DESIGNING OPTIMAL NETWORK TOPOLOGIES UNDER MULTIPLE EFFICIENCY AND ROBUSTNESS CONSTRAINTS

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

by

SANKET PATIL

International Institute of Information Technology–Bangalore

2011

To *Appaji* and *Amma*

Declaration

This is to certify that the thesis titled "Designing Optimal Network Topologies under Multiple Efficiency and Robustness Constraints" being submitted to the International Institute of Information Technology - Bangalore, for the award of the degree of Doctor of Philosophy, is a record of bonafide research work carried out by Sanket Patil.

SRINATH SRINIVASA

17.01.11

INTERNATIONAL INSTITUTE OF INFORMATION TECHNOLOGY BANGALORE

Declaration

This is to certify that the thesis titled "Designing Optimal Network Topologies under Multiple Efficiency and Robustness Constraints" being submitted to the International Institute of Information Technology–Bangalore, for the award of the degree of Doctor of Philosophy, is a record of bona fide research work carried out by me. The results in this thesis have not been submitted to any other university or institute for the award of any degree or diploma.

Date: 28^{th} June, 2011

Sanket Patil
International Institute of Information Technology–Bangalore
India

Acknowledgements

I would like to begin by thanking my advisor Prof. Srinath Srinivasa. Over the last seven years that I have worked with him he has patiently helped in refining my flashes of insights and inarticulate ideas. He has constantly steered my work towards being substantial. I have also enjoyed long hours of discussions with him on various technical and non-technical topics. I look forward to collaborating with him in the future as well.

I have been associated with the Open Systems Laboratory (OSL) since 2003. It has been a great association and one of its kind. Most of what I have learned about the research lifestyle I have learned through OSL. I want to thank all the present and past members of OSL for creating and maintaining a vibrant atmosphere. Specifically, I want to thank my colleagues Mandar Mutalikdesai, Karthik B. R., Saikat Mukherjee and Aditya Ramana Rachakonda for the numerous useful discussions that we have had in the lab; and also for all the fun we had together. I also want to acknowledge two of the Master's students–Reshma Ratnani and Roshni T. Raj–for helping in implementing a tool called TopAZ which I extensively used for my thesis.

I was at Purdue University for about eighteen months during my PhD as a visiting scholar. I was invited by Prof. Venkat Venkatasubramanian who is also my co-advisor. His work in the area of complex networks gave a good direction for my thesis work. A part of the thesis work is an extension of his work. I have learned a lot from working with him and interactions with him have always been lively. I am also grateful to him for hosting me at Purdue and helping me through a difficult phase there.

I was working as part of Laboratory for Intelligent Processing Systems (LIPS) during my stint at Purdue. I am thankful to members of LIPS with whom I had many useful discussions. I had an extended collaboration with Dr. Tanu Malik (who was at Purdue University at the time) and Prof. Amitabh Chaudhary (University of Notre Dame) with whom I coauthored a paper that extended my thesis work. I thank them both.

I want to thank Prof. P. C. P. Bhatt who was a part of my comprehensive exam committee and gave useful advice that helped in addressing some of the shortcomings in the thesis. On the same note, I want to thank Prof. Dinesha and Prof. Shrisha Rao for helping me refine my problem formulation. While I have not had much chance to discuss my thesis work with Prof. Mahabala and Prof. Prabhu, every time I have met them I have learned something new.

I want to thank the members of my thesis review committee who gave detailed insightful comments about my PhD research work. The comments have helped me greatly in preparing this final version of my thesis.

On the same note, I would like to thank the members of my oral examination (PhD Defense) committee for raising a number of interesting points of discussion and possible future research directions. I want to especially mention Prof. S. K. Gupta (IIT-Delhi) and Prof. Rajagopalan (IIIT-Bangalore) in this regard.

I have spent a significant portion of my last eight years in Bangalore at IIIT–Bangalore. I joined IIIT–Bangalore after my under-graduation not wanting to take the usual path. And that has actually made all the difference! I want to extend my gratitude to Prof. Sadagopan for his constant support and encouragement. I want to thank Mr. Abraham who was the registrar during most of my IIIT life for helping in administrative matters and listening patiently to our grievances. I also want to thank the administrative staff: Mr. Murugan, Ms. Chandrika, Mr. Somashekhar, and the new registrar Mr. Ramachandra.

I was an Infosys scholar at IIITB. I am thankful to Infosys for assisting me financially, thus providing me with an opportunity to pursue long term research. I am also grateful to Siemens Corporate Research and Technologies, where I am now an employee after starting there as an intern, for supporting me through the last leg of my PhD.

The hazards of PhD life would be insufferable without the constant support of friends and family. I want to especially thank Pradeep, Mandar, Gopi, Chitra and Dattu for always being there. Thanks for listening to the woes, and thanks for all the cathartic sessions we have had! I want to fondly mention my friend Neelakanth with whom I keep exchanging ideas regarding classical music, literature, philosophy, and of course, my thesis. I cherish the great doses of fun that I have had with my brother Nikshep and my extended family.

I want to thank Sur—first a close friend and now my partner—for the love, care and most importantly, patience. One also needs some external pressure to reach a logical conclusion in research, which she ably provided.

I made the final nitpickings at this thesis through the night of Sunday, the 26^{th} of June, 2011, in a hospital ward while I and Sur awaited our baby! I did manage to compile the final version by Monday morning just before our daughter arrived. At a later date she might claim she had a role to play in my thesis! I would perhaps concede.

Finally, my parents have not only allowed me to make my own choices in my career as well as personal life but also supported me to the hilt. But they would surely frown upon a "thanks" from me! I dedicate this thesis to my parents.

Abstract

Designing optimal network topologies is an important problem across many application domains, such as: distributed information systems, supply networks, content delivery networks and network-centric warfare. The requirements of optimality vary with the purpose for which a network is built. Further, there are conflicting optimality requirements within a network that need to be balanced.

The operational objective of network design is to minimize the "cost of communication" in a network, i.e. to maximize network *efficiency*. However, efficiency has to be achieved under several constraints. The lack of reliability on the part of machines and links poses issues of resilience (or *robustness*) of the network in the face of failures. Since nodes and links can fail, it might be important to have alternate communication paths between pairs of nodes.

The number of links that constitute a network poses infrastructure and maintenance costs. An asymmetry in the distribution of links across nodes poses issues of load balancing and congestion. Congestion in turn can cause high latency, loss of network data and low availability, thus reducing the performance of a network.

These disparate requirements pose conflicting constraints on network design. For example, in a data-centric p2p network, the index to data elements is distributed across all nodes in the network to form a *symmetric* graph. Such an arrangement (topology) makes the distributed index robust to failures and load balanced. It also distributes "bookkeeping" costs uniformly. However, it increases the *lookup cost* when compared to a traditional search-tree like index structure.

In case of supply networks, facilities are susceptible to failures due to disasters such as hurricanes and earthquakes. Recovering from failures of facilities can take several days. Therefore, it is useful to build redundant facilities to minimize operational delays. However, redundancy increases set up and maintenance costs.

In this thesis, we address the kind of network design problems under multiple constraints as described above. We model the problem of network design for different design metrics and trade-offs. Every combination of metrics corresponds to performance requirements of a class of networks. Further, within a class

of networks, the relative emphasis on optimization parameters might vary across specific deployments.

We use three critical system parameters, *efficiency*, *robustness* and *cost* to model performance requirements. Two application dependent environment variables are used to vary the relative importance between the above parameters. Using a genetic algorithm process called *topology breeding*, we evolve optimal topologies under different environmental conditions.

A triple of performance metrics, *(efficiency, robustness, cost)*, is associated with the notion of an *Optimal Topology Space* (OTS)—a space formed by the three performance parameters in which every point represents a trade-off between them. We use topology breeding to "sample" this space to discover optimal topologies at different points. This results in a compendium of topologies from which a designer can choose depending on specific requirements.

We observe some special families of topologies that recur in specific "sub-spaces" in an OTS: trees, hubs–and–spokes, circular hubs and circular symmetric topologies. In general, we observe two prominent sub-spaces or classes of topologies: (1) *Hub–and–Spoke* topologies, with high efficiency, high resilience to random failures and low cost, and (2) *Circular Skip Lists* (CSL), with high robustness to random failures as well as targeted attacks, and high efficiency at moderate cost.

Further, we observe a sharp transition from hub–and–spoke topologies to CSL as emphasis on robustness increases or cost restrictions become less severe. We analyze CSLs to observe that they show structural motifs that are optimal with respect to several metrics: symmetry, low diameter, high connectivity and hamiltonicity. The findings from this work corroborate many recent results concerning optimal network design in applications such as distributed hash tables and communication networks. Based on these results and our empirical observations, we claim that CSLs form potential underpinnings for optimal network design under a variety of performance requirements.

Since CSLs are observed to be important structures for network design, we develop simple deterministic algorithms to construct them. We also provide some theoretical guarantees about the optimality of *regular graphs* (almost all of which are CSLs) under perturbations such as node arrivals and departures. For instance, we develop techniques to extend a given regular graph while preserving its regularity and/or connectivity and/or hamiltonicity.

Contents

1 Introduction **1**

1.1 Network Optimality: Case Studies 3

1.2 Performance Trade-offs in Network Design 7

1.3 Optimal Topology Spaces 9

1.4 Contributions 13

1.5 Organization 15

2 Background **17**

2.1 Special Families of Graphs 21

2.2 Network Design: Application Areas 31

2.2.1 Steiner Trees and Applications 31

2.2.2 Distributed Hash Tables (DHTs) 32

2.2.3 Complex Networks 36

2.3 Network Design: Philosophical Understanding 43

2.3.1 Statistical Physics of Networks 43
2.3.2 Emergence vs Design 45
2.4 Positioning Our Work 47

3 A Genetic Algorithm Framework 49

3.1 Network Design: Optimization Problem 51
3.2 Optimization Framework 55
3.2.1 Performance Metrics 56
3.2.2 Efficiency 57
3.2.3 Robustness 59
3.2.4 Cost 64
3.2.5 Fitness 66
3.3 Genetic Algorithm Formulation 67
3.3.1 Initialization 68
3.3.2 Selection 70
3.3.3 Reproduction 70
3.3.4 Termination 73
3.4 Topology Breeding 74

4 Topology Breeding: Optimal Topology Spaces 75

4.1 Undirected Optimal Topology Spaces (UDOTS) 76
4.1.1 Diameter vs Degree Centrality 76
4.1.2 Diameter vs Edge Connectivity 86
4.1.3 APL vs Functional Robustness 89

4.2 Directed Optimal Topology Spaces (DOTS) 97
4.2.1 Diameter vs Degree Centrality 98
4.2.2 Diameter vs Edge Connectivity 100
4.2.3 APL vs Functional Robustness 101

5 Characterizing Optimal Topology Spaces 106

5.1 Classes of Optimal Topologies 107
5.1.1 Hubs–and–Spokes 107
5.1.2 Circular Hubs 109
5.1.3 Symmetric Multihubs 110
5.1.4 Circular Skip Lists 112
5.2 “Phase Transitions” from Hubs–and–Spokes to CSL 114
5.3 Distribution of Optimal Topologies 118
5.4 Optimal Structural Motifs of CSLs 127

6 Deterministic Techniques 131

6.1 Circular Skip Lists: Definition 133
6.2 Variants of Graph Powers 135
6.2.1 $\mathcal{G} + E_r$–Graphs as Circular Skip Lists 136
6.3 (Balanced) Base–b Networks 144

7 Regular Graphs 152

7.1 Connected Regular Graphs: Basics 154
7.2 Extending Regular Graphs 156

7.3 Regularity, Connectivity and Hamiltonicity 162

7.3.1 Regularity and Connectivity 163

7.3.2 Regularity and Hamiltonicity 168

7.4 Regular Directed Graphs 171

8 Concluding Remarks 175

8.1 Future Directions 178

8.1.1 Structural Signatures and Snapping 178

8.1.2 Adaptable Networks under Varying Performance Requirements 179

A List of Related Publications 181

List of Figures

2.1 Families of Graphs 23

3.1 GA Initialization: Sample Chromosomal Strings 69

3.2 Single Point Equal Edges Crossover 72

4.1 d vs p_{skew} (p, α, κ, $\beta = 0$) 78

4.2 d vs p_{skew} (p, κ, α, $\beta = 0$) 79

4.3 d vs p_{skew} (p, κ, $\beta = 1$) 83

4.4 d vs p_{skew} (p, κ, $\beta = 1$) 84

4.5 λ (p, κ, $\alpha = 0$, $\beta = 1$) 87

4.6 λ (p, κ, $\alpha = 1$, $\beta = 0$) 88

4.7 l vs ρ_{fn} (degree) (p, κ, $\beta = 0$) 90

4.8 l vs ρ_{fn} (degree) (p, κ, $\beta = 1$) 91

4.9 l vs ρ_{fn} (node betweenness) (p, κ, $\beta = 0$) 94

4.10 l vs ρ_{fn} (node betweenness) (p, κ, $\beta = 0.4$) 95

4.11 l vs ρ_{fn} (node betweenness) (p, κ, $\beta = 1$) 96
4.12 d vs ρ_{skew} (p, κ, $\beta = 1$) 98
4.13 λ (p, κ, $\beta = 0$) 100
4.14 l vs ρ_{fn} (degree) (p, κ, $\beta = 1$) 102
4.15 l vs ρ_{fn} (degree) (p, κ, $\beta = 0.6$) 103
4.16 l vs ρ_{fn} (degree) (p, κ, $\beta = 0.6$) 104

5.1 Undirected Hubs–and–Spokes 107
5.2 Undirected Circular Hubs 109
5.3 Undirected Symmetric Multihubs 111
5.4 Undirected Circular Skip Lists 113
5.5 ϕ vs. α vs. p 115
5.6 Fitness vs α 116
5.7 Value vs α 117
5.8 η vs. ρ vs. κ for η_d and ρ_p 119
5.9 ρ vs. η vs. κ for η_d and ρ_p 120
5.10 η_d vs. ρ_p 121
5.11 η_d vs. κ 122
5.12 ρ_p vs. κ 123
5.13 η vs. ρ vs. κ for η_l and ρ_{fn} 124
5.14 ρ vs. η vs. κ for η_l and ρ_{fn} 124
5.15 η_l vs ρ_{fn} 125
5.16 η_l vs κ 125

5.17 ρ_{fn} vs. κ 126

5.18 ρ vs. η vs. κ for η_l and ρ_{fn} 127

5.19 An example directed CSL 128

6.1 r–polygon insertion 137

6.2 gcd–polygon insertion 138

6.3 Hamiltonian insertion 139

6.4 “Cartwheel” like structure 140

6.5 Undirected polygon insertion 142

6.6 Polygon-Halving Insertion 143

6.7 Balanced Ternary Network–Formation 147

6.8 A balanced ternary network with $n = 27$ 147

6.9 Ternary Network–Formation 150

7.1 A regular graph with a cut-edge 156

7.2 Extending even regular graphs 158

7.3 Extending odd regular graphs 159

7.4 Extension by merging 161

7.5 Connectivity preserving extension (even regular) 164

7.6 Connectivity preserving extension (odd regular) 166

7.7 Regular directed graph 172

List of Tables

1.1 Relative emphasis on performance parameters 10

3.1 GA Reproduction: List of crossover operators used in topology breeding 71

5.1 Structural properties of an example directed CSL 128

The White Rabbit put on his spectacles. "Where shall I begin, please your Majesty?" he asked.
"Begin at the beginning," the King said gravely, "and go on till you come to the end: then stop."

Lewis Carrol, *Alice in Wonderland*

1

Introduction

A crucial element of performance in large scale decentralized systems is the network infrastructure. As such, designing optimal network infrastructures is an important problem across many application domains, such as: distributed information systems, content distribution networks (CDN), supply chain networks, transportation networks, and recently, network centric warfare (NCW).

Optimality of a network has different connotations in different application domains determined by the purpose for which the network is built. Further, within an application context, there are conflicting performance requirements that need to be balanced as we describe later on in this section. Thus, designing a network is a multi-objective optimization problem.

The operational objective of network design is to minimize the "cost of communication" in a network (in other words, to maximize network *efficiency*). However, efficiency has to be achieved under several constraints. The lack of reliability on the part of nodes (e.g., machines or facilities) and edges (e.g., communication links) poses issues of resilience (or *robustness*) of the network in the face of failures. Since nodes and links can fail, it might be important to have alternate communication paths between pairs of nodes. The number of links that constitute a network poses infrastructure and maintenance costs. An asymmetry in the distribution of links across nodes poses issues of load balancing and congestion.

These disparate requirements pose conflicting constraints on network design. In this thesis, we address the problem of designing optimal networks under multiple conflicting performance requirements. We model performance requirements in terms of three parameters that are critical to any distributed system: *efficiency*, *robustness*, and *cost*. We propose that network design is governed by the trade-offs between the above three parameters. We use a genetic algorithm process to traverse the space formed by efficiency, robustness and cost, and uncover optimal topologies at various points in this space.

By conducting a series of such experiments, we discover topologies for several classes of performance requirements. We observe and highlight certain recurring patterns or *optimal structural motifs* that potentially form useful guidelines for designing optimal networks under multiple constraints.

There is a lot of work towards addressing the problem of optimal network design in the last five decades. In general, approaches towards addressing this problem in literature are of the following type: (1) designing networks for specific (domain dependent) performance requirements, (2) evaluation of families of graphs in search of specific optimal properties; using this knowledge in network design, and (3) efforts towards explaining the theoretical underpinnings or mechanisms that lead to certain optimal properties or behaviours in networks.

We take an approach that is different from the above. While we consider performance requirements of specific domains as motivating examples, the design process itself is domain independent and relies on optimization parameters abstracted in terms of structural (graph theoretical) properties. Instead of starting from well known families of graphs, we identify families of optimal topologies through a process of evolutionary optimization. Finally, we study some of these families (or classes), analyze properties that lead to their optimality and develop theoretical results concerning these properties.

1.1 Network Optimality: Case Studies

Below we briefly describe examples from different domains to motivate the problem of network design under multiple constraints.

Distributed Indices for p2p Networks: A distributed index (DI) is the core element of any distributed hash table (DHT), which in turn is the basic data structure of a data-centric network. This is an index structure that is spread across the system and used for data lookup. For purposes of load balancing and fault tolerance, such index structures are distributed across all machines in the system (or at least across all the reliable machines in the system). The efficiency of the index, and that of the system itself, is determined by the lookup complexity. This is a bound on the number of lookups or network connections that an application program needs to make, before it finds the required data element.

Unlike conventional index structures that are based on variants of search trees, distributed index structures do not have a single point of entry. A tree-like distributed data structure would be very efficient, with the bound on the number of lookups being equal to the height of the tree, which is small. However, a tree-like structure would also mean that the machine hosting the root node bears a disproportionately high lookup load. Distributed indices are hence typically designed as graphs where lookup requests can start from anywhere in the data structure.

Designing a distributed index for a p2p network also involves constraints such as: handling *churn* [Stutzbach and Rejaie, 2006] (or the frequent arrivals and departures of nodes), unreliable connections, and bandwidth bottlenecks. In addition, there may be restrictions on the storage and computing abilities of the participating nodes.

Thus, the DI topology has to be such that the lookup complexity is optimized, or at least bounded to a reasonable maximum number of lookups per key. The lookup complexity is hence bounded by the *diameter* of the DI topology. The design objective of a DI can be described as: bounding the lookup complexity under constraints on symmetric load distribution and robustness.

Network Centric Warfare (NCW): In recent times, network centric warfare has received a lot of research interest [Alberts et al., 1999]. NCW is a doctrine propounded by the United States Department of Defense. The objective of NCW is to increase performance by basing the activities of a battle around an optimal information network. Such an information network helps in sharing information, which leads to improved quality of *situational awareness* and collaboration, thus increasing the effectiveness of missions [Cebrowski, 2005]. Similar philosophies have emerged elsewhere, viz; network enabled capability (United Kingdom) and network based defense (Sweden).

An important challenge in NCW arises due to the high probability of *targeted attacks* on nodes and links. Therefore, robustness in the event of an attack is the most important design objective. Targeted attacks are different from random node failures that is characteristic of say, p2p networks. Designing for robustness against targeted attacks requires that any node is no more or less "important" than any other node [Dekker and Colbert, 2004]. In other words, the loss of any node (or edge) should not cause (significantly) more damage than the loss of any other.

At the same time, the network should be cost effective. While additional links might increase the efficiency, by creating shorter paths, and robustness, by creating alternate paths, they also increase the infrastructure cost of a network. A high number of links also implies a high maintenance cost. Thus, in NCW, the design objective is to build a symmetric and robust network under constraints on infrastructure and maintenance costs.

Supply Networks: A supply network (SN) design is a classic *facility location* problem [Drezner, 1995]. The primary design objective is to build manufacturing, processing and distribution facilities such that the installation and operational costs are minimized. However, there are constraints posed by dependencies between facilities, environmental uncertainties such as variable demand and disruptions [Santoso et al., 2005].

Installation of manufacturing and processing facilities are largely governed by geographical constraints such as availability of raw material. Warehouses should be installed so as to cover demand from all parts of a state or country, while minimizing the cost of transportation. Disruptions such as hurricanes or floods force huge losses on facilities, also making them unavailable during the time of recovery. Therefore, the design needs to consider adding redundant units. However, redundant units increase costs of set up and maintenance. Thus, designing an SN involves designing a cost effective network while ensuring multiple supply routes at least in parts of the network that are important and/or vulnerable.

Content Delivery Networks (CDN): The important challenge in a content delivery network is congestion in traffic flow. Congestion leads to increased latency hence decreasing the network efficiency. The main source of congestion is imbalanced load distribution. A hub–and–spoke topology achieves a very small average node separation with a small number of communication links. But such a topology is susceptible to congestion. This is because, there is a small set of nodes through which most of the traffic flows. Therefore, building a topology that results in load balancing is important. However, a topology that balances load on nodes might neither be very efficient nor economical.

Hence, the objective is to build a network such that congestion is minimized under constraints on cost.

1.2 Performance Trade-offs in Network Design

This thesis addresses the kind of network design problems that are motivated above. The objective of this work is to look for underlying recurring patterns of design when networks need to satisfy multiple performance requirements. A compendium of such patterns would serve as a handy guide to the network designer, as well as help in future research seeking to probe into underlying principles behind network design.

We extend the work of [Venkatasubramanian et al., 2004] on the emergence of *complex network* topologies under different constraints. They propose a formalism based on the Darwinian theory of evolution: the structure of a network evolves so as to optimize the short term liveness needs (efficiency) of a network as well as its long term survival chances (robustness). A *selection pressure* variable decides the trade-off between short term and long term objectives. Different structures emerge for different trade-offs. Hence, the structure of a network is a result of adaptation under selection pressure.

We extend this formalism to develop the concept of *optimal topology spaces*

(OTS). An OTS is formed by the three critical system parameters that pose constraints: *efficiency*, *robustness* and *cost*. Efficiency measures the cost of communication in a network. The lower the cost of communication, the higher the efficiency of a network. Robustness measures the resilience of a network to node and edge failures, and uncertain traffic patterns. Cost measures the setup and maintenance costs of a network.

We use multiple performance metrics to model different classes of networks. Each problem instance results in an OTS as we explain in Section 1.3. By "sampling" an OTS at different points, we observe optimal network topologies for the corresponding trade-offs.

We model efficiency, robustness and cost in terms of structural or graph theoretic properties. A structural metric models only certain aspects of network performance. Therefore, we use multiple metrics to model network performance. The specific structural metrics that we use depend on a network's performance characteristics.

As we discussed previously, the three critical parameters, efficiency, robustness, and cost, pose conflicting requirements to network design. In order to design a network, we need to consider the trade-offs between the above parameters. The trade-offs vary with applications, depending on their relative importance in each application context. Some examples follow:

1. In a p2p network, failure of nodes and links are random, frequent and inexpensive; whereas in a supply network, failures can be either random or due to a deliberate attack, infrequent yet very expensive. Therefore, robustness

is more critical in the design of supply networks than in the design of p2p networks.

2. In a content distribution network, load balancing is an important challenge in order to reduce congestion and increase efficiency; whereas in a data grid or an airline network, it is efficient and cost effective for a small number of high performance nodes (and/or links) to carry most of the load in the network.

3. In an information network such as a p2p network, the cost of setting up links is inexpensive. So, more links can be added if that improves efficiency and/or robustness. Whereas adding links is expensive in most networks that require physical infrastructure, such as transportation networks or supply networks.

In table 1.1, we note typical emphasis given to design challenges in different networks. Please note that the values we have mentioned are only indicative of the general case; they might vary in special cases.

1.3 Optimal Topology Spaces

We model optimality or performance requirements in terms of trade-offs between efficiency, robustness and cost. The optimization technique we use is genetic algorithms. A fitness function governs the optimization process. Fitness (ϕ) is a

	Cost of Communication	**Resilience to Failures**	**Load Balancing**	**Infrastructure Cost**	**Maintenance Cost**
DHT	Moderate	Moderate	High	Moderate	Moderate
NCW	Moderate	High	High	Moderate	Moderate
SN	High	Moderate	Low	High	High
CDN	Moderate	High	High	Moderate	Moderate
Data Grids	High	High	Low	Low	Low
Airlines	High	Moderate	Low	High	High

Table 1.1: Relative importance for optimizing performance parameters in different application domains: distributed hash tables (DHT), network centric warfare (NCW), supply networks (SN), content distribution networks (CDN), data grids and airline networks. The above values are only indicative of the general case.

function of efficiency (η), robustness (ρ) and cost (κ). Environmental variable, α ($0 \leq \alpha \leq 1$), decides the trade-off between efficiency and robustness. Environmental variable, β ($0 \leq \beta \leq 1$), acts as a cost control parameter.

The generic fitness function is given by the following relation:

$$\phi = \alpha\rho + (1 - \alpha)\,\eta - \beta\kappa$$

We are interested in designing (strongly) connected network topologies. Every topology has a *fitness* value. The generic optimization problem is to find the set of edges, $\mathcal{E}$, in a graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, such that, the fitness, $\phi(\mathcal{G})$, is maximized:

$$\arg\max_{\mathcal{E}} \phi\left(\mathcal{G}\left(\mathcal{V}, \mathcal{E}\right)\right)$$

We let a genetic algorithm process evolve the fittest topology, which we then propose as the optimal topology under a given set of constraints.

Every instance of the above optimization problem corresponds to an optimal topology space. By varying the constraints and trade-offs, we evolve optimal topologies at different points in an OTS. In Chapter 4, we present optimal topology spaces corresponding to different problem instances. Figure 4.3 represents an example OTS uncovered as part of this work. It corresponds to an experiment in which the optimization objective is to minimize diameter, d, (maximize efficiency, η), under different values of maximum permissible degree, p, for nodes, given a limit on the maximum number of undirected edges, e, in the graph. Emphasis on robustness is dictated by the values of α. Further, since the cost control parameter, $\beta = 1$, topologies try and "shed away" as many superfluous edges—edges that do not contribute in increasing the efficiency and/or robustness—as possible.

We can see several patterns in the optimal topology space. Trees emerge as optimal topologies under severe restrictions on cost (for low values of e or κ). A class of topologies that may be termed as "hubs–and–spokes" emerges as optimal when the constraint on the maximum permissible degree is relaxed. Circular hubs, multihubs and "circular skip lists" (circular structures with "chords" of different lengths) emerge as optimal topologies under intermediate values for the limit on the maximum number of edges and maximum permissible degree. Circular skip lists are prevalent everywhere in the optimal topology space except at the lower boundary of maximum edges and the upper boundary of maximum permissible degree (refer to Figure 4.4).

