighting two points:

- Complexity reaches a maximum when states are similar
- The impact of increasing flow on complexity is higher when there is less variation between states.

While this insight is valuable, concentrating purely on complexity would result in misguided judgements. Consider the impact of parallelisation, on a single element which is able to operate either at full capacity or no capacity:

- $Flow = 50\% \times 100\% + 50\% \times 0\% = 50\%$
- $Complexity = -50\% \log(50\%) - 50\% \log(50\%) = 0.6931472$

Dividing capacity equally between two elements would have no impact on flows at the expense of added complexity:

- $Flow = 25\% \times 50\% + 50\% \times 50\% + 25\% \times 0\% = 50\%$
- $Complexity = -25\% \log(25\%) - 50\% \log(50\%) - 25\% \log(25\%) = 1.039721$

Yet, there is still a benefit to parallelisation which is not captured by this analysis: the riskiness of the operational characteristics of the network have decreased. Consider the standard deviation of flows:

- Single Element:  $stdev(0\% \times 50\%, 100\% \times 50\%) = .353$
- Two elements in parallel:  $stdev(0\% \times 25\%, 50\% \times 50\%, 100\% \times 25\%) = 0.144$

**Complexity as a Function of Flow and State Variation**




Figure 4.1: Complexity against Throughput

indicating a tradeoff between complexity and the riskiness of flows.

**Insight 5 (Complexity and Risk)** *Increased complexity can reduce the riskiness of a network and make it more resilient.*

*Slice* further enhances the analysis of this tradeoff by attributing risk to idiosyncratic and systematic factors.

### 4.1.2 Complicated Version

By using Jensen's Inequality and the fact that  $-\sum p \log p$  is a concave function, it is possible to derive a number of interesting properties of the *Architectural Model*.

First, note the following:

**Maximum Value of the Entropy Function** The entropy function takes its maximum value when all values are equal in the summation sign  $H(p) = -\sum_{i=1}^n p_i \log p_i$ . For a proof, see Section B.5.

**Concavity of Dimensions** The entropy function is concave in the number of dimensions of  $\psi$

**Concavity of Entropy** The entropy function is concave in the values of  $\psi$ <sup>1</sup>.

This means that mathematically, to increase complexity:

- The size of the *Operational Matrix* can be increased:
  - Increase *Operational Complexity*: increase the length of the matrix by adding more *Operations*
  - Increase *Element Complexity*: increase the width of the matrix by adding more *Elements*
- *Scheduling Complexity* can be increased, essentially by decreasing the number of zeros in the *Traffic Precedence Matrix*:
  - Decreasing precedence relationships between *Operations*, or

---

<sup>1</sup>In thermodynamics, the concavity of entropy and its dimensions implies that mixing two gases of equal entropy results in a gas with higher entropy.

- Allowing more *Operations* to be performed by each *Element*
- *Traffic Complexity* can be increased, by increasing the entropy of  $\nu$

These are consequences of the concavity of the entropy function and Jensen's Inequality. They are used below to gain valuable insights into network operations.

### 4.1.2.1 Impact of the Network Precedence Matrices

#### **Insight 6 (Operational and Element Similarities Increases Complexity)**

*The Operational and Elements Complexity metrics in Section 3.3.3.2 are maximised when all operations are possible on all elements.*

**Insight 7 (Scheduling Flexibility Increases Complexity)** *Reduced precedence requirements decrease the number of zeros in the Network Precedence Matrix. Again, due to Jensen's Inequality this maximises Architectural Complexity.*

Both insights are a relatively straightforward application of Jensen's Inequality to independent dimensions of the matrices.

Practically, similarity between operations or elements means that the observer has to ask more questions to determine where something went wrong in the systems to determine the erring element and/or operation. Similarly, a lack of operational precedence means that anything could be happening anytime, making it harder to determine state and increasing *Interrogation Costs*.

### 4.1.2.2 Size Matters

Furthermore, all these effects are related to the overall size of the *Traffic Precedence Matrix*: the entropy of a (normalised) smaller matrix will be impacted more by an increase in its size or by a decrease in the number of zeroes (see (CT91)).

**Insight 8 (The Impact of Size on Complexity)** *The complexity of a small system is more sensitive to architectural changes.*

This is related to the point made in Section 2.3 on the magnitude of Entropy Prices for constraints in networks: when Entropy Prices for constraints are high, the likelihood of a configuration change becomes very low.

### 4.1.2.3 Complexity and Architecture

The *Traffic Precedence Matrix* has sub-synergistic effects on *Architectural Complexity*. This means that networks with a low level of complexity will experience a larger increase in complexity from variations in operational requirements. This is an important result when the network architect is constrained and must focus his/her efforts on sub-systems which will have the most impact. This result implies that architectural changes will have more effect when systems are not operated near their theoretical complexity limits.

**Insight 9 (The Impact of Complexity on Architecture)** *Networks operating at lower levels of Architectural Complexity will respond more to changes in their architecture.*

### 4.1.2.4 Traffic Mix and Optimal Control Policy

Traffic can change in two ways:

- The types of traffic can increase
- The entropy of the traffic mix can increase

Either effect would cause an increase in the complexity of a system. However, a smaller system would be impacted more by changes in the traffic mix relative to a larger system.

**Insight 10 (Traffic Mix)** *The complexity of a smaller network is more sensitive to traffic mix changes. A smaller network will undergo a larger increase in overall complexity if the traffic mix is varied.*

Combined with the insight on the Impact of Complexity on Architecture, this implies that it makes more sense to focus engineering efforts on smaller networks where the complexity payoffs will be larger.

#### 4.1.2.5 Scheduling Complexity Insights

The best network performance for a given complexity level would be given by an optimal control policy. Control policies are a series of decisions scheduling *Operations* on *Elements* to maximise throughput. This section considers the impact of complexity on control policies and performance.

In the context of a network, a control policy is equivalent to scheduling traffic flows on the network to achieve a certain performance level. An optimal control policy would be dynamic, and at each instant, would consider the current state of the network as traffic flows arrive. The alternative would be a sub-optimal scheduling policy which did not take into concern the state of the network or traffic dynamics and followed a static control policy.

Performance is defined in terms of waiting times: when a flow arrives at an ingress node, if it has to wait to be served by an *Element*, then performance will degrade. This means that performance can be defined in terms of execution times as follows:

$$p(\text{execution times}) = \left[ 1 - \prod_{i \in m} (1 - p_i(\text{waiting time})) \right] \quad (4.4)$$

It is relatively easy to see that increasing the number of *Elements*, *Operations* would decrease waiting times.

$$\prod_{i \in m} (1 - p_i(\text{waiting time})) \leq \prod_{i \in m'} (1 - p_i(\text{waiting time})) \quad (4.5)$$

where  $m' > m$ .

Furthermore, increasing the number of *Elements* which can perform a certain *Operation* and reducing scheduling requirements also decrease waiting times. The former correspond to an increase in the dimensions of  $\Lambda$  and the latter correspond to a decrease in the number of zeros in  $\Lambda$ : both increase *Architectural Complexity*. This leads to the following insight:

**Insight 11 (Complexity increases Performance)** *Increasing Architectural Complexity will lead to an overall increase in the performance of a network when the network is operating under an optimal control policy.*

In some ways, this is an obvious result: there would be no reason to go for higher complexity in networks if there wasn't some gain. In Section 4.1.1, one benefit to increasing complexity was reduced variability in performance. Another is outright increased performance.

What is not so clear is what happens when the network is not operated under an optimal control policy. There are two key insights as they relate to this issue:

1. First of all, due to the concavity of Equation 4.4, the improvement in performance drops as the complexity of the network increases.
2. More importantly, the difference in the performance of a network operated optimally and under a static control policy is greater the higher is the complexity of the network.

The loss in performance from a static control policy is relatively easy to prove: consider a network with no complexity, i.e. one *Element* processing one *Operation* on one *Flow* at a time. The difference in network performance from any scheduling policy is zero. However, for any other network, the difference in performance will be positive. (a more thorough probabilistic analysis in the context of a manufacturing plant can be found in (Kle75) based on  $M/M/$  queueing systems).

These are important enough that they are worth formalising:

**Insight 12 (Diminishing Returns from Complexity)** *The performance gains from increasing complexity in a network are monotonically decreasing in the complexity of the network.*

**Insight 13 (Deadweight Loss)** *Networks operating at higher levels of complexity will show a higher relative performance degradation if their control policies are not optimal.*

These results have strong implications. In order to improve the performance of a network, the network engineer may choose to increase the complexity of the network. However, as the network becomes more and more complex, an optimal control policy will be harder and harder to calculate. This may also happen because *Traffic Complexity* may increase. In this case, the deadweight losses may be significantly higher.



This is one of the fundamental ideas behind *Entropic Routing*, analysed in Section 6.1.1: if the objective is performance, traffic flows should be segregated based on variability. Predictable flows should be sent down network paths that are less complex and variable flows should be sent down network paths that are more complex. If the objective is efficiency, however, the exact opposite would be appropriate. More variable flows should be sent down less complex paths that are easier to optimise and more predictable flows should be sent down more complex paths that are harder to optimise.

### 4.1.2.6 Architectural Complexity Conclusions

Overall, this implies that small networks are more sensitive to changes in traffic, and need to be scheduled and controlled more optimally. Hence, resources should be spent on increasing the capabilities of the scheduler. Larger networks are less sensitive to changes in traffic and can tolerate sub-optimal control policies.

Note that for any network, *Operational Availability* can be increased by:

- Increasing the number of elements, hence increasing the length of the *Operational Matrix*
- Increasing the number of element on which an operation can be executed, increasing the width of the *Operational Matrix*
- Decreasing precedence requirements between operations, decreasing the number of zeros in the *Precedence Matrix*

All three actions result in an increase in the Operational Availability of the network, at the expense of increased *Architectural Complexity*. For an application of some of these insights in the manufacturing domain, see (BR96).

## 4.2 Measurement Model Insights

The *Measurement Model* yields useful insights in the area of risk-control. The key source of complexity in the *Measurement Model* will be due to a deviation from its simplistic assumptions, which are based on homogeneity of elements. In the context of the *Measurement Model*, there are two sources of heterogeneity:

- Heterogeneity of Trigger Probabilities, i.e.  $a_i \neq a$
- Heterogeneity of Couplings, i.e.  $M_i \neq M$

Increases in idiosyncratic and systematic risks due to changes in these parameters is investigated. This is followed by a discussion of how the *Implied Complexity* of a network increases with changes in the heterogeneity of elements.

### 4.2.1 Heterogeneity of Trigger Probabilities: Impact on *Slice*

Heterogeneity of trigger probabilities increases idiosyncratic risks in a network. To understand this, consider the case when all the elements have a coupling of  $a = 1$ . In this case, the probability of just one element triggering a measurement is equal to the probability of the element with the highest trigger probability. On the other hand, if the coupling was 0, then the probability of just one element triggering a measurement would be equal to the average trigger probability of the network. Hence, it must be that a heterogeneity of trigger probabilities increases the idiosyncratic risks in the network. The multivariate distributions CDF has to add up to 1. If increasing variety increases the likelihood that one fault will happen within a certain period, the likelihood that multiple faults will happen must decrease.

**Insight 14 (Diversity in Elements Increases Idiosyncratic Risks)** *Increasing the Variety of Network Elements Decreases the Likelihood of Systematic Events.*

### 4.2.2 Coupling Heterogeneity

The impact of coupling heterogeneity is not so clear cut: while a heterogeneity of couplings would indicate areas of the one-factor covariance matrix implied in Equation 3.21 which would be highly correlated, indicating an increase in the incidence of clustered triggers, if average coupling were to be kept constant (i.e. links were rearranged, none added and none taken out), then idiosyncratic risks would also increase. Hence, heterogeneity in couplings increases both idiosyncratic and systematic risks.




Figure 4.2: Entropy ag Trigger Probabilities

#### 4.2.2.1 The Causes of *Network Complexity* ?

The root cause of *Network Complexity* is related to the way in which the network is coupled together, and the heterogeneity of its elements. To illustrate this, the entropy of implied model probability distributions is calculated for various scenarios.

**Trigger Probabilities and Entropy** As trigger probabilities increase, the entropy of the network will first decrease, and then increase. This can be seen in Figure 4.2.

As the coupling between elements and the systematic factor increases, the entropy of the network decreases (Figure 4.3). Combined, entropy is maximised when a network has very low trigger probabilities and very low coupling.

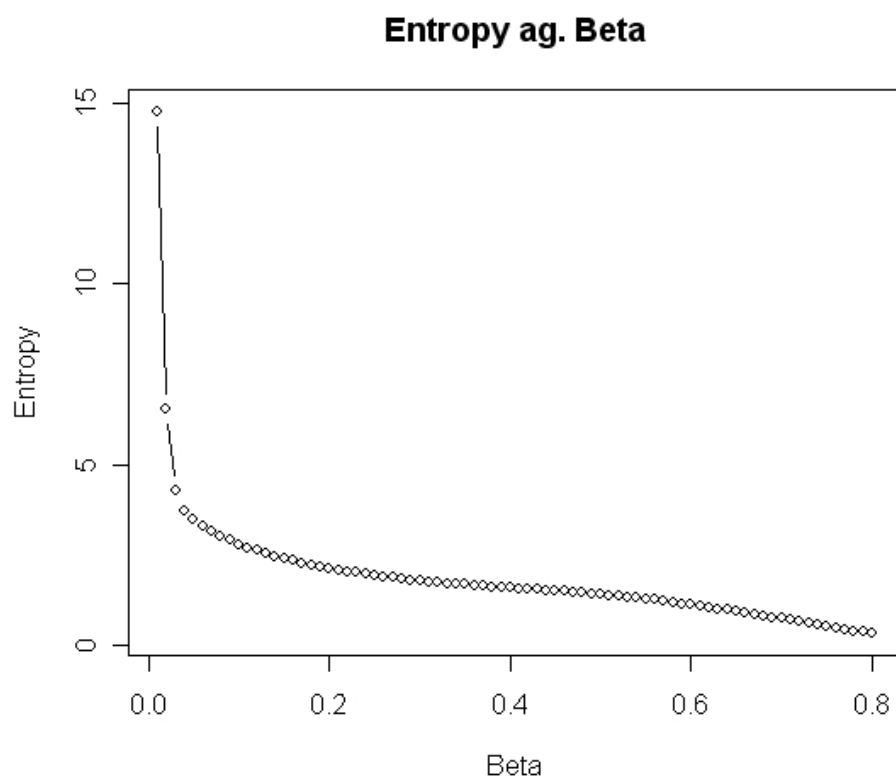


Figure 4.3: Entropy ag Beta

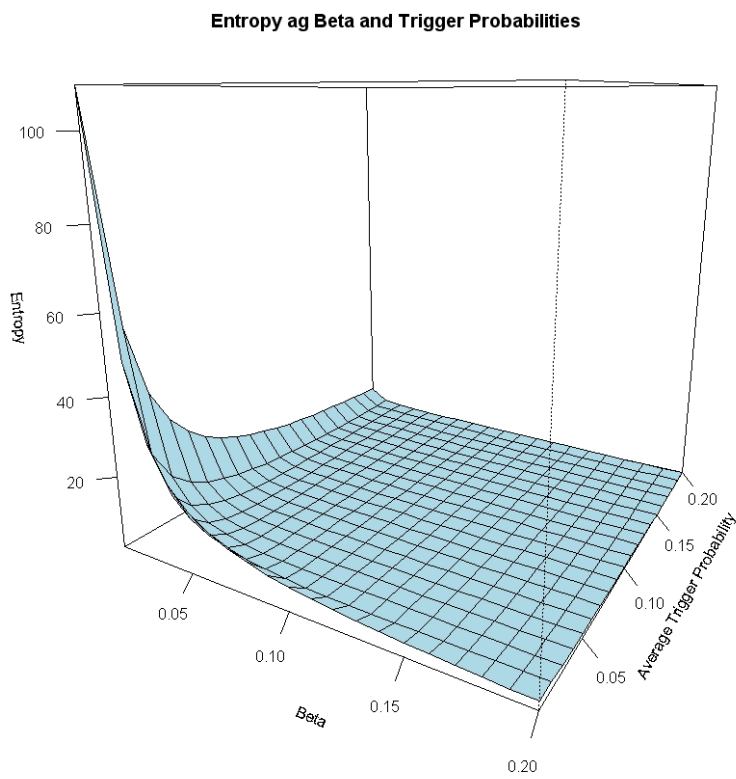


Figure 4.4: Entropy ag Beta and Trigger Probability




Figure 4.5: Entropy ag Beta and Trigger Probability, Small

This coincides with intuition: a network with low *Implied Coupling* corresponds to a network which is highly idiosyncratic. Highly idiosyncratic networks possess high structural complexity. On the flip side, network which are highly coupled possess low structural complexity.

Leading on from the insights of the previous sections, one can conclude that:

- Networks operating under sub-optimal scheduling policies should be highly coupled: this reduces complexity and hence deadweight losses.
- Networks operating under optimal scheduling policies can afford to be lightly coupled.
- Networks experiencing higher traffic variability operated by agents utilising statistical inference will react by becoming decoupled. This will further increase the entropy of the system and drive the network towards higher entropy.

Most importantly, if network architecture capabilities are a constrained resource, they should be directed towards smaller networks, as increasing the scheduling efficiency of smaller networks will yield greater benefits.

## 4.3 Completing the Circle

It is time to complete the circle between the entropic hypothesis of Section 2, the *Architectural Model*, and the *Measurement Model*.

### 4.3.1 Heterogeneity, Complexity and Risk

To do this, extensions in Section 3.4.4 can be analysed to investigate the relationship between heterogeneity in the infrastructure of the network, complexity and risks to performance.

The true nature of the linkage between the *Architectural Model* and the *Measurement Model* lies in the original specification for a complexity measure: what are the *Interrogation Costs* associated with determining state in a network ?

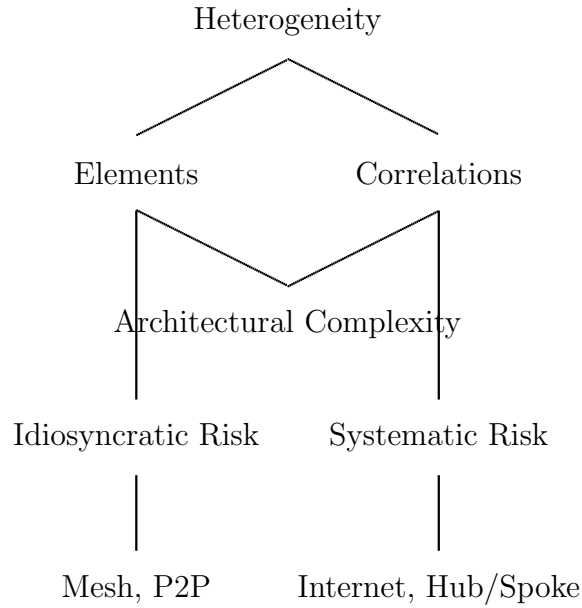


Figure 4.6: Relationship between Heterogeneity and Complexity

The *Architectural Model* was meant to answer this question qualitatively, and to breakdown these costs into their constituent *Operational*, *Element*, and *Scheduling* components. The notion of *Network Complexity* which comes out of the mechanics of the *Measurement Model* is meant to answer this question quantitatively, by positing the complexity of a network against the simplest model imaginable for explaining network behaviour. The link between these two models is very straightforward: networks with higher *Architectural Complexity* should be proxied well by networks which realise high levels of *Network Complexity*.

Once the complexity of a network is determined, the next issue becomes one of risk management: what are the risks that can be tolerated when transferring a certain kind of data and how risky is a network, or a sub-system ? This is answered by the insights in this Section: say that:

- Networks with higher complexity have higher idiosyncratic risks.
- Networks with lower complexity have higher systematic risks.

This can be tied in with the performance analysis of networks:



- Networks with high complexity can tolerate sub-optimal control policies.
- Networks with low complexity have low dead-weight losses
- Small networks are more sensitive to architectural decisions.

### 4.3.2 Practical Issues

How are all these insights and the tools of Chapter 3 useful to an Engineer ? The *Architectural Model* and the *Measurement Model* serve separate purposes in aiding the operations and management of a network, and need to be used in different ways.

#### 4.3.2.1 Making Architectural Decisions

The insights related to the *Architectural Model* can be used in a disciplined manner to guide any architectural decision related to a network in the following manner:

First, consider the impact of a specific design decision on the following factors:

**Operational Complexity** Will the number of operation increase or decrease ?

**Element Complexity** Will the number of elements increase or decrease ?

**Scheduling Complexity** Will the number of operations which an element can perform increase or decrease ?

**Traffic Complexity** Will the variability of the traffic carried increase ?

This will lead to a conclusion about whether the *Architectural Complexity* of a network will increase or not.

Once this is concluded, one can consider the impact on the network as it relates to:

**Size** Smaller networks are more sensitive to changes in *Architectural Complexity*

**Control Policy** More complex networks can tolerate sub-optimal control policies better: the performance loss is lower. On the flipside, the deadweight loss from sub-optimal control is higher for more complex networks: the efficiency loss is higher.

**Traffic** More variable traffic requires more complex networks.

### 4.3.2.2 Operating an Actual Network

This is where the *Measurement Model* can become very useful. First of all, different kinds of traffic can tolerate different kinds of risks. The tools of the *Measurement Model* such as *Network Complexity* and *Implied Coupling* could also be used in real-time (this is the topic of *Entropic Routing*, see Section 6.1.1) to decomposed a network into sub-systems or paths with differing levels of complexity. Different types of traffic could then be routed along different sub-systems depending on what kind of risks could be tolerated.

Consider the tradeoff between quality and quantity in the delivery of data across a network:

**Quantity Constraints** Traffic with real-time, high quantity delivery constraints can tolerate idiosyncratic risks, but not systematic risks: systematic risks effect all components in a network simultaneously and would make it harder to meet timing constraints. The tradeoff would be that certain packets may never be delivered. **Such traffic should be routed down high complexity network with low systematic risks.**

**Quality Constraints** Traffic with quality delivery constraints can tolerate systematic risks as they are not time sensitive but cannot tolerate idiosyncratic risks as they would require frequent resending of large packets, and/or too much re-sequencing due to idiosyncratic networks. **Such traffic should be routed down low complexity networks with low idiosyncratic risks.**

Finally, the *Implied Coupling* measure is an important tool for a network engineer. For example, the *Implied Coupling* of different sub-systems can be

measured using *Reference Measurements*, and the engineer can decide where to focus architectural resources accordingly, in light of the insight in Section [4.1.2.3](#) which says that architectural variations impact networks operating at low levels of complexity more. This can also be used to channel traffic according to its characteristics.

# Chapter 5

## Model Evaluation

This Chapter evaluates the *Architectural* and *Measurement Models* using data from a real network, and against real-life architectural decisions. Due to the differences between the two models, it is not possible to evaluate the two models within the same framework.

The *Architectural Model* is a tool for making systematic and consistent decisions about architectural issues. It needs to be evaluated by considering real architectures and testing it's ability to systematically break-down the issues into meaningful conclusions.

The *Measurement Model* acts on real data, so it is evaluated against a real network. *Reference Measurements* from a real network are used to generate actual entropy measurements, and an actual *Slice* curve is generated, to illustrate the tools made available by this formulation.

### 5.1 Evaluating the *Architectural Model*

The *Architectural Model's* purpose is to facilitate systematic thinking about high-level issues by putting the decision process into a common framework across domains. A number of interesting architectural questions lend themselves to the analysis made possible by the insights of Chapter 4.