Figure 4.13 corresponds to an experiment in which the optimization objective is to maximize edge connectivity, λ, for different values of maximum permissible outdegree, p^{out}, and maximum number of directed edges, e. Here, complete emphasis is placed on robustness ($\alpha = 1$), and none on efficiency. Also, since there is no cost control ($\beta = 0$), no shedding away of edges occurs.

Circular skip lists (CSL) are prevalent throughout the space in Figure 4.13 because: (1) there is complete emphasis on robustness, and (2) cost constraints are not severe. Further, in directed graphs, CSLs are more natural than in undirected graphs even when the emphasis on robustness is not high. This is because, in order to ensure strong connectivity, which is our basic requirement for a valid topology, directed topologies have to be mostly rings or CSLs, unless we are at the upper boundaries of maximum permissible in and out degrees, p^{in} and p^{out}, where strongly connected hubs–and–spokes are possible.

Further, both in undirected and directed topologies, we observe that the trade-off between efficiency and robustness is pronounced only under severe restrictions on cost. In all other cases, CSLs are optimal topologies in terms of balancing efficiency, robustness and cost. The transition from hub–and–spoke topologies to CSLs is sharp with CSLs emerging as soon as the constraint on cost is not severe. CSLs are characterized by several motifs that are optimal under a variety of requirements: low diameter and low average path length; resilience to random failures as well as targeted attacks due to the presence of multiple independent paths; homogeneity with near zero entropy of degree distribution, near zero skew in node and edge betweenness; and hamiltonicity.

In this thesis, we describe the above and other optimal topology spaces uncovered by an exploratory technique based on genetic algorithms. We call this exploration as "topology breeding". Exploration of such spaces was done for both undirected and directed graph topologies.

1.4 Contributions

The contributions of this thesis may be enumerated as follows:

1. [Venkatasubramanian et al., 2004] addressed the problem of network design as an evolutionary optimization problem involving the trade-offs between three system parameters, efficiency, robustness and cost. We extend their work by modelling performance requirements and trade-offs in various classes of networks. We model optimization objectives and constraints occurring in different classes of networks as instances of the above problem. Further, we define a number of new optimality metrics for efficiency, robustness and cost in terms of structural properties.

2. For every instance of the problem, the three critical parameters efficiency, robustness and cost form an *optimal topology space*. Every point in this

space represents a certain trade-off between the parameters. We model these trade-offs using two environmental variables: (1) α, $(0 \leq \alpha \leq 1)$ is a slider that indicates the emphasis on robustness in designing a network, and (2) β, $(0 \leq \beta \leq 1)$, is a slider that indicates the emphasis on cost control. We "sample" the optimal topology space to find the optimal topologies at different points using a genetic algorithm.

3. We observe recurring optimal topology patterns in optimal topology spaces: for example, hubs–and–spokes, circular hubs, symmetric multihubs and circular skip lists. These patterns occur in different *sub-spaces* in an OTS.

4. We study one sub-space, that of circular skip lists, in detail. Circular skip lists are featured with several optimal structural properties such as, high connectivity, low diameter, presence of hamiltonian circuits and high symmetry. Each of these properties is important for network design in various application areas.

   Based on empirical observations in our work and similar results obtained independently by other studies, such as [Guimerá et al., 2002; Valente et al., 2004; Zhao et al., 2005; Donetti et al., 2005, 2006; Mondragón C, 2006], we claim that CSLs have the potential to form theoretical underpinnings for network design under constraints on symmetry, which is an important property in applications such as DHTs, CDNs and NCW.

5. Since CSLs are observed to be important structures for network design, we develop simple deterministic algorithms to construct circular skip lists

with interesting properties. Further, we provide some theoretical guarantees about the optimality of regular graphs (almost all of which are circular skip lists) under perturbations such as node arrivals and departures. For instance, we develop techniques to extend a given regular graph while preserving its regularity and/or connectivity and/or hamiltonicity.

1.5 Organization

The rest of this thesis is organized as follows. We start with an overview of existing literature covering various aspects of network design (Chapter 2). We then proceed to formally model network design as an optimization problem. We propose a design framework for optimal topologies based on evolutionary optimization (Chapter 3). We also describe in brief our experimental setup called "topology breeding". In Chapter 4, we present different optimal topology spaces and describe classes of topologies that recur. Next, we provide detailed analyses (Chapter 5) of topologies that are bred through the topology breeding experiments.

This thesis can be logically divided into two connected parts. The first part, which we described above, concerns with modelling the problem of designing optimal networks under multiple conflicting constraints as an evolutionary optimization problem, and observing patterns of topologies that result from the evolutionary process. In the second part, viz; chapters 5, 6 and 7, we explore a particular subspace of optimal topologies, which is that of circular skip lists.

Circular skip lists (CSL) are structures that not only recur in all optimal topology spaces, but prevail under most circumstances. Here, we analyze structural motifs of CSLs that appear to be optimal for network design across several applications. Since CSLs appear to be useful topologies in many applications, we develop deterministic algorithms to construct CSLs. Finally, we develop theoretical guarantees concerning optimal properties of regular graphs under perturbations. We describe future work in Chapter 8 and conclude.

The difficulty of literature is not to write, but to write what you mean;
not to affect your reader, but to affect him precisely as you wish.

Robert Louis Stevenson

2
Background

Network design has elicited widespread interest in both theory and practice over the last five decades. Several aspects of the problem have been explored. The problem of scheduling flights between cities with minimal stopovers under constraints on capacities of airports is discussed as early as [Erdös et al., 1962]. Infrastructure networks, such as telephone networks and data communication networks, have been studied since the 1970s by modelling them as Steiner tree problems [Garey et al., 1977]. Problems in circuit layout optimization also pertain to network design [Barahona and Reinelt, 1988; Lengauer, 1990].

Graph theoretical properties of networks are studied to establish the extremal properties of families of graphs under several constraints that are useful in network design [Bollobás, 2004]. With the advent of the Internet and the World Wide Web (WWW), information networks such as p2p networks [Lua et al., 2005] and online social networks [Kumar et al., 2010] are hot areas of research.

Recently, [Berners-Lee et al., 2006] have started the process of defining a framework for *web science* towards understanding the web better. A large and expanding body of effort also exists towards developing a philosophical understanding of network design [Barabási and Crandall, 2002; Börner et al., 2007], leading to the emergence of the field *network science*.

The underlying structure or the *topology* of a network has far reaching consequences for the performance of the network. Topology governs the *connectivity* i.e. the availability of communication paths between pairs of nodes in a network, the amount of data/material flow across the network, and the speed of flow. For example, topology influences the diffusion of ideas in a society, the spread of diseases on a global scale and the travelling patterns of people in a transportation network.

Innovation diffusion [Rogers, 1995] explores the spread and adoption of products, practices and ideas in a society. It tries to predict not only the rate of adoption, but also to explain why certain ideas diffuse rapidly while other do not. It has been emphasized that the most important factor in the adoption of new ideas is sociological [Abrahamson and Rosenkopf, 1997; Valente and Davis, 1999]. Several *network models* have been proposed to understand the dynamics of innovation diffusion [Cowan, 2005; Meade and Islam, 2006; Rahmandad and Sterman, 2008].

[Cantono and Silverberg, 2009] develop an agent based percolation model for diffusion. In their model, every agent is willing to pay a certain price to adopt a new technology (for example, renewable energy source), drawn from a log-normal distribution. However, the adoption happens only when both the prices fall below an agent's price threshold, and if at least one of the agent's neighbours has already adopted. Similarly, the number and the kind of people that have adopted an idea influences the probability of adoption for a person [Valente, 2005; Delre et al., 2007].

[Christakis and Fowler, 2008a] study the causes behind the reduction in the prevalence of smoking in the united states over the last 30 years. They model social networks of smokers and non-smokers and observe that cessation of smoking spreads rapidly through close social ties. Also, smokers seem to find themselves increasingly socially isolated while non-smokers find themselves in larger social clusters. [Fowler and Christakis, 2007] also study spreading patterns of obesity and happiness [Fowler and Christakis, 2008], and peer effects on health [Christakis and Fowler, 2008b] over the three decade period, and observe similar effects.

Network topology influences the behaviour of travellers in terms of travel distance, frequency and mode of travel [Bento et al., 2005]. It is in turn important to understand the mobility patterns of humans: this understanding is crucial in several areas such as epidemic prevention, emergency response and urban planning [González et al., 2008].

[Eubank et al., 2004] study the outbreak of smallpox in an urban social network based on traffic patterns. They model the network as bipartite graphs: of individuals and the people they contact, and the locations they frequent. They find that the graph of locations is *scale-free*, which allows to detect and prevent the outbreak of diseases by monitoring crucial hubs. [Tatem et al., 2006] describe the correlation between the volume of international ship and aircraft movements between different parts of the world, and the spread of certain viruses.

Airport networks influence and are influenced by geo-political and trade relations between countries [Guimerá et al., 2005]. Airline networks have, over the years, evolved into one of two broad structures: hub–and–spoke and point–to–point. Studies suggest that major policy decisions such as alliances, mergers and code sharing are heavily dependent on network structures of carriers [Adler and Smilowitz, 2007; Alderighi et al., 2005; Dennis, 2005]. Further, network topology also impacts delays due to congestion, scheduling, frequency and pricing [Brueckner and Zhang, 2001; McAfee and Te Velde, 2006; Forbes, 2008].

The success of information networks such as p2p networks essentially depends on the performance of the underlying network structure. Factors that are crucial for designing a successful information network are: availability of data, efficiency and scalability of processing queries, replication, consistency, and resilience to failures. Studies show that strong relationships exist between the above factors and topological properties of the underlying p2p network [Gummadi et al., 2003;

Christin et al., 2005; Fletcher et al., 2005; Small et al., 2008].

In this chapter, we provide an overview of the research done in these various areas concerning optimal network design. This serves as a background to the work done as a part of this thesis.

2.1 Special Families of Graphs

Degree-Diameter Problem: The degree-diameter problem ((Δ, D)–problem) is the problem of finding the largest graph—in terms of the number of nodes—for a given maximum degree, Δ, in the graph and a given diameter, D. In the digraph (directed graph) version of the problem, a maximum *outdegree* is given. This problem is well studied with respect to several aspects such as, establishing upper bounds, finding as large graphs as possible for given degree–diameter pairs, and relating other structural properties, viz; girth, connectivity and hamiltonicity to this problem.

The *Moore Bound* [Hoffman and Singleton, 1960] defines this upper bound as: the number of nodes, N_{max}, that can be packed in a graph of fixed degree Δ and diameter D. To estimate this, imagine a Δ–ary tree of depth D. The maximum number of nodes, N_{max} can be estimated as:

$$N_{max} = 1 + \Delta + \Delta^2 + \Delta^3 + \ldots + \Delta^D = \frac{\left(\Delta^{D+1} - 1\right)}{\Delta - 1}$$

As a consequence, the lower bound on the diameter of a graph with N nodes and a fixed degree Δ is estimated as:

$$D_{min} = \lceil \log_{\Delta} (N \cdot (\Delta - 1) + 1) \rceil - 1$$

As a corollary, we have the maximum number of nodes within a distance of H, for a fixed degree Δ, estimated as:

$$\frac{\Delta^{H+1} - 1}{\Delta - 1}$$

A *Moore Graph* is a Δ–regular graph that has a diameter D, in which the number of nodes is equal to the Moore Bound. Moore graphs were described by Hoffman and Singleton [Hoffman and Singleton, 1960; Singleton, 1968], who showed that there exists only one Moore graph each with (1) $\Delta = 3$ and $D = 2$ (Petersen graph) [Harary, 1994], and (2) $\Delta = 7$ and $D = 2$ (Hoffman-Singleton graph) [Bondy and Murty, 1976]. There are no other Moore graphs for $\Delta > 2$, except for a possible $\Delta = 57$ and $D = 2$ Moore graph. Thus, although Moore graphs are optimal structures, it is not possible to construct non-trivial Moore graphs. [Bannai and Ito, 1973] prove that there exist no Moore graphs with $D \geq 4$ and $\Delta > 2$. Further impossibility results are provided by [Friedman, 1971]. The possibility of directed Moore graphs was explored by [Bridges and Toueg, 1980] who conclude that unless, $\Delta = 1$ or $D = 1$, directed Moore graphs are impossible.

Studies have been conducted to find families of graphs that are close to the Moore bound, while not quite being Moore graphs. [Chung, 1987] discusses several extremal properties concerning the bounds on the increase/decrease in the diameter of a graph after removing/adding a certain number of edges. A technique to construct large graphs starting from Moore bipartite graphs is discussed in [Comellas and Gómez, 1999]. [Fiol et al., 1983] discuss a technique called line digraph iterations to construct large (Δ, D) graphs. [Tri Baskoro et al., 1998] investigate the structural conditions required to construct directed graphs with order close to that of Moore graphs.

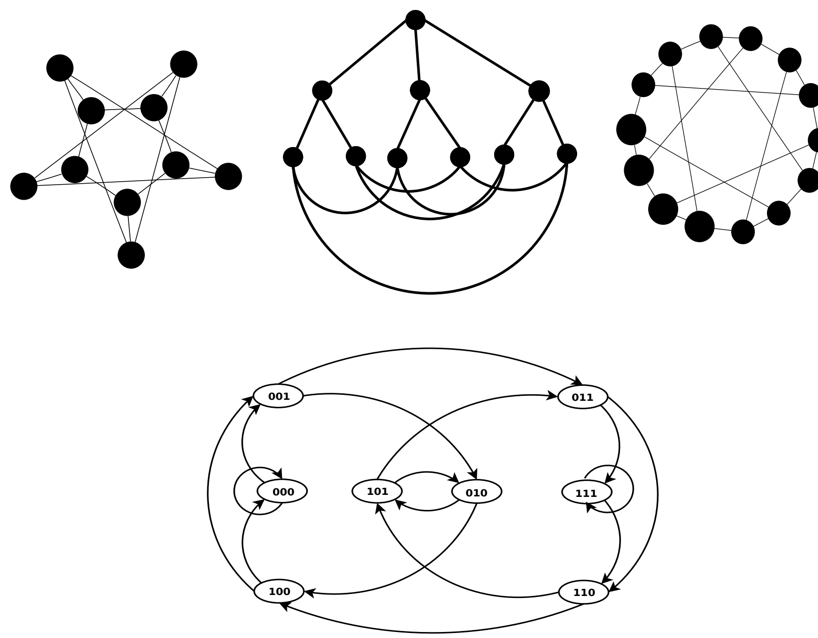

Figure 2.1: Different graph families: (1) Petersen Graph (2) Petersen Graph depicting a Moore Graph (3) Heawood Graph (4) de Bruijn Graph with 8 nodes.

Active research is still on concerning Moore graphs and their relatives. Refer to [Loz and Širán, 2008] for the latest records of large (Δ, D) graphs. Refer to [Miller and Širán, 2005] for a detailed survey of the different approaches in attacking the (Δ, D)–problem: construction techniques; restricted versions of the problem; and, the relationship of the problem with other properties such as girth and connectivity. They also study the (Δ, D)–problem with respect to other families of graphs such as Cayley graphs and *vertex transitive graphs*. We discuss some of these results next.

While Moore graphs are by definition regular graphs, [Singleton, 1968] proved there is no "irregular" Moore graph. Irregular graphs, in the strict sense, wherein every node has a different degree, are impossible. However, the word irregular is used in general to mean "non-regular". [Alon et al., 2002] derive the Moore bound for such irregular graphs. Given a graph with a certain degree sequence or a graph wherein the average degree is non-integral, they derive a strict upper bound for the maximum number of nodes that can be packed in such a graph.

Special families of graphs have been studied in search of optimal structural properties useful in network design. Most of these families are closely related to the degree-diameter problem and hence are recurring.

Cage Graphs: A cage graph [Biggs, 1993] is the smallest graph with a given girth g and a regularity r. Please refer to [Exoo and Jajcay, 2008] for a comprehensive survey of results in cage graphs. Cage graphs are closely related to the (Δ, D)–problem. Cage graphs are obtained by inverting this problem: find the minimum number of nodes, N_{min} to form an r–regular graph of girth g. The following lower bound can be found by considering all vertices whose distance from a given vertex is at most $\left\lfloor \frac{g-1}{2} \right\rfloor$:

$$N_{min}\left(r,g\right) \geq \begin{cases} 1+r+r\left(r-1\right)+\ldots+r\left(r-1\right)^{\frac{g-3}{2}} & : \ if\ g\ is\ odd \\ 2\left(1+r+r\left(r-1\right)+\ldots+r\left(r-1\right)^{\frac{g-2}{2}}\right) & : \ if\ g\ is\ even \end{cases}$$

A Moore graph has a girth of $2D+1$. Figure 2.1 shows the Petersen graph which is a $(3,5)$–cage with 10 nodes. The Heawood graph is a $(3,5)$–cage with 12 nodes.

Cage graphs have several properties that are important in network design. [Jiang and Mubayi, 1998] prove that every cage with regularity, $r \geq 3$, is at least *3-vertex connected*. In other words, at least 3 nodes need to be removed from a $(r \geq 3, g)$–cage in order to disconnect it. Along similar lines, [Wang et al., 2003] prove that an $(r \geq 3, g)$–cage is *k–edge connected* if g is odd. In general, it has been conjectured that cage graphs are maximally connected, i.e., their connectivity is equal to their degree. In network design high connectivity implies greater resilience to failures. With their symmetric structure, cages also offer good load balancing in networks [Decker et al., 2000].

Expander Graphs: Expander graphs, which are sparse structures with high connectivity, find applications in a variety of fields: theoretical computer science, probability theory, coding theory, and, communication networks [Sarnak, 2004; Hoory et al., 2006]. Informally, an expander is a graph, $\mathcal{G}\left(\mathcal{V},\mathcal{E}\right)$, in which every subset of vertices, $S \subset \mathcal{V}$, connects to many vertices in the complementary subset, $\bar{S}$, ($S \cup \bar{S} = \mathcal{V}$ and $S \cap \bar{S} = \emptyset$).

Given an r–regular ($r \geq 2$), n node (multi)graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, we define the edge boundary, δS of S, as the set of all edges, (u, v), such that, $u \in S$, and $v \in \bar{S}$. The *expansion parameter* of $\mathcal{G}$, $h(\mathcal{G})$, is defined as:

$$h(\mathcal{G}) = \min_{1 \leq |S| \leq \frac{n}{2}} \frac{|\delta S|}{|S|}$$

Expansion is closely related to the notion of *conductance* of graphs. Given a subset of vertices, S, and the edge boundary, δS, as above, conductance for $(S, \bar{S})$ is defined as the ratio of δS to the edges within S. Conductance of a graph is the minimum conductance over all its vertex subsets. Conductance is defined in the context of small *mixing time* of random walks. On a graph with high conductance, a random walk that starts within a small subset, S, has a non-zero probability of being on any node in the graph within a few steps. This is because of the high proportion of edges that "cross over" from S to $\bar{S}$. In other words, conductance measures how quickly a random walk on a graph converges to a uniform distribution.

Expanders have several properties that are important to network design such as, regularity, presence of hamiltonian circuits, high connectivity and fast mixing times for random walks. They have been discussed extensively in the context of unstructured p2p networks. [Law and Siu, 2003] discuss a distributed algorithm for the construction of d–regular random expander graphs, which have d hamiltonian circuits. [Gkantsidis et al., 2006] evaluate the effectiveness of random walks for searching and construction of unstructured p2p networks. They discuss several construction techniques for expanders. [Aspnes and Yin, 2008] extend the above efforts to device distributed algorithms to *maintain* the property of expansion in p2p networks in the face of node insertions.

Ramanujan graphs are an instance of expander graphs, known to be good *expanders* [Hoory et al., 2006] leading to high node and edge connectivity despite being sparse. Ramanujan graphs are also used for *key distribution* in wireless sensor networks [Çamtepe et al., 2006]. They argue that using Ramanujan graphs ensures highly resilient and secure key distribution.

de Bruijn Graphs: A de Bruijn graph [de Bruijn, 1946] is a directed graph, where each node is mapped onto an identifier in the identifier space formed by all m–length strings of an alphabet of length n. The number of such strings is n^m, which is the number of nodes in the corresponding de Bruijn graph. Every node has exactly m outgoing edges. The m edges are drawn by right shifting each node identifier by 1 position, and adding each of the n symbols in the alphabet at the end. Figure 2.1 has an example de Bruijn graph.

A de Bruijn graph guarantees a diameter of $\log_m(n^m)$. It is a regular graph with every node having the same indegree and outdegree, m. Networks based on de Bruijn graphs have been claimed to be diameter optimal [Loguinov et al., 2005] and find application in DHT design.

Cayley Graphs: Given a finite group, G, and a set of generators, S, ($S \subset G$), there exists a graph whose vertices correspond to the elements of G, and edges correspond to functions over the generators. Such a graph is called a Cayley graph [Magnus et al., 1976].

Cayley graphs are important in the context of the (Δ, D)–problem, as discussed in [Dinneen and Hafner, 1994; Dougherty and Faber, 2004; Hafner, 1992]. Cayley graphs use concepts from group theory to model symmetric interconnection networks [Akers and Krishnamurthy, 1989; Xiao and Parhami, 2005]. They offer hamiltonian decompositions, thus making them suitable for designing robust and load balanced networks. They have been proposed as a model for DHT design [Qu et al., 2006]. [Hsieh and Hsiao, 2006] study a family of Cayley graphs called k–degree Cayley graphs that are featured by regularity, logarithmic diameter and maximal connectivity.

[Xiao and Parhami, 2006] suggest that Cayley graphs are a good model for small-world networks. They also study the construction of *biswapped networks* starting from Cayley graphs [Xiao et al., 2007]. Biswapped networks are proposed to be load balanced and fault tolerant. [Šiagiová and Širáň, 2005] discuss techniques to construct large Cayley graphs.

Random Regular Graphs: A random r–regular graph is selected from the uniform probability space, $G_{n,r}$, of all n–node, r–regular graphs, where $3 \leq r < n$ and nr is even [Bollobás, 2001]. Generating random regular graphs is a nontrivial problem. [Wormald, 1999] has extensively studied models of random regular graphs and construction techniques. [Meringer, 1999] discusses the problem of generating complete lists of regular graphs up to isomorphism, and proposes an algorithm to generate regular graphs of a given degree and a given number of nodes. The paper also reports some work on generating regular graphs of a given girth (the length of the shortest cycle in a graph). The generator can be used to verify conjectures about regular graphs.

Random regular graphs have several optimal properties such as small diameter, high connectivity and symmetry that can be proved *almost surely*. For example, a random r–regular graph is almost surely r–connected. While one can construct a regular graph whose connectivity is lesser than its regularity, the probability of doing so tends to 0 as the number of nodes, n, increases. If $\epsilon > 0$, is a positive constant, and d is the smallest integer that satisfies the following relation,

$$(r-1)^{d-1} \geq (2+\epsilon)\, rn \, \ln \, n$$

then, almost surely, a random r–regular graph has diameter d [Bollobás, 2001].

[Steger and Wormald, 1999] study properties of large random regular graphs and report several results on the relation between regularity and properties such as hamiltonicity and connectivity. A result shows that almost all regular graphs are hamiltonian [Robinson and Wormald, 1994]. Hamiltonicity is an important problem concerning robustness of a network.

Vertex Transitive Graphs: A graph $\mathcal{G}$ is vertex transitive [Godsil et al., 2001] if given any two of its vertices, v_1 and v_2, there exists an automorphism:

$$f : \mathcal{V}(\mathcal{G}) \rightarrow \mathcal{V}(\mathcal{G})$$

$$f(v_1) = v_2$$

Informally, in a vertex transitive graph, every vertex has the same *local environment*, such that no two vertices can be distinguished based on their neighbourhoods.

A related family of graphs is that of edge transitive graphs. In edge transitive graphs, any pair of edges are equivalent under some automorphism. Again, no two edges can be distinguished based on their neighbourhoods. A graph which is both vertex and edge transitive is called a *symmetric graph*. If a regular graph is only edge transitive but not vertex transitive, it is called *semisymmetric*.

Vertex transitive (as well as, edge transitive) graphs are important in the context of the (Δ, D)–problem, and more generally, in network design. They are highly symmetric, by definition. Most vertex transitive graphs are hamiltonian. In fact, there are only five known non-hamiltonian connected vertex transitive graphs (the Petersen graph is one of them). All Cayley graphs are vertex transitive, though the converse does not hold [McKay and Praeger, 1994].

Vertex transitive graphs offer high node and edge connectivity: their edge connectivity is equal to the degree of the graph, d; and node connectivity is at least $2\left(d+1\right)/3$. Thus, they are ideally suited in network design where chances of targeted attacks are high, viz; network centric warfare [Dekker and Colbert, 2004]. The counterpart family of *edge transitive* graphs are also well suited under similar constraints [McKay et al., 1998].

Hypercube Graph: Hypercube graph topologies offer logarithmic diameters and degree regularity, thus making good candidates for network design. A hypercube graph is an m–regular graph of 2^m nodes, in which each node represents a subset of a set of m elements. In other words, imagine a set of all m–length binary strings. These represent the 2^m nodes in the graph. Each node connects to all other nodes that are at a *Hamming Distance* of 1, forming an m–regular graph. A hypercube graph has a diameter of m, which is the maximum hamming distance between any two nodes. Being a regular graph, it is also highly symmetric.

2.2 Network Design: Application Areas

2.2.1 Steiner Trees and Applications

Given a graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, a length function, L, on the edges, and a subset of vertices, $R \subset \mathcal{V}$, called terminals, a Steiner tree is a subgraph that connects all vertices in R. Further, the tree may include non-terminals, i.e the vertices in $\mathcal{V} \notin R$ as well. Such vertices are called Steiner points. The Steiner tree problem is to find the shortest Steiner tree, i.e. a tree whose total length is minimum. The problem is known to be NP-hard [Garey et al., 1977].

Extensive work exists in literature concerning Steiner tree problems in network design [Hwang, 1993; Du et al., 2001]. Approximate algorithms exist for a class of generalized Steiner tree problems, wherein a minimum cost network needs to be designed under a set of connectivity requirements between pairs of nodes [Agrawal et al., 1995; Gropl et al., 2001; Chlebík and Chlebíková, 2002]. Such algorithms are by and large based on greedy strategies. However, recently, dynamic programming based algorithms [Molle et al., 2006] and randomized algorithms [Garg et al., 2000] are being proposed.

In the directed version of the Steiner tree problem, edges are weighted and directed. Directed Steiner tree problems find application in many network design and routing problems. [Charikar et al., 1999] provide approximation for directed Steiner tree problems. [Garg et al., 2000] discuss the group Steiner tree problem in which we are given several groups of vertices, and the purpose is to construct a minimum weight Steiner tree that has at least one vertex from each group.

2.2.2 Distributed Hash Tables (DHTs)

The topology design problem has received significant interest in the area of data-centric peer-to-peer networks. DHTs such as Chord [Stoica et al., 2003], Pastry [Rowstron and Druschel, 2001], and Koorde [Kaashoek and Karger, 2003] use DHTs to lookup where a given hash bucket is located. DHTs continue to be an area of active research [Fraigniaud and Gauron, 2006; Loguinov et al., 2005; Ratnasamy et al., 2002; Xu et al., 2004].