It is very straightforward to put the *Architectural Model* to work. When considering an architecture, the following need to be considered:

**Impact on the width and variability of the *Operational Matrix*** Wider, less variable rows imply an increase in complexity due to *Element Complexity*, making it harder to figure out which element failed and which operations it was executing when it failed.

**Impact on the height and variability of the *Operational Matrix*** Longer, less variable columns imply an increase in complexity due to *Operational Complexity*, making it harder to figure out which operation failed on which element.

**Impact on the Number of Zeroes in the *Precedence Matrix*** If there are less operational precedence requirements, it will be more difficult to figure out which operation is currently executing.

At it's heart, all this analysis is meant to proxy for *Interrogation Costs*. To illustrate the *Architectural Models* ability to unearth interesting insights about the complexity implications of various design decisions, a variety of network designs are systematically evaluated:

1. Connection-less vs. Connection-oriented paradigms
2. Multi-path routing
3. Different ways of rolling out new features into the network
4. ATM vs. MPLS

The key point of this evaluation is that it is *systematic*. When confronted with any architectural issue, one considers step-by-step the impact on:

- *Operational Complexity*
- *Element Complexity*
- *Scheduling Complexity*
- *Traffic Complexity*

Combined, these lead to an overall conclusion about how the overall *Architectural Complexity* of a network will change as a result of a certain design decision.

### 5.1.1 Connection-oriented vs. Connection-less Paradigms of Networking: OSPF and MPLS for Traffic Engineering

The OSPF protocol is a link-state, hierarchical IGP, and uses Dijkstra's algorithm to calculate the shortest paths, using cost as a routing metric. The link state database of the network topology is identical on all routers in a common area. In the absence of an intermediate protocol like MPLS (MPLS is sometimes called a 2.5 Layer protocol), OSPF achieves traffic engineering objectives by varying the weights associated with paths.

MPLS, on the other hand, emulates properties of a circuit-switched network over packet-switched networks using labels, operating between Layer 2 (data link layer) and Layer 3 (Network Layer). In addition to IP Traffic, MPLS can also carry native ATM, SONET, and Ethernet frames.

MPLS and OSPF cannot be compared directly because its label forwarding table is constructed using IGP routing protocols, of which OSPF is one of the most popular (IS-IS being another alternative). However, MPLS makes significant changes to the means by which traffic engineering requirements are met in an OSPF network, and a comparison is possible between connection-oriented and connection-less paradigms in this context. Traffic engineering requirements in a network using MPLS are met by establishing paths for dedicated traffic, by packet type. In theory, using configurable OSPF link weights, many MPLS paths can be duplicated with OSPF. Besides, there is some empirical evidence indicating that the performance difference between a network which only uses OSPF and one which uses MPLS is not that great for working traffic (see (FT00)).

However, MPLS can distinguish packets towards the same destination IP address and treat them differently. Furthermore, MPLS can distribute traffic unequally on multiple paths. Both of these characteristics makes *Entropic Routing*, based on the notion of routing traffic along available paths as a function of its statistical and risk characteristics, much easier to implement.

Furthermore, a key advantage of MPLS is network resilience. OSPF is a distributed protocol, meaning that information regarding element failure takes some time to reach all forwarding tables. Even if the convergence of this update were

not an issue, any traffic engineering methodology which was based on link metrics would require a subsequent optimisation which is likely to take a significant amount of time. In the case of MPLS, this optimisation can be performed offline for typical failure patterns.

The machinery of the *Architectural Model* can be used to unearth these issues systematically. In a connection-oriented paradigm like MPLS, even though information is transmitted through the same elements as in OSPF, there is a significant decrease in the size of the *Operational Matrix*: this is because the path computation is only performed by source elements, and not by all the intervening elements. This triggers an architectural decrease in the complexity of the network (see Chapter 4) through a reduction in the height of the *Operational Matrix*. This is arguably the main reason why MPLS is preferred as the means to achieve traffic engineering requirements over OSPFs' native infrastructure in most small and medium-size networks: MPLS *Interrogation Costs* are significantly smaller than OSPF. Furthermore, there is also a decrease in the size of the *Precedence Matrices*, as label switching is much easier than prefix matching, which places consistency requirements across the fabric of the network.

Potentially, the MPLS intermediate layer can also resolves routing, admitting a further reduction in entropy by obviating the need for the router function entirely: this has the potential to reduce the width of the Operational Matrix. These devices are called *Label Switch Routers*. However, a Complexity-Risk Tradeoff (see Section 5) makes it harder to adopt *Label Switched Routers*, especially in light of increased restoration functionality required in the IP layer.

**Combinatorial Complexity** In (FT00), the authors describe a heuristic which in practice overcomes the computational intractability issues associated with OSPF weight setting (see Section 1.2.1). They do assume that the traffic demand matrix is known, which one of the objections to their proposal. Another is the requisite increase in *Scheduling Complexity* as the neighbourhood structure of the local search heuristic significantly increases precedence requirements. This is amplified because the requirements apply on a per element basis, implying that while there may be a polynomial increase in the efficiency of the network, there might be a combinatorial increase in it's complexity. Hence, in spite of

the increase in efficiency, it would not be surprising to witness that most operators would stick with the simplest rules (like inverse-capacity) for setting features specific to individual elements.

In concluding, the machinery of the *Architectural Model* is systematically applied to the connection-oriented paradigm (MPLS), in comparison with the connection-less paradigm (OSPF):

**Operational Complexity** Lower because only source elements compute paths

**Element Complexity** Ambiguous.

**Scheduling Complexity** Lower, because label switching is easier than prefix matching.

**Traffic Complexity** Ambiguous.

leading to lower *Architectural Complexity* for connection-oriented networking over connection-less networking.

This means that, the connection-oriented paradigm is more suitable for:

**Size** Small networks over which the operator has total control.

**Control Policy** Networks which can optimally controlled.

**Traffic** Networks which carry predictable traffic.

### 5.1.2 Multi Path Routing

The same arguments apply to multi-path routing: in theory, multiple paths from source-destination pairs have the potential to result in significant efficiencies. However, this would again trigger a combinatorial increase in the entropy of the *Precedence Matrices*, and also further, the *Operational Matrices* if new operations were required of elements. Unless the routing algorithm is a *Pareto Improvement* in the efficiency-complexity frontier, it is unlikely to witness significant take-up with operators who would face significant increases in the architectural complexity of their network as a result.

Systematically applying the machinery of the *Architectural Model* to multi-path routing:



**Operational Complexity** Higher because intervening elements have to carry out more operations.

**Element Complexity** Ambiguous.

**Scheduling Complexity** Ambiguous.

**Traffic Complexity** Higher, because the variety of flows across an element at any point in time will increase.

leading to overall higher *Architectural Complexity*.

This means that multi-path routing is more suitable for:

**Size** Large networks.

**Control Policy** Networks which may be sub-optimally controlled.

**Traffic** Ambiguous.

### 5.1.3 Rolling Features into the Address Space

There is one way to achieve multi-path routing without an increase in the *Operational Complexity* of the network: roll the feature into the addressing space of the router.

The idea of rolling new features into the addressing space resolves an important complexity problem associated with the introduction of features which require coordination with individual elements. From a complexity standpoint, the best way of doing this is to use existing *Operations* and *Elements*. Since routers are already built to route, a new feature which just makes use of that functionality would not trigger a combinatorial increase in the size of the *Operational matrix*. Obviously, the mapping from the feature to the address is an additional operation, but this scales linearly with the number of features and not with the number of elements. The only additional requirements are with the ingress and egress routers.

In some ways, feature introduction through the address space is used in certain contexts:

- GMPLS<sup>1</sup> treats circuits, lambdas, fibres and bundles all as labels
- Circuit Cross connect where the router only looks as far as a layer 2 ID, pushing routing table management to the edge, reducing per-element operational requirements

Once again, evaluating systematically:

**Operational Complexity** Lower: while the ingress and egress nodes would have to do more, any alternative means of introducing a feature by introducing new elements and new operations on those elements would be most costly.

**Element Complexity** Much lower, as the number of elements could decrease significantly.

**Scheduling Complexity** Possibly higher, as the address space will become more crowded.

**Traffic Complexity** Ambiguous.

leading to lower *Architectural Complexity*.

Smaller networks are more susceptible to changes in architecture. This could be an important consideration in deciding whether to admit new features into a network using new elements or through the existing infrastructure of addressing.

### 5.1.4 ATM vs. MPLS: Flexibility vs. Complexity

Most of the analysis of why ATM has failed against MPLS focuses on extremely domain specific issues, such as the complexity of PNNI<sup>2</sup>, the arbitrariness of

---

<sup>1</sup>Generalized MPLS (GMPLS) extends MPLS to provide the control plane (signaling and routing) for devices that switch in any of these domains: packet, time, wavelength, and fibre.

<sup>2</sup>PNNI is a link-state routing protocol for ATM networks that automatically finds paths in the network using neighbour discovery techniques and then assists in setting up SVCs (switched virtual circuits) between end systems.

boundaries between VP<sup>1</sup> and VC Switching, and the size of ATM frames. Typically, practical difficulties associated with overlay inter-operability are also highlighted.

While the literature is rich with technical issues, the analysis lacks transportability, and it is difficult to learn a lesson from the ATM/MPLS experience. This topic is a good example of how the *Architectural* analysis in this thesis succeeds in being systematic and consistent.

The overall bent of the existing research claims that that ATM is actually more complicated than MPLS. The analysis in this section debunks this claim, where it is made. Specifically, complex MPLS label stacking would indicate that the *Precedence matrix* for MPLS is significantly wider than ATM<sup>2</sup>.

There is an alternative *Architectural* reasons for the success of MPLS:

- Sub-optimal control policies: networks operating under sub-optimal control policies need to be operated at higher levels of complexity (see Section

---

<sup>1</sup>There are two major types of ATM connections: VC (Virtual Channel) and VP (Virtual Path). They differ in the way they are identified, and hence also in how they are switched. VCs are uniquely identified on the port of an ATM switch by the VPI and VCI numbers. VPs are uniquely identified on the ATM port by VPI number only. ATM cells that arrive on a VC are switched based on both the VPI (Virtual Path Interconnect) and VCI (Virtual Circuit Interconnect) number in the cell header. On the other hand, ATM cells that arrive on a VP are switched based on the VPI number only. VPs consist of all VCs with a specified VPI number on the ATM port. ATM networks rely on the Virtual Circuit (VC) and Virtual Path (VP) concept to provide unicast connection-oriented services with Quality of Service guarantees.

<sup>2</sup>**Experts call MPLS bad for 'Net VPNs based on Multi-protocol Label Switching said to be risky. Backbone management challenges also cited. By CAROLYN DUFFY MARSAN, Network World, 08/06/01** “ Two prominent Internet researchers from AT&T Labs are among a growing number of experts raising red flags about Multi-protocol Label Switching, a next-generation traffic engineering technology backed by network industry leaders such as Cisco, Juniper Networks and AT&T itself.

The researchers - security guru Steve Bellovin and network operations expert Randy Bush - say MPLS create serious network management challenges for Internet backbone providers. Even more dire are their warnings about potential security and privacy problems for companies that deploy MPLS-based VPNs.

MPLS VPNs are a "great way to sell routers, but they greatly complicate the core of the Internet," Bush says. ...

Critics Bush and Bellovin claim MPLS is unnecessary because carriers can run frame relay or ATM traffic directly over an Internet backbone."

4.1.2.5).

- Reduction in the *Operational Matrix* as the control plane is merged with IP

So, while MPLS results in a Precedence Matrix which has much higher entropy than ATM as there are less precedence requirements, its operational efficiency in the presence of sub-optimal control policies (especially sub-optimal in the area of schedulability) is an overriding tradeoff.

Along these lines, the same analysis probably also applies to Lambda switching:

- Increases *Scheduling Complexity*,
- Is ambiguous with respect to *Operational Complexity*
- Probably increases overall complexity
- Enables the network to operate at higher levels of complexity

As a system response and at a higher-level of analysis, it is possible to interpret both MPLS and Lambda Switching in the context of Maximum Entropy Production: they are the networking communities reaction to increasing the entropy of networks *Traffic Vector* through higher structural complexity. See Section 6.1.2.

Keeping to the discipline, analysing MPLS vs. ATM:

**Operational Complexity** Higher, as intervening elements have to carry out more operations.

**Element Complexity** Slightly lower, as there is likely to be a higher requirement for dedicated ATM hardware over MPLS.

**Scheduling Complexity** Higher, as MPLS will have less precedence requirement over ATM.

**Traffic Complexity** Ambiguous.

probably leading to higher *Architectural Complexity* for MPLS over ATM.

This means that, MPLS is more suitable for networks that are:

---

## 5.2 Evaluating the Measurement Model

**Size** Larger than ATM networks.

**Control Policy** Networks which are harder to control optimally.

**Traffic** Networks which will carry variable traffic.

### 5.1.5 Concluding the Evaluation of the *Architectural Model*

One of the main goals of this thesis was to come up with a systematic way of thinking about architectural issues that would be easy to apply step-by-step, and would also lead to concrete proposals.

The *Architectural Model* seems to achieve this purpose. The existing body of literature which compares architectures is either too domain-specific, too technical, or does not use a systematic methodology which is easily exportable across domains. It is difficult to learn lessons from the analyses which can be reused.

This makes it hard for a network engineer to build his know-how of how networks behave, as experiences with different technologies cannot be translated into useful artefacts for making decisions in the future.

The step-by-step procedure by which *Architectural Model* breaks down a networking decision into its components and translates them into specific proposals is meant to overcome this shortcoming.

## 5.2 Evaluating the Measurement Model

The *Measurement Model* is a tool for quantifying the complexity and risk characteristics of a network. It uses *Reference Measurements* as a building block.

The methods of the *Measurement Model* relate to closed networks over which engineers have architectural control. Hence, the Internet, or other multi-homed networks which are under multiple control domains are not considered.

### 5.2.1 Source Data

The data used to illustrate the machinery in this thesis came from the corporate Wide Area Networks of Credit Suisse, an international Investment Bank. An overlay reporting functionality was built to collect and report daily data on the

---

## 5.2 Evaluating the Measurement Model

the categories above using sampling routers in five locations: London, New York, Singapore, Hong Kong and Tokyo. SNMP<sup>1</sup> was used to collect data at these servers and the data was stored in a central database. The data sampling rate was 144 samples/24 hours. Over a two year period, in every 24 hours epoch, reports were generated for:

**Performance** Consisting of Availability, Latency, Average Latency, and the Standard Deviation of Latency

**Utilisation** Consisting of Average Inbound and Outbound, and Peak Inbound and Outbound Traffic Measurements

**Errors** Consisting of Availability, Drops, Errors, Inbound and Outbound Discarded Packet

Data was collected on:

- an average of 240 links between 282 sites,
- a minimum of 98 links on some days
- a maximum of 282 links, and
- with a standard deviation of 25 sites.

For confidentiality and security reasons, data has been anonymised and data collection ends at least three months before submission of this thesis.

### 5.2.2 The Creation and Selection of Reference Measurements

The basis for utilising the *Measurement Model* in a statistical context are Bernoulli-(0,1) measurements. An element creates a trigger Bernoulli-(1) measurement when a threshold is breached. Hence, the native data sets above are not suitable

---

<sup>1</sup>The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. It is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite.

## 5.2 Evaluating the Measurement Model

Domain	Variable	Lower Threshold	Upper Threshold
WAN	Average Inbound Bandwidth	5	40
	Average Outbound Bandwidth	5	40
	Inbound Utilisation		70
	Outbound Utilisation		70
	Availability	99.9	
	Average Latency		50
	Standard Deviation of Latency		200
Firewall	Average Inbound Bandwidth	20	80
	Peak CPU Utilisation		90
	Average Availability	99	
	Average Latency		5

Table 5.1: Typical Thresholds for Network Measurement Data

for use as Reference Measurements, and need to go through a conversion process. For each data set, a threshold is chosen and values above the threshold for each element constitute a trigger.

Furthermore, the key variables in the context of the *Measurement Model* are probabilities. As the model is a 1-period model (in Section 7.1), each epoch in the data set, i.e. measurements over 24 hours, constitutes a sample of the percentage of elements in a hypothetically infinite-sized network which trigger a Reference Measurement: the number of elements above the threshold are counted and divided by the number of elements in total. This constitutes one sample, and the series of samples constitute the time series.

This lead to the issue of how to select the right:

- Network engineers typically determine thresholds to satisfy a quality constraint, either in the context of an interactive transfer of data between agents or in the context of bulk data transfer. Some typical values are:
- The threshold can be chosen using Sample Entropy, the negative natural logarithm of an estimate of the conditional probability that sub-series (epochs) of length  $m$  that match pointwise within a tolerance  $r$  also match at the next point.

---

## 5.2 Evaluating the Measurement Model

Sample entropy is a potential idea for selecting a threshold as it relates Reference Measurements to the “health” of the network. In the study of physiological time series, it has been observed that healthy heartbeats have high entropy and sample entropy is the technique used to measure such biological entropies (see (RM00)). As an analogy, a threshold with the lowest Sample Entropy could be the most informative about an inefficiency in the network.

In any case, the most reasonable choice is the former, values typically chosen by engineers: this relates the entropic discussion around values that matter. This is the methodology used in this thesis.

### 5.2.2.1 Using the Model

The analysis that follows uses Average Latency in the network with a threshold of 50 millisecond: hence, the time series consists of the percentage of all elements reporting average latency above 50 milliseconds over two years. Figure 5.1 and Figure 5.2 display the empirical Cumulative Distribution and Density of the time series.

The *Implied Complexity* of the network is calculated over the empirical density of the time series data using a 100-day sliding window. The implied density shows smooth tendencies for the networks complexity to gradually increase and decrease over time as the network adjusts to traffic patterns and is reconfigured post capacity and quality reviews. This can be seen in Figure 5.3.

The Complexity of the Network is measured as the Kullback-Leibler distance between the networks empirical density and model-implied density. To calculate this measure, Powell’s method (see (PTVF92)) is used to calculate the best fit *Measurement Model* coupling parameter to the empirical density using Kullback-Leibler distance as a minimising objective. The complexity is then the remaining distance between the two densities. Over a 100-day moving window, this can be seen in Section 3.2.

The minimisation procedure is remarkably efficient. Furthermore, it can be seen that the model does a fairly good job of explaining most of the networks behaviour. Once again, note that model entropy is higher than *Implied Complexity* (see Section 3.4.5:



## 5.2 Evaluating the Measurement Model

---

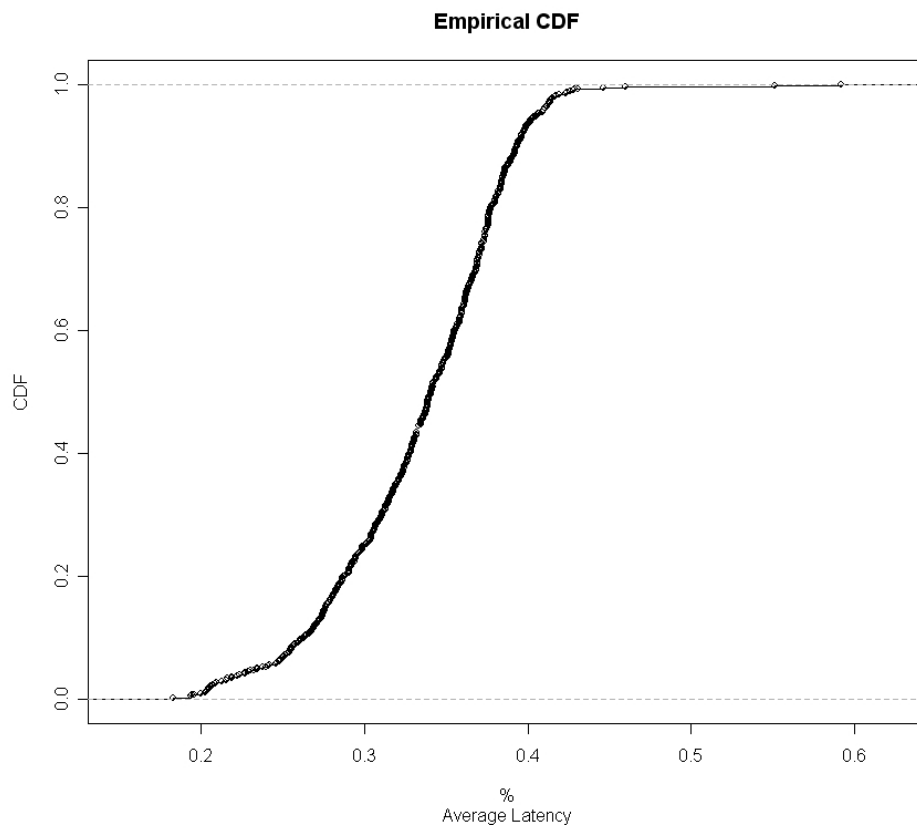


Figure 5.1: Empirical CDF, Average Latency  $50ms$

## 5.2 Evaluating the Measurement Model

---

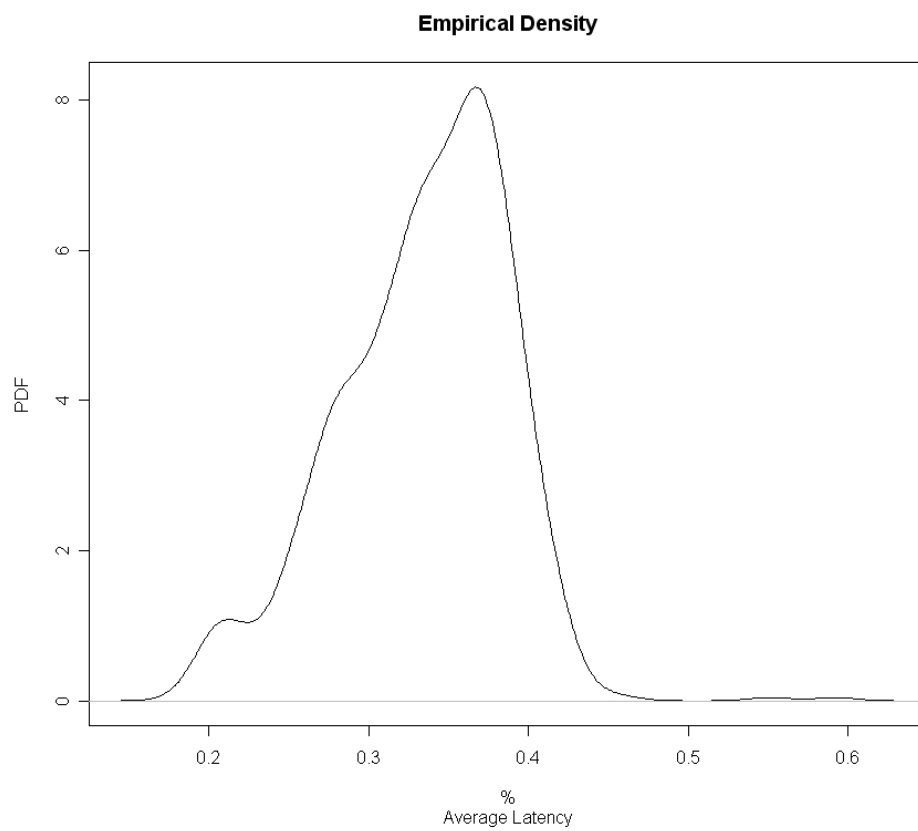


Figure 5.2: Empirical PDF, Average Latency  $50ms$

## 5.2 Evaluating the Measurement Model

---

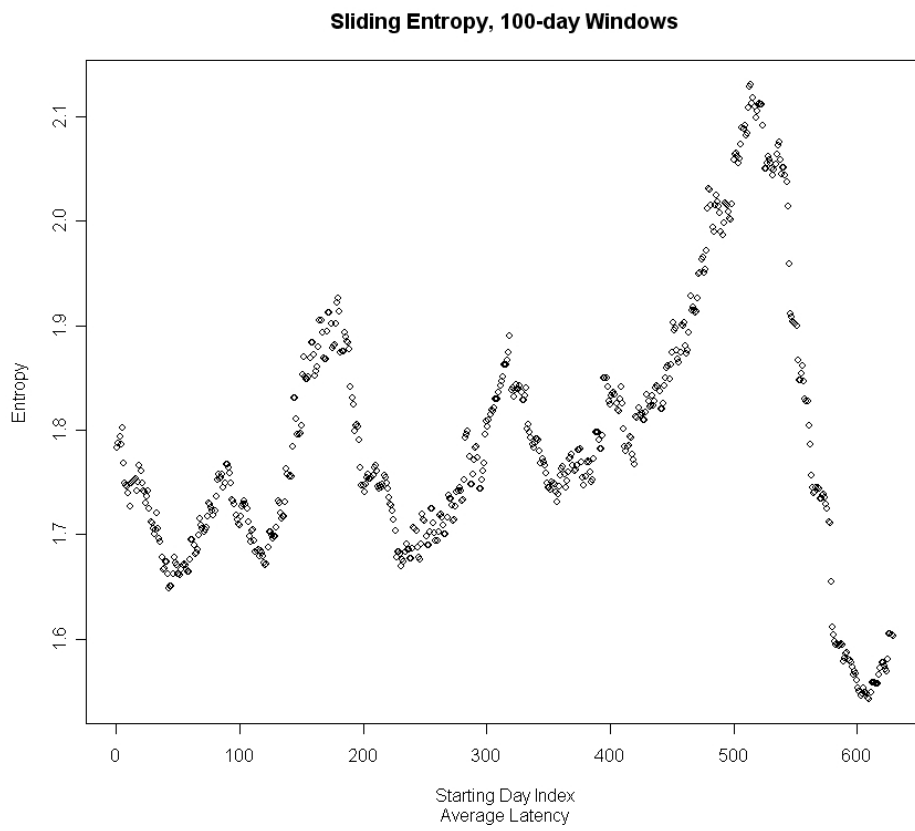


Figure 5.3: Sliding Entropy, Average Latency 50ms

## 5.2 Evaluating the Measurement Model

---

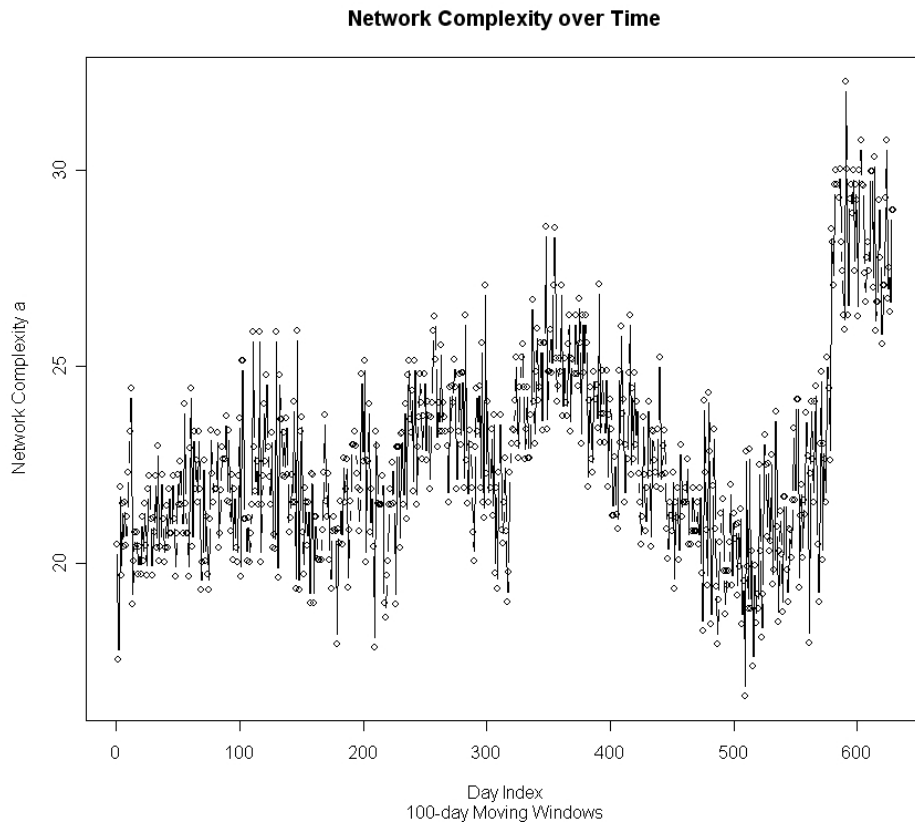


Figure 5.4: Network Complexity, Average Latency 50ms




Figure 5.5: Empirical and Model Entropy, Average Latency  $50ms$

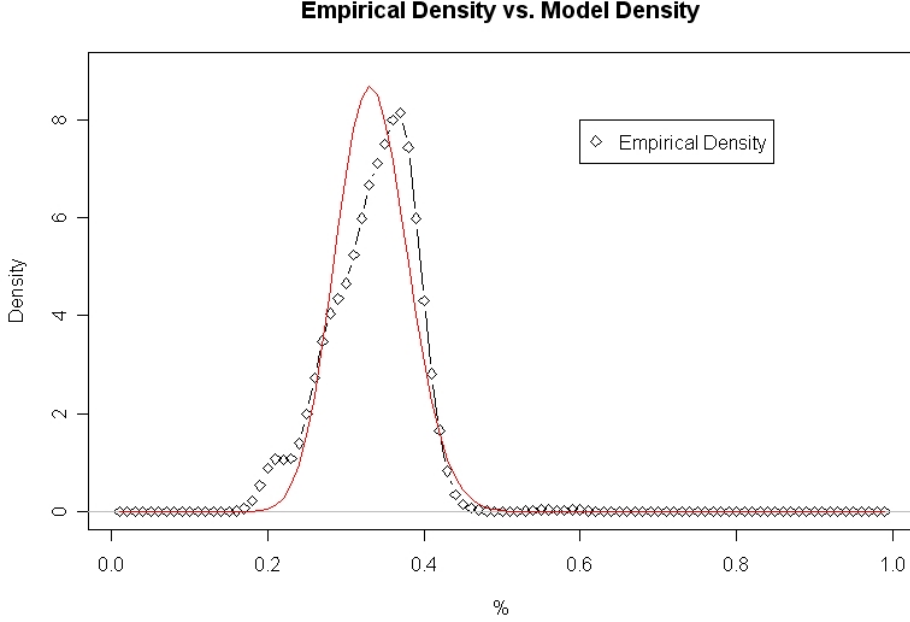


Figure 5.6: Empirical and Model Density Average Latency 50ms

An interesting parameter to observe over time is the *Implied Coupling* in the network. This is the parameter  $a$  in Equation 3.21 in Section 3.4.1, also defined in Section 3.5.4.

$$A_i = a_i M + \sqrt{1 - a_i^2} X_i \quad (5.1)$$

One potential validation of the underlying structural model proposed in Section 3.4 is to see if different *Reference Measurement* data imply the same coupling. This is illustrated in Figure 5.8.

### 5.2.2.2 Calculating the *Slice Curve*

Using Powell's Method (see (PTVF92)) on the sum-of-squared differences between model density from Section 3.49 and the empirical density of the *Reference Measurement*, it possible to solve for the shape of the *Slice Curve*.

What is more interesting is to take first derivatives of this curve, which is the instantaneous coupling rate for a very thin *Slice*, i.e. the coupling rate as it

## 5.2 Evaluating the Measurement Model

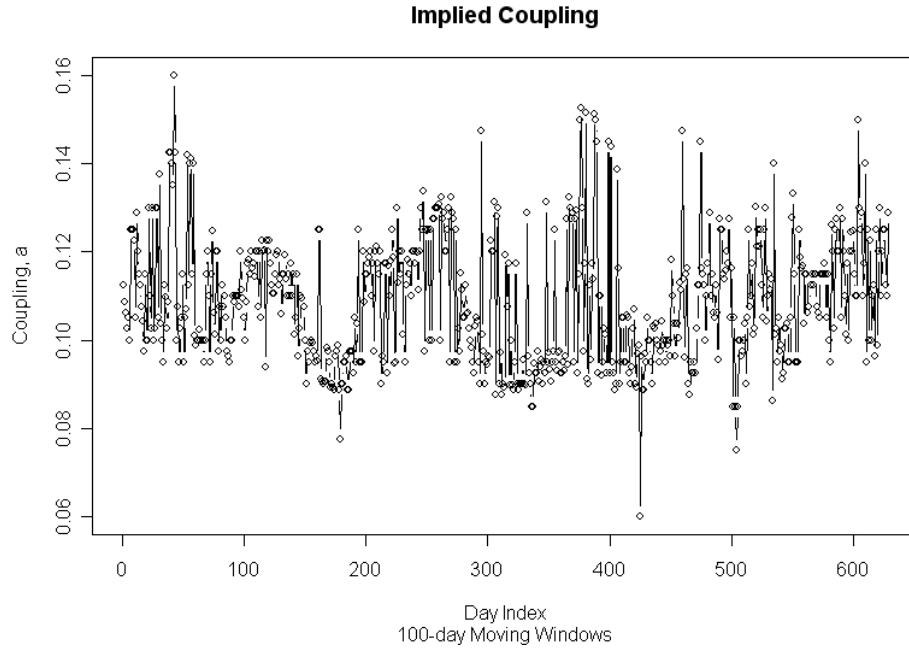


Figure 5.7: Implied Coupling Average Latency 50ms




Figure 5.8: Implied Coupling across Reference Measurements

## 5.2 Evaluating the Measurement Model

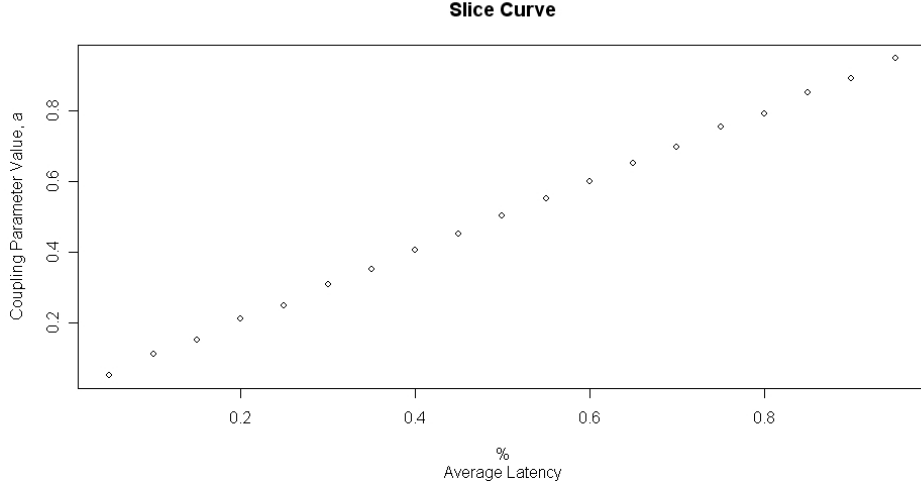


Figure 5.9: Slice Curve Average Latency 50ms

applies to  $Slice_{K_1, K_1+\epsilon}$ .

One possible assertion is that for traffic which can tolerate a medium level of risk, the *Implied Coupling* of the network drops: the network operates at lower levels of *Architectural Complexity*. For traffic with hard idiosyncratic constraints or hard systematic constraints the network does not need as much structural complexity.

### 5.2.3 Concluding the Evaluation of the *Measurement Model*

The evaluations carried out in this Chapters are not a conclusive and full evaluation of the tools designed in this thesis. A full analysis would cross-check interactions between the two models over time and across different network configurations.

For example, the *Architectural Model* would be used to predict whether one configuration is more complex over another, and *Network Complexity* measures could be used to test this assertion.

Furthermore, it would be interesting to see whether a single networks *Complexity* changes as it's architecture evolves over time to allow for new protocols and elements, such as firewalls, VPNs, and so on.

This was outside the scope of this thesis for the following reasons:



## 5.2 Evaluating the Measurement Model

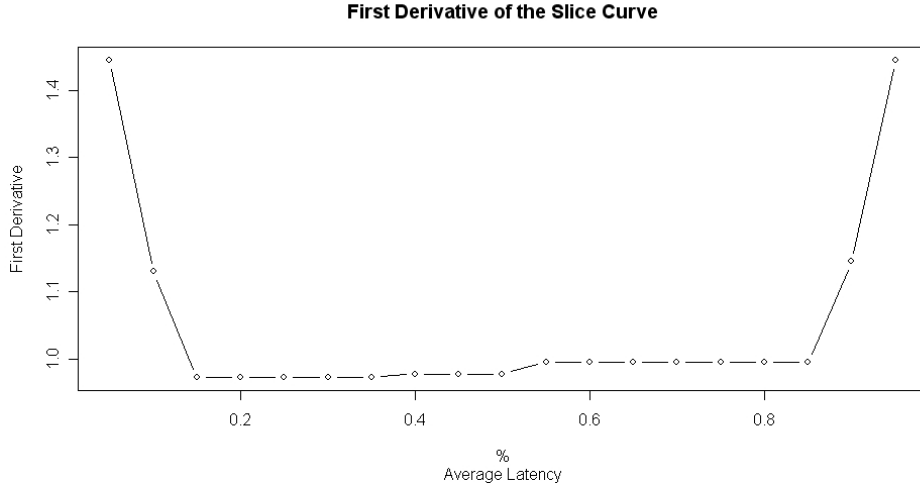


Figure 5.10: Slice Curve, First Derivative Average Latency 50ms

**Time** To consider the impact of architectural changes on a network, one needs significant time series relating to that network over all periods during which the architecture evolves. A typical medium-sized network will go through a major revamp once every two years at the earliest.

**Data** The techniques in this thesis, especially as they relate to the *Measurement Model*, apply to closed networks. It is very hard to get access to such data as most corporates are excessively protective.

Nonetheless, it has been shown that the complexity measures of Section 3.4 are easily calculatable and they exhibit interesting aspects of a networks behaviour over time. Specifically, the *Implied Coupling* parameters and the *Slice* curve relay significant insights about the networks behaviour over time and it's adaptability to different kinds of traffic.

# Chapter 6

## Implications

Network entropy will increase. This is the key implication of the entropic proposal of Section 2.2.1 and the trends outlined in Section 1.10. The increase in entropy will materialise primarily at the IP layer.

This can be seen as a steady progression: starting with the predictable laws of voice networking in the late 1950's, to the power law relationships and exponential, but predictable growth of the late 90's, to today's unpredictable data growth rates and bursty, fractal traffic associated with peer-to-peer and wireless networks<sup>1</sup>.

The rest of this Chapter analyses some of the far reaching implications for networks, networking research, and the financial marketing of telecommunications services from the *Entropic Formulation*.

### 6.1 Networking Implications

The *Entropic Formulation* of Chapter 2 and the mathematical tools developed in Section 3.4 have significant implications for the future of networking.

---

<sup>1</sup>Nonetheless, it must be noted that the key reason for utilising entropy is not the thermodynamic implications but the information theoretic requirement to account for interrogation costs. Increasing entropy is driven by increased efficiency and diversity due to the utilisation of rules of statistical inference by agents with informational constraints.

### 6.1.1 Entropic Routing

**Definition 23 (Entropic Routing)** *The use of a direct measure of the networks complexity and risk characteristics to make routing decisions.*

The arguments of Section 2.2.1 and the machinery of Section 3.4 indicate that information theoretic measures of a networks static complexity are invariants, and hence potentially a source of feedback in flow control: this leads to the idea of *Entropic Routing*.

Flow control<sup>1</sup> offers traffic to the network based on some information about the network state. For these purposes, *Slice* is a fast, easy to calculate parameter that can be fine-tuned as an objective in a feedback flow control system.

*Slice* has the potential to relay a rich characterisation of the state of the network, and can be used to match differing traffic patterns to different paths depending on their risk characterisations. For example, in the case of two traffic flows sending an equal amount of traffic, one bursty and time-insensitive, and the other a steady stream, the router would send the former down routes with *Slice* characteristics indicating low idiosyncratic risk, and the latter down routes indicating low systematic risk. Networks built to service bulk customers would prefer the former. Networks built to service retail would prefer the latter.

Minimising *Slice* with a high  $K_1$  would reduce the systematic risks in a network at the cost of increasing idiosyncratic risks, making the network more suitable for time-sensitive traffic. Minimising *Slice* with low thresholds would have the opposite effect, making the network more suitable for high-volume traffic.

Of specific interest in this context is the notion of *Complexity Neutral* routing decisions: in certain contexts, the network engineer may be interested in control decisions which do not impact the overall complexity of a network.

Here are some results an entropic ingress router could make in the presence of information about the shape of the *Slice Curve*:

- If a specific traffic requirement necessitated low idiosyncratic risks, a router may set-up a virtual path through the network in the absence of *Slice* measurements below a certain threshold (i.e. if the network was too risky from an idiosyncratic perspective).

---

<sup>1</sup>Also, sometimes termed feedback flow control.

- If a specific traffic requirement necessitated low systematic risks but could tolerate high idiosyncratic risk, the router may prefer to route the traffic without any dedicated resources.
- Traffic under long-term constraints would be routed using *Complexity Neutral Slice* parameters, to prevent degradations due to configurational changes.
- Long-term traffic under peak constraints would be routed along *Low Idiosyncratic Paths*.

*Complexity Neutrality* can be interpreted in two ways:

**Architectural Neutrality** These are configuration decisions which either do not effect *Architectural Complexity*. This can happen either through offsets between *Operational*, *Element* or *Scheduling Complexity*, or by making architectural decisions outside this domain.

**Network Complexity** These are routing decisions which do not effect *Network Complexity*: for small changes to the coupling parameter in the *Measurement Model* in the context of the calibration process, all decisions will be Complexity Neutral. This happens because the coupling parameter is the result of an optimisation minimising the KL distance between the *Measurement Model* and the realised distribution of the network<sup>1</sup>

*Complexity Neutral* routes and routing decisions are important for network control. For example, in the presence of sub-optimal control policies, entropic routers could direct their flows down Complexity Neutral routes.

Finally, the sub-systems extensions in Section 3.4.4 admits exciting possibilities for *Slice* as a candidate for feedback flow in BGP. This was already mentioned in Section 4.1.2.5 when considering the insights from the *Architectural Model* in the presence of sub-optimal control policies.

---

<sup>1</sup>This is not entirely correct: underlying this assumption is that Network Complexity is not a function of  $a$ .

### 6.1.2 The Red Queen Principle and Maximum Entropy Production

The Red Queen principle<sup>1</sup> is the term used to explain the low level competition between agents which resembles an arms race in a system with increasing entropy. The main assertion is that opposing forces created by innovation and competition bring macro stability to the system. Competition between agents increases the entropy of a system.

A typical response of a system out of equilibrium is structural complexity. This is a direct consequence of the *Principle of Maximum Entropy Production* mentioned in Chapter 2. Red Queen type competition is one way in which innovative forms and more complex organisms are created in large-scale systems. In the Schumpeterian economic setting, this usually surfaces in the form of monopolistic firms which innovate to create technological barriers as the competition draws nearer. This increases the rate of innovation in the economy in the form of structural complexity.

The increase in *Architectural Complexity* emanating from an increase *Element Complexity* could be seen in this light: by increasing the structural complexity of elements and operations in a network, vendors are reacting to increased competition amongst themselves, and amongst elements in a network for resources. MPLS, a stateful protocol which introduces significant control overheads, and feature creep in routers are symptoms of this predicament.

Feature creep can also be seen in light of the analysis in Section 4.1.2.5. According to this analysis, feature creep is equivalent to decreasing the number of zeros in the *Network Precedence Matrix*, leading to increased *Architectural Complexity* in the network. In this way, feature creep is a systems response to increased variability in the traffic mix, as posited by the *Principle of Maximum Entropy Production* for dealing with non-equilibrium states. Additional features

---

<sup>1</sup>Lewis Carol has the following passage in his book *Alice Through the Looking Glass*, which has coined the term *Red Queen principle* : *Well in our country said Alice, still panting a little youd generally get to somewhere else if you ran very fast for a long time as we been doing. A slow kind of country! said the Red Queen. Now here, you see, it takes all the running you can do, to be in the same place.*” See (CM95)

in a router increase the dimensionality of the network and make it possible for a network to attain higher levels of entropy sooner.

Increasing static complexity can also be seen as a control response to increasing variability in the *Traffic Vector*: under the assumptions of increasing entropy in network traffic, increasing structural complexity is the natural response of an operator with fixed efficiency targets.

As before, this is a cycle: as the underlying infrastructure of the network, such as Lambda switching, increases the entropy of the network, existing devices built around assumptions based on predictable traffic (such as limited buffer size) will become inadequate. This leads to *an arms race*, increasing the segmentation and structural complexity of the system.

### 6.1.2.1 Separating Routing from Routers

A potential implication of increased entropy will be that routing control and packet level processing functions will be forced to operate at lower and lower levels of granularity. This trend is at odds with future speed requirement which leave roughly  $8ns$  for each packet lookup in a router with current technologies. The computation and speed requirements of optically switched networks can only be met if the routing function and the forwarding function are separated and become highly specialised, meaning that their complexity will decrease. In (FBR<sup>+</sup>04), the authors argue “*the Case for Separating Routing from Routers*” based also on increasing complexity, even though theirs is not a holistic view as the one espoused here.

### 6.1.2.2 End of the End-to-End Argument

The *End-to-End* argument is a design principle: it suggests that the actual code for an application and services required to execute it should be moved upward in a layered design. This has the effect of increasing the flexibility and autonomy of the application designer.

While the added flexibility is welcome, especially during periods of high-growth when there is significant ambiguity about which applications will be needed, the mechanics of *Architectural Complexity* imply that the *End-to-End*

argument implies networks which are much more complex: the size of the *operational* and *network precedence matrices* must be much larger than otherwise due to the layering of the network and the requirement that network elements are not able to specialise.

This means that the *End-to-End* argument brings with it significant increases in the complexity of the lower layers of the network.

What is the impact of this design decision on network performance ? Systems with high complexity do not operate well in the presence of sub-optimal control policies (see Section 4.1.2.5). Furthermore, it is increasingly difficult to calculate optimal control policies when network entropy is increasing, an eventuality for all networks, whether due to competitive pressures, or Red Queen dynamics. This implies that there is a key weakness with the *End-to-End* argument, and that it is not a design decision suitable for all contexts.

The end of the *End-to-End* argument has been announced many times in the past. While this is not the stated goal of this thesis, a systematic analysis of this design principle with respect to network performance can cast a shadow over its suitability in contexts where manageability is important and sub-optimal performance can be costly. The re-emergence of connection-oriented protocols in the form of MPLS, the de-layering of network protocols, and feature creep in routers can be seen as one manifestation of this reality,

### 6.1.3 Network Planning

Current networking practices separate network planning from traffic engineering (Section 1.7). Time scale differences between capacity planning and facility location is the main reason for this. Increasing entropy means traffic patterns will become fractal, and differences between time scales will disappear. This could mean that in the future, the network planning and traffic engineering functions in telecommunication firms will merge.

### 6.1.4 Element Design

In Section 2.3.1, the bijection between entropy and log utility maximisation was pointed out. Agents which maximise entropy also maximise their long-term

growth rate. In a closed system, they will dominate with probability one. This means that only network elements utilising some form of statistical inference based on the principle of maximum entropy (see (Jay57)) will be viable. When designing new elements for a network which must compete for resources, the designer must take this into account: unless elements use some rudimentary form of entropy based statistical inference for competing in non-cooperative environments, they will not be able to secure the resources required to carry out their duties.

### 6.1.5 Infrastructure

Complex infrastructure in real-world markets exists as a reaction to increased entropy and to ease the computational burdens of participating agents:

- *Regulators*: Decrease computational burdens on agents by dictating best practices.
- *Brokers*: Decrease computational burden, and increase entropy through data dissemination.
- *Derivatives*: Increased entropy through security complexity.

In the future, the infrastructure required to maintain and operate networks will also become increasingly complex. This has already happened in the Internet, through the formation of ICANN<sup>1</sup>, its increased pro-activity in regulating the Internet, and the system of 75 registrars who finance its budget. This is all a sign of increased structural complexity as a response to increasing entropy in the Internet. Social, regulatory and technical complex adaptive systems aimed at reigning in the increasing entropy of the Internet should be expected to increase in the future.

---

<sup>1</sup>The Internet Corporation for Assigned Names and Numbers (ICANN) is an internationally organised, non-profit corporation that has responsibility for Internet Protocol (IP) address space allocation, protocol identifier assignment, generic (gTLD) and country code (ccTLD) Top-Level Domain name system management, and root server system management functions.



## 6.2 Financial Implications

### 6.2.1 Insurance

In Section 1.6.1, it was argued that trading point-to-point capacity was a misguided attempt at creating markets in telecommunications, and that this was why the markets for traded bandwidth never materialised.

There were also other problems. In the financial markets, when large-scale risk are identified, risk transfer does not usually begin to take place initially through trading. The first step for risk transfer in the financial markets usually begins with insurance<sup>1</sup>. Insured risk sometimes evolves and becomes tradable, but this is not always guaranteed.

Furthermore, due to the implications of CAPM, insurance markets typically provide of protection for systematic risks, and not idiosyncratic risks (see Section 3.4.5.3).

From this perspective, it can be seen that the markets for telecommunications not only started out in the wrong order, they started out with the wrong risks: the idiosyncratic risk associated with the price of a single point-to-point link matters even less than the overall price of capacity for most large scale telecommunication operators: their sheer size means they are hedged against single point prices for free due to diversity in their networks, and would not pay a premium to hedge risks associated with point-to-point bandwidth prices.

Insurance providers rarely provide coverage for first loss risk: a provider would then be incentivised to take out insurance in parts of the network they know asymmetrically to be weak. this problem is identified as “*asymmetric information*” in economics.

---

<sup>1</sup>A good current example for this procession is the market for variable annuities, which are essentially securities linked to the equity markets that have minimum guarantees. The minimum guarantee creates a “gap” risk for the seller of the annuity, a risk which was initially hedged in the reinsurance market, not in the equity derivatives markets. However, since the beginning of 2005, protection on this risk has moved to the traded market for equity derivatives, as the price of reinsurance has increased. The increase in reinsurance premiums has happened because variable annuity sellers have been forced by regulators to cover a larger portion of their gap risk than was previously mandated.

The true, und hedgeable, systematic risks faced by providers of telecom services are:

- cascading failures<sup>1</sup>
- operational failures to deliver on large scale contracts
- persistent congestion

The solution to this is subordination: insurers are never interested in insuring first-loss risks due to asymmetric information problems, and providers are not interested in buying first loss risk because they can hedge those risks through diversification.

To calculate the right levels of subordination, one needs a way to work efficiently with the multivariate distributions of errors and failures realised in large scale networks. It turns out, *Slice* measurements characterise multivariate distributions directly in terms of subordination, and hence can play the symbolic role of a numéraire in this market. They can be used to value insurance contracts that cover these risks. It is possible that, over time, *Sliced* risks from a pool of providers could become a tradeable asset<sup>2</sup>.

### 6.2.2 Impact on Strategy

If a market in network insurance materialised, the financial structure of most service providers would change substantially. Currently, large service providers have to “insure themselves” through size, as there is not a market for network insurance. A market for network insurance would release significant amounts of capital and could potentially create much smaller firms.

---

<sup>1</sup>Typically leading to losses that can scale to \$ Billions.

<sup>2</sup>Note that, even with a subordinated contract (where the telecom provider takes first losses and the insurance provider covers losses beyond a certain point) is still subject to significant conflicts of interest when the network provider can game the underlying network in his favour. Through element selection (see Section 4.2.1) and network configuration (see Section 4.1.2) the provider can increase systematic risks, borne by the insurer, and reduce idiosyncratic risks. Any insurance market would have to contract out such risks.

## 6.3 Murphy's Law of Networks

Entropy is the result of increased efficiencies in the network, in the form of rapid provisioning and multi-source/destination traffic. Current technical developments driving higher efficiency include multi-path routing <sup>1</sup>, optical switching, and multi-homed applications. They are also increasing the entropy of the network, as routes are set-up and torn down with increasing frequency. This is a cycle: increased efficiency in the network increasingly attracts interactive, bursty, voluminous traffic, leading to a sort of *Murphy's Law* for networks:

**Principle 4 (Murphy's Law of Network Traffic)** *Any predictability in data traffic will decrease over time as network elements increase disorder through increasing efficiencies.*

There is already significant empirical support of this hypothesis. For example, in (XZB05), the authors mention that there is an:

- increase in unwanted traffic
- emergence of disruptive applications
- emergence of new services on traditional ports, and the
- emergence of traditional services on non-standard ports.

In any case, high entropy in a network should probably be interpreted as a good thing. It has been known for a long time that heartbeats characterised by high entropy (see (Iga03)) belong to healthy hearts: similarly, the best way to measure the health of a network is probably to measure its entropy<sup>2</sup>.

---

<sup>1</sup>In (KV05), the authors show that a decentralised sufficient condition for the local stability of end-to-end algorithms for joint routing and rate control is restricted by the round-trip time of that route alone, and not other routes. If such a guarantee could be given, i.e. round-trip time measurements stability, scalable load-sharing across paths would be possible.

<sup>2</sup>“Internet performance is improving each year with packet losses typically improving by 40-50% per year and Round Trip Times (RTTs) by 10-20% and, for some regions such as S. E. Europe, even more. Geosynchronous satellite connections are still important to countries with poor telecommunications infrastructure. In general for HEP countries satellite links are being replaced with land-line links with improved performance (in particular for RTT).” [January 2005 Report of the ICFA-SCIC Monitoring Working Group](#)

# Chapter 7

## Conclusions

This thesis is an attempt at changing the way we reason about networks in a fundamental way. So far, the existing body of networking research has focused on the network as the side-effect of interacting elements. This thesis proposes that one should instead focus directly on the macroscopic features of networks directly.

When a market is not complete, there is no single valuation measure for making decisions, and agents risk preferences matter. How does one decide which valuation measure to use when there is a continuum of alternatives ? There are two alternative solutions to this problem:

**Utility-based** Choose a representative agent, and assume the market sets prices to maximise the utility of that agent.

**Pricing Kernels-based** Choose a valuation measure which is not directly related to the utility preferences of individual agents, but is related to the statistical properties of asset prices<sup>1</sup>.

One is left with the dual problem of which utility function to choose in the case of the representative agent whose utility is to be maximised, or which ad hoc criteria to choose as a valuation measure for making decisions. The neoclassical set-up uses the former methodology. This thesis defends the use of the latter methodology, argues why the entropic criteria amongst others is the right one for networks, and proposes various ways in which this criteria could be adapted to

---

<sup>1</sup>A pricing kernel is the stochastic discount factor in a model of asset prices.

---

networks. It then pushes the boundaries of this assumption to arrive at various insights about the complexity of networks using two practical models.

Why the latter methodology ? The assumptions underlying the neoclassical paradigm are very unrealistic, and this is one reason. However, many models make unrealistic assumptions, and these can be corrected in later extensions: the *Measurement Model* is one good example.

There is a much more fundamental objection to the neoclassical paradigm which is based on the enormous information requirements associated with the equilibrium it proposes: as the entropy of consumption bundles is zero (all agents of the same type consume the same bundle), the neoclassical paradigm would require that all the information available in the network would be required at all times for equilibrium to be maintained. This is not only a difficult task, but also means that the neoclassical paradigm has a vanishingly small probability associated with it, amongst all possible equilibriums. This information objection holds for any network model which is based on the core assumptions of the neoclassical paradigm, even if other issues are addresses. This argument is a core contribution of this thesis.

Hence, the objection is not just in the small details, but is of a very fundamental nature. This also makes clear why the alternative proposed in this thesis is perhaps much more defensible: by comparison (and by construction) the maximum entropy equilibrium is the most likely equilibrium amongst all equilibriums.

That one should focus directly on the statistical properties of network characteristics as opposed to the side-effects of toy models based on utility maximising agents is the basis for arguing that the research agenda needs ReReshaping.

The two models developed to analyse real-life networks and architectural decisions are the key mathematical contributions of this thesis. They:

- illustrate that some of the key emergent characteristics of the complexity of a network could be tied to notions of entropy, and
- show that highly tractable models could be built to proxy both the complexity and the riskiness of a network.

The goal of this thesis was to create a systematic methodology for considering networking issues. The result is a number of insights, and the implications of

---

the methodology that has been built. The key results of this analysis can be summarised as:

**The Entropic Formulation** Researchers should ensconce their modelling around an *Entropic Formulation* of the networks configuration, on the grounds that:

1. The Micro-foundations of a zero-sum game between competing agents implies a maximisation of a systems entropy.
2. Any system which tries to achieve maximum informational efficiency will trend towards higher entropy.
3. A probabilistic analysis of likely configurations of a network make available the rich toolkit of statistical mechanics, which also imply that the most likely configuration will be that which has the highest entropy.

**Important Complexity Tradeoffs** The *Architectural Model* is a rich source of insights about the complexity of networks.

1. When the number of *Elements* or *Operations* increase, the complexity of a network increases.
2. When the *Scheduling Requirements* of a network decrease, its complexity increases. This could happen either because more *Elements* can carry out more *Operations* or the sequencing requirements are decreased.
3. Smaller networks are more sensitive to changes in the complexity of a network. These changes can be *Operational*, related to *Elements* or to the *Scheduling* requirements in the network.
4. Increasing complexity can increase resilience and/or performance.
5. Architectural changes will impact networks operating at lower levels of complexity more than networks operating at maximum complexity. This means that architects should direct their resources towards sub-systems with low complexity for maximum engineering impact.

---

This last point perhaps also relates to the direction that future networking research should take: large, complex existing systems, like the Internet, will not show large benefits from engineering. They have already reached a high-level of efficiency, perhaps evidenced by their complexity. The academic community may be better served if research were targetted at smaller, less complex networking models, where the benefits from engineering and new techniques may reap larger rewards.

**Performance and Complexity** While some of the insights around complexity tradeoffs are perhaps obvious, placing them into the networking context and seeing how they could be applied to reasoning about architectural issues is a core contribution of this thesis. However, what is not so obvious are the two key insights related to network performance and complexity:

1. The increase in network performance from increasing complexity shows diminishing returns.
2. The performance degradation from sub-optimal control policies are higher for more complex networks.

This means that the network engineer must evaluate carefully the tradeoffs between network performance and complexity.

**Risk** Associated with the complexity of a network are risks. Large systems are open to two kinds of risks, *idiosyncratic* and *systematic*, and they relate to the complexity of a network in different ways.

1. Highly complex networks face idiosyncratic risks.
2. Networks with low complexity face systematic risk.

**Implications for Networks** The *Entropic Formulation* posits a number of important propositions for the future evolution of networks and proposes alternative explanations for why certain complex phenomena emerge:

1. The Entropy of networks, in the form of network traffic and structural complexity, can be expected to increase over time as they become more efficient.

2. Entropy maximisation reveals that there are *Entropy prices* associated with the constraints in a network. The analysis indicates that configuration changes which cause large changes in the entropy of a network are less likely.
3. Networks may display structural complexity either in the form of emergent phenomena, or in the institutions they maintain in order to function smoothly. These structure allow networks to attain Maximum Entropy states faster.

As a final point, we point out an alternative interpretation of the *Implied Coupling* metrix. Every experienced engineer knows that how tightly a network is coupled together impacts performance, resilience and complexity. However, it is generally very difficult to characterise and measure the actual coupling in a network, either because the information is simply not available due to data limitations, or it is hard to interpret the various alternative graph-theoretic measures, like clustering and transivity, directly in terms of performance and complexity. As an alternative, the *Implied Coupling* measure in this thesis can be interpreted as the *effective coupling* of the network, rolling all the information about the networks performance including linkage, resource sharing, co-dependent traffic flows into one, simple figure which is comparable across networks.

## 7.1 Future Work

There is significant room for expanding and validating the methodology in this thesis, in the form of interesting and practical work for the future.

### 7.1.1 Empirical

While the data in this thesis illustrates that the *Measurement Model* is useful for characterising network complexity, it would be useful to correlate complexity measurements with actual network changes.



Another obvious area for further research is correlations between network decisions and rises and falls in idiosyncratic and systematic risks. For this, *Slice* measurements need to be regressed against network decisions. Such data is also generally hard to come by.

Further model validation is an ambitious research agenda. For full validation of the proposals in this thesis, significant data collection in the setting of a large autonomous system over significant periods of time would be required. Furthermore, as this is a thesis about architecture, these measurements would have to be crossed with architectural changes. This data requirements was outside the scope of the resources available during the writing of this thesis. Small laboratory networks are not suitable for this kind of work: the large, homogeneous approximations underlying the mathematical relationships in this thesis require relatively large networks for the numbers to have significance.<sup>1</sup>

Of all the findings of Section 5, the most interesting one probably relates to the shape of the first derivative of the *Slice Curve*. While this is just an assertion (and hence, potential future work), it could be that the shape of this curve has something to do with how flexible a network is. It would be interesting empirical research to relate the shape of this measure to network configurations and technology which enabled automatic reconfigurability.

### 7.1.2 Routing

There is obvious research in the area of control related to the use of entropy to make routing decisions. The issues that would require further analysis include:

**Entropy Flooding** A few routers using entropic routing may cause instability in a network in the presence of other routers which do not. This is because

---

<sup>1</sup>One key problem hindering this is security: most corporate networks abide by strict security and confidentiality requirements making it practically impossible to cross network configuration data with actual performance, error and utilisation data in the context of an academic article. It would be interesting to see if an avenue for this research could be formulated. Even when there is permission for the data to be utilised, there is a limit requiring that data used not be current (hence, the data in this thesis is at least three months old).

the behaviours of the routers which do use entropic routing would speed the ascent up the entropy gradient. While the entropy-enabled routers would be able to adapt to the increase in the entropy of the network, other elements in the network would be “driven to extinction”.<sup>1</sup>

**Interaction Between Complexity and Risk** It is possible that entropic routers routing on the basis of an objective function which aims to minimise idiosyncratic risks in a system coupled with low complexity could cause instability.

Furthermore, data collected for the purposes of this thesis was relatively coarse. This thesis makes the claim that it is possible for the mathematical techniques in this thesis to be the basis for even fine grain routing of data flows in a network. That claim would require testing with higher frequency data, more than the 144 samples collected for use in this thesis. This is another interesting area of research in the future.

### 7.1.3 Modelling

There is always useful work to be done in the direction of extending the *Measurement Model* further. The obvious ones are:

1. Adding a time dimension to the model: the model in Section 3.4 is a one period model.
2. Modelling traffic variability explicitly
3. Changing the coupling specification in the model from a Gaussian to something more elaborate

The final point relates to the parameter called *Implied Coupling* In Section 3.5.4 and in Figure 5.7. This parameter is an indication of how tightly the network is coupled together over the many resources that are shared, and the operations that are executed on multiple elements. In the context of the Measurement Model,

---

<sup>1</sup>This relates to the point made in Section 2.3.1, that in a closed system with, entropic agents will dominate the actions of all agents.

this parameter has a very simple interpretation: it is the flat correlation in a multivariate Gaussian. Extending the model to incorporate non-linear dependencies would be an interesting avenue of research,

While these are interesting extensions, they must be noted with the caveat that one strength of the methodology is the use of a simple, yet structural model (see Section 6). Extreme extensions to the model risk squandering the gains from simplicity to “feature creep” in modelling requirements.

### 7.1.4 Applying the Entropic Formulation

In Sections 2.2.1.1 and 2.2.2.2, an outline was given for how the *Entropic Formulation* could be applied to actual networking problems. This is perhaps the most interesting area for future research. The Kelly and Gibbens papers apply the minimax methodology to networks, but the entropic implications are not really studied. It is easy to imagine many other contexts where the minimax, zero-sum game theoretic formulation makes sense for networks, such as in denial of service attacks. Again, it appears that the entropic angle has been missed by most authors.

It is surprising that more has not been made of the statistical mechanics approach to networking evidenced in this thesis, especially as there has been significant crossover between the two fields of telecommunications research and the physical sciences in many adjacent sub-fields. Could the diversion provided by the neoclassical paradigm be to blame? In any case, as networks grow in size, become reconfigurable and the elements on networks acquire more “intelligence” through the use of statistical inference to make control decisions, it is inevitable that network researchers will be forced to turn to statistical techniques to make predictions about how their networks will behave. This is the most interesting avenue for future research into networking as a result of the work in this thesis.

### 7.1.5 Looking Back over the Fence

There is also interesting research in other fields which could reverse the direction of information flow in this thesis. This thesis was built on work from neighbouring

fields: manufacturing and finance. The key machinery used from finance was the actuarial practice of making homogeneity assumptions for credit risk and measuring risk in terms of idiosyncratic and systematic risks.

A key limitation of those models, already highlighted, is the assumption of static trigger probabilities. Networking research using queueing in stochastic environments is one potential way to deal with this. Current approaches to market microstructure research such as herding behaviour, order clustering, and long-range dependencies in stock prices could also benefit from models based on queueing theory. It would be interesting to take existing networking models to see how they could be applied in return to the fields of finance and manufacturing.

# Appendix A

## Rabbits and Traps

Work inspired by this thesis led to this essay titled “Rabbits and traps”, which won second place in the The John Rose Prize, for “the best explanation of a scientific principle of general interest”.

### Rabbits and traps

Imagine you are a rabbit, living a luxurious life at the edge of a carrot farm. After raiding the farm for many months, the farmer has finally had it: he is not taking any more chances. He is going to lay some traps.

What would you prefer ? That he lay all the traps densely in one section of the farm, or that he spread them out evenly across ?

If you are one single rabbit, you prefer that he lay the traps bunched up together in one section. Most of the time, you will probably be OK. If you hit one trap, you hit them all. If you hit none in one section of the farm, chances are there are none anywhere around.

Things change if you are a large herd of rabbits. In this scenario, you would prefer it if he laid them out randomly. However small, there is a chance that as you go carrot hunting, none of you will hit any traps. However, if all the traps are bunched up together in one section of the farm, it is almost certain that some of the rabbits will end up caught in the section with the traps..

---

This simple analogy describes an important concept which is increasingly important in our lives: the impact of correlations on large complex systems.

First things first: what is a large complex system ? You get one of these whenever there is more than one trap, rabbit-related questions need answers, and simply counting the number of traps and rabbits does not give a good enough answer. Just a simple headcount is not enough to answer how one can be sure at least a certain number of rabbits gets caught every month: one also needs layout information. The difference between large complex systems and large simple systems is whether basic reductionist techniques suffice when answering important questions, or whether one also needs information about the system as a whole.

Correlation itself is more subtle. Its about the likelihood of a rabbit being caught if another rabbit has been caught. In the example above, when all the traps are side-by-side, correlation is high. If one rabbit has been caught, chances are that other rabbits in the vicinity are also suffering. On the other hand, low correlation happens when the traps are randomly dispersed. Even if you are dancing with the reaper, the likelihood of a a rabbit happily munching away next door may not change then.

### **The Subtleties: Correlation neutral rabbits**

Lets complicate things for Mr. Wabbit. First of all, it is clear that there is an inflection point at which the dispersion of traps and the likelihood of a rabbit being caught turns. At one extreme, traps are entirely randomly dispersed and the likelihood of being caught by a trap is unaffected by whether the rabbit next door is suffering in the clutches of a manacle. The other extreme makes it certain that if your buddy has been caught, you are doomed. Most likely, no matter which direction you hop, there is a trap.

This is the case for rabbits hopping around independently of each other. However, what about small groups of rabbits ? As the rabbits munch in groups of two or more, interestingly enough, there is a point at which how correlated the traps are makes no difference: the dispersion of traps offsets the correlation of rabbits ! We can call these, correlation neutral rabbits !

This brings us to the first lesson to be learnt from rabbits: if one is facing an adversary who has a limited number of traps but can vary their dispersion

---

up and down to increase his chances of baiting , there is an optimum, medium group size for munching which makes one invariant to how the traps are layed out. Small groups are better then being alone and being together.

### **A common misunderstanding about rabbits and traps (and correlation) !**

Most people who have ever read anything about major large scale mishaps have some intuition which tells them that very high correlations are bad. Whether it is the 1987 market crash, the 1998 Long Term Capital Crisis, the electricity blackouts in California, or pandemics, they know that when lots of things become highly correlated with each other, bad things can happen simultaneously to everyone. You dont need a silly story about rabbits to tell you that !

However, the fact is, in the real world, there are usually only a limited number of traps. Herein lies the main misconception. While it is true that all the traps may set simultaneously one day, this would mean that on other days, they were not setting at all. This brings us to the second lesson: if one concentrates on trying to prevent large scale losses, one is probably paying the price through small losses which add up over time: a rabbit here and a rabbit there.

The focus on correlation is partially skewed by the reporting of extreme events. No-one likes to report that one small rabbit got caught here and another rabbit caught there. Much juicier if there is a farm where 30 rabbit all got caught in the same day. Books are written about tulip crazes, not about the seemingly random (but regular) losses by small flower growers occurring independently of each other over long periods of time. The reporters focus on one end of the extreme and ignore the other. While it is true that the events surrounding the Russian default in 1998 were disastrous, they were preceded by many years when things went fabulously for most market participants.

In rabbit-speak, this is what happens when the farmer puts all his traps in one place. There will be large swathes of time when no rabbit get caught. Then, one day, all the rabbits get caught together. The opposite would be a case where the farmer caught one or two rabbits with certainty almost every month, but never all of them.

### **Who cares about rabbits ?**

---

We should all care about rabbits. We lead lives in complicated circumstances where the likelihood of something important happening to us is a function of correlation and large numbers.

If you are a sole ADSL user and connectivity matters, you want to ensure that the traps are bunched up together at BT. You want to be sure that the smart engineers there have spent all their money on ensuring that single users are served well and the likelihood that you will hit a spurious fault which effects your line and only your line is low. This of course will mean that they spent less money on ensuring that the whole system never ever goes down, as they have limited resources.

On the other hand, if you are a large corporation, you want the exact opposite: you prefer it if BT spent their money to ensure large scale outages which effects many users are unlikely. If one of your many offices is losing connectivity, chances are the other will be up and running and can take some of the excess load from loss of connectivity to a small portion of other users.

This is the third lesson to be learnt from rabbits: While being small and simple is nice, in certain cases there are advantages to being large and complicated when one has no control over small, random mishaps.