There are a number of important performance measures based on which distributed indexes are evaluated: lookup cost, cost of maintaining routing tables in the face of changes in topology, robustness in the face of failures and symmetry in load distribution. More often than not these measures are conflicting. In the context of peer-to-peer networks, a high emphasis is placed on the "symmetric" nature of the data structure. All peers are assumed to have nearly equal degree in forming the DHT.

DHTs primarily optimize the cost of lookup or the diameter of the DHT topology. In Chord [Stoica et al., 2003], an identifier space is formed as a logical circle, to which both nodes and keys are mapped using a hash function. The Chord topology is a skip list, wherein each node connects to $\log n$ other nodes on the circle to achieve a diameter of $\log n$ and symmetric degree distribution.

Koorde [Kaashoek and Karger, 2003] implements a de Bruijn graph (which we discussed earlier) based network, in which a $\log n$ diameter is achieved with a fixed degree of 2. The paper also describes an extension to a degree–m de Bruijn graph, using which a diameter of $\log_m n$ can be achieved. Another implementation of de Bruijn graphs is D2B [Fraigniaud and Gauron, 2006], which is similar to Koorde. DHT topologies based on de Bruijn graphs continue to be proposed [Awerbuch and Scheideler, 2009].

[Loguinov et al., 2005] provide a detailed graph-theoretic analysis of peer-to-peer networks with respect to routing distances and resilience to faults. The paper argues that de Bruijn graphs offer the optimal diameter topology among the class of practically useful graphs. They offer a low diameter. Also, they come close to satisfying the Moore Bound. Being directed graphs, de Bruijn graphs suit very well for applications like DHTs. However, it can be shown that an undirected de Bruijn graph is not the best possible topology in terms of diameter optimality. Also, the topology of de Bruijn graphs are fixed and cannot be altered to have better diameters by compromising on the degree distribution. Therefore, de Bruijn graphs cannot be claimed to be optimal in a general sense.

In HyperCuP [Schlosser et al., 2003], a hypercube graph is constructed in a distributed manner by assuming that each node in an evolving hypercube takes more than one position in the hypercube. That is, the topology of the next dimensional hypercube is implicitly present in the present hypercube, with some of the nodes also acting as “virtual” nodes to complete the hypercube graph. Similarly, when nodes go away, some of the existing nodes take their positions along with their own.

Viceroy [Malkhi et al., 2002] is an implementation of an approximate butterfly network. Nodes are arranged in $\log n$ levels, with the nodes at each level forming a ring topology with each node having an outgoing link to a successor and a predecessor. Apart from the "neighbours" on the ring, each node has long range outgoing links to five other nodes across the $\log n$ levels. These levels and nodes are chosen by a randomized process. Viceroy claims to achieve a $\log n$ diameter with a fixed degree of 7. Ulysses [Kumar et al., 2004] is another implementation of the butterfly, though with its $\log n$ neighbours, it is not a fixed degree graph.

[Freedman and Vingralek, 2002] propose a distributed trie based approach. It is based on [Plaxton et al., 1999]'s prefix based routing, wherein a k–ary prefix tree is maintained in a distributed manner. The maximum degree of a node is $k + 1$ and the diameter is $2\lceil \log_k n \rceil$. [Ratnasamy et al., 2001] discuss a *Content Addressable Network* (CAN), which forms an identifier space over an approximation of an m–dimensional torus. CAN has a fixed degree of $2m$ and provides a diameter of $\frac{mn^{\frac{1}{m}}}{2}$.

There are several random graph based DHTs. [Pandurangan et al., 2002] propose an algorithm to construct a DHT dynamically under a probabilistic model of node arrivals and departures. The DHT has a fixed low degree and achieves a logarithmic diameter almost certainly. However, their algorithm is partially centralized since new nodes register with a central server and the server selects a set of random nodes to connect to, rather than the node doing it. [Law and Siu, 2003] take up the problem of distributed construction of random regular graphs that have m Hamiltonian circuits. The graph is $2m$–regular with each node having a degree of $2m$, and m hamiltonian circuits passing though each node. With a high probability these graphs will achieve a diameter of $\log_m n$.

Kleinberg discusses a *small-world* [Kleinberg, 2000] where the nodes are arranged in a metric space having a distance function defined across any pair of nodes. Each node is modelled as having short range links to its neighbours till a certain range, plus some random long range links to nodes that are far away. Kleinberg argues that although we can find several families of such small-world topologies that have short routing paths, it is very difficult to find short paths between pairs of nodes based on local information. He goes on to prove that there is only a unique family of small-worlds for which a decentralized navigation algorithm can be constructed. Symphony [Manku et al., 2003] extends Kleinberg's work. The Symphony topology is a ring with a set of long range links, which are randomly drawn from a family of continuous harmonic distributions.

[Ratnasamy et al., 2003] identify the trade-offs between routing table length (degree) and length of lookup (diameter). They argue that DHTs with a $\log n$ degree achieve a $\log n$ diameter, whereas fixed degree DHTs, say the degree is m, achieve a diameter of $n^{\frac{1}{m}}$. They ask if these are the optimal trade-offs, and if they are not, then will the topologies that improve this trade-off do so at the cost of some other desirable properties. [Xu et al., 2004] take off from here. They argue that congestion (load) in the network is an important parameter that will be affected badly by trying to improve the above trade-offs. They conjecture that for a uniform load distribution, the above trade-offs are asymptotically optimal. However, when uniform load distribution is not one of the desired properties, we can easily develop DHTs with better degree and efficiency performance. [Gummadi et al., 2003] consider *flexibility* in neighbour and route selection that different DHT topologies offer. They conclude that the ring structure performs the best.

2.2.3 Complex Networks

Complex networks are networks—either naturally occurring or man made—that are featured with non-trivial structural properties such as power law degree distributions, small average path length and ubiquitous community structures. Unlike random graphs, complex networks have apparent *design* in them, for performance, or, as in case of naturally occurring networks, they seem to have evolved to survive. Example complex networks include, the World Wide Web, the brain, protein interaction networks, trade networks and social networks.

Structural properties of complex networks have been studied extensively, especially following the seminal work by [Strogatz, 2001], [Albert and Barabási, 2002] and [Barabási and Bonabeau, 2003]. Disparate classes of complex networks have been analyzed. There are attempts to relate their structural properties such as diameter, degree distribution, betweenness distribution, clustering coefficient and modularity, to their performance. Below we briefly review a representative sample of the work done in complex networks over the last decade. For a detailed coverage, refer to [Albert and Barabási, 2002].

Internet and the WWW

In their seminal work, [Faloutsos et al., 1999] explore the degree distribution of the Internet topology at both router and autonomous system levels using a network map obtained from *traceroute*. They conclude that the degree distributions follow a power law for both networks. [Govindan and Tangmunarunkit, 2002] developed a system called *Mercator* to model the router level network to study the Internet topology. Their results corroborate the existence of power law degree distribution. [Magoni and Pansiot, 2001] study the autonomous system level topology based on six instances of BGP data. They study the distributions of degree, distance, number of pairwise shortest paths etc. They extend the work of [Faloutsos et al., 1999] by finding five new power law relations.

[Medina et al., 2000] explore the phenomenon that lead to power laws in Internet topology. They conjecture that the following are the main factors leading to power laws: (1) preferential attachment of new nodes to existing nodes, and (2) incremental growth of the network. [Zhou and Mondragón, 2004] propose the *rich-club model*, which is a measure of "mixing" of nodes in a network. A network is *assortative* if nodes connect to similar nodes (ex: high degree nodes connecting to other high degree nodes), and *disassortative* if nodes connect to dissimilar nodes. A rich club is a set of high degree nodes in a network. The rich-club connectivity measures how well the club members connect to each other. The authors extend this model to explain the autonomous system level topology.

While there is a lot of work on understanding the Internet topology, a primary criticism of these efforts is the lack of validation of the data as well as the models. [Chen et al., 2002] argue that the BGP data that was used to study the Internet topology, which was also later used as a basis to develop the preferential attachment model, do not generate an accurate map of the Internet as it misses 20-50% of the connections. They suggest that the growth mechanism of the network topology might be entirely different from preferential attachment. Work still continues concerning improved ways of mapping the topology. [Mahadevan et al., 2006] propose gathering data from three different sources—BGP, traceroute and WHOIS—and using a measure called *joint degree distribution*. Newer approaches using percolation theory and fractal geometry are proposed [Carmi et al., 2007] that study properties such as connectivity and self-similarity.

[Albert et al., 1999] started a series of research efforts toward understanding the WWW topology. They constructed a model of the WWW topology by using local connectivity of web pages to observe that the indegree and outdegree distributions of the WWW follow a power law. This is remarkably different from Poisson distributions that are observed in Erdös-Rényi random graphs, indicating that the underlying growth mechanism of the web is unlike that of a random graph. This lead to the Barabàsi-Albert model of preferential attachment on the web [Barabási and Albert, 1999; Albert and Barabási, 2002]. [Adamic et al., 2000] used a different dataset than [Albert et al., 1999]: they consider a node to be a domain and a link only if it is across domains. They too observe power law degree distributions.

Transportation Networks

[Latora and Marchiori, 2002] study the Boston subway network to analyze properties such as, the average path length (APL) and clustering coefficient. They conclude that the Boston subway is a small-world. More interestingly, they compute the efficiency of the network in terms of APL and observe a cost vs efficiency trade-off: if bus routes are also considered in the network along with the subway routes, the efficiency increases. [Sienkiewicz and Hołyst, 2005] analyze 22 metropolitan transport networks in Poland to observe that the networks are small-worlds with clear hierarchical organizations. The degree distributions are either power laws or exponential. The smaller transport networks are disassortative while they tend to become assortative as they consider larger ($n > 500$) networks.

[Sen et al., 2003] study the Indian railway network to observe an exponential degree distribution. [Rosvall et al., 2005] discuss a road network from an information theoretic perspective. They compute the *entropy* of a city's road network in terms of the amount of information needed to navigate through it. They compare the complexity of road networks in cities based on this measure.

There are several efforts in the last few years towards analyzing airline networks. [Guimerá and Amaral, 2004] study the world-wide airport network (WAN) and report "anomalous" high betweenness centrality values for nodes that have relatively low degrees. They argue that power-law generating models such as *preferential attachment* (PA) cannot explain this observation. Instead, they propose that the anomalous centrality is a result of a network designed under several geo-political constraints leading to well-defined, geographically separated community structures [Guimerá et al., 2005]. Further, they classify cities based on their intra and inter-community "participation" into roles such as hubs, provincial hubs, connectors and peripheral nodes.

[Barrat et al., 2004] consider a weighted network model of WAN, and analyze the degree and betweenness distributions, clustering coefficient and the *mixing patterns* of nodes. Here, weights on edges are expressed in terms of the number of available passenger seats between pairs of airports. They report that considering weights allows a richer characterization of a network. In a related work, [Barthélemy et al., 2005] propose a model of network evolution based not only on the network topology, but also on the strengths of interaction between nodes represented by edge weights.

[Colizza et al., 2006] study the role of air transportation networks in the spreading pattern of diseases. They report a correlation between the heterogeneity in the structure of the air transportation network and the seemingly erratic, global scale spreading of diseases.

Domestic airline networks have also been studied in an effort to better understand the dynamics of transportation networks at a lower scale. [Li and Cai, 2004] observe that the degree distribution of the airport network of China (ANC) follows a double Pareto distribution. [Bagler, 2008] reports that the airport network of India (ANI) is a small-world like ANC and WAN, however with a much smaller clustering coefficient. Further, ANI suggests a *disassortative mixing* with hubs having a lot more low-degree neighbours than high-degree neighbours. This is in contrast with WAN wherein *assortative mixing* can be observed [Barrat et al., 2004]. Similar analysis on the Italian airport network shows a small-world network characterized by a fractal nature [Guida and Maria, 2007].

Biological Networks

Foodwebs model the flow of nutrition among species. Species form the nodes in a foodweb. The edges are *trophic links*, which represent “who eats whom”. [Williams and Martinez, 2000] analyze several foodwebs across different habitats for properties such as average degree and link density. Based on these properties, they propose a simple model that generate foodwebs. [Camacho et al., 2002a] extend the above work to develop analytical expressions for features of ecological interest such as predator and prey distributions. In a related work, they observe robust patterns in foodwebs that obey concepts in statistical physics such as universality and scale [Camacho et al., 2002b]. [Stouffer et al., 2007] observe and classify *motifs* in foodwebs in terms of their link structure.

[Smart et al., 2008] study metabolic networks, wherein they represent the cellular metabolism as a directed bipartite graph with two types of nodes—*reaction nodes* and *metabolite nodes*. A directed edge exists from a metabolite node to a reaction node if the metabolite participates in the reaction by being an enzyme. By subjecting it to a cascading failure model they conclude that metabolic networks are exceptionally resilient to failures. [Guimerá et al., 2007] classify edges, i.e. interactions in a metabolic network in the order of importance.

[Jeong et al., 2001] study the correlation between the number of interactions (degree) of a protein and the likelihood of its deletion proving to be lethal (importance) in an yeast protein interaction network. They observe a positive correlation. They also observe that the yeast protein interaction network is scale-free.

Communication Networks

In the last few years, communication networks are being studied from a complex networks perspective. [Valente et al., 2004] report that the most robust structure in the face of random failures or targeted attacks is one with at most three distinct node degrees in the network. Unless the load handling capacity of a node is directly proportional to its degree, random-regular graphs and Cayley trees are shown to be better suited for designing traffic flow in networks [Zhao et al., 2005].

[Guimerá et al., 2002] study the problem of "searchability", i.e. the ability to find short paths based on local information, in the presence of congestion, and report that only two classes of optimal topologies exist: highly polarized (like a star topology) networks, when the load is low; and *homogeneous isotropic networks* with symmetric node betweenness, as the load increases.

[Donetti et al., 2005, 2006] study optimal structures for high synchronizability and low first-passage times for random walkers, and propose *entangled networks* with highly homogeneous structures. [Mondragón C, 2006] reports a sharp transition from hub–and–spoke to decentralized as the load on the network increases.

2.3 Network Design: Philosophical Understanding

There is a lot of recent interest in developing a philosophical understanding of the optimal network design problem. [Motter and Toroczkai, 2007] provide a philosophical overview of network optimization, and identify four classes of network optimization problems in the increasing order of complexity: (1) *Structural Optimization* – optimizing a property (e.g., diameter), given a static property (e.g., degree sequence) (2) *Dynamics Optimization on a Static Graph* – optimizing flow on the edges in a static graph to minimize a global property (e.g., congestion) (3) *Structural Optimization for Dynamics* – finding the optimal structure given a property (e.g., flow dynamics), and finally, (4) *dynamics-driven network optimization* – designing an optimal structure in situations where dynamics, such as failures, can affect performance.

Several models have been proposed to capture the mechanisms that generate topologies. The role of *emergence* of networks as opposed to *design* is a major topic of debate in the last few years. Below, we give a brief overview of these topics.

2.3.1 Statistical Physics of Networks

A large body of literature exists that attempts to characterize complex networks inspired from ideas in statistical physics such as *emergence* and *self-organization*. The most popular model of self-organization is *self-organized criticality* (SOC) proposed by [Bak, 1996]. The essential intuition behind SOC is that systems with a large number of interacting components will "self-organize" to a critical point at which occurrence of events follows a power law. Since this behaviour "emerges" without the need for tuning parameters, system design has only a secondary role.

Several network generation models have been developed as efforts towards explaining seemingly universal motifs in complex networks such as power law degree distribution, community structures and small diameters. They largely use ideas similar to self organized criticality. Perhaps the most popular model is the Barabási-Albert model. [Barabási and Bonabeau, 2003] propose *preferential attachment* as the underlying mechanism for the emergence of power laws in natural and man-made networks. The model states the following. When a new node joins a network, it makes a number of connections. The probability of the new node connecting to an existing node is biased to be proportional to the degree of the latter. That is, the probability, p_i, of a new node connecting to an existing node, i, with degree, k_i, is,

$$p_i = \frac{k_i}{\sum_j k_j}$$

This model generates networks that have power law degree distributions and small average path lengths. Such networks are robust against random failures; but fragile when high degree nodes are targeted.

[Watts and Strogatz, 1998] propose a model to generate networks that exhibit "small-world" properties with small average path lengths and high clustering coefficients leading to the formation of "communities". The Watts-Strogatz model uses a parameter, β, where, $0 \leq \beta \leq 1$, that determines the amount of randomness in the network. Given the number of nodes, N, parameter β, and mean degree, K, such that, $N \gg K \gg \ln n \gg 1$, the model constructs an undirected graph with $\frac{NK}{2}$ edges. The model starts by constructing a ring lattice, and then "rewires" $\beta\frac{NK}{2}$ number of edges. Rewiring is done by replacing an edge (i, j) by an edge (i, k), where k is randomly chosen from the N nodes.

The Watts-Strogatz model generates networks that have more or less homogeneous degree centrality. They have small average path lengths and high clustering coefficients.

[Kleinberg, 2000] provides a theoretical framework for small-world networks by developing exact conditions under which small-worlds occur. [Newman, 2006] reports the existence of community structures in disparate complex networks.

2.3.2 Emergence vs Design

In contrast to the statistical mechanical explanations of phenomena like power laws as emerging due to simple underlying forces like preferential attachment, there are recent efforts which propose that the occurrence of complex phenomena are a result of rigorous design for performance. A network's functional requirements are defined in terms of optimality objectives. The occurrence of interesting motifs is thus the result of trade-offs between objectives. [Carlson and Doyle, 1999] propose *highly optimized tolerance* (HOT), a mechanism to generate power laws as a result of trade-offs between efficiency and robustness in a network. They also challenge the validity of models such as scale-free networks since those models ignore domain dependent trade-offs [Alderson et al., 2005]. Further, the scale-free model attempts to characterize networks based solely on degree distribution, which might be inadequate [Willinger et al., 2009]. For example, a study of the world-wide airport network [Guimerá and Amaral, 2004] shows that the nodes with the highest betweenness centrality values do not necessarily correspond to the nodes with the highest degrees, despite the degree distribution being a power law.

[Venkatasubramanian et al., 2004] study the emergence of network topologies under varying efficiency and robustness requirements. They propose a formalism based on the Darwinian evolutionary theory: networks evolve their structures so as to optimize their short term liveness (efficiency) as well as long term survival chances (robustness). A selection pressure variable decides the trade-off between short term and long term survival goals. The star topology emerges as the most efficient and least robust, when the selection pressure is completely on efficiency. The circle topology emerges as the most robust and least efficient, when the selection pressure is completely on robustness. Hub–and–spoke structures of different types emerge for intermediate selection pressures. Thus, they propose that the structure of a network is the result of adaptation under selection pressure.

In a related study [Venkatasubramanian, 2007], suggests that power laws occur because a power law network has the maximum *adaptability* or chance of survival in a variety of environmental conditions. In other words, a power law is a result of a network trying to maximize uncertainty or information entropy about its environment.

2.4 Positioning Our Work

In general, the following types of research efforts can be seen in literature towards addressing the problem of optimal network design:

- designing for specific (domain dependent) performance requirements
- evaluation of families of graphs in search of specific optimal properties; using this knowledge in network design
- efforts towards explaining the theoretical underpinnings or mechanisms that lead to certain optimal properties or behaviours of networks.

In this work, we model the problem of network design for a variety of performance metrics and trade-offs. Each combination of performance metrics corresponds to a class of network design problems. For instance, in case of DHT design, minimizing diameter (lookup cost) is the objective, which is to be achieved under constraints on maximum node degree (bookkeeping cost) and symmetric lookup load on nodes (uniform degree centrality). In case of CDNs, the average communication cost (average path length) should be minimized under the constraints on symmetric load distribution (uniform betweenness centrality) to reduce congestion. Therefore, we build a design framework that models multiple performance parameters.

Further, within a class of networks, the emphasis on performance parameters might vary based on specific deployment requirements. Therefore, rather than using predefined trade-off values, we evolve optimal topologies for several trade-offs using a genetic algorithm process. This leads to a compendium of topologies from which a designer can chose depending on specific requirements.

Rather than mapping known families of graphs such as Cage graphs and de Bruijn graphs to application contexts, we start by modelling the requirements of applications, and let an evolutionary process decide the families of topologies that are optimal for those applications. Also, our focus is towards observing optimal structural motifs in topologies that emerge rather than explaining the mechanisms that lead to those structures.

The findings from this work corroborate several recent results. [Mondragón C, 2006] reports that the transition from hub–and–spoke networks to symmetric networks is sharp as the load in a network increases. We observe that this holds under other metrics of robustness as well. [Bermond et al., 1998; Gummadi et al., 2003] show that graphs with hamiltonian decompositions are optimal for DHTs. [Valente et al., 2004] report networks with a nearly symmetric degree centrality as most robust against failures as well as attacks. [Guimerá et al., 2002] propose networks with symmetric node betweenness as most optimal to handle congestion in traffic flow networks. Regular graphs with *optimal connectivity* are proposed [Dekker and Colbert, 2004] to be ideal under targeted attacks. We observe that circular skip lists, which emerge as the most prominent class of topologies in our work, are characterized by all these properties which are optimal under different circumstances.

Fast, good, cheap. Pick any two.

The Project Triangle

3

A Genetic Algorithm Framework

In this chapter, we develop a genetic algorithm based framework to address the problem of designing optimal networks under multiple efficiency and robustness constraints. We use this framework to evolve optimal topologies that form optimal topology spaces (OTS). We start by discussing the features of the problems we address in this thesis.

Consider a network, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where the number of nodes is, $|\mathcal{V}| = n$. The problem of optimal network design is to find the set of edges or connections, $\mathcal{E}$, by adding which an objective function, $\Omega(\mathcal{G}(\mathcal{V}, \mathcal{E}))$ is maximized, while satisfying a set of constraints, $\Gamma(\mathcal{G}(\mathcal{V}, \mathcal{E}))$.

Below we list the assumptions under which we address the network design problem in this thesis:

1. We design both directed and undirected networks. However, for an instance of the design problem, edges are exclusively either directed or undirected. We do not allow *mixed networks*.

2. We design unlabelled directed and undirected networks. The ordering of nodes and edges is not important.

3. We address designing connected undirected networks and strongly connected directed networks. We do not deal with designing disconnected networks.

4. We address designing unweighted networks. That is, we do not consider node and/or edge weights.

5. Finally, in our networks, there exists at most one edge between a pair of nodes. And an edge connects exactly two distinct nodes. In other words, we do not allow *multigraphs*, *self-loops* and *hypergraphs*.

Some of the above assumptions are not limitations of our framework: the framework supports mixed networks, disconnected components and weighted networks. However, a limitation of the framework is that it cannot model a real world weighted network without substantial extensions. For instance, the framework does not support heterogeneous nodes and links. Also, it does not model domain specific requirements such as turn restrictions in road networks.

3.1 Network Design: Optimization Problem

Given a graph, $\mathcal{G}(\mathcal{V},\mathcal{E})$, there are $n(n-1)$ potential directed edges or connections that can be added to $\mathcal{G}$. The maximum number of edges in a directed graph, $e_{max}^{directed} = n(n-1)$. If the edges are undirected, then there are $n\frac{n-1}{2}$ potential connections. The maximum number of edges in an undirected graph is, $e_{max}^{undirected} = n\frac{n-1}{2}$. The minimum number of directed edges required to design a *strongly connected* directed network, $e_{min}^{directed} = n$. The minimum number of undirected edges required to design a connected undirected network, $e_{min}^{undirected} = n-1$.

Given this, the network design problem is to chose a set of edges, $\mathcal{E}$, where,

$$e_{min}^{directed} \leq |\mathcal{E}| \leq e_{max}^{directed} \quad : \; if\; \mathcal{G}\; is\; directed$$

$$e_{min}^{undirected} \leq |\mathcal{E}| \leq e_{max}^{undirected} \quad : \; if\; \mathcal{G}\; is\; undirected,$$

to maximize an objective function, $\Omega(\mathcal{G}(\mathcal{V},\mathcal{E}))$, such that a set of constraints, $\Gamma(\mathcal{G}(\mathcal{V},\mathcal{E}))$, are satisfied.

In the unconstrained version of the above problem, every connected graph, $\mathcal{G}'$, with $|\mathcal{E}|$ edges is a candidate. The $\mathcal{G}'$ that maximizes Ω is the optimal network.

The problem of counting all graphs of a certain kind—in this case, all graphs with a given number of vertices and edges—falls in the field of *graph enumeration* [Harary and Palmer, 1973]. Graphs can be *vertex labelled*, i.e., every vertex in the graph has a label that distinguishes it from other vertices for the purpose of enumeration, or *unlabelled*. The number of directed labelled graphs with n nodes and m edges is $\binom{n(n-1)}{m}$; the number of undirected labelled graphs with n nodes and m edges is $\binom{n(n-1)/2}{m}$.

We are interested in unlabelled graphs, since it is the structure or topology of the network that is important to us and not the ordering of the nodes. In general, unlabelled enumeration problems are tougher to solve than labelled enumeration problems. *Pòlya's Enumeration Theorem* [Harary and Palmer, 1973] gives the number of unlabelled non-isomorphic graphs with a given number of nodes and edges. Generating all non-isomorphic unlabelled graphs is a different problem altogether. There are no known polynomial time algorithms to resolve isomorphism of two graphs.

Further, as we discussed previously, one or more constraints exist for a given network design problem. In the constrained version of the above problem, the candidates are an enumeration of all graphs that satisfy a set of constraints. An immediate constraint is that all graphs should be connected. In addition, there may be constraints on the degree of nodes. While the number of candidate solutions comes down as a result of constraints, there are no known ways of enumerating and/or generating arbitrarily large graphs that satisfy a given set of constraints to the best of our knowledge.

Below, we define an instance of the optimal network design problem under multiple constraints.

Problem Instance 1. *Given a network,* $\mathcal{G}(\mathcal{V}, \mathcal{E})$*, of* $\mathbf{n}$ *nodes, each having a maximum permissible degree of* $\mathbf{p}$*, where,* $2 \leq p \leq n-1$*, add a maximum of* $\mathbf{e}$ *undirected edges to minimize the diameter* $\mathbf{d}$ *of the network. Ensure that the degree skew in the resulting network,* $\mathbf{p_{skew}}$*, i.e. the difference between the maximum degree in the network,* $\hat{\mathbf{p}}$*, and the average degree of the network,* $\bar{\mathbf{p}}$*, is no more than* δ*, where* $\mathbf{0} \leq \delta < \mathbf{n} - \mathbf{2}$.