Tolga Uzuner  
PhD Student  
Computer Lab  
Kings College  
Box 441  
itu20@cl.cam.ac.uk



# Appendix B

## Entropy

The entropy of prices,  $H(g) \equiv H(F_T|g)$  of the random price  $F_t$  under  $g$  is defined as:

$$H(g) \equiv - \sum_i g_i \log(g_i) \tag{B.1}$$

in (CT91).

Broadly, entropy is a measure of the information content in a probability distribution. Probability distributions with higher entropy have lower information content, the meaning of which is formalised in a number of settings broadly as to mean the number of bits required to transmit a certain message through a communication channel when words occur with a certain probability. A language with low entropy can be transmitted with a smaller number of bits<sup>1</sup>.

Entropy is not randomness: seemingly chaotic messages can have very high or very low entropy, and there is a vast literature dedicated to this topic. (GM94) book covers these issues in detail, and in Chapter 2, we will also provide a brief discussion of the relevant issues.

The estimation of the distribution of an underlying asset from a set of market prices using the principle of maximum entropy is least committal with respect to unknown or missing information and is hence the least prejudiced. As a matter

---

<sup>1</sup>In this context,  $\log$  has base 2.

of fact, the maximum entropy distribution is the only information about assets that can be inferred from price data alone.

Relative entropy and the Principle of Maximum Entropy admit a dual interpretation: the relative entropy between the Uniform Distribution is the implicit objective function one minimises when one maximises entropy. This is also why maximising entropy is sometimes referred to as a smoothing procedure.

## B.1 Maximum Entropy and Relative Entropy

The Kullback-Leibler (see (KL51)) cross-entropy measure relates distributions through a closeness metric:

$$I_{p-\dot{p}} = \sum_j P_j \log \frac{P_j}{\dot{P}_j} \quad (\text{B.2})$$

If two distributions are identical, this measure returns 0. In many contexts, the entropic approach to economics is posed in this measure, as opposed to maximising entropy. For example, (AFHS97) minimise cross-entropy to an a priori distribution, allowing the inclusion of prior knowledge of the asset distribution. The a priori is typically a moment constraint in the form of past realised volatility, skewness and kurtosis in prices.

However, it is easy to see that maximising entropy in the absence of a prior is equivalent to relative entropy methods when the prior is set equal to the uniform distribution, which has the highest entropy amongst all distributions.

Because the  $-\log$  function is convex,  $\text{Avg}(-\log(\frac{p_i}{q_i})) \geq -\log\left(\text{Avg}\left(\frac{p_i}{q_i}\right)\right)$  due to Jensen's Inequality. This implies:

$$I_{p-\dot{p}} \geq -\log \sum_i \left( q_i \frac{p_i}{q_i} \right) = -\log \sum_i p_i = -\log 1 = 0 \quad (\text{B.3})$$

meaning that the Relative Entropy is strictly non-negative and equal to zero if and only if  $p = \dot{p}$ .

## B.2 Axiomatic Formulations of Entropy

In his original text, Shannon used an axiomatic justification for the *Entropic Formulation*. His axioms were (see (CT91)):

- Normalisation:  $H_2(\frac{1}{2}, \frac{1}{2}) = 1$
- Continuity:  $H_2(p, 1 - p)$  is a continuous function of  $p$ .
- Grouping:  $H_m = (p_1, p_2, \dots, p_m) = H_{m-1}(p_1 + p_2, p_3, \dots, p_m) + (p_1 + p_2)H_2(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2})$

This implies that  $H_m = (p_1, p_2, \dots, p_m) = -\sum_{i=1}^m p_i \log_2 p_i$  for  $m = 2, 3, \dots$ . A sketch of this proof from the perspective of *Information* can be found in Section 3.3.2.1.

## B.3 Principle of Maximum Entropy

The principle of maximum entropy is based on the observation that information entropy can be seen as a numerical measure which describes how uninformative a particular probability distribution is from zero (completely informative) to  $\log(m)$  (completely uninformative). Of all distributions, the most uninformative distribution is the Uniform distribution. Choosing a distribution with lower entropy results in the assumption of information which one does not possess. Thus, the maximum entropy distribution, given a priori information, is the only reasonable distribution.

Formally, the principle of maximum entropy is a method for analyzing available information in order to determine a unique probability distribution:

**Definition 24 (Principle of Maximum Entropy)** *Choose the probability model  $\tilde{f}$  which maximises the entropy of  $x$ :*

$$H(x|\pi_f) = H(\{x = x_n\}_1^N|\pi_f) = -K \sum_{n=1}^N f(x_n) \log_2 f(x_n) \quad (\text{B.4})$$

*as this is the least biased distribution that encodes certain given information while remaining consistent with the given information.*

Using the *Wallis* derivation, it is also possible to arrive at the Principle of Maximum Entropy making no reference to information entropy as “uncertainty” or some other subjective measure. In this derivation, the entropy function is not assumed a priori, is found in the course of using strictly combinatorial arguments. It is the same argument used for the derivation of partition functions in statistical mechanics. It concludes with the finding that to maximise entropy under the constraints of testable information is the most probable of all “fair” random epistemic distributions, in the limit as the probability levels go from discrete to continuous.

## B.4 Jensen's Inequality

For a real continuous concave function:

$$\frac{\sum f(x_i)}{n} \leq f\left(\frac{\sum x_i}{n}\right) \quad (\text{B.5})$$

In Economics, Jensen's Inequality is used to show that an investor with a concave utility function prefers a certain return to the same expected return with uncertainty. In statistics, it is used to show that  $E[X^2] \geq (E[X])^2$ . As  $E[X^2] - (E[X])^2$  is the variance, this also implies that any random variable for which  $E[X^2]$  is finite has a variance and a mean.

In Information Theory, Jensen's Inequality makes use of the fact that  $-\sum_i g_i \log(g_i)$  is a concave function, to make statements about how entropy behaves as a function the properties of  $g_i$ .

## B.5 Properties of Shannon's Entropy

The Shannon entropy satisfies the following properties:

- For any  $n$ ,  $H_n(p_1, \dots, p_n)$  is a continuous and symmetric function on variables  $p_1, p_2, \dots, p_n$ .
- Event of probability zero does not contribute to the entropy, i.e. for any  $n$ ,  $H_{n+1}(p_1, \dots, p_n, 0) = H_n(p_1, \dots, p_n)$ .

## B.5 Properties of Shannon's Entropy

---

- Entropy is maximised when the probability distribution is uniform. For all  $n$ ,  $H_n(p_1, \dots, p_n) \leq H_n\left(\frac{1}{n}, \dots, \frac{1}{n}\right)$ .

Following from the Jensen inequality,  $H(X) = E\left[\log_b\left(\frac{1}{p(X)}\right)\right] \leq \log_b\left(E\left[\frac{1}{p(X)}\right]\right) = \log_b(n)$ .

- If  $p_{ij}$ ,  $1 \leq i \leq m$ ,  $1 \leq j \leq n$  are non-negative real numbers summing up to one, and  $q_i = \sum_{j=1}^n p_{ij}$ , then  $H_{mn}(p_{11}, \dots, p_{mn}) = H_m(q_1, \dots, q_m) + \sum_{i=1}^m q_i H_n\left(\frac{p_{i1}}{q_i}, \dots, \frac{p_{in}}{q_i}\right)$ .

The entropy  $H_n\left(\frac{p_{i1}}{q_i}, \dots, \frac{p_{in}}{q_i}\right)$  is the entropy of the probability distribution conditioned on group  $i$ . This property means that the total information is the sum of the information gained in the first step,  $H_m(q_1, \dots, q_m)$ , and a weighted sum of the entropies conditioned on each group.

It can be shown that the only function satisfying the above assumptions is of the form  $H_n(p_1, \dots, p_n) = -k \sum_{i=1}^n p_i \log p_i$ , where  $k$  is a positive constant representing the desired unit of measurement.

### B.5.1 Kraft Inequality

A particular application of Information Theory is to the practical implementation of codes. Kraft's Inequality is condition which specifies the necessary and sufficient conditions for the existence of prefix codes. Prefix codes are instantaneously decodable codes, instantaneous because no code word is a prefix of any other, and the decoding process requires a single pass over the code.

Formally, let each source symbol from the alphabet

$$S = \{s_1, s_2, \dots, s_n\} \tag{B.6}$$

be encoded into a uniquely decodable code over an alphabet of size  $r$  with codeword lengths

$$\ell_1, \ell_2, \dots, \ell_n \tag{B.7}$$

Then, Kraft's Inequality states that:

$$\sum_{i=1}^n \left(\frac{1}{r}\right)^{\ell_i} \leq 1 \tag{B.8}$$

The converse also implies that a uniquely decodable code over an alphabet of size  $r$  exists if  $\ell_1, \ell_2, \dots, \ell_n$  satisfies the Kraft Inequality.

In the context of this thesis, the prefix code would be a representation of results for a binary poll, as it relates to Section 2.2.1. The Kraft Inequality is then used to prove the lower bound on the length of these codes as the Entropy of the codes.

## B.6 Renyi Entropy

Renyi Entropy generalises Shannon measure of information:

$$H_\alpha(x) = \frac{1}{1-\alpha} \log_2 \left( \sum_{i=1}^n p_i^\alpha \right) \quad (\text{B.9})$$

for  $\alpha \neq 1$ .

In the limit, when  $\alpha = 1$ , one gets back Shannon's Entropy. It has been applied to measure complexity in dynamical systems in relation to “subjective” observers. When  $\alpha = 1$ , it becomes possible to separate out variables in a multivariate distribution:

$$H(A, B) = H_\alpha(A) + H_\alpha(B) \quad (\text{B.10})$$

Hence, Renyi Entropy can be seen as a relaxation of the additivity requirements in Section B.2.

Renyi entropies are used in various diversity indices and uncertainty measures. For example, when  $\alpha = 2$ , one gets a measure of information which is slightly different from that of Shannon's. Under Shannon, information measures the reduction in uncertainty that results from the occurrence of an event. When  $\alpha = 2$ , a numerical value is assigned to not a reduction in uncertainty, but to a transition in the probability distribution. This generalises information from the mere occurrence of an event to the change in the probability distribution arising from an entire ensemble of events. This special case of Renyi Entropy is dubbed *Variational Information*.

## B.7 Connections to Other Areas

In B.2, entropy is measured with Base 2. By changing the logarithmic base to  $e$ , one gets the Maxwell-Boltzmann-Gibbs entropy, frequently  $S$  in the physics nomenclature. The Maxwell-Boltzmann-Gibbs entropy measures the number of ways  $N$  atoms can be arranged in  $m$  cells. Let  $N_i$  be the number of atoms in cell  $i$ , where  $i = 1, \dots, m$ . Letting  $W$  be the number of arrangements, and  $S = \log_e W$ :

$$W = \left[ \frac{N!}{\prod_{1 \leq i \leq m} N_i!} \right]^{\frac{1}{N}} \quad (\text{B.11})$$

Let  $p_i = \frac{N_i}{N}$ , then, using Stirling's formula:

$$S \approx - \sum_{1 \leq i \leq m} p_i \log_e p_i \quad (\text{B.12})$$

Another connection is to *ergodic theory*: if entropies for two spaces are the same, there is a measure preserving map that also preserves Bernoulli shifts.

# Appendix C

## Brief Overview of Statistical Mechanics

In physics, one assumes that the system is composed of  $N$  molecules with a total energy  $E$ . Instead of attempting the impossible by focusing on the energy of each individual molecule, one focuses on the populations of *energetic states*: this can be graphically described by imagining that one one buckets energy states into discrete components and then tries to determine how many molecules are in a particular bucket. In equilibrium, the number of molecules in each bucket is constant, on average: there will be  $n_0$  molecules in the state with energy  $E_0$ ,  $n_1$  with  $E_1$ , and so on.

Due to the fact that

$$N = \sum n_i \tag{C.1}$$

Each of the different ways in which the energy buckets could be populated constitute a configuration. One such configuration is  $N, \dots, 0$ , and there are many more.

The number of arrangements,  $W$ , corresponding to a given configuration  $n_0, n_1, \dots$  is given by:

$$W = \frac{N!}{n_0!n_1!n_2!\dots} \tag{C.2}$$



---

corresponding to the number of distinguishable ways  $N$  objects can be sorted into bins with  $n_i$  objects in bin  $i$ .

Applying Stirling's approximation,

$$\ln n! \approx n \ln n - n \quad (\text{C.3})$$

and taking logs, one gets

$$\ln W = N \ln N - \sum (n_i \ln n_i) \quad (\text{C.4})$$

where  $-\frac{\sum (n_i \ln n_i)}{N}$  is the entropy of a specific configuration.

A configuration dominates another if there are more ways to obtain it. The most dominant configuration will be the one with the largest value of  $W$  with respect to  $n_i$ . As can be seen from Equation C.4, this is equivalent to finding the configuration with maximum entropy.

In statistical mechanics, the partition function is an important quantity that encodes the statistical properties of a system. Most of the thermodynamic variables of the system, such as the total energy, free energy, entropy, and pressure, can be expressed in terms of the partition function or its derivatives. The partition function encodes all the information of a system e.g. total energy/utility, fluctuation of energy, correlation functions, etc., without the need for considering microscopic quantities like the exact energy state of a molecule, or in the case of a network, certain degrees of specific nodes in the network. The forms of partition functions of statistical physics is analogous to a probability distribution obtained from maximizing informational entropy when only average values are known/given.

In order to understand the partition function, how it can be derived, and why it works, it is important to recognize that these macroscopic properties reflect the average behavior of the components. The partition function may be either calculated analytically, or approximated by computer simulations.

Another important feature of statistical mechanics is the constraints operating on the system, such as volume. Typically, the key constraints in operation in physics are the total number of molecules and the total energy of the systems

$$N = \sum n_i \quad (\text{C.5})$$

$$E = \sum n_i E_i \quad (\text{C.6})$$

---

where there are  $n_i$  molecules in energy state  $E_i$ .

In order to find the way in which molecules are distributed by energy states, which is equivalent to finding the number of molecules  $n_i$  in energy state  $E_i$ , the constraints are usually factored in through Lagrangian multipliers:

$$\begin{aligned} f(n_i) &= \ln(W) + \alpha \left( N - \sum n_i \right) + \beta \left( E - \sum n_i E_i \right) \\ &= N \ln N - \sum n_i \ln n_i - \alpha' \left( N - \sum n_i \right) + \beta \left( E - \sum n_i E_i \right) \end{aligned} \quad (\text{C.7})$$

# Appendix D

## Perfect Markets

The use of market metaphors to solve technical problems in telecommunication networks have been based on the results of Neoclassical Economics for the last two decades. While there is nothing wrong with turning to market mathematics, the reliance on Neoclassical Economic theory is arguably misguided. A large body of research in the last two decades has debunked the applicability of the neoclassical paradigm to real-life financial networks, and many of these arguments are equally applicable in the realm of networks.

The crux of this argument is based on intractability: in the hypothetical setting, agents face an insurmountable task. In what is known as the Arrow-Debreu framework, one needs a complete market for each good spanning all contingent states of an economy. Agents then have to keep track of prices and transact in an unrealistically large number of markets. This casts a pall over the neoclassical paradigm, equivalent in many ways to the incalculability criticism levied on the socialist agenda: this time, it is the economic agent, rather than the central planner, who faces an equally intractable problem in the neoclassical paradigm. This does not even account for the fact that, in the real world, market participation is costly: agents would face unbearable transaction costs if they were required to participate in all markets all the time.

Markets in the real world do solve problems relevant to telecommunication networks: efficient allocation of resources, enabling transactions between agents,

---

centralising information. They just don't solve them using the machinery of neoclassical economics.

### D.0.1 Neoclassical Economics

General equilibrium theory underlies all the major developments that have come to be known as Neoclassical Economics. Leon Walras (1834-1910), known generally as the father of General Equilibrium Theory, formalised the building blocks of this theory in a market system in which the prices, production and consumption of all goods are interrelated. Through the effects of complementarity and substitution, a change in the price of one good effects another: for example, a change in the price of oil changes the price of airline tickets. This, in turn, effects tourism.

In theory, this means that calculating the equilibrium price of just one good, requires information on all goods and knowledge of all demand and supply functions in the system. It is assumed that access to this information is cost-free and the ensuing calculations are instantaneous (see (HK88)).

The economics profession has witnessed a progression of increasingly sophisticated tools before arriving at the Arrow-Debreu formalisation. These are:

1. Adam Smith's invisible hand
2. Walras' auctioneer
3. Marshall's partial equilibrium, culminating with
4. Arrow and Debreu's proof of the existence and uniqueness of a General Equilibrium under certain (pretty stringent) conditions.

The central conclusion of this analysis is the idea of *laissez-faire*: markets allocate resources optimally in an unregulated, free economy.

---

### D.0.1.1 Walrasian Auctions

Walrasian auctions, are a form of a simultaneous auction in which each participant is able to calculate a demand function for each good in an economy, submit this to an auctioneer (Walrasian) who is then able to calculate total demand across all agents. The auctioneer then proceeds to match perfectly total supply and demand in the economy.

Notions of equilibrium are used to determine prices for goods across single or multiple markets. In “partial equilibrium” aggregate demand and supply are equalised under demand and supply functions which takes price as input. The key simplification is independence between goods: the price of one good does not affect another.

Walras’ key contribution was to expand partial equilibrium to general equilibrium, when there is complementarity and substitution between goods and the price of one good effects the prices of all other goods in the economy. It is this calculation which the Walrasian Auctioneer performs, above and beyond that of simple “partial equilibrium models”. However, there is a strong implications: the solution to the system of equations which solve Walras’ problem requires market clearing.

### D.0.1.2 Arrow-Debreu

Kenneth Arrow and Gerard Debreu formulated the first set of results around the dynamics of an equilibrium in the Neoclassical setting in their classic paper in 1954 (([Arr65](#); [Arr71](#); [AD54](#))). They were able to prove the existence of a multi-market equilibrium in which no excess demand or supply exists. This proof is based on the following key assumptions:

- Convexity of the Agents Utility Function
- Perfect Competition
- Demand Independence

---

The implication of their result is that there will be a set of prices at which aggregate demand and supply are matched for every good in the economy. Furthermore, there will be a spot and forward price for every good, for all time periods. This defines Maximally Complete Markets.

**Arrow-Debreu Security** To make their argument, Arrow and Debreu((AD54)) made use of a hypothetical construct dubbed an Arrow-Debreu Security. This is a security which pays one unit of numéraire<sup>1</sup> upon the realisation of a specific, identifiable state in the economy, and otherwise pays zero. They are used to analyse derivative contracts, whose payoffs are specified in terms of a linear combination of Arrow-Debreu securities. These securities, by arbitrage, sum to 1, and like probabilities, are exhaustive and mutually exclusive.

The price paid today for a state price, i.e. an Arrow-Debreu security, resembles a density. This is just a resemblance, because the probability has undergone what is called a “change of measure”<sup>2</sup>, due to arbitrage. It is important to see that they are not quite probabilities, even when translated into their continuous price and time limit.

The Arrow-Debreu method advanced General Equilibrium theory by giving it consistent dynamics and spatiality. Under Arrow-Debreu, equilibria is not achieved via an auction, but via a game in which choices for each agent are specified a priori, and there exist real-valued functions which specify the agents payoff. Each agent then acts selfishly to maximise their payoff, and the Nash Equilibrium of this game is that from which no agent would want to deviate unilaterally. In this context, Convexity is needed to prove existence, taking into account inconsistency and lack of independence in agent choices (non-transitivity). Arrow-Debreu markets attain *Pareto Optimality* if they are *complete*.

---

<sup>1</sup>French for measuring/counting. The numéraire is the money unit of measure within an abstract macroeconomic model in which there is no actual money or currency. A standard use is to define one unit of some kind of goods output as the money unit of measure for wages, measuring the worth of goods and services relative to another.

<sup>2</sup>A technique to transform one stochastic process to another to simplify calculation.

---

## D.0.2 Pareto Optimality

Defined as “the best that could be achieved without disadvantaging at least one group,” a Pareto Optimal allocation amongst a set of individuals is one in which it would not be possible to better any single individual without making someone worse off. It is a key characterisation of equilibrium conditions. The notion of Pareto Optimality led to the two Fundamental Theorems of Welfare Economics which are the key contribution of Neoclassical Economics.

### D.0.2.1 First Fundamental Theorem of Welfare Economics

The First Fundamental theorem makes a strong statement about free markets. It proves that a system of free markets will lead to a Pareto efficient outcome. The key assumptions underlying the first theorem are:

- Markets exist for all goods
- Markets are perfectly competitive
- Transaction costs are negligible

The assumptions of the First theorem are relatively weak. No convexity is needed.

### D.0.2.2 Second Fundamental Theorem of Welfare Economics

The Second Theorem reverses the First by adding *equilibrium*, albeit at the expense of a stronger assumption. The Second Theorem claims the inverse of the first by proving that a set of consistent prices will exist for any efficient allocation and an equilibrium will hold. This is because any initial distribution of goods will eventually lead to a Pareto Optimal outcome. By definition, no deviations will follow.

However, the second theorem does require an additional assumption: convexity in Consumer preferences. This roughly implies diminishing marginal utility.

One of the important implications of the second theorem is that efficiency and equality can be separated and do not involve a tradeoff. Convexity brings with

---

it the benefit of existence, but not uniqueness. However, in a complete market, risk neutral probabilities, and hence prices, are unique and there are no arbitrage.

### D.0.3 Lack of Arbitrage and Market Completeness

An arbitrage exists if and only if either:

1. Two portfolios can be created that have identical payoffs in every state but have different costs; or
2. Two portfolios can be created with equal costs, but where the first portfolio has at least the same payoff as the second in all states, but has a higher payoff in at least one state; or
3. A portfolio can be created with zero cost, but which has a non-negative payoff in all states and a positive payoff in at least one state.

The lack of arbitrage is a key condition in the Neoclassical setting, as it is the basis for consistency and uniqueness in the General Equilibrium.

Formally, assume an economy can take one of  $n$  states at time  $T$ . Let  $A = [A_{ij}]$  be the payoff matrix at time  $T$  for a set of Arrow-Debreu Securities, that of asset  $j$  in state  $i$ . The columns of  $A$  represent the payoffs of asset  $j$  and the rows represent state  $i$ .

If  $\text{rank}(A) = m$ , i.e. the number of linearly independent assets is equal to the number of states, which means that  $m$  linearly independent assets can be combined to generate any payoff at time  $T$ . This is what is meant by market completeness.

(Ros76) proves the following:

**Lack of Arbitrage 1** *There is no arbitrage if and only if there exists a row vector of strictly positive state prices  $q = [q_1 \ q_2 \ \dots \ q_m]$  such that  $p = q \cdot A$*

The role of arbitrage should also not be overlooked in the real world. Arbitrage pricing is the bedrock of derivatives pricing theory. If the payoffs of an instrument can be replicated, statically or dynamically, by trading in a portfolio of securities,



---

the value of the portfolio and the contingent must be equal. This is known as the *replication argument*. Because the combination of the option and a short position in the replicating portfolio is risk-free, the return must also be risk-free, which gives rise to the *risk-neutral*<sup>1</sup> argument<sup>2</sup>. If the prices of the primary assets are known, the price of the derivative is a simple consequence.

Because Arrow-Debreu prices are risk-neutral probabilities and any derivative contract can be replicated through a linear combination of Arrow securities, the Second Theorem rules out arbitrage.

More importantly, arbitrage is the glue which holds market together, by acting as a transmission mechanism for prices from one market to another. This is an oft-criticised, yet important function that arbitrage plays in the world economy. If interest rates in one country rise while falling in another, arbitrage determines new rates of exchange of one currency for another in the forward markets<sup>3</sup>. If arbitrage markets did not operate, interest rates and foreign exchange markets would not function and entire economies would not be able to balance their trade flows. The inability to balance trade flows in the long run can and do have disastrous results that can lead to decades of economic malaise(SS01a).

#### D.0.4 The Premises of Neoclassical Economics

The premises of the neoclassical set-up are an early indication of its weakness:

1. Equilibrium is attained for all prices under perfect competition
2. Movement from one equilibrium to the next is instantaneous
3. Equilibrium is “normal”
4. The optimum outcome is achieved at equilibrium

---

<sup>1</sup>Risk-neutral probabilities refer colloquially to the probability measure which replicates the price (present value) of a derivatives contract when applied to the payoff profile at maturity, even in markets known to be incomplete.

<sup>2</sup>The primary asset replicating the derivative may need to be dynamically traded to replicate the payoff. Black-Scholes is the personification of this result.

<sup>3</sup> A result originally due to Keynes.

- 
5. There is no adverse selection, no intermediation, no transaction costs and no information.

Indeed, the Arrow-Debreu General Equilibrium is in one sense, a reverse-engineered solution: the assumptions are requirements for a pre-defined adequate solution. However, none of these assumptions bear any resemblance to reality, especially those relating to transaction costs and information. This criticism applies not only to the application of this framework to financial markets, but also to the telecommunication systems. Markets are not perfect.

# Appendix E

## Further Results

In Chapter 5, the *Measurement Model* was evaluated using time series data on the average latency of the Credit Suisse network. Table 5.2.2 provides a number of alternative *Reference Measurements* which could also be used as the basis for analysing the complexity of a network.

There are a number of interesting issues in the right choice of *Reference Measurement*:

- the suitability of the measure for the purposes of a certain analysis
- the relationship between various measurements
- the impact of the window used to calculate *Network Complexity* and *Implied Coupling* trends are calculated

Some of these measures may be clearly correlated. *Network Complexity* measures based on the standard deviation of latency measures, for example Figure E.1, and discarded packets, Figure E.3, seem to be negatively correlated. This may indicate a tradeoff between the two measures of performance.

In depth analysis of these results is left as future work. A cursory exposition of the results is graphically illustrated below.