We need to chose an $\mathcal{E}$, where $|\mathcal{E}| \leq e$, from the set of all possible undirected edges, whose size is, $e_{max}^{undirected} = \frac{n(n-1)}{2}$. Let x^T be an $\frac{n(n-1)}{2}$–vector, representing all undirected edges in $\mathcal{G}$:

$$x^T = \{x_{ij} : \ i, \ j \in \mathcal{V}\}$$

$$x_{ij} = \begin{cases} 0 & : edge \ \ (i,j) \notin \mathcal{V} \\ 1 & : edge \ \ (i,j) \in \mathcal{V} \end{cases}$$

Since the diameter is determined by the arrangement of the edges, $d = f\left(x^T\right)$, the following is the mathematical programming model for Problem Instance 1.

$$\arg\min_{x^T} \mathbf{f}\left(\mathbf{x^T}\right) \tag{Ω-1}$$

subject to,

$$x_{ii} = 0,\ x_{ij} + x_{ji} = 1 \tag{Γ–1.1}$$

$$\sum_{i,j} x_{ij} \leq e \tag{Γ–1.2}$$

$$\sum_{j} x_{ji} + \sum_{j} x_{ij},\ \forall\, i \in \mathcal{V} \leq p \tag{Γ–1.3}$$

The diameter of a network is defined as the longest of the all pairs shortest paths. Therefore, we need to define our objective function further. Let P_{ij} denote the shortest path between nodes i and j.

$$f\left(x^T\right) = d = \max_{i,j} P_{ij}$$

$$P_{ij} = \min(P_{ij}, \min_{ij}(P_{ik} + x_{kj}))$$

Further, the constraint on degree skew is defined as follows. Let p_i be the degree of the i^{th} node.

$$\hat{p} = \max_{i}\ p_i \qquad \bar{p} = \frac{\sum_{i,j}\ p_i}{n}$$

$$p_{skew} = \hat{p} - \bar{p} \leq \delta \tag{Γ–1.4}$$

3.2 Optimization Framework

What we described in the previous section is an instance of a combinatorial optimization problem—we need to find an optimal object (optimal set of edges) in a large but finite set of objects. In other words, an optimization process has to search a large discrete solution space to find the optimal solution. The solution space can be formed by graphical enumeration under a given set of constraints. Each solution is then evaluated to find the optimal solution. However, this approach is not feasible as we discussed earlier.

Therefore, we propose an evolutionary optimization based framework to address the kind of network design problems defined previously. Our approach for the exploration of optimal topology spaces is called *topology breeding*. It relies on genetic algorithms as the optimization technique to aid in this exploration. In this section, we formally define the topology breeding setup.

[Venkatasubramanian et al., 2004] propose that performance of a network depends on three critical system parameters: efficiency, robustness and cost. In their work, efficiency is defined in terms of the average path length of the resulting graph topology and robustness in terms of the number (and size) of connected components a node deletion causes in the graph. A selection pressure variable decides the relative importance of efficiency and robustness. Using these measures, they let topologies evolve under different trade-offs through a genetic algorithm process.

We extend the above formalism to accommodate multiple constraints and optimality objectives to study different classes of networks.

- We define efficiency (η) in terms of properties such as diameter, average path length and closeness centrality.
- We define robustness (ρ) in terms of properties such as degree centrality, node betweenness centrality and edge connectivity.
- We design topologies with an arbitrary number of edges. Edges in a network can be directed or undirected.
- The trade-off between robustness and efficiency is decided by an environmental variable, α.
- A variable, β, is used as a cost control parameter.

3.2.1 Performance Metrics

Different combinations of design metrics are applicable in different scenarios. In DHTs, minimizing lookup complexity (minimal diameter) while maintaining small and symmetric finger tables across machines (degree symmetry) is a design objective. To handle traffic flow, designing networks with low average path lengths while balancing load on the nodes (node betweenness symmetry) to avoid congestion is important. In case of NCW and supply chains, having alternate paths when a communication link fails (optimal edge connectivity) is a design requirement.

We use a genetic algorithm optimization process called "topology breeding" to address the problems such as above. Topology breeding is governed by a "fitness function". Fitness of a graph, $\phi(\mathcal{G}(\mathcal{V}, \mathcal{E}))$, is defined in terms of the optimization parameters: efficiency (η), robustness (ρ) and infrastructure cost (κ). The maximum permissible degree (p) is also a design constraint.

We model efficiency, robustness and cost in terms of structural or graph theoretic properties. A structural metric models only certain aspects of network performance. Therefore, we use multiple metrics to model performance. The specific structural metrics that we use to model a network depends on the network's performance characteristics. Below we note some useful metrics in network design and analysis. For a comprehensive account of metrics used in network design and analysis, please refer to [Costa et al., 2007].

3.2.2 Efficiency

Efficiency measures the cost of communication in a network. There are a number of ways in which efficiency can be defined. The *diameter* of a network is the upper bound on the communication cost in the network. The *average path length* (APL) gives the communication cost on average. Maximizing the symmetry in the distribution of distances between pairs of nodes can also serve as a useful measure of efficiency. Therefore, per node *eccentricity* (longest of all shortest paths from a node) and *closeness centrality* values can also be used to define efficiency. In this thesis, we use efficiency based on two metrics in the design process—diameter and average path length.

Efficiency based on Diameter

The worst diameter (dia_{max}) for a **connected undirected graph** of n nodes is $n-1$, which is the diameter of a straight line graph. The best diameter (dia_{min}) is 1, which is the diameter of a clique (complete graph). In other words, a topology is most efficient if the diameter is 1, and least efficient if it is $n-1$. We map a diameter, d, that falls in the interval, $[n-1, 1]$, to a value of efficiency, η_d, in the interval, $[0, 1]$, as:

$$\eta_d = 1 - \frac{d-1}{n-2} \tag{3.1}$$

In case of a **strongly connected directed graph** with n nodes, the worst diameter is $n-1$, which occurs for a circle topology. The best diameter is 1, which is the diameter of a directed complete graph. Since both the upper and lower bounds of diameter are the same as in undirected graphs, the expression for η_d in case of directed graphs is the same as above.

Expression (3.1) is due to the following mapping function. For a graph $\mathcal{G}$, with a diameter, $dia_{\mathcal{G}}$:

$$\eta_d(\mathcal{G}) = 1 - \frac{dia_{\mathcal{G}} - dia_{min}}{dia_{max} - dia_{min}} \tag{3.2}$$

By this definition, an undirected straight line topology has an η_d of 0, a directed circular topology has an η_d of 0, a clique has an η_d of 1 and so on.

Efficiency based on APL

It can be shown that, the worst APL, l_{max} for a **connected undirected graph** of n nodes, which occurs again for a straight line, is $\frac{n+1}{3}$. The best APL, l_{min} is 1, which occurs for a clique. We map an APL, l, that falls in the interval, $\left[\frac{n+1}{3}, 1\right]$, to a value of efficiency, η_l, in the interval, $[0, 1]$, as:

$$\eta_l = 1 - 3\frac{l-1}{n-2} \tag{3.3}$$

Similarly, for **strongly connected directed graphs**, the worst case APL, the APL of the directed circle, is $\frac{n}{2}$. The best case APL is again 1, for the directed clique. Thus, the expression for efficiency of directed graphs, in terms of APL is:

$$\eta_l = 1 - 2\frac{l-1}{n-2} \tag{3.4}$$

Expressions (3.3) and (3.4) are due to the following mapping function. For a graph $\mathcal{G}$, with an APL, $l_{\mathcal{G}}$:

$$\eta_l(\mathcal{G}) = 1 - \frac{l_{\mathcal{G}} - l_{min}}{l_{max} - l_{min}} \tag{3.5}$$

3.2.3 Robustness

Robustness measures the resilience of a network in the face of perturbations such as: node and edge failures; and, variable network load. Robustness is often defined in terms of the *skew* in the *importance* of nodes and edges. As such, centrality measures, viz; degree, node betweenness and edge betweenness sequences can be used. When there is a skew in the centrality measures, a small number of nodes and/or edges are more important than the others. Thus, their failure affects the network's performance much more than failures in the rest of the network. On the other hand, a symmetric centrality sequence ensures robustness to random failures as well as targeted attacks of nodes/edges.

Connectivity of a network is also indicative of robustness. If a network has a *vertex connectivity* of m, it implies there are m *vertex independent paths* between any pair of nodes in the network. Similarly, there is a notion of *edge connectivity*. Menger's theorem states that the size of the minimum vertex (or edge) cut of a graph is equal to the maximum number of pairwise vertex (or edge) independent paths in the graph [Harary, 1994]. Thus, connectivity has an important role in networks prone to failures.

In this work, we use several definitions of robustness to cover different perspectives:

1. *Degree Centrality* (ρ_p): Measure ρ_p is based on the skew in degree centrality, to cover the symmetric load perspective (as in DHTs).

2. *Node Betweenness* (ρ_{nb}): Measure ρ_{nb} is based on node betweenness, to cover the perspectives of targeted attacks as well as congestion (as in CDNs).

3. *Edge Connectivity* (ρ_λ): Measure ρ_λ is based on edge connectivity, to cover the targeted attack perspective (as in NCW).

4. *Node Deletions* (ρ_{fn}): Measure ρ_{fn} is based on the effect of random or targeted node deletions.

Robustness based on Degree Centrality

For **connected undirected graphs**, we define *skew in degree centrality* as the difference between the *maximum degree* in the graph ($\hat{p}$) and the *mean degree* of the nodes ($\bar{p}$). For a connected undirected graph of n nodes, the worst skew ($skew_{max}$) occurs for the star topology. The central node has a degree of $n-1$ and all the nodes surrounding it have a degree of 1. Therefore, the worst skew is $\frac{(n-1)(n-2)}{n}$. The best skew ($skew_{min}$) is 0, which is when all the nodes have the same degree. This occurs when the topologies are *regular* graph topologies as in a circular topology or a clique.

$$\rho_p = 1 - \frac{n(\hat{p} - \bar{p})}{(n-1)(n-2)} \tag{3.6}$$

In case of **strongly connected directed graphs**, we consider outdegrees. Again, the worst skew occurs for a directed graph in which the central node has an outdegree of $n-1$. All other nodes have an outdegree of 1. Like undirected graphs, the worst skew in degree centrality is $\frac{(n-1)(n-2)}{n}$. The best skew, which is 0, occurs for circular topologies or for a clique. As a result, the above relation continues to hold for directed graphs.

Expression (3.6) is due to the following mapping function. For a graph $\mathcal{G}$, with a degree skew, $skew_{\mathcal{G}}$:

$$\rho_p(\mathcal{G}) = 1 - \frac{skew_{\mathcal{G}} - skew_{min}}{skew_{max} - skew_{min}} \tag{3.7}$$

Robustness based on Node Betweenness Centrality

Similar to the skew in degree centrality, we define the skew in node betweenness centrality ($nbskew$) as the difference between the maximum node betweenness ($\hat{nb}$) and the average node betweenness ($\bar{nb}$) in a graph. For a **connected undirected graph**, the worst skew ($nbskew_{max}$) occurs for a star topology. In this case, all the shortest paths pass through the central hub of the star topology, and none through the rest of the nodes. Therefore, $nbskew_{max} = \frac{n-1}{n}$. The best skew occurs when an equal number of shortest paths pass through each node in the graph, $nbskew_{min} = 0$, as in a circle. Thus:

$$\rho_{nb} = 1 - n\frac{\hat{nb} - \bar{nb}}{n-1} \tag{3.8}$$

In case of **strongly connected directed graphs**, the above relation continues to hold. Expression (3.8) is due to the following mapping function. For a graph $\mathcal{G}$, with a node betweenness skew, $nbskew_{\mathcal{G}}$:

$$\rho_{nb}(\mathcal{G}) = 1 - \frac{nbskew_{\mathcal{G}} - nbskew_{min}}{nbskew_{max} - nbskew_{min}} \tag{3.9}$$

Robustness based on Edge Connectivity

A third way to measure robustness, ρ_{λ}, is in terms of *edge connectivity* (λ). Edge connectivity is the minimum number of edges whose removal renders a network disconnected. In case of an **undirected graph**, the tree topologies have the worst connectivity ($connectivity_{min}$) of 1; the clique has the best connectivity

($connectivity_{max}$), $n-1$. Thus, robustness, when defined in terms of connectivity is:

$$\rho_\lambda = \frac{\lambda - 1}{n - 2} \tag{3.10}$$

In case of **directed graphs**, the circle has the worst connectivity of 1. Again, the directed clique has the best connectivity, $n - 1$. Hence, the above result continues to hold.

The expression for robustness with respect to connectivity, ρ_λ, is due to the following mapping function. For a graph $\mathcal{G}$, with connectivity, $connectivity_{\mathcal{G}}$:

$$\rho_\lambda(\mathcal{G}) = \frac{connectivity_{\mathcal{G}} - connectivity_{min}}{connectivity_{max} - connectivity_{min}} \tag{3.11}$$

Functional Robustness based on Node Deletions

We also develop certain measures that we call *functional robustness*. Functional robustness indicates how well a network is doing after one or more nodes/edges stop working, either due to random failures or targeted attacks.

When a node v is deleted from a (strongly) connected graph, $\mathcal{G}$, the remaining network, $\mathcal{G} - v$ has one or more (strongly) connected components, $\mathcal{C}$, where $1 \leq |\mathcal{C}| \leq n - 1$. Let $\hat{\mathcal{C}}$ be the largest (strongly) connected component after a single node deletion.

We define the following quantities: (1) $nwsize = \frac{|\hat{\mathcal{C}}|}{n-1}$, the biggest fraction of the total remaining network that remains (strongly) connected after a node deletion, (2) $failure = 1.0 - nwsize$, fraction of the network that has failed after a node deletion.

We derive a robustness measure based on the above that holds for both directed and undirected graphs. The minimum failure occurs when a node deletion does not add any connected components to the network, that is the rest of the network remains intact. Therefore, $failure_{min} = 0$. The maximum failure occurs in case of a star network when its central node is deleted, resulting in $n-1$ components of size 1 each. Therefore, $failure_{max} = \frac{n-2}{n-1}$. The functional robustness of a graph, $\mathcal{G}$ is defined as:

$$\rho_{fn}(\mathcal{G}) = 1.0 - \frac{failure_{\mathcal{G}} - failure_{min}}{failure_{max} - failure_{min}} \tag{3.12}$$

We use this measure in our experiments to measure the effect of random node deletions as well as targeted attacks.

3.2.4 Cost

We divide cost into two components: (1) infrastructure cost as a function of the number of edges, e, in the network, and (2) node level maintenance/“bookkeeping” cost, as a function of the node’s degree, p (in case of directed graphs, outdegree, p^{out}). We place upper bounds on both these in our topology design.

Infrastructure Cost (κ): The minimum number of edges ($e_{min}^{undirected}$) required to have a **connected undirected graph** is $n-1$. We associate a cost, $\kappa = 0$, to a minimally connected graph. Any “extra” edge has an associated cost. All extra edges cost the same. An undirected clique has the highest cost, with $e_{max}^{undirected} =$

$\frac{n(n-1)}{2}$ number of edges.

We define the cost of an undirected topology with e edges as the ratio of the number of extra edges in a topology to the number of extra edges in the clique with the same number of nodes:

$$\kappa = \frac{e - e_{min}^{undirected}}{e_{max}^{undirected} - e_{min}^{undirected}} \tag{3.13}$$

In case of a **directed graph**, the minimum number of edges ($e_{min}^{directed}$) required to have a strongly connected graph is n, corresponding to the circle topology. And the maximum number of edges in a directed graph is, $e_{max}^{directed} = n(n-1)$, for a directed clique.

The cost of a directed topology with e edges is:

$$\kappa = \frac{e - e_{min}^{directed}}{e_{max}^{directed} - e_{min}^{directed}} \tag{3.14}$$

Maximum Permissible Degree (p): The Maximum Permissible Degree, p, is an upper limit on the number of edges that can be incident on a node. In case of a directed graph, p^{out}, is the upper limit on the number of outgoing edges. Degree is a measure of the local "bookkeeping cost". It can also be thought of as the amount of load a node is handling through edges incident on it.

3.2.5 Fitness

The general fitness function is as follows:

$$\phi = \alpha\rho + (1-\alpha)\eta - \beta\kappa \tag{3.15}$$

Here, $0 \leq \alpha \leq 1$, is an application dependent parameter that acts as a slider between efficiency and robustness. A high value of α indicates that a high emphasis should be placed on the robustness of topologies during topology breeding. The parameter $\beta,\ 0 \leq \beta \leq 1$, is used for additional cost control (in addition to the upper bound, κ). When set a high value, β helps the evolutionary process to "squeeze out" the most cost-effective topology that achieves a certain efficiency and robustness (controlled by α) by removing as many superfluous edges as possible.

Thus, the global optimization objective is to find the set of edges to construct the fittest graph:

$$\arg\max_{\mathcal{E}} \phi\left(\mathcal{G}\left(\mathcal{V}, \mathcal{E}\right)\right) \tag{3.16}$$

The rationale for defining fitness as a linear function over η, ρ and κ is that a linear combination is the least biased expression of the trade-offs. There is no clear way of correlating efficiency, robustness and cost to the best of our knowledge. While efficiency is non-decreasing with increasing cost, we can also find highly efficient ($\eta \approx 1$) topologies such as the star topology at low ($\kappa = 0$) costs. In case of robustness, the effect of cost depends on the metric used to measure robustness.

If degree symmetry is the robustness measure, topologies such as the straight line and the circle have high robustness ($\rho_p \approx 1$) even at low costs ($\kappa = 0$). In fact, addition of a few edges might reduce the robustness in such networks, due to degree asymmetry. If edge connectivity is the measure of robustness, robustness is non-decreasing with increasing cost. Similarly, the trade-offs between efficiency and robustness are also not clear. As a result, we use a simple linear function for fitness.

We present results from our experiments using several combinations of efficiency, robustness and cost measures (Chapter 4). Below, we give two examples in order to illustrate the fitness function:

1. *Parameters*: efficiency in terms of diameter (η_d), robustness in terms of degree centrality (ρ_p), cost, κ

   Fitness Function: $\phi_{d,p} = \alpha\rho_p + (1-\alpha)\eta_d - \beta\kappa$

2. *Parameters*: efficiency in terms of diameter (η_d), robustness in terms of edge connectivity (ρ_λ), cost, κ

   Fitness Function: $\phi_{d,\lambda} = \alpha\rho_\lambda + (1-\alpha)\eta_d - \beta\kappa$.

3.3 Genetic Algorithm Formulation

Genetic algorithm (GA) is a search technique developed to imitate the process of evolution through natural selection. We employ a GA process in our work to evolve or “breed” optimal topologies that are the “fittest” in terms of predefined performance requirements starting from random graphs with arbitrary fitness.

A genetic algorithm process runs through the following stages: *initialization*, *selection*, *reproduction* and *termination*. Typically, there is only one stage of initialization followed by several iterations of selection and reproduction until a "convergence" is reached, at which stage the process enters termination. Below, we discuss this process with respect to our problem.

3.3.1 Initialization

We start with a large number of *seed graphs* (typically 1000s). The seed graphs are generated randomly under the p and e constraints. These seed graphs form the *chromosomal strings* for the genetic algorithm. Thus, a chromosome consists of edges under the p and e constraints. In other words, a chromosome (or a seed graph) is a combination over the set of all possible edges in a n node network, such that the the number of edges in the combination is at most e, and each node gets a degree of at most p.

Table 3.1 shows a sample of chromosomal strings corresponding to undirected topologies. We represent chromosomes as edge lists instead of binary strings for the sake of convenience. These seed graphs were randomly generated with the following parameters: $n = 10$, $e \leq 14$ and $p \leq 5$.

Chromosome	Edge List	Topology
Seed Graph-1	[10 5, 1 6, 4 10, 6 9, 3 9, 5 8, 7 10, 8 7, 10 3, 4 3, 2 1, 5 7, 4 9, 6 4]	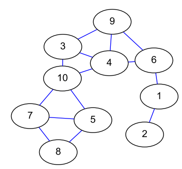
Seed Graph-2	[2 5, 2 7, 8 2, 8 1, 6 5, 8 9, 2 10, 3 10, 9 5, 1 7, 2 1, 7 3]	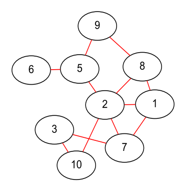

Figure 3.1: GA Initialization: Sample Chromosomal Strings

3.3.2 Selection

We retain only the connected seed graphs. After retaining the connected seed graphs, we sort them based on their fitness. Fitness is computed using the fitness function corresponding to a particular experiment.

A *mating pool* of a certain size is formed such that 80% of the *mates* are chosen based on fitness using the *fitness proportionate selection* or *roulette wheel selection* [Bäck, 1996]. In a fitness proportionate selection method, the fitness of an individual chromosome (seed graph) determines the probability of its being selected. Given a population of N, the probability, P_i, of a seed graph with fitness, ϕ_i, being selected is: $P_i = \frac{\phi_i}{\sum_{j=1}^{N} \phi_j}$

The remaining 20% of the mating pool is formed by choosing randomly from amongst the lesser fit seed graphs. Mixing some less fit offspring in the population is an effort towards reducing the chances of local minima.

3.3.3 Reproduction

We do an "all-pairs cross-over" over the mating pool. We use a number of crossover (CO) operators, some of them standard and some defined by us. Table 3.1 lists the crossover operators we use along with their descriptions.

The crossover operators are chosen randomly for every pair of mates. During crossover, it is ensured that the p and e constraints are not breached. Also, duplicate edges are removed. Figure 3.2 shows an example offspring resulting from the crossover of Seed Graph-1 and Seed Graph 2 when the *single point equal edges*

Operator	**Description**
Clone	Offspring are produced by simply cloning the parents
One Point with Equal Edges	Both parents contribute half their edges to produce two offspring
Cut–and–Splice	A variant of the above wherein the parents have different crossover points, resulting in offspring of different lengths
One Point with Equal Nodes	Both parents contribute half their nodes, along with edges incident on them, to produce two offspring
Two Point	Parents' edge lists are cut at two points and two offspring created using the six portions
Copy Edges	A random number of nodes are chosen from both the parents and offspring are formed by copying the edges incident on those nodes
Transfer Edges	A variant of the above in which offspring are formed by replacing randomly chosen edges from one parent with edges from another parent
Best Fit Edges	Offspring are created by choosing a greater fraction of edges from the fitter parent
Best Fit Nodes	Offspring are created by choosing a greater fraction of nodes, along with the edges incident on them, from the fitter parent

Table 3.1: GA Reproduction: List of crossover operators used in topology breeding

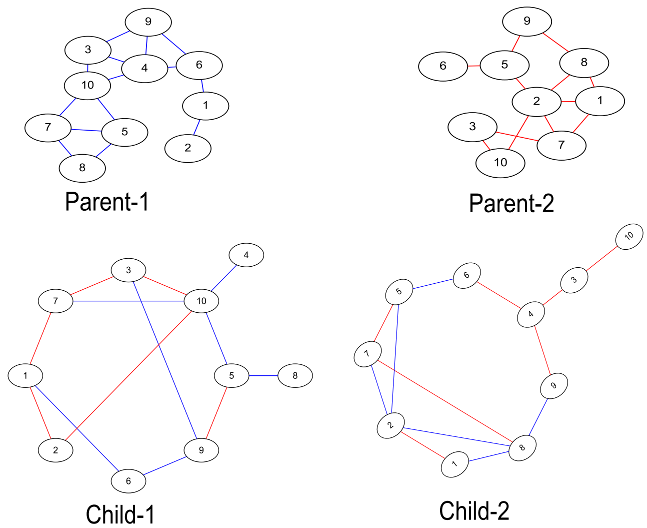

Figure 3.2: *Single Point Equal Edges* crossover of Seed Graph-1 (Parent-1) and Seed Graph-2 (Parent-2) to form two offspring

crossover operator is used.

At the end of the crossover phase, we have a large number of offspring, $(2 \cdot (mating\ pool\ size)^2)$. A small percentage (typically, 5%) of the offspring generated by crossover undergo random mutations. Mutations include addition, deletion or replacement of edges. It is ensured that the constraints on maximum degree and the maximum number of edges is not breached during mutation.

The offspring, some of which are mutated, form the seed graphs for the next round of selection. The whole process is repeated over several generations until convergence is reached. At the end of the process, we get an optimal topology for a given experiment.

3.3.4 Termination

We determine convergence using the following two parameters: (1) $iter$, which is the minimum number of iterations of the GA, and (2) $convergence$, which is the difference between the maximum fitness and minimum fitness in a generation of offspring. The GA proceeds until the number of iterations is $iter$ or $convergence \leq \epsilon$, where ϵ is a small positive value, whichever is later.

When $convergence \leq \epsilon$, essentially all offspring have the same fitness (most of them have the same topology). A topology with the highest fitness is then chosen as the optimal topology.

We observe that convergence is reached fast, within 20 iterations. The rate and quality of convergence depends on the number of seed graphs we use initially. For a small number of seed graphs, we observe that convergence is neither fast nor are the topologies generated very fit. Therefore, for our experiments, we have used a large number of seed graphs (1000s of seed graphs).

There is a possibility that a GA converges by ending up in a local minimum. We have tried to reduce the chances of local minima by inducing as much randomness in the GA as possible. Specifically, we take the following steps: (1) implement a large number of crossover operators and chose them randomly, (2) mutate a small percentage of a generation of offspring by random addition, deletion and replacement of edges, and (3) during the selection stage, choose 80% of the next mating pool by roulette wheel selection and the remaining 20% uniformly randomly from the least fit offspring.

3.4 Topology Breeding

To conduct experiments, we have developed a tool called Topology Analysis and visualiZation (TopAZ) as part of this work. TopAZ is a set of graph design, analysis and visualization libraries. Specifically, the genetic algorithm is implemented in a library called *TopBreed*. TopAZ is available in Python and Java versions.

Our experimental setup is called topology breeding: we "breed" or evolve optimal topologies under various constraints. We conducted topology breeding experiments using different design metrics for networks with up to 200 nodes.

Different types of structures emerge at different points in spaces of optimal topologies defined by the optimization parameters. In the next chapter, we present our main findings.

It is the pervading law of all things organic and inorganic,
Of all things physical and metaphysical,
Of all things human and all things super-human,
Of all true manifestations of the head,
Of the heart, of the soul,
That the life is recognizable in its expression,
That form ever follows function. This is the law.

Louis H. Sullivan. *The Tall Office Building Artistically Considered*

4

Topology Breeding: Optimal Topology Spaces

In this chapter, we address instances of the problem of network design under multiple constraints, which we described in Chapter 3. For each instance, we present an Optimal Topology Space (OTS) corresponding to the instance.

An OTS is formed by the different optimization parameters such as maximum degree (p), maximum number of edges (e), and the environmental variables, α and β. Every point in an OTS corresponds to an instance of a topology breeding experiment, and results in an optimal topology. We conduct several instances of topology breeding over an OTS to uncover optimal topologies at different points. This is akin to *sampling* an OTS.

We conducted experiments using different design metrics for networks with up to 200 nodes. Different types of structures emerge at different points in an OTS. Below we present a summary of our main findings. We consider undirected and directed topologies separately. We present topologies with a small number of nodes ($n = 20, 30, 50$) so as to illustrate the structures clearly. However, our observations hold for larger networks as well.

4.1 Undirected Optimal Topology Spaces (UDOTS)

4.1.1 Minimize Diameter under Constraints on Degree Centrality and Cost

Consider Problem Instance 1 from Chapter 3: *Given a network, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, of $\mathbf{n}$ nodes, each having a maximum permissible degree of $\mathbf{p}$, add a maximum of $\mathbf{e}$ undirected edges to minimize the diameter $\mathbf{d}$ of the network (Expression (3.1)). Ensure that the degree skew in the resulting network, $\mathbf{p_{skew}}$, i.e. the difference between the maximum degree in the network, $\hat{\mathbf{p}}$, and the average degree of the* *network,* $\bar{\mathbf{p}}$*, is no more than* δ*, where* $\mathbf{0} \leq \delta \leq \mathbf{n} - \mathbf{2}$ (Expression (3.6)).