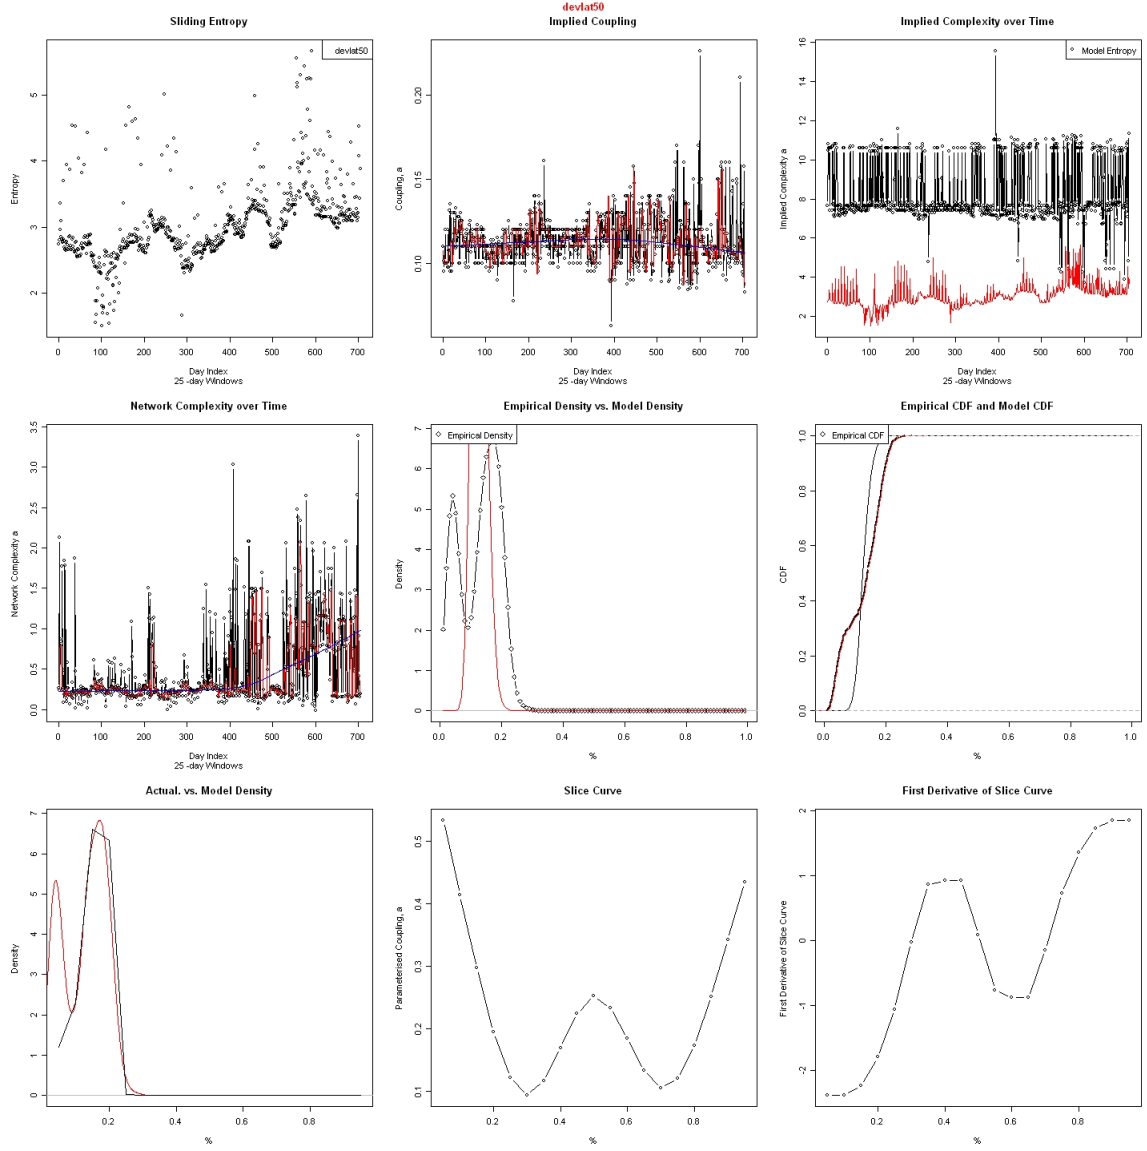


Figure E.1: Latency Standard Deviation, 25-day Moving Windows, 50-ms Triggers

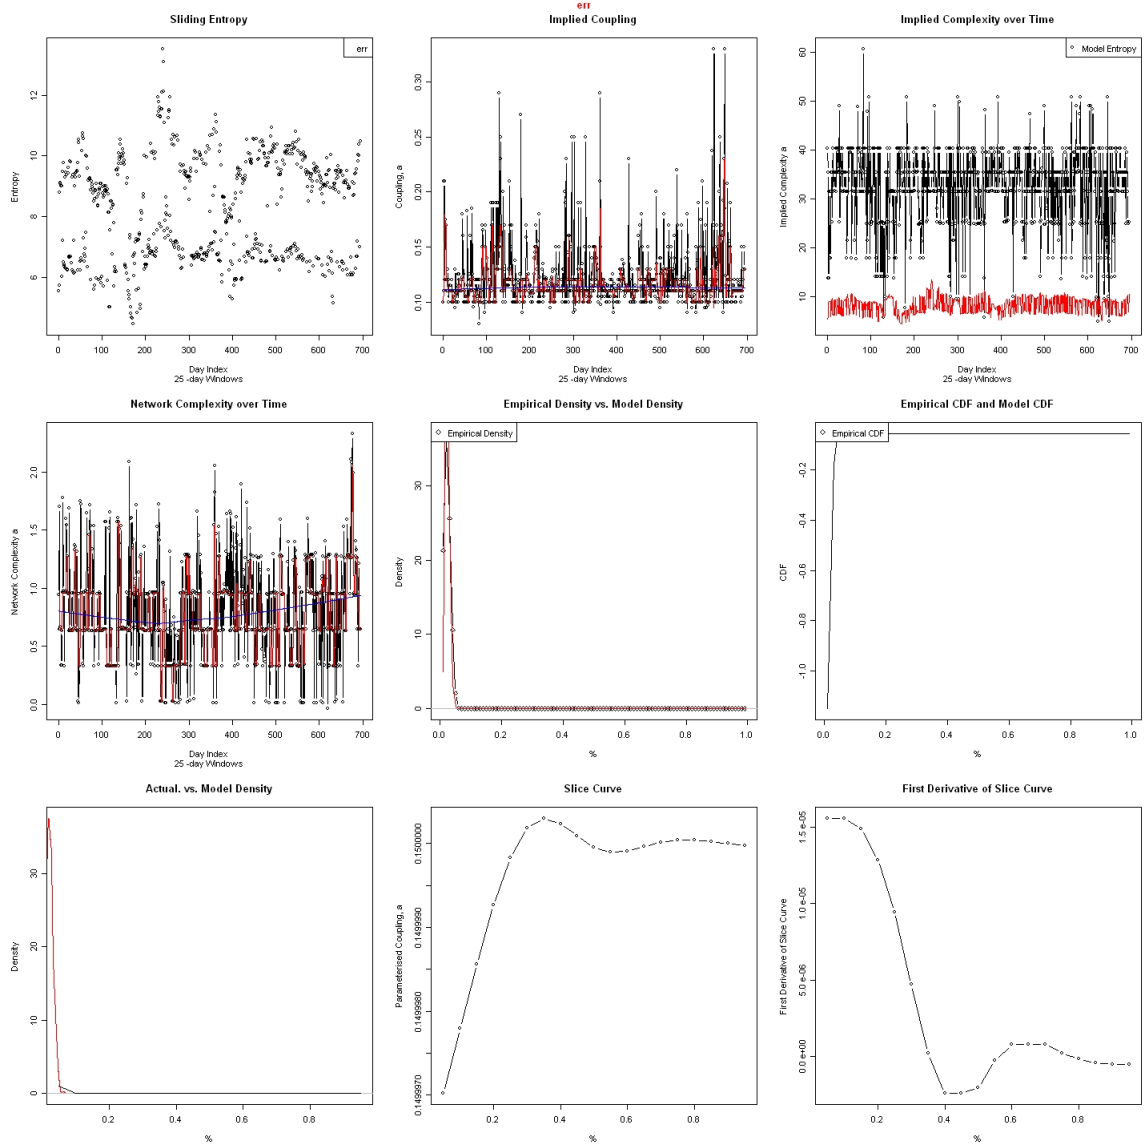


Figure E.2: Error Rates, 25-day Moving Windows

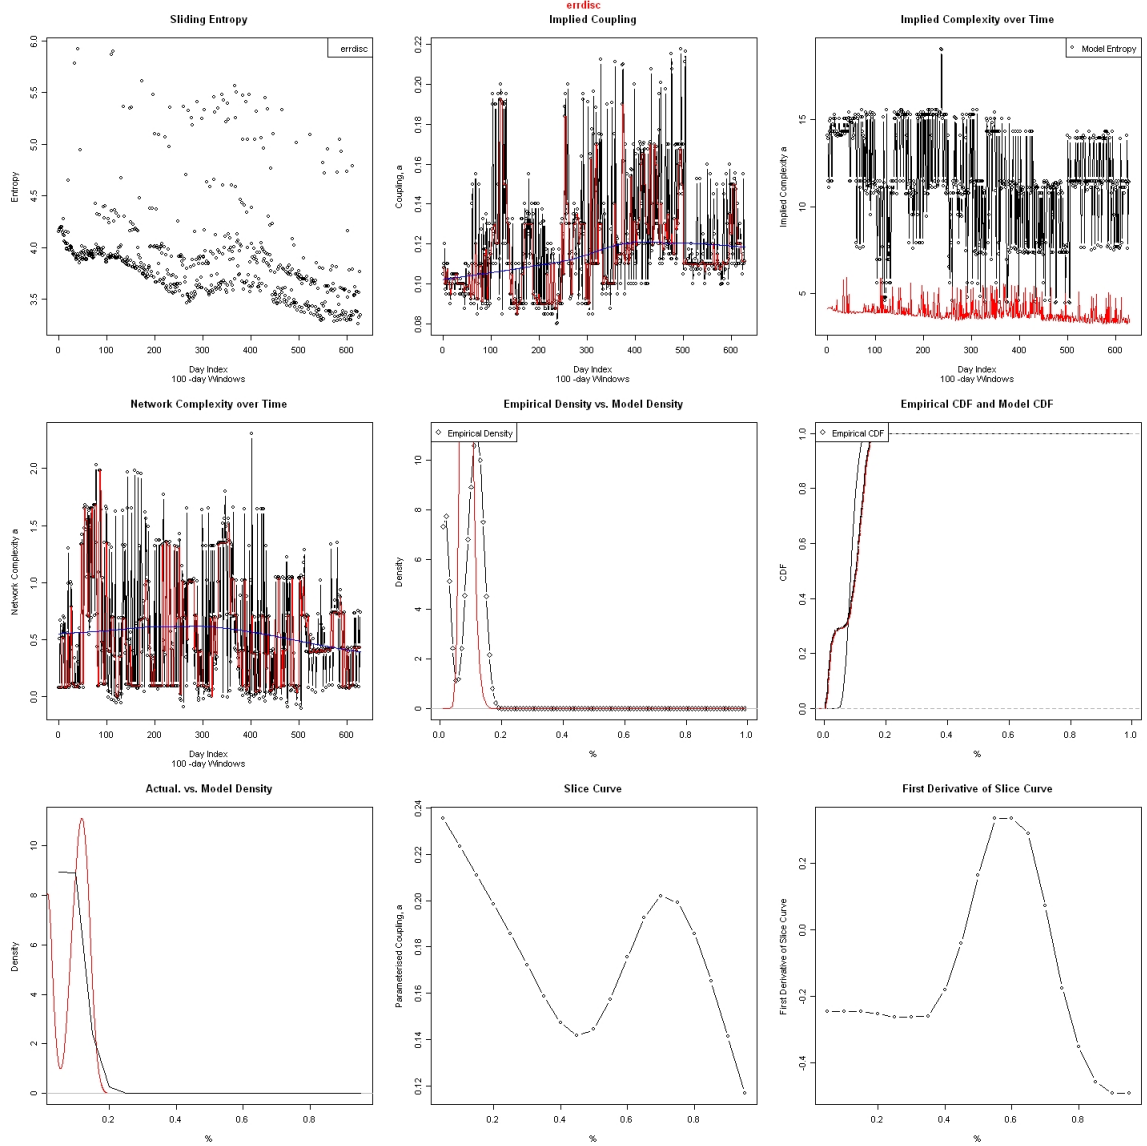


Figure E.3: Discarded Packets, 100-day Windows

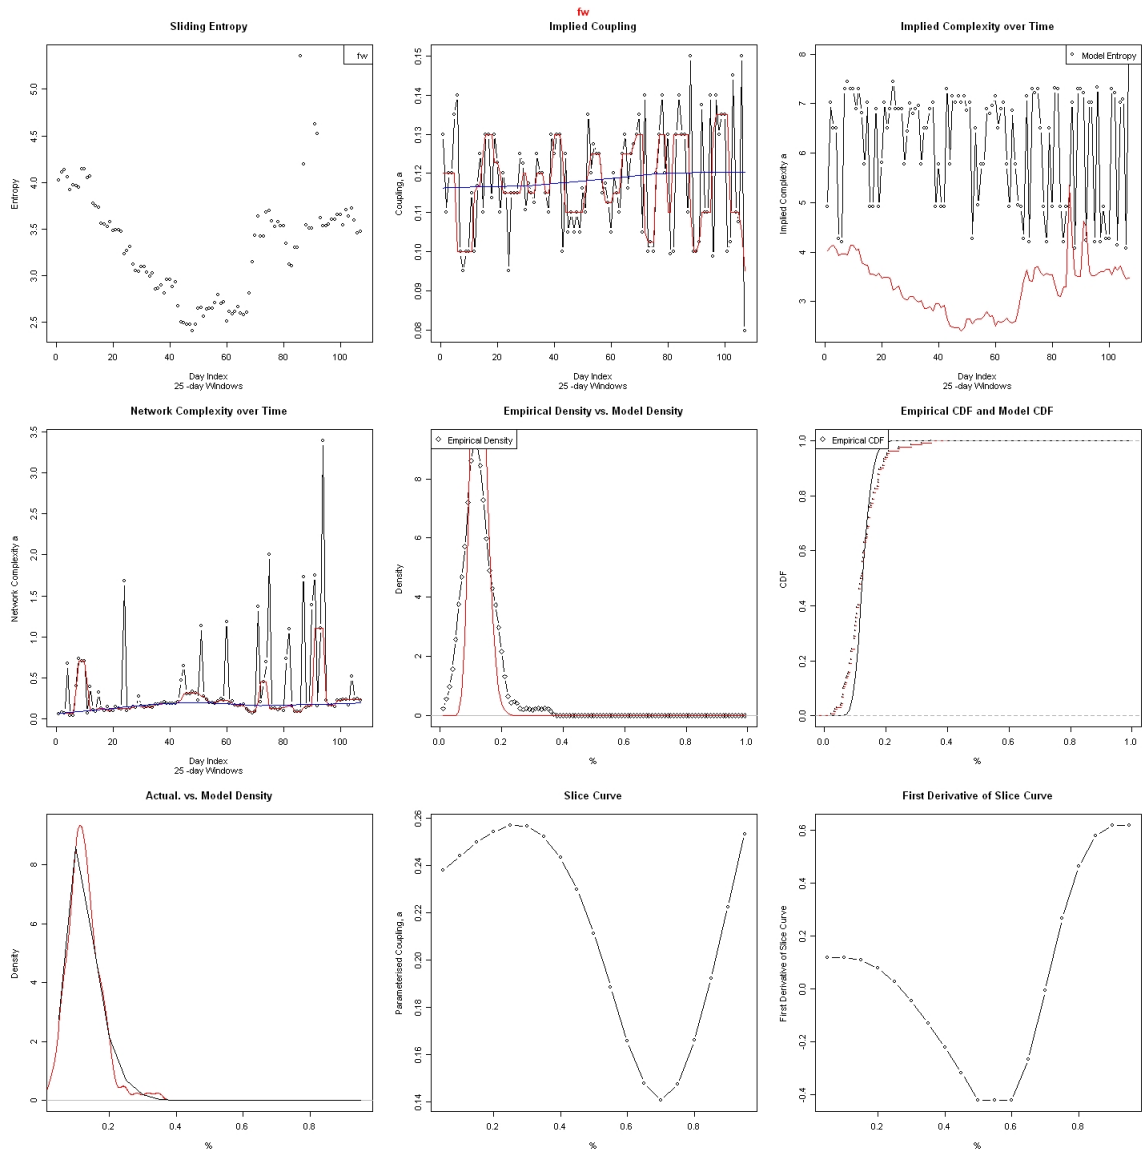


Figure E.4: Firewall Utilisation, 25-day Moving Windows

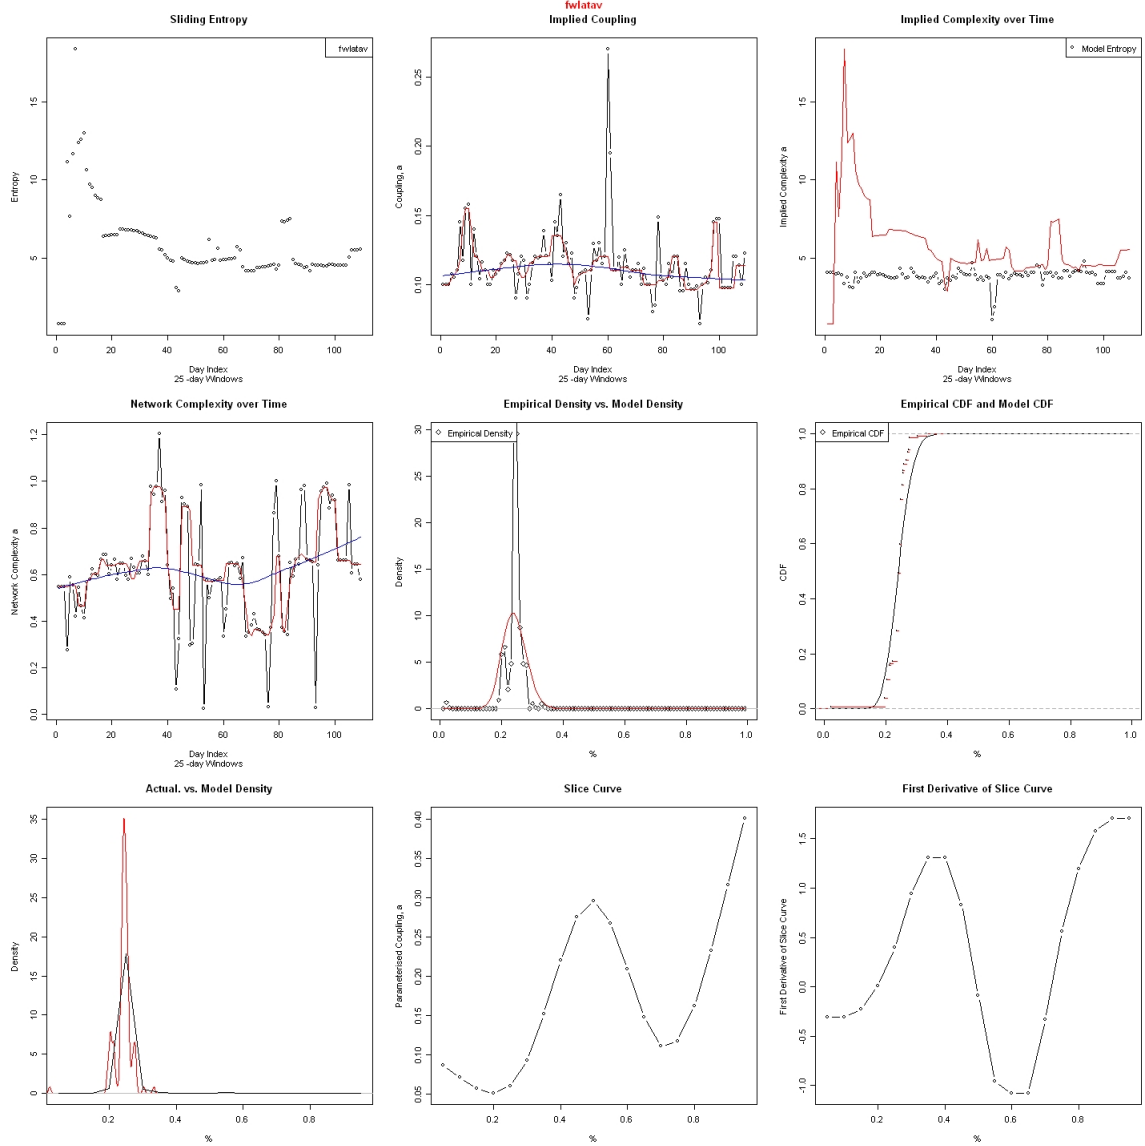


Figure E.5: Firewall Average Latency, 25-Day Moving Windows



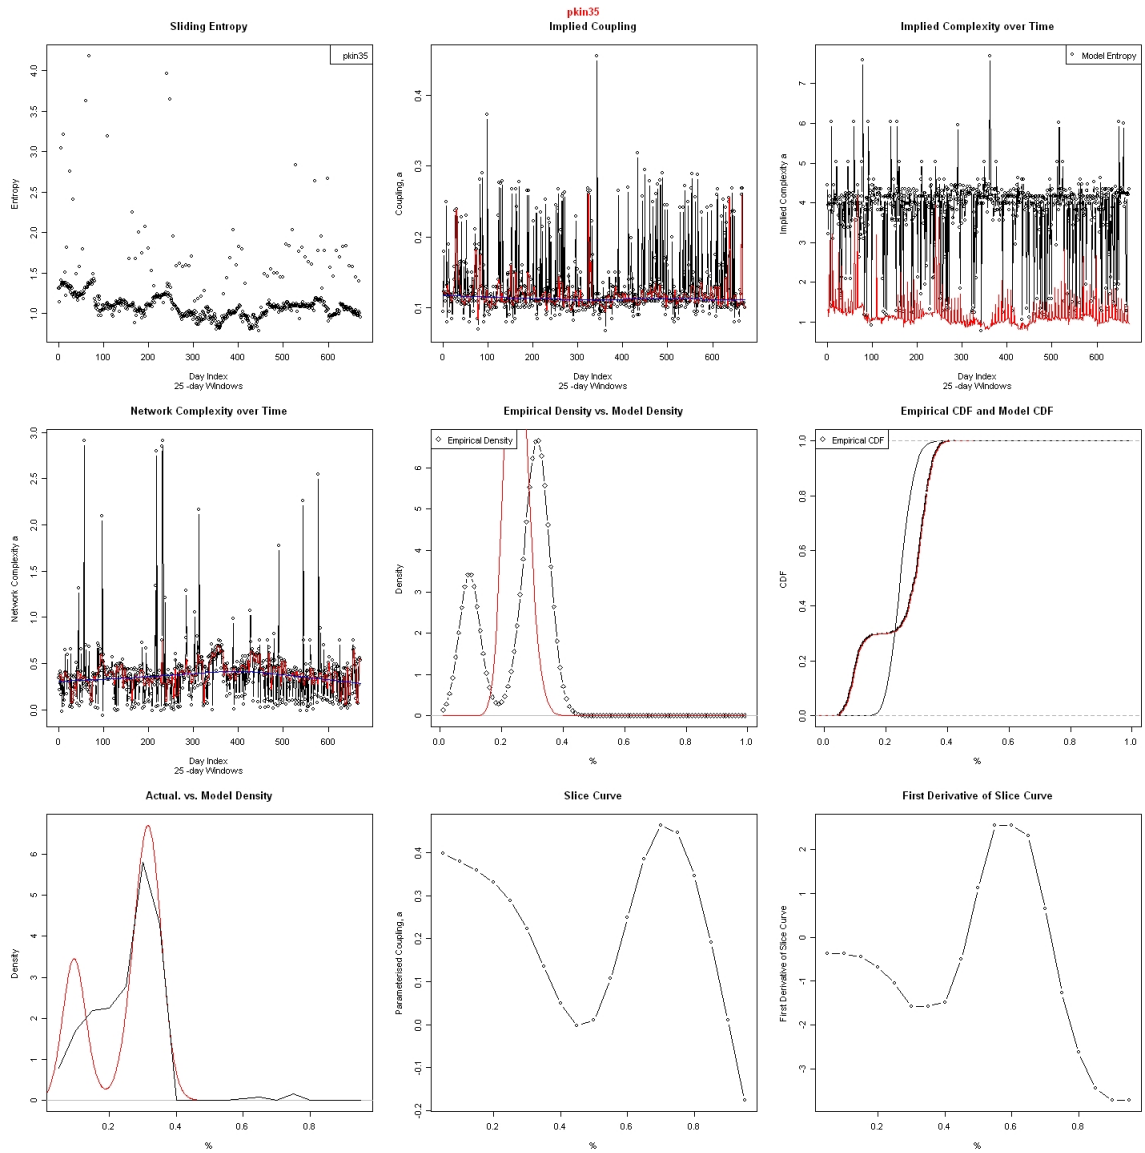


Figure E.6: Peak Inbound Traffic, 25-day Moving Windows, 35-ms Triggers

# Glossary

## Architectural Complexity

A measure of the static complexity in a network arising from the way in which various *Elements* perform certain *Operations* on a *Traffic Vector*. It is meant to be a proxy for the *Interrogations Costs* associated with determining the state of a network at a certain point in time. It is measured using the *Architectural Model*. Formally:

$$H = - \sum_{u=1}^q \sum_{v=1}^q \sum_{z=1}^m \sum_{l=1}^b \tilde{\psi}_{uvzl} \log \tilde{\psi}_{uvzl} \quad (\text{E.1})$$

## Ashby's Law of Requisite Variety

, [121](#)

The larger the variety of actions available to a control system, the larger the variety of perturbations it is able to compensate. This implies that the more variety there is in a system the more information has to extract from it in order to understand and control it., [22](#)

## Complexity Neutral

*Complexity Neutral* decisions refer to control decisions in a network which do not impact the *Implied Complexity* of the networks behaviour., [155](#)

<b>Element Complexity</b>	A sub-measure of <i>Architectural Complexity</i> , meant to quantify the complexity in a network associated with the variety of <i>Elements</i> performing <i>Operations</i> on a network. The <i>Operational Matrix</i> is used to quantify this complexity., <a href="#">122</a>
<b>Entropic Formulation</b>	The use of the Entropy of a network as the key macroscopic feature for the purposes of control and behaviour prediction., <a href="#">22</a>
<b>Entropic Routing</b>	The use of measures of the networks complexity and risk characteristics to make routing decisions., <a href="#">202</a>
<b>Entropy Flooding</b>	In closed systems, the behavioural rules of agents maximising entropy will dictate the dynamical behaviour of the system over time., <a href="#">94</a>
<b>Hakanssons Paradox</b>	If markets are complete options are not needed; if they are incomplete then according to financial theory, they cannot be priced., <a href="#">62</a>
<b>Idiosyncratic Risk</b>	In finance, price changes due to circumstances specific to a certain asset, not related to the overall market., <a href="#">16</a>
<b>Implied Complexity</b>	The complexity of a network in the absence of any model. Formally, the entropy implied by the histogram of actual network <i>Reference Measurement</i> data., <a href="#">104</a>
<b>Implied Coupling</b>	<i>Implied Coupling</i> measures how tightly the elements in a network are coupled together., <a href="#">152</a>
<b>Information, <math>I(p)</math></b>	A quantitative measure of information, representing the physical occurrence of a an event $X = x$ , which a priori is assumed to have probability $p$ . Under the assumption that $I(p)$ is always non-negative and decreases with increasing $p$ , it can be shown that $I(p) = -\log(p)$ ., <a href="#">115</a>
<b>Investors Problem</b>	Make it as hard as possible for Market Makersmarket maker to infer the probability distribution of asset prices., <a href="#">72</a>

<b>Law of Constraints</b>	A characterisations of the 1st Law of Thermodynamics in the language of networks and markets. It is a preservation law, which posits that goods are not created through transactions., <a href="#">75</a>
<b>Law of Encapsulation</b>	A characterisation of the 0th Law of Thermodynamics in the language of networks and markets. It essentially posits that in a system of networks controlled by agents, no transactions will take place if their valuation is the same for all agents., <a href="#">75</a>
<b>Law of Preferences</b>	A characterisation of the 2nd Law of Thermodynamics in the language of networks and markets. Agents will only do transactions that increase their utility., <a href="#">75</a>
<b>Market-Makers Problem</b>	Find the least number of questions to ask investors to discover their collective belief about future prices., <a href="#">71</a>
<b>Measurement Model</b>	A parsimonious model which uses <i>Reference Measurements</i> , describes as much of the networks behaviour as possible, and acts as an easy to use proxy for the <i>Architectural Model</i> ., <a href="#">126</a>
<b>Memoisation</b>	A common technique from functional programming, in which a value is stored once it has been computed to avoid re-evaluating the expression., <a href="#">98</a>
<b>Model Complexity</b>	Model Complexity is meant to quantify the complexity associated with extending a base model for a network to account for known details in a networks configuration which deviate from the base case assumptions. It is an information theoretic distance between the simplest possible model of a system and extensions which aim to increase its explanatory power. In this context, it can be measured as the Kullback-LeiblerKullback-Leibler distance between the CDF implied in Equation <a href="#">3.29</a> and various extensions in Section <a href="#">3.4.4.</a> , <a href="#">141</a>

**Neoclassical Economics**

A school of economics which focuses on individual decision making based on preference relationships amongst consumption bundles. The maximisation of utility or profit is usually posited as an objective function. Most mainstream microeconomics is based on the neoclassical paradigm., [14](#)

**Network Complexity**

*Network Complexity* is meant to measure what is not explained by the model. This is the essence of *Effective Complexity*. Hence, it is defined as that part of a networks *Implied Complexity* which deviates from the predictions of the model the engineer has built of the network. Using the *Measurement Model*, it is quantified as the Kullback-Leibler between the empirical behaviour of the network (proxied through *Reference Measurements* and the behaviour expected by the *Measurement Model*., [129](#)

**Network Precedence Matrix**

Represents the precedence relationships between operations., [121](#)

**Operational Complexity**

A sub-measure of *Architectural Complexity*, meant to quantify the complexity in a network associated with the variety of *Operations* being performed by *Elements* on a network. The *Operational Matrix* is used to quantify this complexity., [122](#)

**Operational Matrix**

Represents the complexity arising from the way operations are carried out by different elements on the network., [121](#)

**Reference Measurement**

An observation on *Elements* which network engineers care about. They are not continuous, but discrete. They can have more than two values, but in the context of Section [3.4](#), they are a Bernoulli-(0,1) Bernoulli. Examples are breaches of capacity utilisation thresholds indicated by (1), and (0) otherwise., [118](#)

<b>Reference Measurements</b>	Measurements which proxy the complexity of the network and are common to all elements. These measurements should be easy to calculate and not specific to individual elements or technologies. Reference Measurements are Bernoulli-(0,1) and triggered when Capacity, determined as a threshold over some characteristic of the <i>Element</i> , is exceeded in usage., <a href="#">125</a>
<b>Service Level Agreement</b>	A formal agreement made between a service provider and a customer which defines a specified level of service, support requirements, and penalties., <a href="#">47</a>
<b>shadow prices</b>	Shadow prices are usually defined as the change in the objective value of an optimal solution when the constraints are relaxed by a unit. The Lagrangian multipliers usually correspond to shadow prices., <a href="#">67</a>
<b>Slice</b>	A risk metric which measures the probability that a certain <i>Reference Measurement</i> of importance associated with a network will take on a value between two thresholds, a lower threshold $K_1$ and an upper threshold $K_2$ . It is meant to proxy <i>Idiosyncratic Risks</i> when the <i>Slice</i> between $K_1$ and $K_2$ on the multivariate distribution associated with the <i>Reference Measurement</i> is located at low values and <i>Systematic risks</i> when the <i>Slice</i> is located at higher values., <a href="#">16</a>
<b>Slice Curve</b>	A graphical depiction of <i>Slice</i> measurements on a network. The lower threshold is always set to zero, the upper threshold runs along the x-axis, with the y-axis displays the corresponding <i>Slice</i> measurement., <a href="#">151</a>
<b>Systematic Risks</b>	In finance, refers to risks emanating from market-wide circumstances. Such risks cannot be hedged, and hence pay investors a premium., <a href="#">16</a>

**Traffic Complexity**

A measure of the complexity of *Traffic Flows* as they relate to overall *Architectural Complexity*:

$$H(\nu) = - \sum \nu \log(\nu) \quad (\text{E.2})$$

**Traffic Precedence Matrix**

., [122](#)

Represents the complexity arising from similarities between sequences of operations on actual traffic flows and the sharing of resources., [121](#)

# References

- [ABEA<sup>+</sup>06] E. Altman, T. Boulogne, R. El-Azouzi, T. Jimenez, and L. Wynter. A survey on networking games in telecommunications. *Comput. Oper. Res.*, 33(2):286–311, 2006. [45](#)
- [AD54] Kenneth J. Arrow and Gérard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954. [236](#), [237](#)
- [AFHS97] M. Avellaneda, C. Friedman, R. Holmes, and D. Samperi. Calibrating volatility surfaces via relative-entropy minimization. *Applied Math Finance*, 4:37–64, 1997. [225](#)
- [AK02] M. Angermann and J. Kammann. Cost metrics for decision problems in wireless ad hoc networking. Technical report, German Aerospace Center (DLR), Institute of Communications and Navigation, 2002. [109](#)
- [AL95] S. Arora and C. Lund. Hardness of approximations. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*. PWS, 1995. [12](#)
- [Arr53] Kenneth J. Arrow. Le rôle des valeurs boursières pour la répartition la meilleure des risques. *Econométrie*, Colloques Internationaux du Centre National de la Recherche Scientifique, Paris, 11:41–47, 1953. Published in English as “The Role of Securities in the Optimal Allocation of Risk-Bearing” in the *Review of Economic Studies*, April 1964, 31 (2), 91–96. [28](#), [100](#)



## REFERENCES

---

- [Arr65] Kenneth J. Arrow. *Aspects of the Theory of Risk-Bearing*. Yrjö Jahnsson Foundation, Helsinki, 1965. 236
- [Arr71] Kenneth J. Arrow. The theory of risk aversion. In *Essays in the Theory of Risk-Bearing*, pages 90–120. North-Holland, Amsterdam, 1971. 236
- [Asson] Semiconduer Industry Association. *The International Technology Roadmap for Semiconductors*. Semiconductor Industry Association, 2001 Edition. 52
- [Awd99] D. O. Awduche. Mpls and traffic engineering in ip networks. *Communications Magazine, IEEE*, 37(12):42–47, 1999. 37
- [Awe99] J. Aweya. Ip router architectures: An overview. Technical report, Nortel Networks, 1999. 50
- [Axt99] Rob Axtell. The complexity of exchange. Computing in Economics and Finance 1999 211, Society for Computational Economics, March 1999. available at <http://ideas.repec.org/p/sce/scecf9/211.html>. 98
- [BBM<sup>+</sup>93] Simona Brugnioni, Guido Bruno, Roberto Manione, Enrico Montariolo, Elio Paschetta, and Luisella Sisto. An expert system for real time fault diagnosis of the italian telecommunications network. In *Proceedings of the IFIP TC6/WG6.6 Third International Symposium on Integrated Network Management with participation of the IEEE Communications Society CNOM and with support from the Institute for Educational Services*, pages 617–628. North-Holland, 1993. 45
- [BCFK95] A. Bouloutas, S. Calo, A. Finkel, and I. Katzela. Distributed fault identification in telecommunication networks. *Journal of Network and Systems Management*, 3:295–312, 1995. 45
- [Ben84] J.-P. Benassy. *The Economics of Market Disequilibrium*. Academic Press, New York, 1984. 61

## REFERENCES

---

- [BFL<sup>+</sup>00] Vasken Bohossian, Charles C. Fan, Paul S. LeMahieu, Marc D. Riedel, Lihao Xu, and Jehoshua Bruck. Computing in the RAIN: A Reliable Array of Independent Nodes. *Lecture Notes in Computer Science*, 1800, 2000. [47](#)
- [BG05] X. Burtschell and J. Gregory. A comparative analysis of cdo pricing models. Technical report, BNP-Paribas, 2005. [152](#)
- [Bin92] K. Binmore. *Fun and Games: a Text on Game Theory*. D.C. Heath, Lexington, 1992. [63](#)
- [BL78] Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, October 1978. [150](#)
- [BLHM<sup>+</sup>96] N. Bjorkman, A. Latour-Henner, A. Miah, Simon Crosby, Ian M. Leslie, M. Davey, Raymond Russell, and Fergal Toomey. Exploring the queueing behaviour of ATM switches. *Performance Evaluation*, 27/28(4):89–98, 1996. [58](#)
- [BR96] S. Benjaafar and R. Ramakrishnan. Modeling, measurement, and evaluation of sequencing flexibility in manufacturing systems. *International Journal of Production Research*, 34:1195–1220, 1996. [168](#)
- [BT89] J. T. Blake and K. S. Trivedi. Multistage interconnection network reliability. *IEEE Trans. Comput.*, 38(11):1600–1604, November 1989. [5](#)
- [BT00] Vincent D. Blondel and John N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000. [65](#)
- [BTW88] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Phys. Rev. A*, 38(1):364–374, Jul 1988. [95](#)
- [Cat01] Ariel Caticha. Entropic dynamics. *MaxEnt 2001, the 21th International Workshop on Bayesian Inference and Maximum Entropy Methods*, Sep 2001. [82](#)

## REFERENCES

---

- [Cat04] Ariel Caticha. Questions, relevance and relative entropy. *MaxEnt 2004, the 24th International Workshop on Bayesian Inference and Maximum Entropy Methods*, September 2004. [82](#)
- [CB96] Mark Crovella and Azer Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proceedings of SIGMETRICS'96: The ACM International Conference on Measurement and Modeling of Computer Systems.*, Philadelphia, Pennsylvania, May 1996. Also, in Performance evaluation review, May 1996, 24(1):160-169. [8](#)
- [CH05] Shu-Heng Chen and Ya-Chi Huang. On the role of risk preference in survivability. In Lipo Wang, Ke Chen, and Yew-Soon Ong, editors, *ICNC (3)*, volume 3612 of *Lecture Notes in Computer Science*, pages 612–621. Springer, 2005. [94](#)
- [CKT03] S. Chakraborty, S. Kunzli, and L. Thiele. A general framework for analysing system properties in platform-based embedded system design. *Design, Automation and Test in Europe (DATE)*, March 2003. [109](#)
- [CM95] Dave Cliff and Geoffrey F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *European Conference on Artificial Life*, pages 200–218, 1995. [204](#)
- [CO98] J. Crowcroft and P. Oechslin. Differentiated end-to-end internet services using a weighted proportional fair sharing tcp. *Computer Communications Review*, 3:53–67, 1998. [46](#)
- [Coo67] J.L.B. Cooper. The foundations of thermodynamics. *Journal of Mathematical Analysis and Applications*, 17:172–193, 1967. [87](#)
- [CT91] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, USA, 1991. [23](#), [72](#), [105](#), [114](#), [118](#), [164](#), [224](#), [226](#)

## REFERENCES

---

- [CW99] Fabián A. Chudak and David P. Williamson. Improved approximation algorithms for capacitated facility location problems. *Lecture Notes in Computer Science*, 1610, 1999. [11](#)
- [Deb54] G. Debreu. Representation of a preference ordering by a numerical function. In R. Thrall, C.H. Coombs, and R. Davies, editors, *Decision Processes*, pages 159–175. Wiley, New York, 1954. [87](#)
- [Deb64] G. Debreu. Continuity properties of Paretian utility. *International Economic Review*, 5(3):285–293, September 1964. [87](#)
- [Dew03] Roderick C. Dewar. Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states. *MATH.GEN.*, 36:631, 2003. [95](#)
- [DFT03] C. S. Daw, C. E. A. Finney, and E. R. Tracy. A review of symbolic analysis of experimental data. *Review of Scientific Instruments*, 74(2):915–930, 2003. [104](#)
- [DG04] Alessio D’Ignazio and Emanuele Giovannetti. From exogenous to endogenous networks: Internet applications. Cambridge Working Papers in Economics 0445, Faculty of Economics (formerly DAE), University of Cambridge, September 2004. available at <http://ideas.repec.org/p/cam/camdae/0445.html>. [86](#)
- [DLO<sup>+</sup>95] Nick G. Duffield, J. T. Lewis, Neil O’Connell, Raymond Russell, and Fergal Toomey. Entropy of ATM traffic streams: A tool for estimating qos parameters. *IEEE Journal of Selected Areas in Communications*, 13(6):981–990, 1995. [58](#)
- [DS86] Darrell Duffie and Wayne Shafer. Equilibrium in incomplete markets ii: Generic existence in stochastic economies. *Journal of Mathematical Economics*, 15(3):199–216, 1986. [61](#)
- [Edm99a] B. Edmonds. *Syntactic Measures of Complexity*. PhD thesis, University of Manchester, 1999. [26](#), [104](#)

## REFERENCES

---

- [Edm99b] Jeff Edmonds. Scheduling in the dark. In *STOC '99: Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 179–188, New York, NY, USA, 1999. ACM Press. [42](#)
- [EJK<sup>+</sup>01] Joseph E. Eggleston, Sugih Jamin, Terence P. Kelly, Jeffrey K. MacKie-Mason, William E. Walsh, and Michael P. Wellman. Survivability through market-based adaptivity: The marx project. Technical report, University of Michigan, 2001. [46](#), [47](#)
- [Fan00] G. Fankhauser. *A Network Architecture Based on Market Principles*. PhD thesis, ETH, 2000. [47](#)
- [FBR<sup>+</sup>04] Nick Feamster, Hari Balakrishnan, Jennifer Rexford, Aman Shaikh, and Kobus van der Merwe. The Case for Separating Routing from Routers. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, Portland, OR, September 2004. [52](#), [205](#)
- [Fer03] Pablo Molinero Fernandez. *Circuit Switching In The Internet*. PhD thesis, Stanford University, 2003. [52](#)
- [FGHW99] Anja Feldmann, Anna C. Gilbert, Polly Huang, and Walter Willinger. Dynamics of IP traffic: A study of the role of variability and the impact of control. In *SIGCOMM*, pages 301–313, 1999. [8](#)
- [FLBS99] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 548–553. Morgan Kaufmann Publishers Inc., 1999. [65](#)
- [FO95] Eric J Friedman and Shmuel S Oren. The complexity of resource allocation and price mechanisms under bounded rationality. *Economic Theory*, 6(2):225–50, July 1995. available at <http://ideas.repec.org/a/spr/joecth/v6y1995i2p225-50.html>. [65](#)

## REFERENCES

---

- [FORR98] E. Fulp, M. Ott, D. Reininger, and D. Reeves. Congestion pricing flow control for computer networks. Technical report, Center for Advanced Computing and Communication, April 1998. [29](#)
- [FS03] Craig Friedman and Sven Sandow. Learning probabilistic models: An expected utility maximization approach. *Journal of Machine Learning Research*, 4:257–291, 2003. [93](#)
- [FT91] Drew Fudenberg and Jean Tirole. *Game Theory*. Mit Press, Cambridge, MA, 1991. [72](#)
- [FT00] Bernard Fortz and Mikkel Thorup. Internet traffic engineering by optimizing OSPF weights. In *INFOCOM (2)*, pages 519–528, 2000. [9](#), [37](#), [181](#), [182](#)
- [GA98] E. Giovanetti and E. Agliardi. Morphogenesis of an institution on a lattice game. *Discrete Dynamics in Nature and Society*, 2:209–213, 1998. [96](#)
- [Gal77] R. G. Gallager. A minimum delay routing algorithm using distributed computation. *IEEE Transactions on Communications*, COM-25:73–85, 1977. [9](#)
- [Gio02] E. Giovanetti. Interconnection, differentiation, and bottlenecks in the internet. *Information Economics and Policy*, 14(3):385–404, 2002. [96](#)
- [GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability : A Guide to the Theory of NP-Completeness (Series of Books in the Mathematical Sciences)*. W. H. Freeman, January 1979. [9](#), [11](#), [63](#)
- [GK98] R. Gibbens and F. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35, 1998. [23](#), [46](#), [68](#), [92](#)
- [GM94] M. Gell-Mann. *The Quark and the Jaguar*. Freeman and Company, 1994. [25](#), [26](#), [107](#), [109](#), [143](#), [224](#)
- [GM95] M. Gell-Mann. What is complexity. *Complexity*, 1(1), 1995. [14](#), [26](#)