Below we discuss the topologies that evolved by varying the optimization parameters. Figure 4.1 shows a sample breed of optimal undirected topologies under different *maximum permissible degree*, p and *emphasis on robustness*, α, for small values of *maximum number of edges*, e . Here, the cost control parameter, $\beta = 0$. Thus, topology breeding finds the most efficient (i.e. diameter optimal) topology using a maximum of e edges without additional cost control, under constraints on node degree, p. The fitness function used is:

$$\phi_{d,p} = \alpha\rho_p + (1 - \alpha)\eta_d$$

When $e = n - 1$ **(**$\kappa = 0$**) and** $p = 2$: Trees emerge as optimal structures at the lower boundary of cost in the OTS. When the maximum number of edges, $e = n - 1$ (i.e., cost, $\kappa = 0$), the only possible connected networks are trees. When $e = n - 1$, and the maximum permissible degree, $p = 2$, the only possible network is a straight line, with diameter, $d = n - 1$, which corresponds to an efficiency, $\eta_d = 0.0$. The degree centrality of a straight line is nearly symmetric, thus it has $\rho_p \approx 1$, for large n, owing to its near regular topology.

When $e = n - 1$ **(**$\kappa = 0$**) and** $p > 2$: As p increases, the diameter decreases, and different kinds of trees emerge. Consider the plane $e = 19$. When $\alpha = 0$, there is no emphasis on robustness. The breeding process evolves topologies to minimize the diameter as constraint p is relaxed. As we move along the p–axis, trees are formed in which the number of "hubs" decreases, while the sizes of the

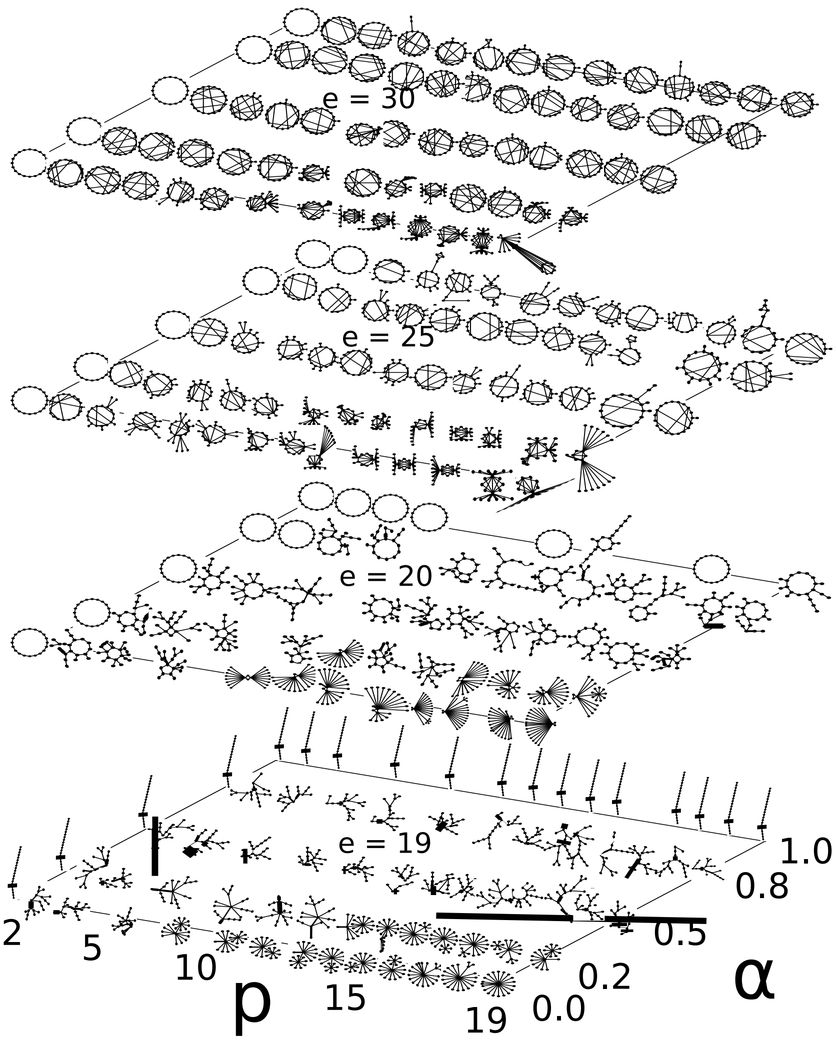

Figure 4.1: An undirected OTS formed by *maximum permissible degree*, p, and, *emphasis on robustness*, α, for different small values of maximum number of edges, e (or cost κ). The objective is to minimize diameter, d, (maximize η_d) under constraints on p, κ, and a bound on degree skew (ρ_p). Here, *cost control parameter*, $\beta = 0$. The no. of nodes, $n = 20$.

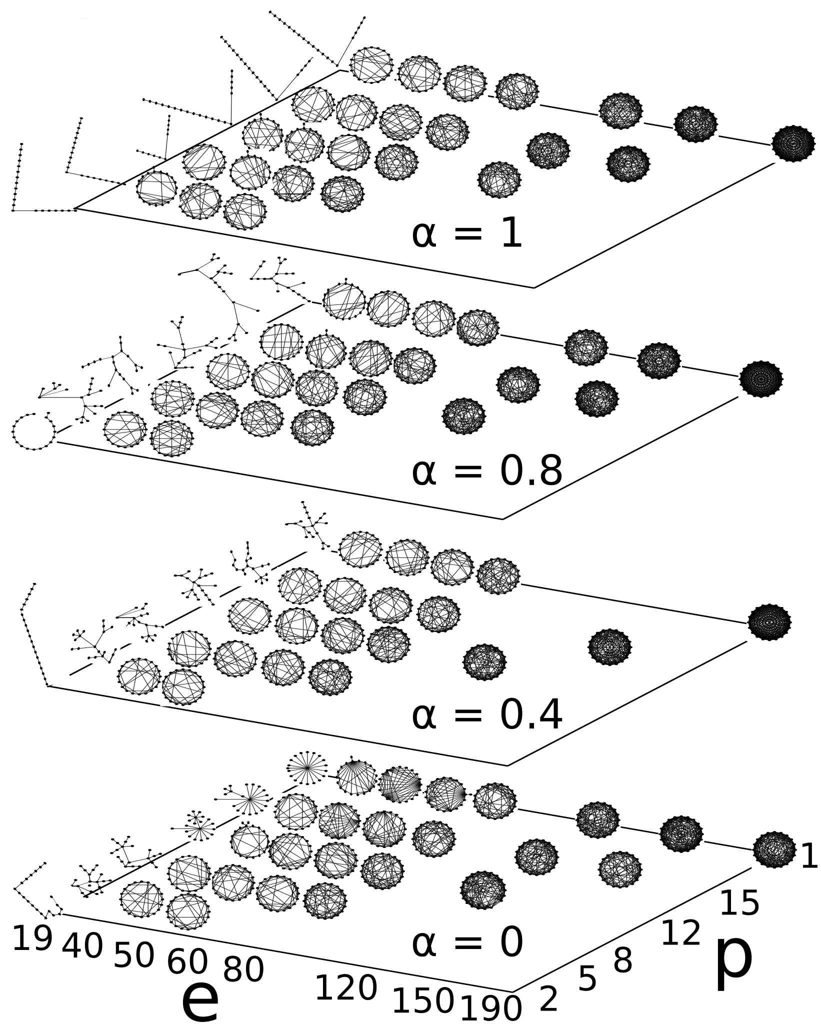

Figure 4.2: An undirected OTS formed by *cost*, κ, *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize diameter, d, (maximize η_d) under constraints on p, κ, and a bound on degree skew (ρ_p). Here, *cost control parameter*, $\beta = 0$. The no. of nodes, $n = 20$.

hubs increase. We call these topologies *hubs–and–spokes*. This process converges at the star topology, which has one central hub connected to $n-1$ spokes. A star is the most efficient topology when $(n-1) \leq e < \frac{n(n-1)}{2}$, with a diameter 2. However, it is the least robust topology in terms of degree centrality.

If we move along the α–axis (in the plane $e = 19$), we see that trees start "folding" into increasingly symmetric topologies in terms of degree centrality. This is due to the increasing emphasis on robustness. When $\alpha = 1$, we see straight lines, which are the most robust (and least efficient) topologies under the given conditions. In a straight line, every node has a degree of 2 except the two end nodes, which have a degree of 1. For a large n, the skew in degree centrality for a straight line is nearly 0.

When $e = n$ **and** $p = 2$: A circle emerges as optimal, in terms of both efficiency and robustness, regardless of the value of α. An undirected circle has a diameter of $\lceil \frac{n}{2} \rceil$. It is 2–regular. Being a hamiltonian circuit, it has an edge connectivity of 2. A circle also has a node connectivity of 2.

When $e = n$ **and** $p > 2$: In this case, *circular hubs* emerge as optimal, which are "hybrid" structures with a central circular core some of whose nodes act as hubs for the non-core nodes. When $\alpha = 0$, as we move along the p–axis, the circular cores grow increasingly smaller such that the diameter of the topology decreases. At $p = n-1$, the topology is a star with an additional edge that is used to form a triangular core. Since $\beta = 0$, the extra edge which does not contribute to diameter reduction is not removed.

As α increases, we observe circular hubs with increasingly large cores. When

$\alpha = 1$, we get circles or near circles as optimal structures.

When $e > n$: We start observing the "folding" of topologies. When $\alpha = 0$, for small values of p, we observe circular hubs. If we move along the p–axis, we observe a different class of topologies which we call "symmetric multihubs with spokes". These are topologies that have the following structure: (1) a small number of high degree nodes (hubs) form a clique, (2) a subset of the non-hub nodes have an edge to each of the hubs, but very few amongst themselves, and (3) the rest of the non-hub nodes are distributed among the hubs as spokes. As e increases further, we observe symmetric multihubs which are similar to the above structures, except there are no spokes. Also, all the hubs have comparable degrees. Please refer to Figure 4.2. In the plane $\alpha = 0$, we can observe such structures as the value of p increases.

As α increases, we observe a class of topologies that we call "circular skip lists" (CSL) (Figure 4.2). A circular skip list has the following structure: (1) every node has an edge to a distinct "successor" node, thus forming a logical circle or a cycle covering all nodes, and (2) every node has zero or more edges that connect to other nodes at random distances ("skips") on the logical circle. A CSL looks like a circle with zero or more random chords. A circle is the minimal CSL, whereas a complete graph (clique) is the maximal CSL.

Circular skip lists are *symmetric* topologies: every node in the topology has nearly the same degree. Hence, they are very robust. They also achieve low diameters at moderate cost. In Figure 4.2, as we move along the e–axis, we observe only CSLs. Since $\beta = 0$, the genetic algorithm does not make any effort towards deleting superfluous edges (edges that do not contribute towards increasing efficiency and/or robustness). Therefore, the CSLs are increasingly dense. At the upper boundary of the e–axis ($e = 190$), we find complete graphs.

Figure 4.4 shows a sample breed of optimal undirected topologies under different *maximum permissible degree*, p, *maximum number of edges*, e, and *emphasis on robustness*, α. Here, the cost control parameter, $\beta = 1$. Thus, topology breeding finds the most efficient (i.e., diameter optimal) as well as the most cost-effective topologies for the given upper bounds on degree, p, cost, κ, robustness, ρ_p (indicated by α). The fitness function used is:

$$\phi_{d,p} = \alpha\rho_p + (1 - \alpha)\eta_d - \kappa$$

The topologies that evolve in this experiment are comparable to the ones in Figures 4.1 and 4.2. The major difference is due to the parameter $\beta = 1$.

1. The plane $e = 19$ in Figure 4.3 is identical to that in Figure 4.1. Since $e = n - 1$, which is the minimum number of edges needed to construct a connected graph, the value of β has no significance.

2. When $e = n$ and $p = 2$, circles are formed. If we increase p while keeping $\alpha = 0$, circular hubs and trees are formed, which give way to hubs–and–spokes, eventually converging to the star topology.

3. In the plane $e = 20$, as we move along the α axis, we find that circular hubs are prevalent. Towards the upper boundaries of α, there seems to be

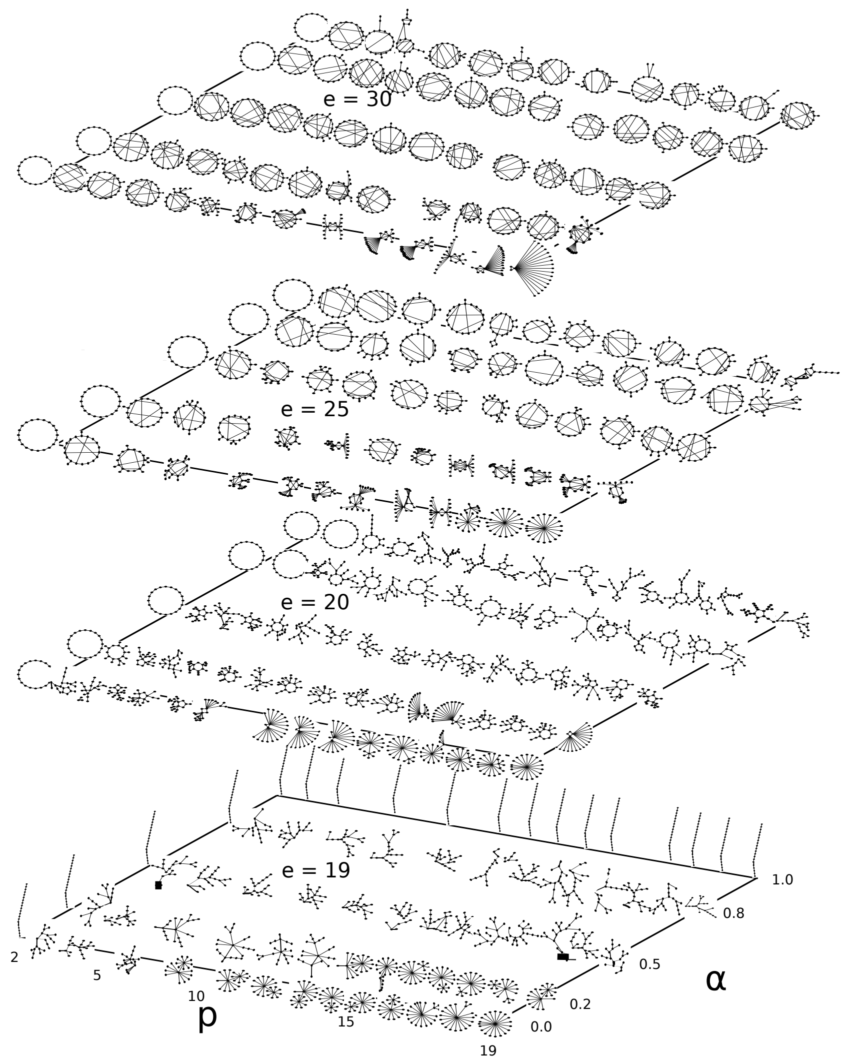

Figure 4.3: An undirected OTS formed by *cost*, κ, *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize diameter, d, (maximize η_d) under constraints on p, e, and a bound on degree skew (ρ_p). Here, *cost control parameter*, $\beta = 1$. The no. of nodes, $n = 20$.

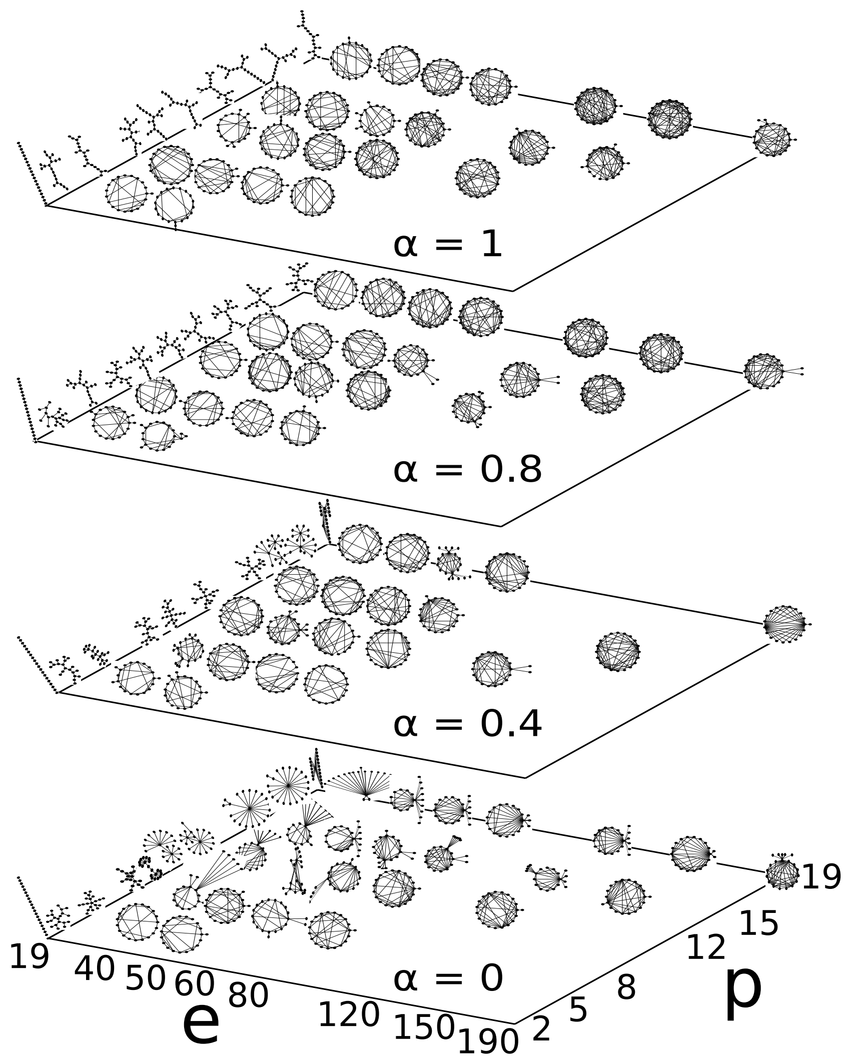

Figure 4.4: An undirected OTS formed by *cost*, κ, *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize diameter, d, (maximize η_d) under constraints on p, e, and a bound on degree skew (ρ_p). Here, *cost control parameter*, $\beta = 1$. The no. of nodes, $n = 20$.

a conflict between two tendencies: use n edges to form a circle, which has symmetric degree centrality, but an extra edge reducing its fitness slightly; and, to lose the one extra edge and become a straight line whose robustness is comparable to that of the circle. This likely leads to the formation of hybrids between a circle and a straight line.

4. As the value of e increases, the topologies that emerge are similar to the corresponding topologies in Figure 4.1. The only major difference arises due to the parameter β. In Figure 4.3, we observe topologies that shed as many superfluous edges as possible.

As in the case of Figure 4.1, for larger values of e, we observe circular skip lists. When the maximum permissible degree, p, is low, CSLs emerge as optimal. As p increases, CSLs start giving way to topologies that are a "hybrid" between CSLs and trees, with a large central loop and small peripheral hubs, as can be seen in figure 4.4. For higher p, the central loops start growing smaller and hubs larger; the topologies start "shedding" redundant edges (owing to a high value of β); eventually hub–and–spoke networks emerge.

With increasing e, the onset of hybrid topologies depends on the values of α and p. For low values of α, we observe that CSLs prevail until the following relation holds, $p \approx \frac{2e}{n}$. In fact, the best CSLs occur at the boundary, $p = \frac{2e}{n}$, where topologies utilize nearly all the edges to achieve low diameters. They are also robust with symmetric degree centrality despite not being optimized for robustness. For p beyond this, CSLs start unfolding and shedding edges to form hybrid structures. As α increases, the unfolding is delayed, and occurs for increasingly higher values of p.

For low values of α, we see networks that belong to the classes *symmetric multihubs* and *symmetric multihubs with spokes* towards the far p end in Figure 4.4. Under the conditions ($\alpha = 0$ and $\beta = 1$), ideally the optimization should converge to the star. However, the evolutionary process is unable to throw away all redundant edges, possibly due to local minima.

4.1.2 Minimize Diameter under Constraints on Edge Connectivity

Problem Instance 2. *Given a network, $\mathcal{G}\left(\mathcal{V}, \mathcal{E}\right)$, of n nodes, each having a maximum permissible degree of p, add a maximum of e undirected edges to maximize the edge connectivity λ of the network (section 3.2.3, Expression (3.10)).*

Figures 4.5 and 4.6 correspond to Problem Instance 2. When $\alpha = 0$, the robustness metric is irrelevant. Hence, Figure 4.5 is similar to the $\alpha = 0$ planes we encountered earlier.

Figure 4.6 shows undirected optimal topologies when the emphasis on robustness is maximum ($\alpha = 1$) and the cost control parameter is relaxed ($\beta = 0$). Edge connectivity is used as the metric of robustness. Thus, the breeding process evolves topologies with as high an edge connectivity, λ, as possible, constrained only by the maximum permissible degree, p, and the number of edges, e. The fitness function used is:

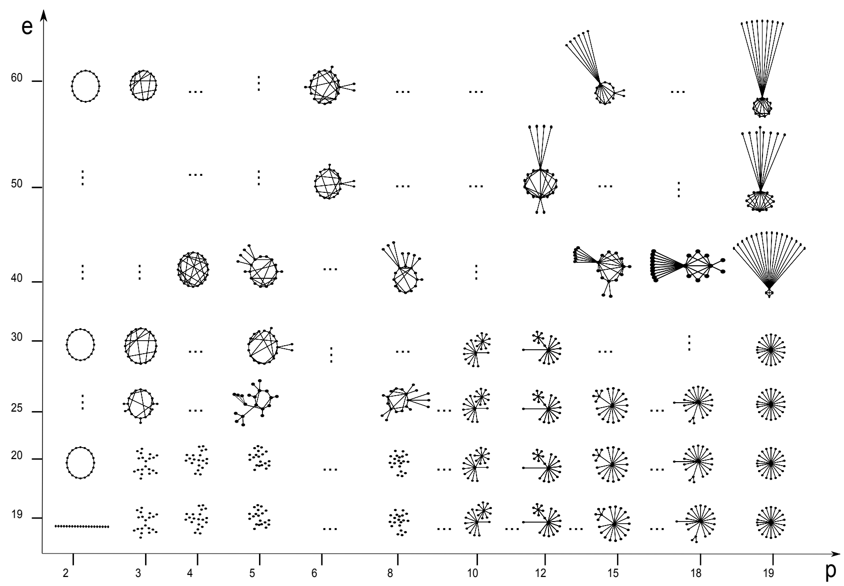

Figure 4.5: Undirected Optimal Topology Space with emphasis on robustness, $\alpha = 0$ and cost control parameter, $\beta = 1$, for different no. of edges, e and maximum permissible degree, p. The no. of nodes, $n = 20$.

$$\phi_{d,\lambda} = \alpha\rho_{\lambda} + (1 - \alpha)\eta_d$$

When $e = n - 1$, no useful topologies emerge. This is because all trees have an edge connectivity of 0. When $e = n$, circles emerge. A circle is 2–edge connected and 2–vertex connected. If a single edge or a node is deleted, the rest of the network still remains connected. Only when any two edges or two non-adjacent nodes are deleted does the network become disconnected.

As e increases, we observe that CSL are prevalent regardless of the value of p. This is because, for trees, hubs–and–spokes and symmetric multihubs with spokes, $\lambda = 1$. Moreover, since $\beta = 0$, there is no evolutionary pressure to shed

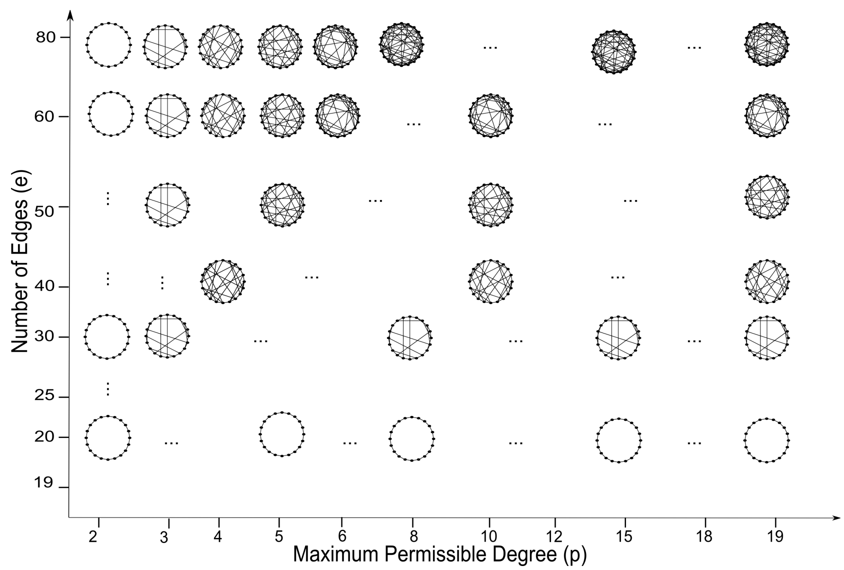

Figure 4.6: Undirected Optimal Topology Space with emphasis on robustness, $\alpha = 1$ and cost control parameter, $\beta = 0$, for different no. of edges, e and maximum permissible degree, p. The no. of nodes, $n = 20$. Robustness metric used is edge connectivity, ρ_λ.

redundant edges.

When $e >= \frac{np}{2}$, we observe that CSLs emerge with λ of $p-1$. Further, only a few node pairs have $p-1$ edge independent paths between them, with most node pairs having p edge independent paths, which is the theoretical upper bound (the maximal connectivity of a graph is equal to the minimal node degree in the graph).

4.1.3 Minimize Average Path Length under Constraints on Functional Robustness

In this section, we consider average path length (APL) as the metric for efficiency (η). Average path length of a network represents the cost of communication in the network on average. We consider three measures of functional robustness based on node deletions (as defined in Section 3.2.3).

Problem Instance 3. *Given a network, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, of* n *nodes, each having a maximum permissible degree of* p*, add a maximum of* e *undirected edges to minimize the average path length, l, of the network. Ensure that the size of the largest connected component, $\hat{\mathcal{C}}$, is no less than $n-k$, where, $1 \leq k \leq n-1$, upon deletion of a node with the highest degree in the network (refer to* $\mathbf{functional\ robustness}$ *in Section 3.2.3, and also Expression (3.12)).*

Figure 4.7 shows a sample breed of optimal undirected topologies under different values of *maximum number of edges*, e, and *maximum permissible degree*, p, for different emphases on functional robustness, ρ_{fn}, with respect to degree. Here, the cost control parameter, $\beta = 0$.

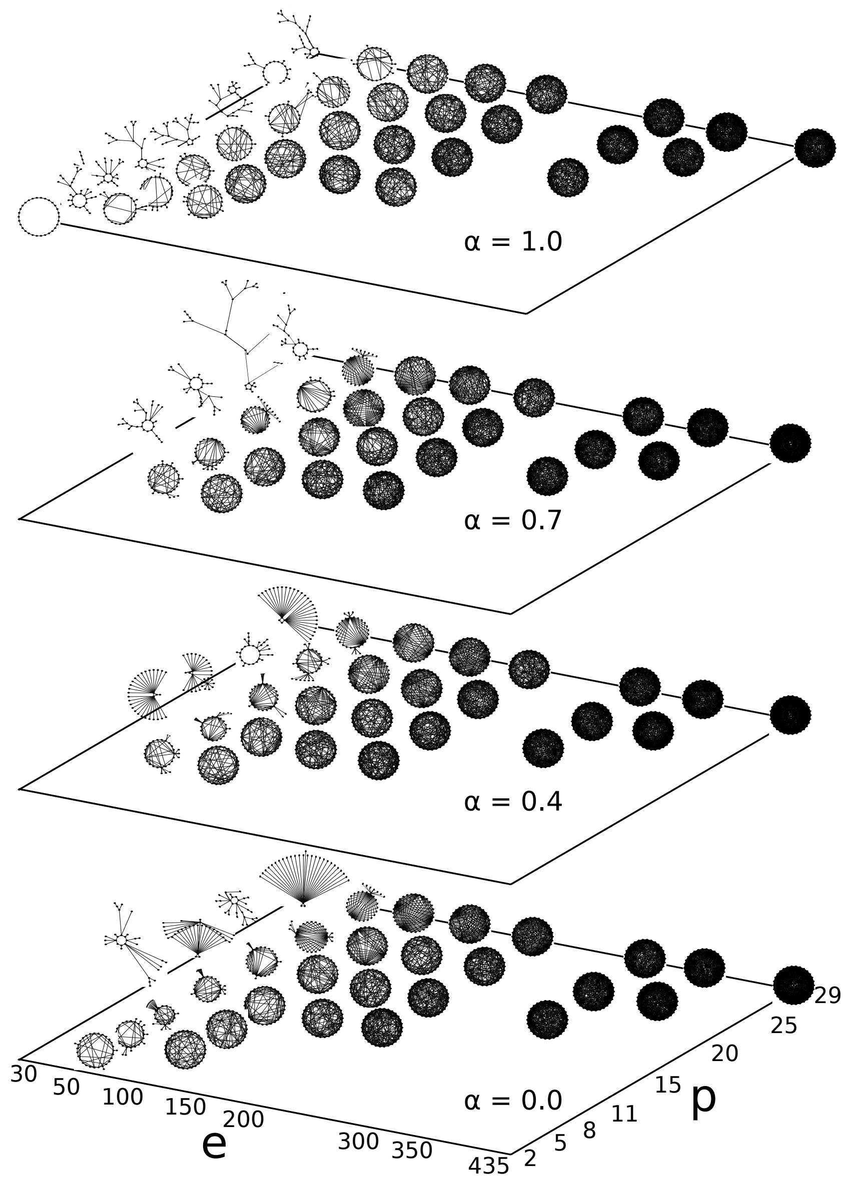

Figure 4.7: An undirected OTS formed by *maximum number of edges*, e, *maximum permissible degree*, p, and, *emphasis on robustness*, α and β. The objective is to minimize APL, under constraints on p, κ and α. Functional robustness with respect to degree is used. Here, $\beta = 0$. The no. of nodes, $n = 30$.

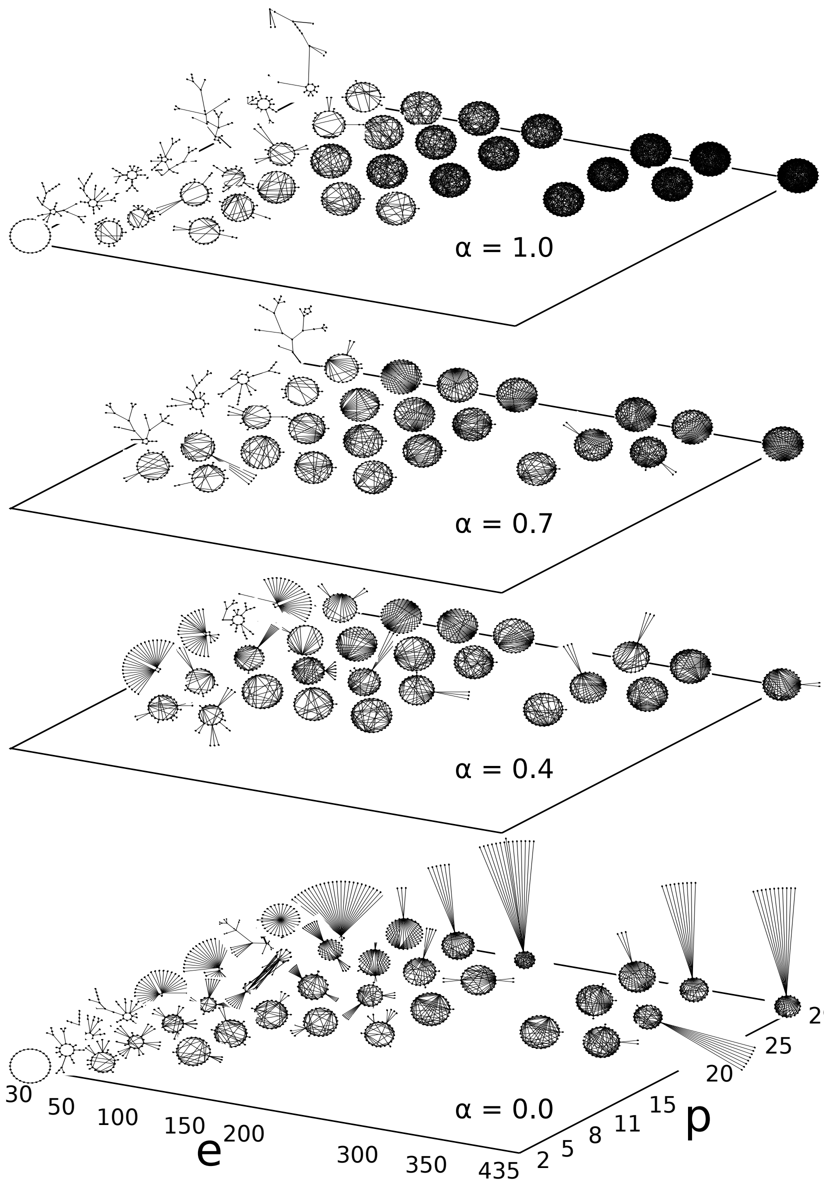

Figure 4.8: An undirected OTS formed by *maximum number of edges*, e, *maximum permissible degree*, p, and, *emphasis on robustness*, α, and cost control parameter, β. The objective is to minimize APL, under constraints on p, κ and α. Functional robustness with respect to degree is used. Here, $\beta = 1$. The no. of nodes, $n = 30$.

The fitness function used is:

$$\phi_{d,p} = \alpha\rho_{fn} + (1-\alpha)\eta_l$$

When α is low, most of the emphasis is one minimizing the average path length. Under restrictions on cost, circular hubs emerge as optimal topologies. As the constraint on cost is relaxed, topologies progress to symmetric multihubs with spokes, symmetric multihubs and CSL.

As α increases, again circular hubs emerge under severe restrictions on cost (refer to the plane $\alpha = 0.4$ towards the $e = 30$ boundary in Figure 4.7). Circular hubs with small cores with a large number of spokes are optimal since the higher emphasis is still on efficiency. However, the spokes are well distributed among the core nodes. Therefore, even if the node with the highest degree goes down, a significant majority of the rest of the network remains connected.

If we move along the α axis, while keeping the cost restriction, we observe topologies that are a hybrid between circular hubs and trees. No node in the network has a significantly high degree. So, the failure of the highest degree node does not affect the remaining network significantly.

Once we start relaxing the cost constraint, we observe increasingly symmetric topologies which eventually converge to dense circular skip lists. We can observe that intermediate topologies such as *symmetric multihubs with spokes* have very few spokes. The number of spokes start reducing further as we move along the α–axis.

Figure 4.8 shows a sample breed of optimal undirected topologies for the same optimization parameters, except the cost control parameter, β. Here, $\beta = 1$.

The fitness function used is:

$$\phi_{d,p} = \alpha\rho_{fn} + (1 - \alpha)\eta_l - \kappa$$

Most of the observations that we made for the previous case ($\beta = 0.0$) hold. However, owing to the cost control parameter, $\beta = 1$, the resulting topologies are less dense as we move along the e–axis.

Problem Instance 4. *Given a network, $\mathcal{G}\,(\mathcal{V}, \mathcal{E})$, of* $\mathbf{n}$ *nodes, each having a maximum permissible degree of* $\mathbf{p}$*, add a maximum of* $\mathbf{e}$ *undirected edges to minimize the average path length, l, of the network. Ensure that the size of the largest connected component, $\hat{\mathcal{C}}$, is no less than $n - k$, where, $1 \leq k \leq n - 1$, upon deletion of a node with the highest betweenness centrality in the network.*

Figures 4.9, 4.10 and 4.11 show OTS formed by *maximum number of edges*, e, and *maximum permissible degree*, p, for different α. Functional robustness, ρ_{fn}, with respect to node betweenness is used. The cost control parameter, β, is 0, 0.4 and 1, respectively.

We make the following observations:

1. As α increases, circular hubs emerge when the maximum number of edges, e is low. Spokes are well distributed among the nodes in the core such that, there are multiple hubs with comparable betweenness values.

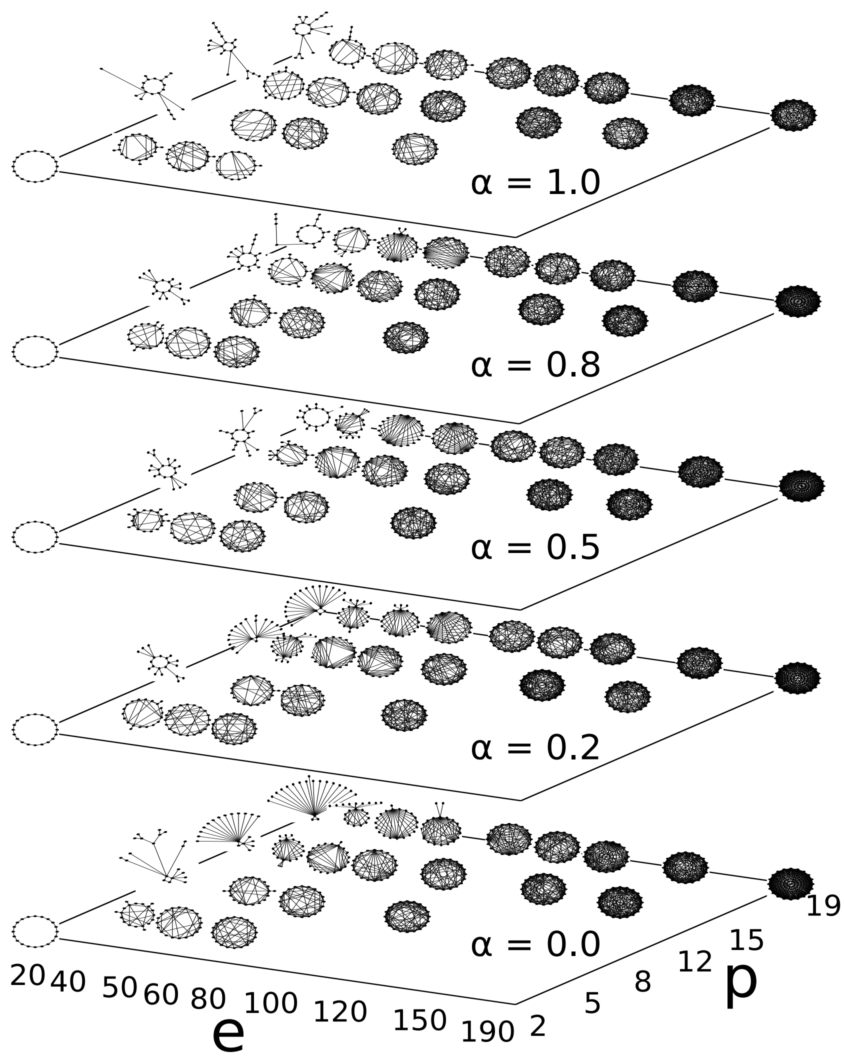

Figure 4.9: An undirected OTS formed by p, κ and α. Functional robustness w.r.t. node betweenness is used. Here, $\beta = 0$. The objective is to minimize APL, under the given constraints. The no. of nodes, $n = 20$.

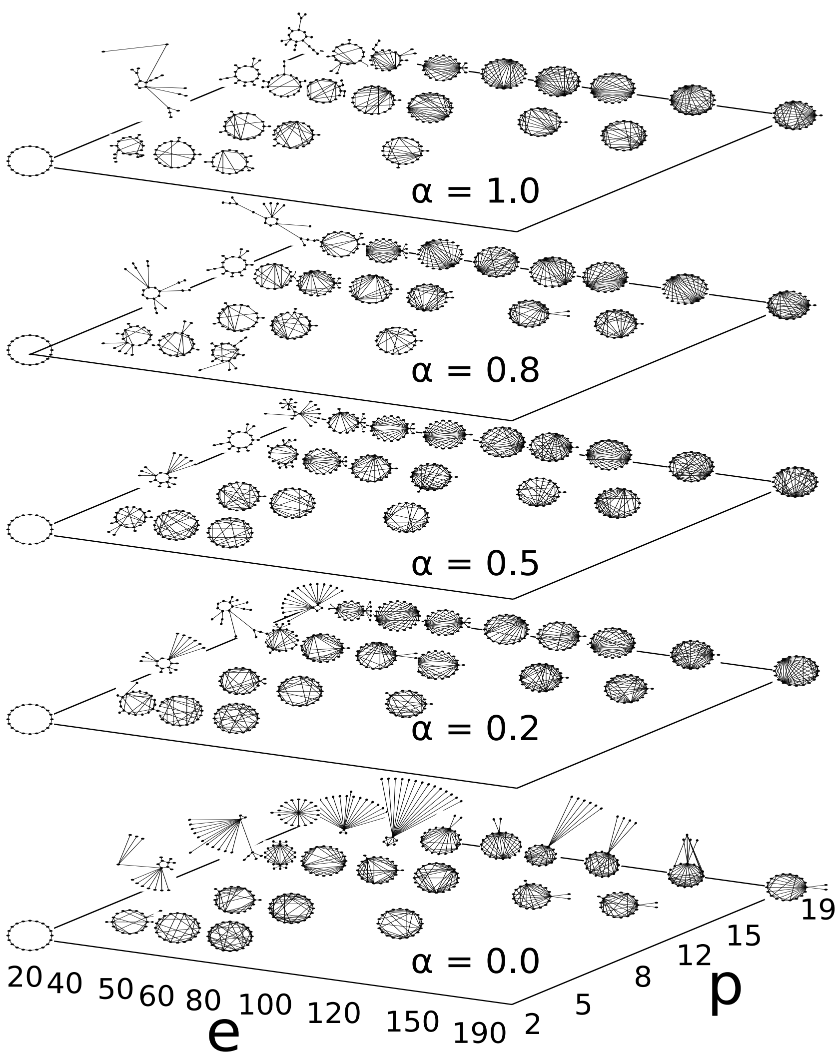

Figure 4.10: An undirected OTS formed by p, κ and α. Functional robustness w.r.t. node betweenness is used. Here, $\beta = 0.4$. The objective is to minimize APL, under the given constraints. The no. of nodes, $n = 20$.

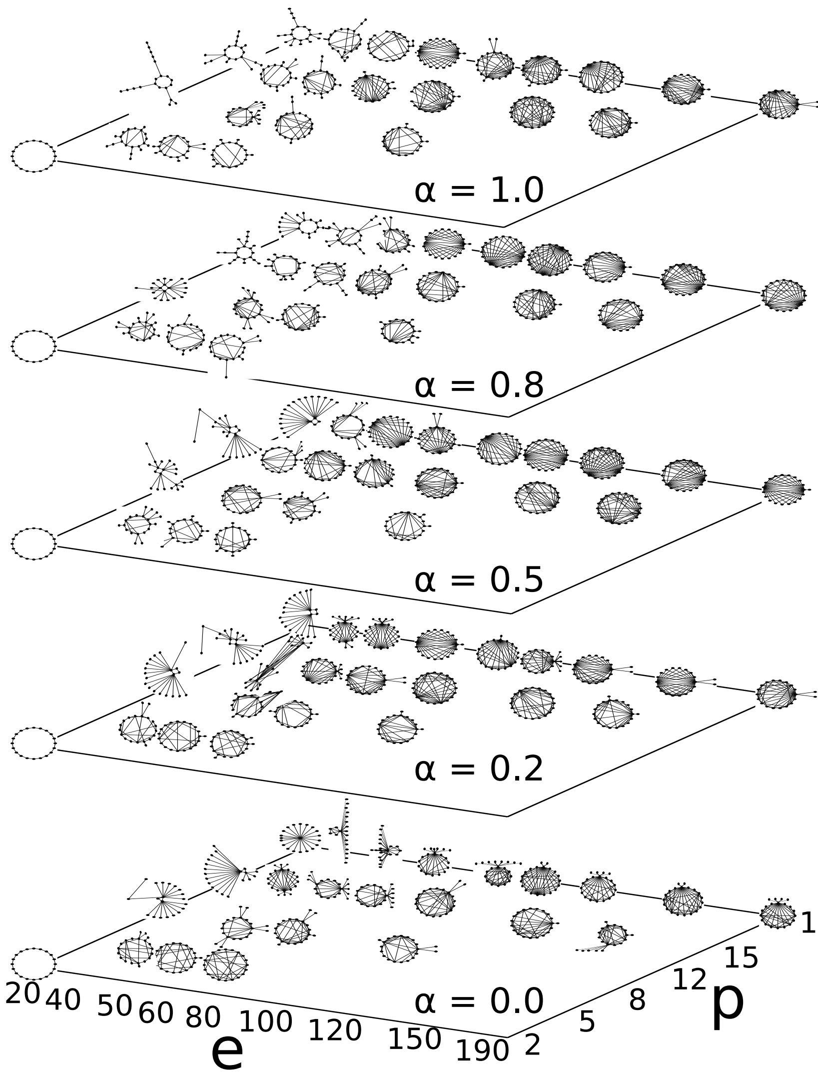

Figure 4.11: An undirected OTS formed by *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize APL, under constraints on p, κ and α. Here, $\beta = 1$. Robustness is defined in terms of node betweenness. The no. of nodes, $n = 20$.

2. As we move along the e–axis, circular hubs with increasingly large cores appear for low values of p. If β has a low value, even at low p values circular hubs have very few spokes.

3. If we move up the p–axis, we observe symmetric multihubs with spokes when α is low. As α increases, we observe symmetric multihubs.

4. Similarly, as e increases further, we observe symmetric multihubs. In symmetric multihubs, nodes have paths through multiple hubs. Failure of a single hub does not bring down the network.

5. At the upper boundaries of α and/or e, we observe circular skip lists whose nodes have comparable betweenness values. As the value of β increases the onset of CSLs is delayed. This is because the evolutionary process tends to settle at symmetric multihubs that are robust to the failure of a single node with high betweenness.

4.2 Directed Optimal Topology Spaces (DOTS)

We conducted similar experiments for directed graphs. Below we present our major findings.

4.2.1 Minimize Diameter under Constraints on Degree Centrality and Cost

Problem Instance 5. *Given a network, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, of* $\mathbf{n}$ *nodes, each having a maximum permissible indegree of* $\mathbf{p^{in}}$ *and a maximum permissible outdegree of* $\mathbf{p^{out}}$*, where,* $1 \leq p^{in} = p^{out} \leq n - 1$*, add a maximum of* $\mathbf{e}$ *directed edges to minimize the diameter* $\mathbf{d}$ *of the network. Ensure that both the indegree skew and the outdegree skew in the resulting network is no more than* δ*, where* $\mathbf{0 \leq \delta < n - 2}$.

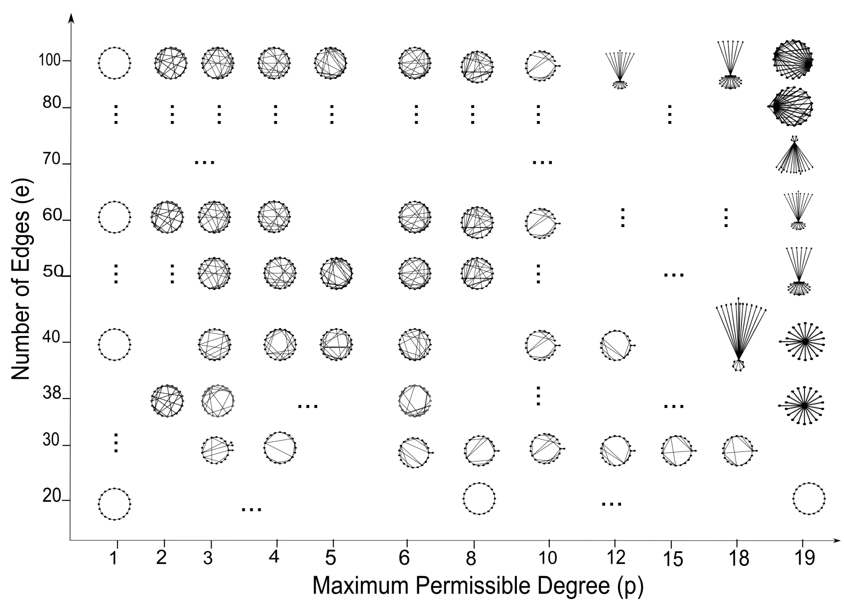

Figure 4.12: Directed Optimal Topology Space with emphasis on robustness, $\alpha = 0$ and cost control parameter, $\beta = 1$, for different no. of edges, e, and maximum permissible degree, p. The no. of nodes, $n = 20$.

Figure 4.12 shows the optimal directed graph topologies for $n = 20$, $\alpha = 0.0$ and $\beta = 1.0$. It corresponds to Problem Instance 5.

The minimally strongly connected directed topology is a circle. A directed circle is the only possible strongly connected topology when the cost, $\kappa = 0$ ($e = n$), regardless of the other parameters. A directed circle has a diameter of $n-1$; therefore, its efficiency, $\eta_d = 0$. However, since every node has an indegree as well as an outdegree of 1, it is symmetric leading to high robustness in this case, $\rho_p = 1$.

We observe that circular skip lists occur more naturally in directed graphs even when the emphasis on robustness, $\alpha = 0$, unlike the case of undirected topologies. This is especially true when there are restrictions on the maximum number of edges, e, and/or the maximum permissible degrees, p^{in} and p^{out}. Directed graphs need to achieve *strong connectivity*, whereas undirected graphs satisfy only *weak connectivity*. Circular skip lists help achieve strong connectivity. This is because, the underlying directed circle is minimally strongly connected. The chords in the skip list help in reducing diameter.

Similar to undirected topologies, hubs–and–spokes achieve better diameter reduction than CSL, which is the objective when $\alpha = 0$. The first directed hubs occur around $e = 2(n-1)$, when the constraint on p is relaxed. These are the equivalents of undirected star and hub-and-spoke networks.

Also, similar to the case of undirected graphs, large hubs with highly clustered cores can be seen for high values of e and p. These are directed equivalents of multihubs with spokes. Towards the upper boundaries of e and κ, we also observe directed symmetric multihubs. Owing to the value of the cost control parameter, $\beta = 1$, topologies try and shed as many superfluous edges as possible.

We observe that except the region near the right hand side of Figure 4.12, CSLs are prevalent everywhere else. This effect is even more complete as α increases, when there is more emphasis on symmetry of degree centrality.

4.2.2 Minimize Diameter under Constraints on Edge Connectivity

Problem Instance 6. *Given a network, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, of $\mathbf{n}$ nodes, each having a maximum permissible indegree of $\mathbf{p^{in}}$ and a maximum permissible outdegree of $\mathbf{p^{out}}$, where, $1 \leq p^{in} = p^{out} \leq n-1$, add a maximum of $\mathbf{e}$ directed edges to maximize the edge connectivity λ of the network (section 3.2.3, Expression (3.10)).*

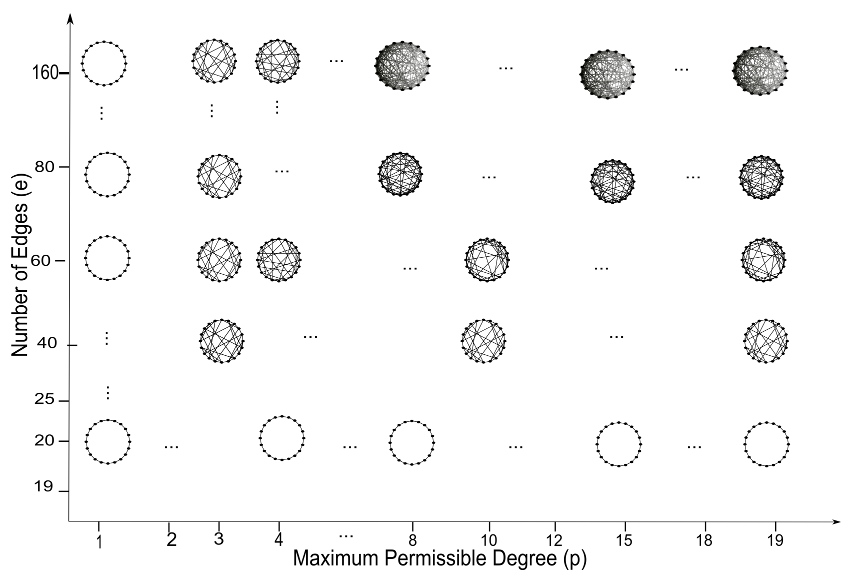

Figure 4.13: Directed Optimal Topology Space with emphasis on robustness, $\alpha = 1$ and cost control parameter, $\beta = 0$, for different no. of edges, e, and maximum permissible degree, p. The no. of nodes, $n = 20$. Robustness metric used is edge connectivity, ρ_λ.

Figure 4.13 shows directed optimal topologies when edge connectivity, λ, is used as the robustness metric. When $e = n$, the only possible strongly connected topologies are directed circles. Directed circles have an edge connectivity of 0, i.e., the deletion of an edge disconnects the network into multiple components (n in this case).

Since the emphasis on robustness is high ($\alpha = 1$), no useful topologies emerge under severe cost restrictions. When we move along the e–axis, with the cost control parameter also relaxed ($\beta = 0$), we see topologies very similar to the corresponding case of undirected graphs (Figure 4.6). CSLs with high edge connectivity prevail. Similarly, when $e > np$, CSLs with λ of $p - 1$ emerge.

In fact, CSLs are the dominant topologies even for higher values of β as well as lower values of α. Only when $\alpha = 0$ (and for high values of β) do we observe other classes of topologies such as hubs–and–spokes (as in the case of Figure 4.12) since the complete emphasis there is on efficiency and the metric for robustness is irrelevant.

4.2.3 Minimize Average Path Length under Constraints on Functional Robustness

In this section, we use average path length as the metric for efficiency. We consider two metrics of functional robustness in our experiments.