## REFERENCES

---

- [GP02] Amos Golan and Jeffrey M. Perloff. Comparison of maximum entropy and higher-order entropy estimators. *Journal of Econometrics*, 107(1-2):195–211, March 2002. available at <http://ideas.repec.org/a/eee/econom/v107y2002i1-2p195-211.html>. 107
- [Gru98] B. Gruschke. Integrated event management: Event correlation using dependency graphs. *DSOM '98*, 1998. 45
- [GY02] Cigdem Gündüz and Bülent Yener. Accuracy and sampling trade-offs for inferring internet router graphs. Technical report, Rensselaer Polytechnic Institute, 2002. 8, 43
- [Hay45] F.A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, September 1945. 98
- [HCF95] K. Houck, S. Calo, and A. Finkel. Towards a practical alarm correlation system. In *Proceedings of the fourth international symposium on Integrated network management IV*, pages 226–237, London, UK, UK, 1995. Chapman & Hall, Ltd. 45
- [HK88] Werner Hildenbrand and Alan P. Kirman. *Equilibrium Analysis — Variations on Themes by Edgeworth and Walras*. Advanced Textbooks in Economics. North-Holland, Amsterdam, 1988. 235
- [Iga03] A.U. Igamberdiev. Living systems are dynamically stable by computing themselves at the quantum level. *Entropy*, 5:76–87, 2003. 210
- [Jay57] Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106:620–630, 1957. 80, 207
- [JP93] Joachim F. Jordaan and Martin Paterok. Event correlation in heterogeneous networks using the osi management framework. In *Proceedings of the IFIP TC6/WG6.6 Third International Symposium on Integrated Network Management with participation of the IEEE*

## REFERENCES

---

- Communications Society CNOM and with support from the Institute for Educational Services*, pages 683–695. North-Holland, 1993. 45
- [JR99] E. Jondeau and M. Rockinger. Estimating gram-charlier expansions with positivity constraints. Papers 56, Banque de France - Direction Generale des Etudes, 1999. available at <http://ideas.repec.org/p/fth/banfra/56.html>. 148, 149
- [Kel56] J. Kelly. A new interpretation of information rate. *Bell Sys. Tech. Journal*, 35:917–926, 1956. 91
- [Kep02] Jussi Keppo. Pricing of bandwidth derivatives under network arbitrage conditions. In *INFORMS International, Hawaii 2001*, 2002. 28
- [KL51] S. Kullback and R.A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951. 225
- [KL92] Y. Korilis and A. Lazar. Why is flow control hard: Optimality, fairness, partial and delayed information. In *Proc. 2nd ORSA Telecommunications Conference*, March 1992. 10
- [Kle75] L. Kleinrock. *Queueing Systems Volume 1: Theory*. Wiley, 1975. 167
- [KM99] P. Key and D. McAuley. Differential qos and pricing in networks: where flow control meets game theory. In *IEE Proc Software*, 1999. 46
- [KMT98] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, volume 49, 1998. 46, 89, 156
- [KP95] B. Kalyanasundaram and Kirk Pruhs. Speed is as powerful as clairvoyance [scheduling problems]. In *FOCS '95: Proceedings of the 36th Annual Symposium on Foundations of Computer Science*



## REFERENCES

---

- (*FOCS'95*), page 214, Washington, DC, USA, 1995. IEEE Computer Society. 42
- [KP97] S. Kätker and M. Paterok. Fault isolation and event correlation for integrated fault management. In *Proceedings of the fifth IFIP/IEEE international symposium on Integrated network management V : integrated management in a virtual world*, pages 583–596, London, UK, UK, 1997. Chapman & Hall, Ltd. 45
- [KS95] Katzela and Schwartz. Schemes for fault identification in communication networks. *IEEE TNWKG: IEEE/ACM Transactions on Networking* IEEE Communications Society, IEEE Computer Society and the ACM with its Special Interest Group on Data Communication (SIGCOMM), ACM Press, 3, 1995. 45
- [KS02] Yuichi Kitamura and Michael Stutzer. Connections between entropic and linear projections in asset pricing estimation. *Journal of Econometrics*, 107(1-2):159–174, March 2002. available at <http://ideas.repec.org/a/eee/econom/v107y2002i1-2p159-174.html>. 92
- [KV05] Frank Kelly and Thomas Voice. Stability of end-to-end algorithms for joint routing and rate control. *SIGCOMM Comput. Commun. Rev.*, 35(2):5–12, 2005. 210
- [KYY<sup>+</sup>95] S. Klinger, S. Yemini, Y. Yemini, D. Ohsie, and S. Stolfo. A coding approach to event correlation. In *Proceedings of the fourth international symposium on Integrated network management IV*, pages 266–277, London, UK, UK, 1995. Chapman & Hall, Ltd. 45
- [LBCX02] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in ip topology measurements. Technical report, Boston University, Computer Science, 2002. 30
- [Lew93] Lundy M. Lewis. A case-based reasoning approach to the resolution of faults in communication networks. In *Proceedings of the IFIP*

## REFERENCES

---

- TC6/WG6.6 Third International Symposium on Integrated Network Management with participation of the IEEE Communications Society CNOM and with support from the Institute for Educational Services*, pages 671–682. North-Holland, 1993. 45
- [Lio04] Panagis Liossatos. Statistical entropy in general equilibrium theory. Working Papers 0414, Florida International University, Department of Economics, July 2004. available at <http://ideas.repec.org/p/fu/wpaper/0414.html>. 75
- [LKSS00] Andre Lucas, Pieter Klaassen, Peter Spreij, and Stefan Straetmans. An analytic approach to credit risk of large corporate bond and loan portfolios. Technical report, Abn Amro Bank, December 07 2000. 140
- [LOR<sup>+</sup>01] Dean H. Lorenz, Ariel Orda, Danny Raz, , and Yuval Shavitt. How good can IP routing be? Technical Report 2001-17, Lucent technologies, 2001. 40
- [LS04] Kate Larson and Tuomas Sandholm. Strategic deliberation and truthful revelation: an impossibility result. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 264–265, New York, NY, USA, 2004. ACM Press. 66
- [Man01] Costas Courcoubetis Manos. An auction mechanism for bandwidth allocation over paths\*. Technical report, Department of Informatics, Athens University of Economics and Business, 2001. 47
- [MFMZ02] Pablo Molinero-Fernández, Nick McKeown, and Hui Zhang. Is ip going to take over the world (of communications? In *ACM HotNets-I*, Princeton, NJ, October 2002. 40
- [MHC03] O. Madani, S. Hanks, and A. Condon. The undecidability of probabilistic planning and related stochastic optimization problems. *Artificial Intelligence*, 2003. 33

## REFERENCES

---

- [MMIC01] Juan R. De Miguel, Ghanshyam B. Mehta, Esteban Indurin, and Juan C. Candeal. exposita notes : Utility and entropy. *Economic Theory*, 17(1):233–238, 2001. available at <http://ideas.repec.org/a/spr/joecth/v17y2001i1p233-238.html>. 87
- [MMV94] Jeffrey K. MacKie-Mason and Hal R. Varian. Pricing the internet. Computational Economics 9401002, the former EconWPA, January 1994. available at <http://ideas.repec.org/p/wpa/wuwpco/9401002.html>. 47
- [Mor02] Richard Mortier. Internet traffic engineering. Technical Report UCAM-CL-TR-532, University of Cambridge, Computer Laboratory, JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom, phone 44 1223 763500, April 2002. 33
- [MP97] A. E. Mackay and E.Z. Prisman. From utility maximization to arbitrage pricing, and back. Working paper, University of Toronto, 1997. 88
- [MS06] L. M. Martyushev and V. D. Seleznev. Maximum entropy production principle in physics, chemistry and biology. *Physics Reports*, 426(1):1–45, April 2006. 95
- [OyG27] Jose Ortega y Gasset. *Mirabeau and Politics*. CUADERNOS DE LA GACETA, 1927. 101
- [PF95] Vern Paxson and Sally Floyd. Wide area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, 1995. 8
- [PS84] I Prigogine and I Stengers. *Order out of Chaos - Man's new dialogue with nature*. Bantam Books, Toronto, 1984. 95
- [PSF<sup>+</sup>01] C. Palmer, G. Siganos, M. Faloutsos, C. Faloutsos, and P. Gibbons. The connectivity and fault-tolerance of the internet topology. In *Workshop on Network-Related Data Management*, 2001. 6

## REFERENCES

---

- [PTVF92] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992. [149](#), [191](#), [197](#)
- [Rab57] M. O. Rabin. Effective computability of winning strategies. *Ann. of Math. Stud.* (Contributions to the Theory of Games), Princeton, 3(39):147–157, 1957. [63](#)
- [RGR02] Zinovi Rabinovich, Claudia V. Goldman, and Jeffrey S. Rosenschein. Non-approximability of decentralized control. Technical Report 2002-29, Leibniz Center for Computer Science, Hebrew University, Jerusalem, 2002. [32](#)
- [RM00] Joshua S. Richman and J. Randall Moorman. Physiological time-series analysis using approximate entropy and sample entropy. *Am J Physiol Heart Circ Physiol*, 278(6):H2039–2049, 2000. [191](#)
- [RO04] Ebbe Rogge and J. Sch Onbucher. Modelling dynamic portfolio credit risk. Technical report, Abn Amro Bank, June 30 2004. [128](#)
- [Ros73] Arnold L. Rosenberg. On the time required to recognize properties of graphs: a problem. *SIGACT News*, 5(4):15–16, 1973. [30](#)
- [Ros76] Stephen A. Ross. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13:341–360, December 1976. [88](#), [239](#)
- [RV76] Rivest and Vuillemin. On recognizing graph properties from adjacency matrices. *TCS: Theoretical Computer Science*, 3, 1976. [30](#)
- [SB01] National Academy of Sciences Computer Science and Telecommunications Board. *Looking Over the Fence at Networks: A Neighbor's View of Networking Research*. National Academies Press, Washington, D.C., 2001. [20](#)
- [SCEH96] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM Computer Communication Review*, 26:19–43, April 1996. [55](#)

## REFERENCES

---

- [Sch04] E. Schlogl. Some useful results on hermite polynomials under linear coordinate transforms. Technical report, Quantitative Finance Research Centre, University of Technology, 2004. [149](#)
- [Sha64] W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442, 1964. [144](#)
- [She95] Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE TNWKG: IEEE/ACM Transactions on Networking* IEEE Communications Society, IEEE Computer Society and the ACM with its Special Interest Group on Data Communication (SIGCOMM), ACM Press, 3, 1995. [9](#), [46](#)
- [SIG03] SIGCOMM. Sigcomm '03: Proceedings of the 2003 conference on applications, technologies, architectures, and protocols for computer communications. In *SIGCOMM Comput. Commun. Rev.*, New York, NY, USA, 2003. ACM Press. General Chair-Anja Feldmann and General Chair-Martina Zitterbart and Program Chair-Jon Crowcroft and Program Chair-David Wetherall. [21](#)
- [SM99] Neil Stratford and Richard Mortier. An economic approach to adaptive resource management. In *Workshop on Hot Topics in Operating Systems*, pages 142–147, 1999. [47](#)
- [SS01a] Makoto Saito and Shigenori Shiratsuka. Financial crises as the failure of arbitrage: Implications for monetary policy. *Monetary and Economic Studies, Special Edition*, 2001. [240](#)
- [SS01b] M. Steinder and A. Sethi. The present and future of event correlation: A need for end-to-end service fault localization. In *IIIS SCI: World Multi-Conf. Systemics Cybernetics Informatics*, 2001. [44](#)
- [Swe89] R. Swenson. Emergent attractors and the law of maximum entropy production: Foundations to a theory of general evolution. *Systems Research*, 6:183–197, 1989. [96](#)

## REFERENCES

---

- [SZ99] Ion Stoica and Hui Zhang. Providing guaranteed services without per flow management. In *SIGCOMM*, pages 81–94, 1999. [42](#)
- [TCGK02] L. Thiele, S. Chakraborty, M. Gries, and S. Knzli. Design space exploration of network processor architectures. In *First Workshop on Network Processors at the 8th International Symposium on High Performance Computer Architecture*, February 2002. [109](#)
- [TG00] M. E. Torres and L. G. Gamero. Relative complexity changes in time series using information measures. *Physica A Statistical Mechanics and its Applications*, 286:457–473, November 2000. [107](#)
- [TKS<sup>+</sup>02] M. Tsai, C. Kulkarni, C. Sauer, N. Shah, and K. Keutzer. A benchmarking methodology for network processors, February 2002. [52](#), [108](#)
- [Tra88] J. F. Traub. Introduction to information-based complexity. In Y. S. Abu-Mostafa, editor, *Complexity in Information Theory*, pages 62–76. Springer-Verlag, New York, 1988. [64](#)
- [Vel05] K. Vela Velupillai. The unreasonable ineffectiveness of mathematics in economics. *Cambridge Journal of Economics*, 29(6):849–872, November 2005. available at <http://ideas.repec.org/a/oup/cambje/v29y2005i6p849-872.html>. [62](#)
- [Wec97] Warren Weckesser. Symbolic dynamics in mathematics, physics, and engineering. In *IMA Industrial Problems Seminar*, 1997. [104](#)
- [WKZL96] Dallas E. Wrege, Edward W. Knightly, Hui Zhang, and Jörg Liebeherr. Deterministic delay bounds for VBR video in packet-switching networks: fundamental limits and practical trade-offs. *IEEE/ACM Transactions on Networking*, 4(3):352–362, 1996. [6](#)
- [Woo00] M. Wooldridge. The computational complexity of agent design problems. In E. Durfee, editor, *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*. IEEE Press, 2000. [32](#)

## REFERENCES

---

- [WP98] W. Willinger and V. Paxson. Where mathematics meets the internet. *Notices of the American Mathematical Society*, 45(8):961–970, 1998. [8](#), [155](#)
- [XZB05] Kuai Xu, Zhi-Li Zhang, and Supratik Bhattacharyya. Profiling internet backbone traffic: behavior models and applications. In *SIGCOMM '05: Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 169–180, New York, NY, USA, 2005. ACM Press. [210](#)
- [ZF93] Hui Zhang and Domenico Ferrari. Rate-controlled static-priority queueing. In *INFOCOM (1)*, pages 227–236, 1993. [42](#)
- [ZHS99] T. Zhou, X. Hu, and E. H.-M. Sha. A probabilistic performance metric for real-time system design. In *A Probabilistic Performance Metric for Real-Time System Design*, pages 90–94, 1999. [108](#)

# Index

- access control, 10
- active networks, 31, 32
  - definition, 31
- Architectural Model, 124
- Architectural Complexity, 15, 121–123, 157, 168, 176, 180, 203, 205
  - definition, 121
- Architectural Model, 24–26, 34, 54, 103, 104, 109, 117, 123, 124, 126, 132, 156, 163, 174, 179, 182, 188
- arms race, 204, 205
- Arrow-Debreu, 60, 234–237, 240
  - security, 237, 239, 240
- Ashby’s Law of Requisite Variety, 22
- ATM, 31, 49–51, 120, 180, 181, 185–187
- autonomous systems, 38, 140
- axiomatic, 115
- Bernoulli, 11, 38, 118, 125, 127, 137, 145, 189, 230
- Brownian, 83, 139
- Capital Asset Pricing Model, 94, 144, 208
- cascading failures, 29, 209
- Cholesky Decomposition, 137
- complete markets, 28, 88, 237
- complex adaptive systems, 100, 101, 146
- Complexity Neutral, 152
  - definition, 155
- complexity neutral, 203
- complexity plane, 104, 147, 152
  - definition, 152
- conditional independence, 11, 13, 20, 22, 145
- Deadweight Loss, 167, 174, 177
- Deliberation Costs, 111
- Deviations from Simplicity, 14, 108, 110, 112
  - definition, 27
- Dijkstra’s algorithm, 181
- duality, 89, 92
- Dynamic Completeness, 100
- Effective Complexity, vi, 15, 26, 103, 108, 110, 123, 124, 129, 157, 158
  - definition, 106
- Efficient Market Hypothesis, 83
- Element Complexity, 14, 15, 120, 121, 123, 157, 163, 176, 180
  - definition, 122
- End-to-End, 205
- Entropic Formulation, 13, 42, 44, 46, 53, 54, 58, 68, 77, 85, 86, 111, 116, 201, 213, 214, 218, 226



- applying, 218
  - definition, 22
  - dissecting, 83
- Entropic Routing, 4, 18, 36, 108, 168, 177, 181, 202, 216, 217
  - definition, 202
- entropy, 68, 70, 86, 210
  - definition, 116
  - Flooding, 94
  - Maxwell-Boltzmann-Gibbs, 230
  - Renyi, 229
- Entropy Representation Problem, 87
- ergodic, 230
- evolution, 99, 101
- extropy, 95
- factor model, 140, 145
- feature creep, 3, 50, 51, 204, 218
- flow control, 10, 18, 103, 202
- Fluctuation Theory, 95
- fractal, 23, 83, 201, 206
- free energy, 94
- free lunch, 68
- gamma path, 100
- Gaussian, 93, 136, 141, 147, 217
  - coupling, 152
  - Distribution, 141, 218
  - noise, 94
- Gell-Mann, 26, 110
- Gram-Charlier Expansion, 26, 147–149
- Hakanssons Paradox
  - definition, 62
- heartbeat, 210
- ICANN, 97, 207
- idiosyncratic risk, 16, 19, 112, 132, 133, 144, 147, 152, 156, 169, 202, 203, 208, 209, 217
  - definition, 16, 132
- IGP, 181
- Implied Complexity, 15, 104, 129, 155, 158, 191
  - definition, 104
- Implied Coupling, 18, 54, 147, 152, 155, 158, 177, 197–200, 217, 242
  - definition, 152
- implied volatility, 19, 146
- Information, 115, 116, 226
  - definition, 115
- Information Based Complexity, 64
- Information Theory, 18, 22–24, 26, 57, 58, 68, 72, 82, 116, 117
- Interrogation Costs, 23, 72, 104, 111, 116, 124, 164, 174, 180, 182
  - definition, 5
- Investors Problem
  - definition, 72
- irreversible transformation, 74, 86
- Jensen’s Inequality, 114, 118, 160, 163, 164, 225, 227
- Kelly Criterion, 91
- Kolmogorov Complexity, 14, 109, 110
- Kullback-Leibler, 76, 116, 129, 130, 141, 149, 152, 191, 225

- kurtosis, 149
- Lagrangian, 92, 233
  - multipliers, 83, 84, 89, 90
  - shadow prices, 23, 84
- Laws of Thermodynamics, 74
- market clearing, 59–61, 236
- market maetaphors, 46
- market maker, 70–72, 86
- market metaphors, 20, 23–25, 57, 58, 155, 234
- Market-Makers Problem
  - definition, 71
- MaxEnt, 95
- Maximum Entropy Production, 94, 95
- Measurement Model, 26, 33, 54, 103, 104, 108, 112, 125, 126, 132, 147, 151, 152, 155, 156, 168, 179, 188–191, 199, 203, 215, 217
  - definition, 126
- memoisation, 98
- minimax, 72
- Model Complexity, 135
  - definition, 140, 141
- MPLS, 31, 39, 40, 49, 50, 52, 103, 180, 181, 185–187, 204, 206
- Murphy’s Law of Networks, 210
- neoclassical, 14, 27, 29, 32, 53, 55, 56, 58–62, 64, 65, 67, 69, 70, 85, 99, 100, 234–236, 239, 240
- Network Complexity, 18, 103, 158, 170, 203, 242
  - definition, 129
  - quantitative definition, 129
- network precedence matrix, 120, 125
- Occam’s Razor, 126, 140
- OMT, 5
- Operational Availability, 168
- Operational Complexity, 111, 119, 123, 157, 176, 180
  - definition, 122
- Operational Matrix, 14, 30, 120, 163, 168, 182, 187
- OSPF, 38–40, 103, 181, 182
- Pareto
  - Efficiency, 76, 238
  - Improvement, 183
  - Optimal, 237, 238
- Pareto Optimal, 238
- PNNI, 185
- principle of maximum entropy, 225–227
  - definition, 226
- Proportional Fairness, 90
- QoS, 80
- Queueing Theory, 57, 145, 167, 219
- random graph, 78
- Red Queen, 204
- Reference Measurement, 124–126, 128, 129, 143, 144, 189–191
  - definition, 118, 124
- Reference Measurements, 178
- replication argument, 62, 100
  - definition, 240

- ReReshaping the Agenda, [vi](#), [55](#), [56](#),  
[101](#), [212](#)
- reversible transformation, [86](#)
- RIP, [40](#)
- risk-neutral, [61](#), [62](#), [240](#)  
definition, [240](#)
- Scheduling Complexity, [15](#), [123](#), [157](#), [176](#),  
[180](#)  
index, [122](#)
- SDH, [48](#)
- SDL, [124](#)
- Second Law of Thermodynamics, [74](#), [87](#),  
[95](#)  
definition, [74](#)
- Self-Organised Criticality, [95](#)
- semantic, [116](#)
- shadow prices, [23](#), [67](#), [84](#), [89–91](#)
- skewness, [149](#)
- SLA, [47](#)
- Slice, [15](#), [16](#), [112](#), [132](#), [133](#), [150](#), [169](#)  
Curve, [151](#)  
definition, [16](#)
- Slice Curve  
definition, [151](#)
- slice curve, [152](#), [154](#), [197](#), [199](#), [200](#), [202](#),  
[216](#)
- Small-world effect, [78](#)
- SNMP, [189](#)
- SONET, [48](#), [181](#)
- statistical inference, [69](#), [80](#), [82](#), [116](#), [174](#),  
[207](#)
- statistical Mechanics, [74](#)
- statistical mechanics, [25](#), [44](#), [57](#), [58](#), [68](#),  
[69](#), [75–77](#), [79](#), [86](#), [110](#), [213](#), [232](#)
- Stirling’s Approximation, [232](#)
- sub-synergistic, [114](#)
- subjective, [14](#), [26](#), [44](#), [108](#), [109](#), [143](#), [227](#)
- SUBSET SUM, [9](#), [11](#)
- Surprisal, [116](#)
- systematic, [180](#)
- systematic risk, [vi](#), [16](#), [112](#), [132](#), [133](#),  
[147](#), [152](#), [156](#), [158](#), [169](#), [202](#), [208](#),  
[209](#), [219](#)  
definition, [16](#), [132](#)
- tátonnement, [28](#), [59](#), [73](#), [77](#)
- thermodynamics, [86](#)
- Traffic Precedence Matrix, [163](#)
- Traffic Complexity, [157](#), [164](#), [176](#), [180](#)  
definition, [122](#)
- Traffic Precedence Matrix, [121](#), [165](#)
- Traffic Vector, [120](#), [121](#), [123](#), [187](#), [205](#)
- trigger probability, [127–129](#), [135](#), [136](#),  
[169](#), [172](#), [173](#)
- UML, [5](#)
- Uniform Distribution, [84](#)
- utility, [32](#), [46](#), [60](#), [64](#), [75](#), [86](#), [90](#), [91](#),  
[206](#), [236](#), [238](#)  
log, [23](#), [46](#), [68](#), [89](#), [90](#), [92](#), [94](#), [156](#)
- Utility Representation Problem, [87](#)
- Variation Information, [229](#)
- VP and VC Switching, [186](#)
- Walras, [59](#), [60](#), [67](#), [73](#), [99](#), [235](#), [236](#)  
Auction, [59](#), [76](#), [77](#), [236](#)

Equilibrium, [63](#), [64](#), [66](#), [76](#), [77](#)  
World Wide Web, [25](#)  
  
zero complexity, [147](#), [150](#)  
zero-sum game, [70–72](#), [213](#)