Problem Instance 7. *Given a network, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, of* n *nodes, each having a maximum permissible indegree of* $\mathrm{p^{in}}$ *and a maximum permissible outdegree of*

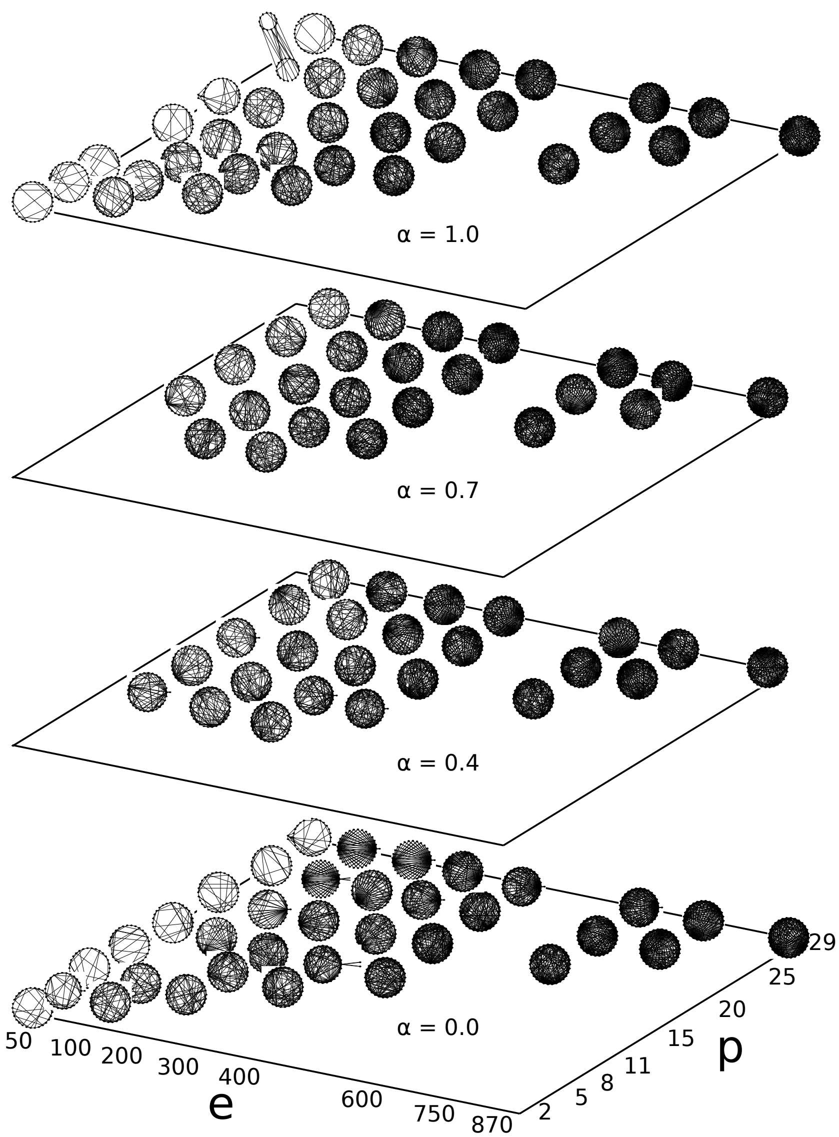

Figure 4.14: A directed OTS formed by *maximum number of edges*, e, *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize APL, under constraints on functional robustness with respect to degree. Here, $\beta = 1$. The no. of nodes, $n = 30$.

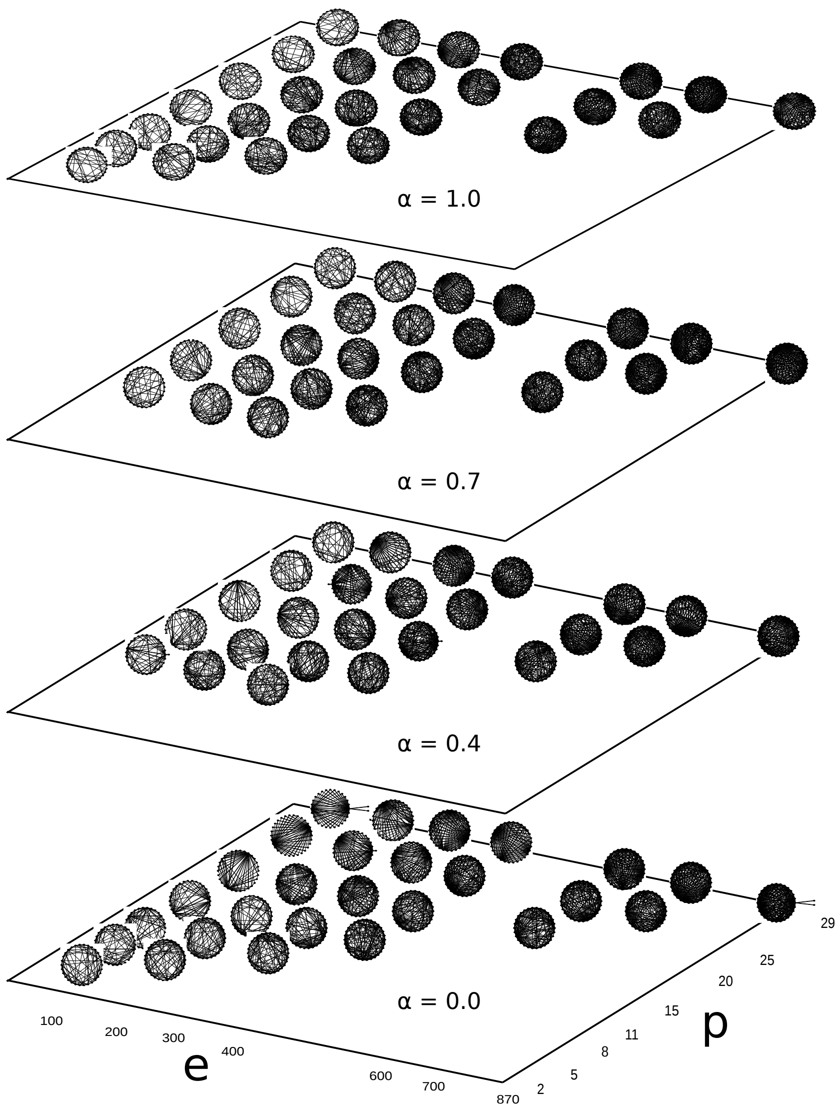

Figure 4.15: A directed OTS formed by *maximum number of edges*, e, *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize APL, under constraints on functional robustness with respect to degree. Here, $\beta = 0.6$. The no. of nodes, $n = 30$.

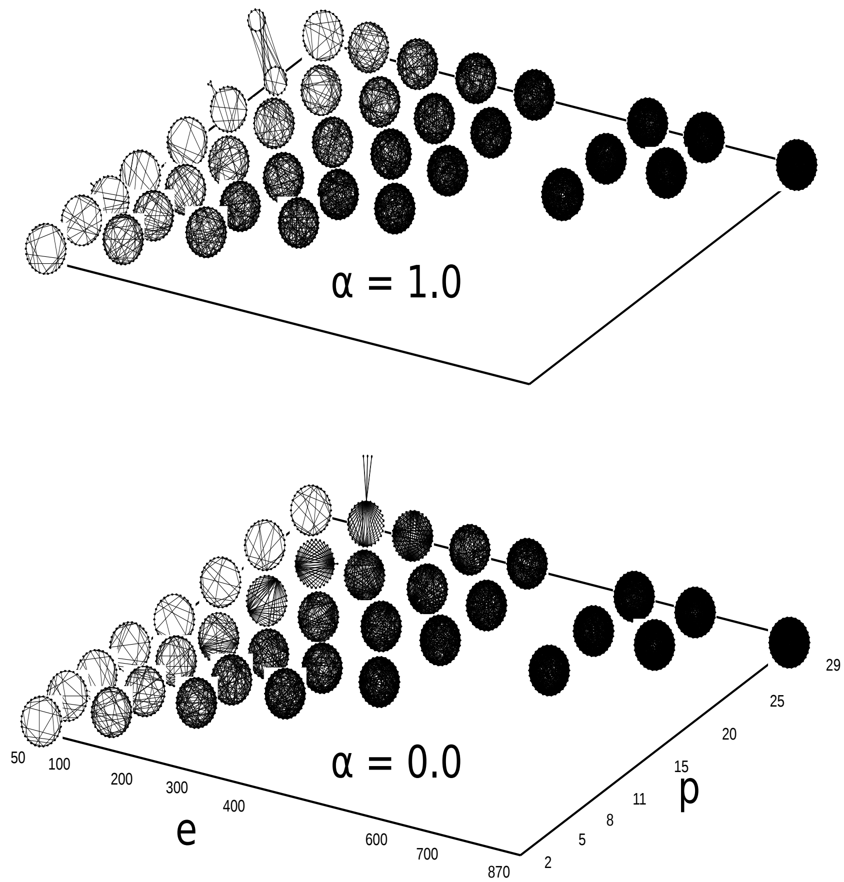

Figure 4.16: A directed OTS formed by *maximum number of edges*, e, *maximum permissible degree*, p, and, *emphasis on robustness*, α. The objective is to minimize APL, under constraints on functional robustness with respect to degree. Here, $\beta = 0$. The no. of nodes, $n = 30$.

$\mathbf{p^{out}}$, *where,* $1 \leq p^{in} = p^{out} \leq n-1$*, add a maximum of* $\mathbf{e}$ *directed edges to minimize the average path length,* l*, of the network. Ensure that the size of the largest strongly connected component,* $\hat{\mathcal{C}}$*, is no less than* $n-k$*, where,* $1 \leq k \leq n-1$*, upon deletion of a node with the highest degree in the network (refer to* **functional robustness** *in Section 3.2.3, and also Expression* (3.12)*).*

Figures 4.14, 4.15 and 4.16 show sample breeds of directed topologies corresponding to Problem Instance 7. Optimal topology spaces are formed by parameters e, p and α. The values of β are 1, 0.6 and 0, respectively.

When $\alpha = 0$, circular skip lists emerge for low values of p out of necessity to form strongly connected networks. When $e = n$, the only possible network is a directed circle which is neither efficient nor robust. For low values of p, CSLs continue to emerge, since the constraint on p does not allow the formation of large hubs. Towards the far end of the p–axis, we observe symmetric multihubs. They are efficient due to the large hubs. Symmetric multihubs give way to CSL as soon as the value of β is lowered as the pressure on shedding redundant edges reduces.

As α increases, CSLs continue to emerge as optimal for low values of p and e. However, instead of symmetric multihubs at the upper p boundary, we observe CSL which have a small number of large hubs. The hubs improve the efficiency of the network, whereas the “closed” structure of CSLs make the topologies robust against hub failures. As we move further along the p and e axes, CSLs become increasingly more symmetric. Their density varies based on the value of β.

Nature uses only the longest threads to weave her patterns, so that each small piece of her fabric reveals the organization of the entire tapestry.

Richard P. Feynman

5

Characterizing Optimal Topology Spaces

In Chapter 4, we conducted different instances of topology breeding experiments to evolve optimal topologies. We observe recurring patterns among topologies that emerge to be optimal such that they can be classified into a set of well-defined classes. These classes occur in specific subspaces as we will show subsequently.

The prominent classes that we observe are: (1) hubs–and–spokes, (2) circular hubs, (3) symmetric multihubs, and (4) circular skip lists. In this chapter, we analyze these topology classes to better understand their properties. In particular, we are interested in the class of circular skip lists.

5.1 Classes of Optimal Topologies

5.1.1 Hubs–and–Spokes

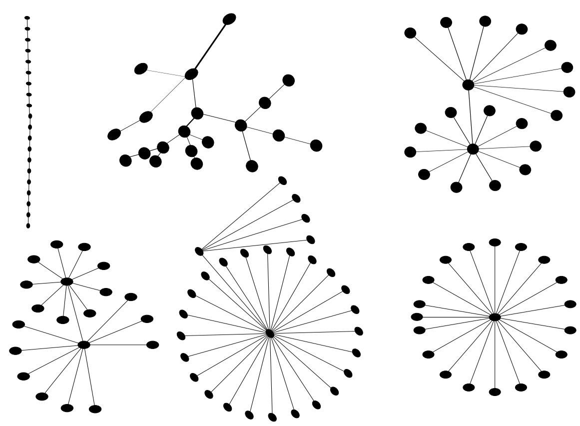

Figure 5.1: Undirected Hubs–and–Spokes

Figure 5.1 show a sample of undirected hubs–and–spokes. Undirected hubs–and–spokes are trees, which occur when conditions 1 and 2 below hold. Condition 3 determines the size of the hub; the bigger the hub, the more efficient the network.

1. Severe restrictions on cost, (e is close to $n-1$ and β is close to 1).

2. Emphasis on efficiency is very high (α is close to 0); in other words, emphasis on robustness is very low.

3. There is no restriction on the maximum permissible degree, p.

In case of directed networks, hubs–and–spokes are rare. They occur under conditions similar to the above, when p^{in} and p^{out} are both unrestricted and $e \approx 2(n-1)$.

We make the following observations about this class:

1. A straight line is the minimal hub–and–spoke topology with $n-2$ hubs each having a degree of 2. A star is the maximal hub–and–spoke with a single central hub connected to $n-1$ spokes.

2. A straight line has the least efficiency in terms of diameter or average path length. However, it has high symmetry in terms of centrality measures.

3. Efficiency increases as hubs become larger, i.e., when the number of hubs reduces and the number of spokes increases.

4. On the other hand, both symmetry and resilience to failures decrease as hubs become larger.

5. A star is the most efficient topology in terms of diameter when, $n-1 \leq e < \frac{n(n-1)}{2}$. It is also the least robust both in terms of symmetry and resilience to targeted attacks as every path in the network goes through the single hub. It is resilient to random node failures.

5.1.2 Circular Hubs

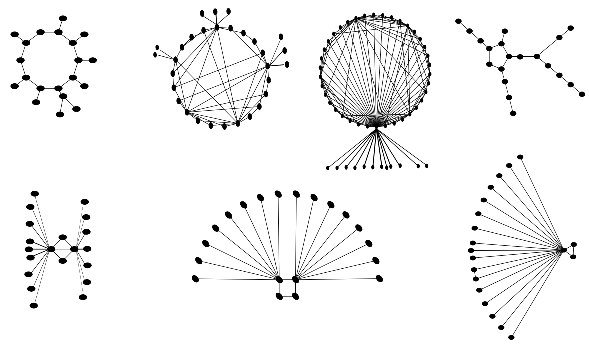

Figure 5.2: Undirected Circular Hubs

Circular hubs are hybrid topologies with a circular core whose nodes have zero or more spokes. Both undirected and directed circular hubs occur when one of the following conditions holds:

1. The maximum number of edges allowed, $e = n$, maximum permissible degree, $p > 2$,

   (a) emphasis on efficiency, $\alpha = 0$, and cost control parameter, $\beta = 0$—since $\beta = 0$, the algorithm does not put effort in removing the one superfluous edge; we observe a small circular hub (a triangle or a square) with the rest of the nodes forming spokes.

   (b) emphasis on efficiency, α, has an intermediate value and cost control parameter, $\beta = 0$; we observe a relatively bigger circular core (formed in an effort to satisfy the emphasis on robustness) with the rest of the nodes forming spokes around it.

2. The maximum number of edges allowed, $n < e < \frac{np}{2}$, where $p > 2$, emphasis on robustness α has an intermediate value, and cost control parameter, β has a high value.

We make the following observations about this class:

1. Circular hubs with small cores are close in structure to hubs and spokes. Therefore, they have low diameters roughly equal to the size of the core. Also, they have low symmetry and resilience to failures.

2. The diameter of circular hubs with large cores depends on the size of the core and the presence of "chords" (edges that are neither part of the circle nor spokes). In the absence of chords, the diameter is roughly equal to half the size of the core. Presence of chords brings down the diameter by forming short cuts between pairs of nodes.

3. Symmetry and resilience to failures of circular hubs with large cores depends on the size of the core and the distribution of the spokes across hub nodes. If the spokes are distributed evenly across a large number of hub nodes, symmetry and resilience to failures is relatively high.

5.1.3 Symmetric Multihubs

Figure 5.3 shows instances of a class of topologies that we call symmetric multihubs. These are topologies that have the following structure: (1) a small number of high degree nodes (hubs) form a clique, (2) a subset of the non-hub nodes have

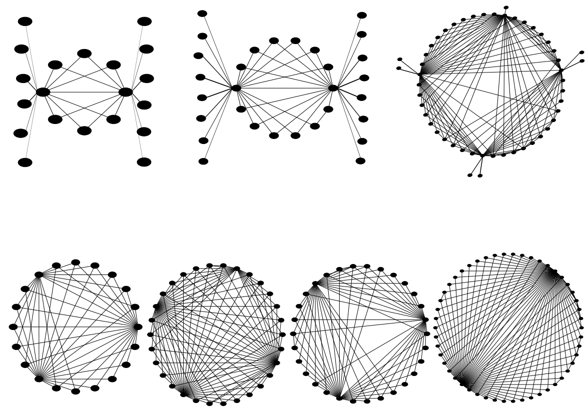

Figure 5.3: Undirected Symmetric Multihubs

an edge to each of the hubs, but very few amongst themselves, and optionally, (3) the rest of the non-hub nodes are distributed among the hubs as spokes. Also, all the hubs have comparable degrees.

Symmetric multihubs occur under the following conditions:

1. When $e > n$, we start observing the "folding" of topologies. When $\alpha = 0$, for small values of p, we observe circular hubs. However, for larger values of p, we observe symmetric multihubs with spokes.

2. Symmetric multihubs occur especially when the constraint on robustness is defined in terms of resilience to node deletions. Under this constraint, symmetric multihubs occur for low and intermediate values of α, low values of β and high values of p. If β is high, then they occur even for high values of α.

Symmetric multihubs have the following properties:

1. The hub nodes form a clique among themselves indicating a *rich club* phenomenon [Zhou and Mondragón, 2004]. Also, the nodes that are neither hubs nor spokes connect to the hub nodes but not to other similar nodes. Thus, we observe a *disassortative mixing* [Newman and Girvan, 2003].

2. Diameters are low since routes can be established through hub nodes in the manner of hubs–and–spokes. Further, the number of hubs does not impact the diameter since the hub nodes form a clique among themselves; a hub node can be reached from another in a single step.

3. Symmetry values are intermediate and improve with the number of hubs. Robustness of symmetric multihubs without spokes is high. Since the non-hub nodes connect to multiple hubs, alternate paths can be found in case of failure of hub nodes. However, the presence of spokes reduces the robustness.

5.1.4 Circular Skip Lists

Circular skip lists (CSL) are the most prominent topologies across different OTS. They have the following structure: (1) a cycle that encompasses all the nodes in the network, i.e., a hamiltonian circuit, and (2) each node has zero or more random chords—edges to nodes at different skips on the hamiltonian circuit.

Figure 5.4 shows a number of undirected CSL topologies. CSLs start emerging as soon as severe constraints on cost are relaxed and/or emphasis on robustness is increased. Specifically, an undirected CSL occurs under the following condi-

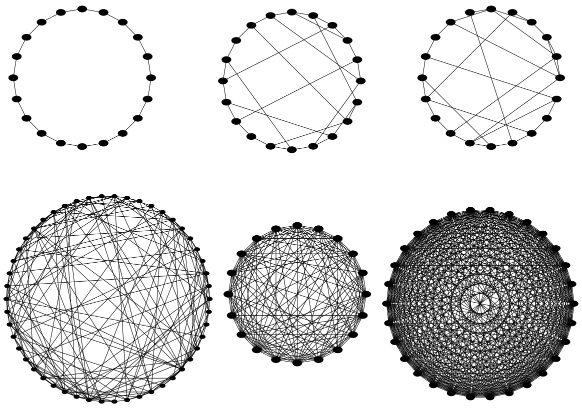

Figure 5.4: Undirected Circular Skip Lists

tions:

1. When $e = n$ and $p = 2$, a circle emerges as the optimal topology, which is the minimal CSL, regardless of the value of α and β. When $e > n$, the cost constraint is less severe. For a value of α higher than 0, CSLs start emerging. As we increase e and reduce β, CSLs with increasing density emerge. When the value of $e = \frac{n(n-1)}{2}$ and $\beta = 0$, we observe complete graphs, which are the maximal CSL.

2. When values of α are low, CSLs prevail until the following relation holds, $p \leq \frac{2e}{n}$. The best CSLs occur at the boundary, $p = \frac{2e}{n}$, where topologies utilize nearly all the edges to achieve high efficiency. They are also robust as a side effect even though the emphasis on robustness is low. For p beyond this, CSLs start unfolding and shedding edges to form hybrid structures such as circular hubs and multihubs. However, as α increases, the unfolding is delayed, and occurs for increasingly higher values of p.

As we have argued before, in case of directed topologies, CSLs are more natural. The minimally strongly connected directed topology is a directed circle, which is the minimal directed CSL. As we can observe across directed OTS, except under rare conditions, CSLs are prevalent. Both undirected and directed CSL topologies are featured with many properties that are optimal as we will discuss in Section 5.4.

5.2 "Phase Transitions" from Hubs–and–Spokes to CSL

In this section, we characterize the spaces of optimal topologies. We show how the classes of topologies we discussed above are distributed in an OTS. We also make some observations that hold across different OTS, i.e., properties about OTS that hold regardless of the optimization parameters that form them.

Figure 5.5 shows the sharp transition from hubs–and–spokes to CSLs as an effect of α on the breeding process, when the maximum number of edges, e is small. Here, $n = 20$ and $\beta = 1.0$. The efficiency metric used is diameter, η_d (Expression (3.1)). The robustness metric used is that of symmetry in degree centrality, ρ_p (Expression (3.6)).

Only when $e = 19$ do we observe hub–and–spoke networks. There are only two types of topologies here that show high fitness: (1) the star topology, which occurs when $\alpha = 0$ and $p = n - 1$, has a diameter of 2; hence it is very efficient, although its robustness is zero, (2) on the other hand, the straight line, which

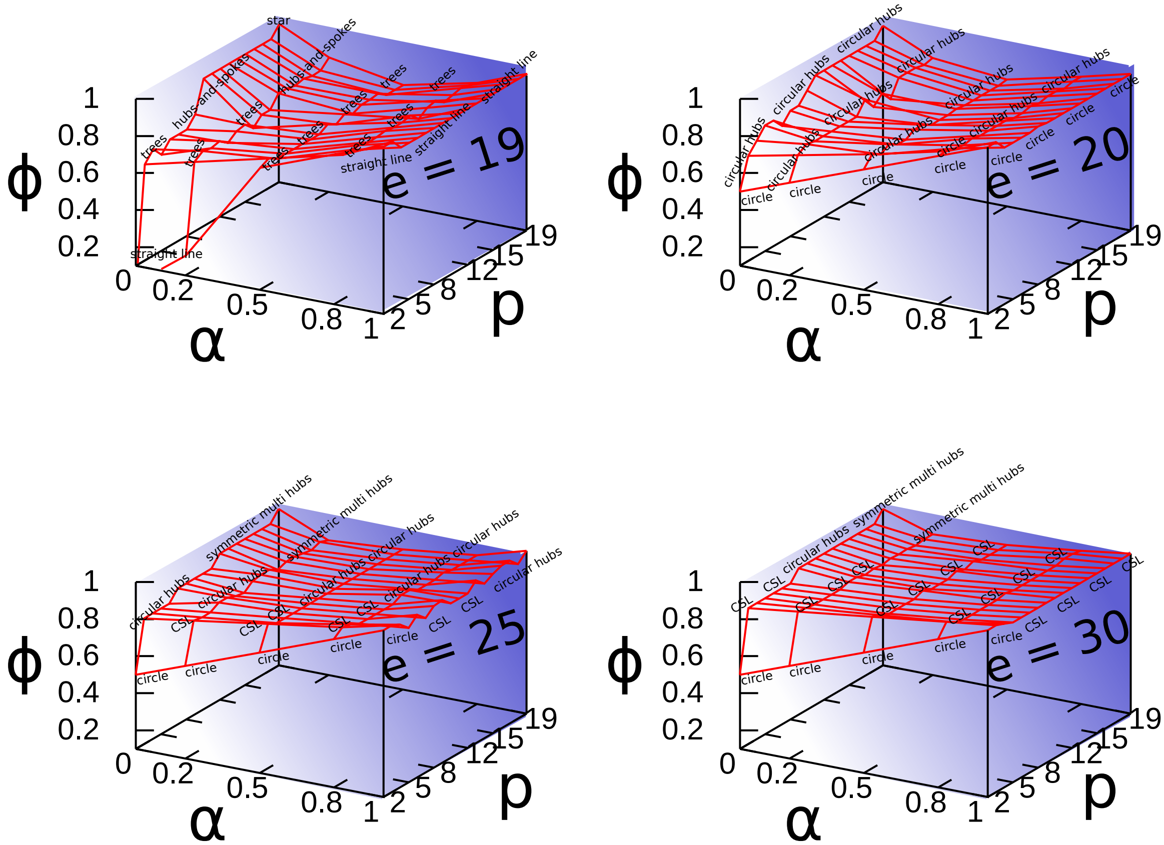

Figure 5.5: *Fitness* (ϕ) *vs.* α *vs.* p. No. of nodes, $n = 20$ and cost control parameter, $\beta = 1.0$. CSLs start emerging as soon as the number of edges, $e > n$.

occurs when $\alpha = 1$, has a nearly symmetric degree centrality; hence it is very robust, although its efficiency is zero. The plot shows the two extremes and the "dip" in fitness for intermediate topologies.

As e increases even slightly, hubs–and–spokes give way to hybrid networks such as circular hubs and symmetric multihubs. With a further increase in e and/or α, CSL topologies emerge. When $e = 30$, we see that CSL are prevalent throughout. We also observe that the surface representing α vs. p vs. ϕ becomes increasingly flat and smooth. This indicates that the trade-off between efficiency and robustness is pronounced only under severe constraints on cost.

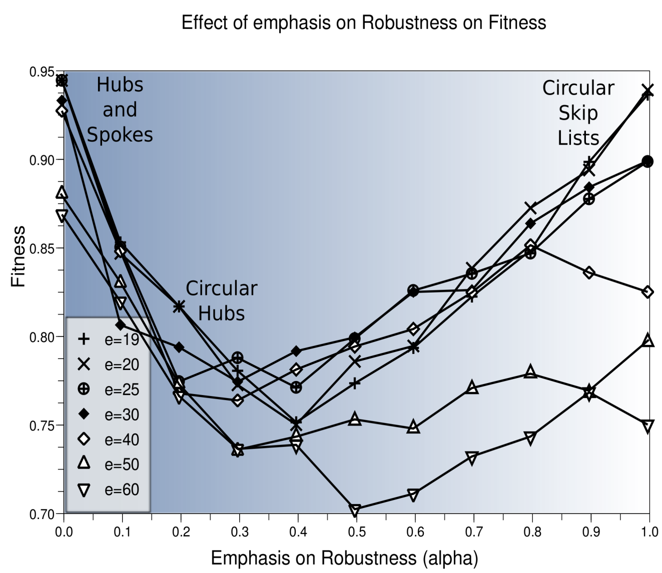

Figure 5.6: *Fitness vs* α. Except for very small no. of edges (e), hub–and–spoke networks start "folding" into CSLs even at a low emphasis on robustness (α), despite the maximum permissible degree, $p = n - 1$. No. of nodes, $n = 20$ and cost control parameter, $\beta = 1.0$.

Figure 5.6 shows a cross section of ϕ vs. α vs. p, with $p = n - 1$ to emphasize the trade-off between efficiency and robustness under severe cost constraints. We can observe that as we increase α slightly, "folding" of topologies starts happening even when the maximum permissible degree, p, is unconstrained.

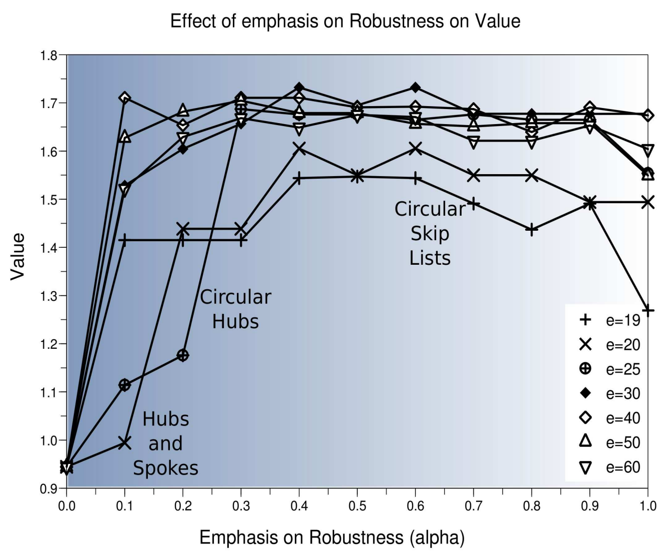

Figure 5.7: *Value vs* α. The *value* of a network, $v = \eta + \rho - \kappa$ rises sharply. Despite $p = n - 1$ and $\beta = 1$, we observe that CSLs set in even for small values of α. Hence, the "flattening" of the value curves.

We define a measure called *value* to which is defined as: $v = \eta + \rho - \kappa$. Value measures what we get back in terms of efficiency and robustness when we invest a certain cost. In Figure 5.7, we plot *value* vs α and observe that the curve rises sharply, again indicating the onset of CSLs even for small increases in α.

Emergence of topologies is dependent on multiple parameters: in this case, the emphasis on robustness, α, the maximum allowed edges, e, and the maximum permissible node degree, p. Here, we have focussed on understanding the behaviour of the breeding process as we increase e in small steps. Figure 5.5 shows the phase transitions from hubs–and–spokes to CSL as we increase the maximum number of edges allowed, e. The order of phase transitions is as follows: hubs–and–spokes $\Longrightarrow$ circle $\Longrightarrow$ circular hubs and/or symmetric multihubs $\Longrightarrow$ CSL.

We did similar analysis for other metrics of efficiency (e.g., average path length) and robustness (e.g., edge connectivity) for both undirected and directed graph topologies that were evolved through topology breeding. We observe similar transitions.

5.3 Distribution of Optimal Topologies

Through topology breeding experiments we obtain a lot of samples of optimal topologies by varying the number of nodes (n), maximum number of edges (e), maximum degree (p), emphases on robustness (α) and cost control (β). In the next set of results, we compute the efficiency (η), robustness (ρ) and cost (κ) for each sample regardless of the parameters under which it emerged; and show the distribution of these samples along these three axes.

Figure 5.8 shows such a plot for undirected topologies. We considered 1000s of *optimal topologies* each of which was generated under a specific combination of n, e, p, α and β. Here, the measure of efficiency is diameter, η_d (Expression (3.1)). The measure of robustness is in terms of skew in degree centrality, ρ_p (Expression (3.6)). Cost is measured in terms of the number of edges (κ).

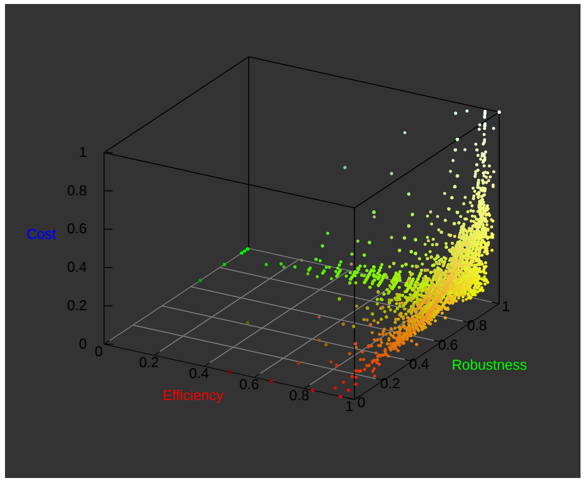

Figure 5.8: *Efficiency vs. Robustness vs. Cost for* η_d *and* ρ_p. Distribution of topologies that evolved as optimal under varying conditions w.r.t. n, e, p, α and β. Here efficiency is computed using diameter (d) and robustness is computed using degree symmetry.

Thus, we compute the diameter, skew in degree centrality and cost of each sample optimal topology.

1. In Figure 5.8, each point corresponds to a sample topology that we obtained through topology breeding. The colour of a point is an RGB value obtained by adding the values of its efficiency (represented by an R value), robustness (represented by a G value) and cost (represented by a B value).

2. We can observe that there is a small number of points that have only R values with very little G and B values. These are topologies that have high efficiency and low cost; but they have highly skewed degree centrality. They correspond to hubs–and–spokes. The red points close to $\eta = 1$ correspond to the star topology.

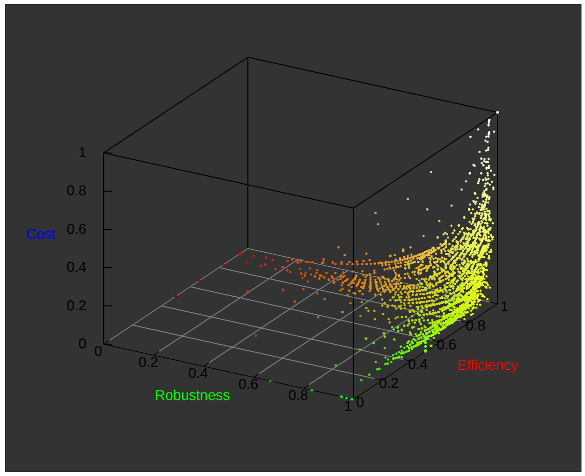

Figure 5.9: *Robustness vs. Efficiency vs. Cost for* η_d *and* ρ_p. The distribution of optimal topologies from the perspective of robustness.

3. Figure 5.9 shows the same plot from the perspective of robustness. Again, we see a small number of points that have high G values, and negligible R and B values. These are highly symmetric topologies corresponding to straight lines and circles.

4. As we can observe, there are very few topologies on the $\kappa = 0$ plane except near the $\eta_d = 1$ and $\rho_p = 1$ corners. Thus, the topologies that cost very little optimize either on diameter or degree symmetry.

5. The topologies that evolved under less severe cost constraints are not subject to a rigorous trade-off between efficiency and robustness. Therefore, as the cost of topologies increases (the "lift" in the plot away from $\kappa = 0$ indicates the increase in the cost of topologies), we observe increasingly more topologies concentrated near the region of high robustness and high

efficiency. The colours of the points indicates that the topologies have both good efficiency and robustness values; and they have moderate costs. These topologies correspond to circular skip lists.

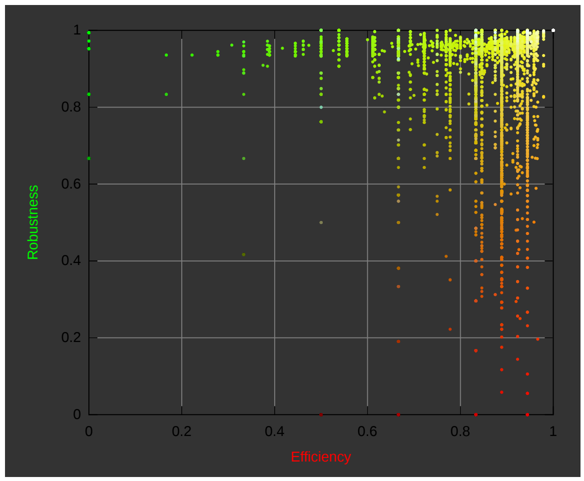

Figure 5.10: *Efficiency (η_d) vs. Robustness (ρ_p).* The projection efficiency vs. robustness shows that most topologies are both efficient and robust; very few are either efficient or robust.

Figure 5.10 shows a projection of the previous plots to show efficiency vs. robustness of the evolved topologies. This projection corroborates our argument that the trade-off between efficiency and robustness is pronounced only under severe cost constraints. As we can observe, most topologies that are efficient are also robust, and vice versa. Invariably these topologies are circular skip lists.

This implies that, under relaxed cost constraints, regardless of whether the emphasis is on efficiency or robustness, the breeding process evolves topologies that balance both requirements. For instance, even when the value of $\alpha = 1$, the process evolves a topology that is not only robust, but also efficient. In this case, high efficiency is a side effect of the topology being a circular skip list, with its long range skips reducing the diameter of the topology. In other words, the influence of α on the optimization process reduces as the constraint on cost is relaxed.

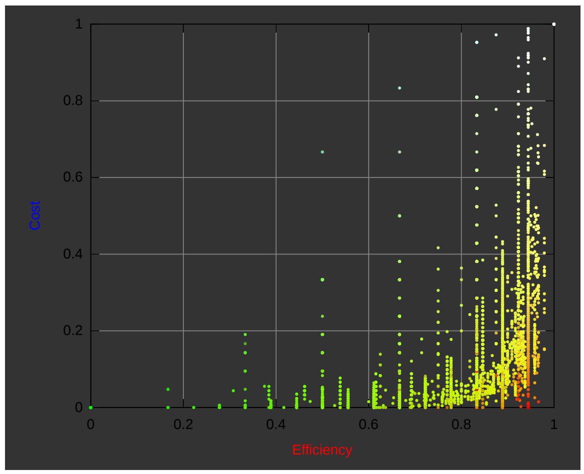

Figure 5.11: *Efficiency (η_d) vs. Cost (κ).* The projection efficiency vs. cost shows that most of the efficient topologies have moderate cost.

Figures 5.11 and 5.12 show projections of efficiency and robustness against cost, respectively. We can observe that most topologies achieve high efficiency and/or robustness under moderate costs.

The shapes formed by the distribution of points are interesting. Figures 5.10 and 5.11 show a series of “pillars”. This is due to the efficiency measure we have chosen: diameter changes in discrete steps; hence the discrete pillars. If we plot optimal topologies for a given value of n alone, the pillars will be equispaced

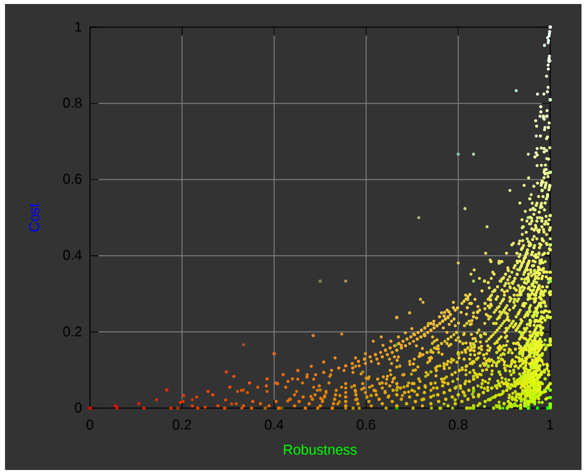

Figure 5.12: *Robustness (ρ_p) vs. Cost (κ)*. The projection robustness vs. cost shows that most of the robust topologies have moderate cost.

(diameter changing by 1).

On the other hand, the measure of robustness—skew in degree symmetry—is a continuous quantity. Hence, we observe the layered curves in Figure 5.12.

We can observe a small number of white points in the above plots, which are topologies that have the highest efficiency and robustness, and they also cost the highest. These correspond to complete graphs.

Figures 5.13 and 5.14 show a distribution of the topologies when we consider APL as the metric for efficiency (Expression (3.3)) against functional robustness with respect to deletion of the node with the highest betweenness (Expression (3.12)).

We make the same observation as before: most topologies are efficient as well as robust regardless of parameters n, p, α and β, as long as the constraint on the

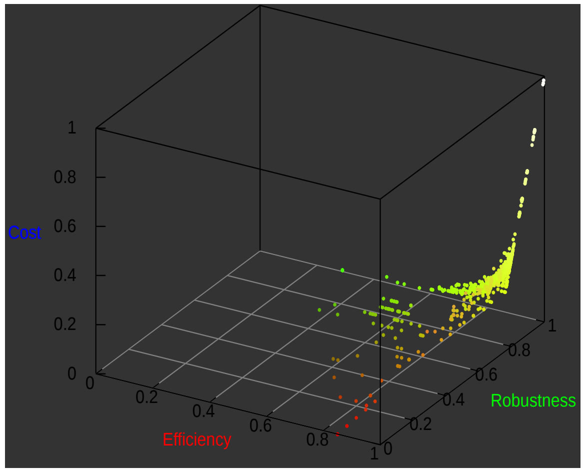

Figure 5.13: *Efficiency vs Functional Robustness vs Cost for η_l and ρ_{fn}.* Distribution of topologies that evolved as optimal under varying conditions w. r. t. n, e, p, α and β. Here efficiency is computed using APL (l) and robustness is computed using functional robustness w. r. t. deletion of the node with the highest betweenness.

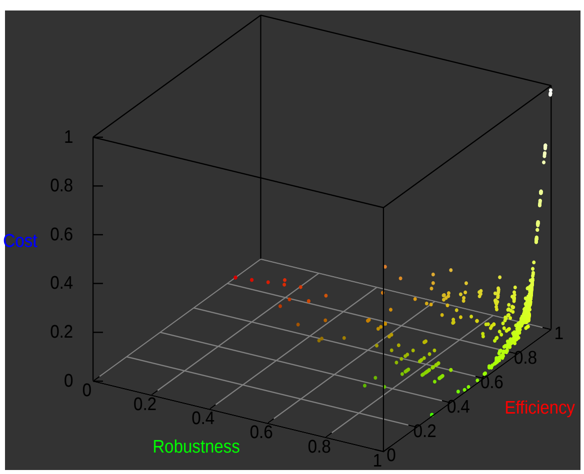

Figure 5.14: *Functional Robustness vs. Efficiency vs. Cost for η_l and ρ_{fn}.* The distribution of optimal topologies from the perspective of robustness.

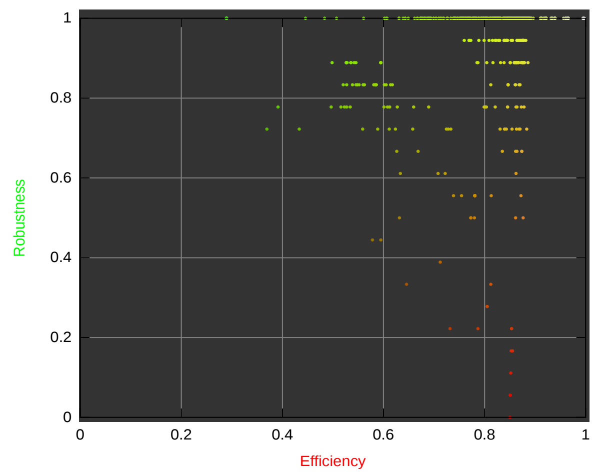

Figure 5.15: *Efficiency (η_l) vs Functional Robustness (ρ_{fn})*. The projection efficiency vs. robustness shows that most topologies are both efficient and robust; very few are either efficient or robust.

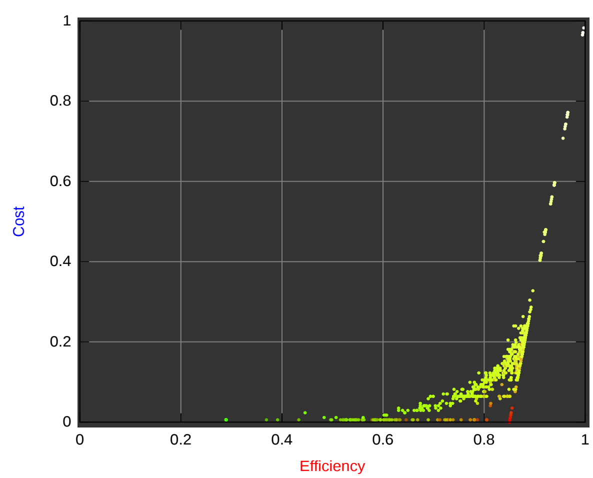

Figure 5.16: *Efficiency (η_l) vs Cost (κ)*. The projection efficiency vs. cost shows that most of the robust topologies have moderate cost.

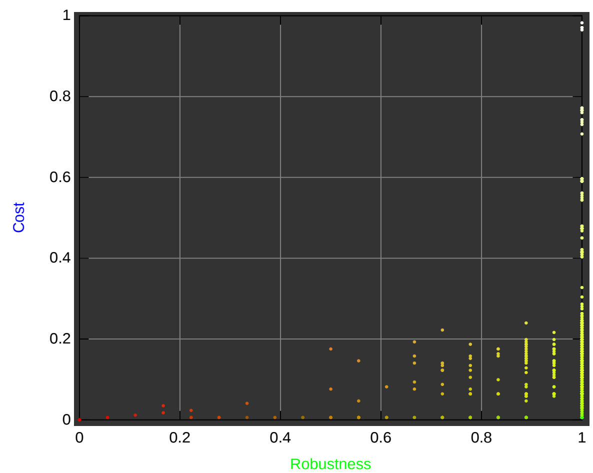

Figure 5.17: *Functional Robustness (ρ_{fn}) vs Cost (κ).* The projection functional robustness vs. cost shows that most of the robust topologies have moderate cost.

cost is not severe. Further, circular skip lists are the most prominent topologies. They achieve low APL and high functional robustness under moderate costs.

Since the metric for efficiency is APL, which is not a discrete quantity (this is unlike diameter), we observe that the plot efficiency vs. cost has a different shape (Figure 5.16). On the other hand, functional robustness with respect to node deletions is a discrete quantity—it is measured as a function of the size of the largest connected component after a node deletion. This quantity changes in discrete steps. This explains the shape of Figure 5.17.

As a part of this work, we have analyzed the behaviour of topologies evolved under a few other metrics. While the shape of the plot varies slightly due to the different types of metrics, the general observations remain the same. We also analyzed directed topologies that evolved through topology breeding to make the same observations (Figure 5.18). Since CSL are more natural in directed net-

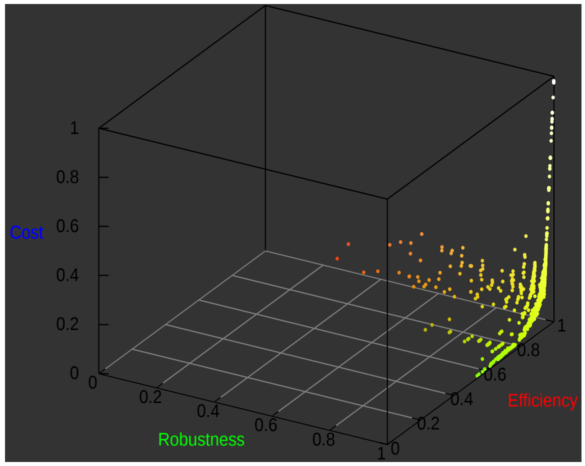

Figure 5.18: *Functional Robustness vs. Efficiency vs. Cost for* η_l *and* ρ_{fn}. The distribution of optimal directed topologies from the perspective of robustness. Here efficiency is computed using APL (l) and robustness is computed using functional robustness w. r. t. deletion of the node with the highest degree.

works, the trade-off between efficiency and robustness is even less pronounced. Most directed topologies that emerges through our breeding process are CSL.

5.4 Optimal Structural Motifs of CSLs

Both directed and undirected circular skip lists show several structural features that are potentially optimal under varied applications. We discuss some of them using an example directed CSL of 32 nodes, shown in figure 5.19.

Hamiltonian Circuits: CSLs contain multiple hamiltonian circuits. Presence of hamiltonian circuits is an optimal feature. Hamiltonians are at least biconnected (both in terms of nodes and edges). They are also considered important in applications such as DHTs [Bermond et al., 1998; Gummadi et al., 2003] as they

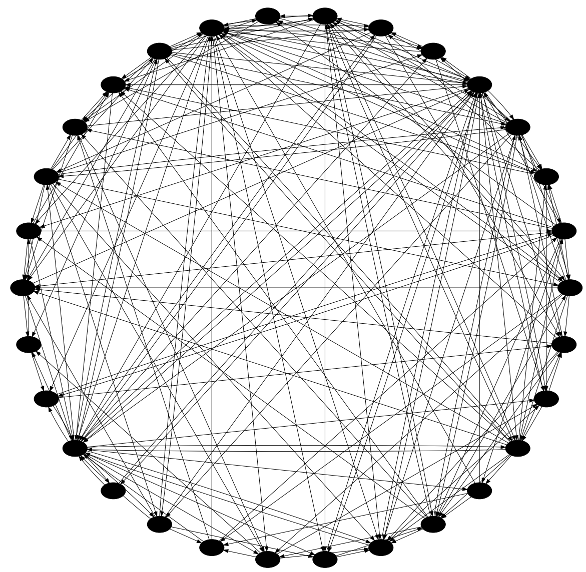

Figure 5.19: A directed circular skip list with $n = 32$, $p = 5$ and $e = 160$. CSLs are highly symmetric and possess several optimal properties.

Table 5.1: Structural properties of an example directed CSL, $n = 32$, $p = 5$, $e = 160$.

Property	Value
Diameter	3
APL	2.53
Eccentricity (min, max)	(3, 3)
Radius	3
Skew in Closeness	0.0016
Entropy of Indegree and Outdegree	0.0 and 0.0
Skew in Node Betweenness	0.005
Skew in Edge Betweenness	0.0025
Edge Connectivity (min, max)	(4, 5)

facilitate design of dynamic routing schemes and load balancing.

Symmetric Centrality Measures: Centralities measure the importance of nodes and edges in a graph. If a small number of nodes/edges have a significantly higher centrality value than the rest, a network is susceptible to congestion. Also, since such nodes/edges are presumably more important than the rest of the network, they are susceptible to targeted attacks.

We measure the skew in node (or edge) betweenness as the difference between the maximum node (or edge) betweenness in the network and the average betweenness. A (normalized) value of 1.0 for the skew indicates that there is only one node (or edge) in the network through which all the traffic flows; whereas a value of 0.0 indicates a uniform load distribution. Table 5.1 shows typical values of betweenness skews for CSLs, which are close to 0. Networks with such homogeneity are not easily congested under heavy traffic. Also, such networks are resilient to both targeted attacks and random failures.

Entropy of Degree Distribution: Entropy of degree distribution is indicative of the heterogeneity of a network. It is measured as, $H = -\sum_p P(p)\ log\ P(p)$, where $P(p)$ is the fraction of nodes with a degree p. Typically, CSLs are regular or nearly regular graphs. Thus, they have an entropy of nearly 0.

Symmetric Distances: CSLs have low diameters and APLs. The distances between node pairs are also homogeneous, leading to a very low skew in closeness centrality (which measures the per node APL). Node eccentricities (which measures the greatest separation a node suffers from another node in the network) are also symmetric. Radius (which is the smallest eccentricity) is typically very close to the diameter. Thus there are very few "peripheral" nodes which incur more communication cost than the rest of the nodes.

Multiple Independent Paths: CSLs are nearly *optimally connected* [Dekker and Colbert, 2004]. Given a CSL with minimum degree, p_{min}, its edge connectivity, λ is at least, $p_{min} - 1$, when $k >= \frac{np}{2}$ (and $k >= np$, for directed graphs).

Recent studies have reported optimal networks with features such as: homogeneous degree [Valente et al., 2004; Zhao et al., 2005], betweenness [Guimerá et al., 2002; Mondragón C, 2006] and optimal connectivity [Dekker and Colbert, 2004]. We observe that CSLs are characterized by all these optimal features.

[...] That is when I understood the magical meaning of the circle. If you go away from a row, you can still come back into it. A row is an open formation. But a circle closes up, and if you go away from it, there is no way back. It is not by chance that the planets move in circles and that a rock coming loose from one of them goes inexorably away, carried off by the centrifugal force. Like a meteorite broken off from a planet, I left the circle and have not yet stopped falling. [...]

Milan Kundera, *The Book of Laughter and Forgetting*

6

Deterministic Techniques for Constructing CSLs

In Chapters 4 and 5, we presented *circular skip lists*, which emerged as optimal topologies in terms of balancing efficiency, robustness and infrastructure cost. We allowed topologies to evolve under various constraints, using genetic algorithm based experiments. Circular skip lists with regular or nearly regular degree centrality, symmetric closeness, high connectivity, low infrastructure cost and low diameter emerge as optimal structures. In this chapter, we discuss simple algorithmic techniques to deterministically construct circular skip lists.

Apart from the work in this thesis, recent findings in many application domains report that "ring"–like or "circular" topologies are ideal for designing networks under requirements of symmetry. Many classes of topologies reported in literature to be optimal are variants of ring–like structures. In other words, they are variants of circular skip lists.

For instance, [Gummadi et al., 2003] suggest that ring topologies are the best in terms of flexibility in DHT routing and resilience. [Qu et al., 2006] propose Cayley graphs [Magnus et al., 1976], which are featured by hamiltonian decompositions and high connectivity as a model for design and analysis of optimal DHTs. They classify existing DHTs as Cayley based and non-Cayley based, and conclude that Cayley based DHTs are better than non-Cayley based ones. [Loguinov et al., 2005] analyze different DHT schemes in terms of lookup cost and resilience to perturbations. DHTs based on variants of ring structures such as Chord [Stoica et al., 2003], hypercube [Schlosser et al., 2003] and Koorde [Kaashoek and Karger, 2003] are shown to do better than other topologies.

Perfect Difference Networks (PDN) [Parhami and Rakov, 2005] are proposed to construct communication networks with asymptotically maximal number of nodes for a given fixed degree and diameter 2. PDNs are symmetric topologies. [Dekker and Colbert, 2004] report that vertex transitive graphs, which are necessarily circular, are robust structures in the face of failures.

[Guimerá et al., 2002] report *homogeneous isotropic networks* with symmetric betweenness as optimal structures to handle traffic load in networks. Other studies concerning networks prone to congestion [Mondragón C, 2006; Zhao et al., 2005] also find symmetric graphs to be the most useful topologies. [Valente et al., 2004] show that topologies with a symmetric degree centrality are the most robust in the face of node and link failures. [Donetti et al., 2005, 2006] introduce *entangled networks*, which have a highly uniform degree and betweenness centrality. They are also characterized by optimal properties in many other respects.

Thus, circular skip lists and their variants emerge as optimal topologies in several independent studies. Our work not only corroborates the above results but also generalizes them. This is because circular skip lists are the most prominent topologies across optimal topology spaces corresponding to a variety of different performance requirements.

6.1 Circular Skip Lists: Definition

We define *circular skip lists* (CSL) to represent a class of network topologies in which (1) the network has one or more logical circles covering all nodes, viz; a hamiltonian cycle, and (2) within such a circle, there exist "chords" or long range connections of various lengths or "skips".

Circular skip lists can be both directed and undirected. The circle topology is the minimal CSL; and the clique is the maximal CSL. In an intermediate CSL, every node connects to a distinct "successor" node so as to form a hamiltonian cycle or a logical circle covering all nodes. In addition, every node connects to zero or more non-successor nodes at different distances or "skips" on the logical circle.

Most of the existing ring-like structures such as Chord [Stoica et al., 2003], de Bruijn graphs [Kaashoek and Karger, 2003] and PDNs [Parhami and Rakov, 2005] can be modelled as circular skip lists. In general, circular skip lists show structural features that are potentially optimal under varied applications: presence of hamiltonian circuits leading to better load balancing; symmetric centrality measures leading to reduced congestion; low diameters at low infrastructure and bookkeeping costs; and, multiple independent paths making a network robust to failures and attacks. In the rest of this chapter we present simple algorithms to construct CSL topologies.

Circular skip lists are similar to the probabilistic skip lists proposed by [Pugh, 1990] and Skip Graphs proposed by [Aspnes and Shah, 2007] in some ways. In Pugh's skip lists, a linked list is built in layers, with the bottom most layer being a normal sorted linked list, and at every subsequent layer, a few "long range" links are created in a randomized fashion. That is, an initial straight line of nodes is augmented with a set of random long range links. In our circular skip lists the equivalent of the bottom most layer is a circle instead of a straight line. This increases the resilience of the system. Also, unlike Pugh's skip lists, all edges are added in a random fashion, including the basic circle.

Skip Graphs based on Pugh's skip lists and address the issue of resilience. In

Skip Graphs, nodes at higher levels can be part of multiple lists instead of a single list. This is similar to a node connecting to other nodes at different skips. This reduces single points of failures. Skip Graphs are modelled as a family of doubly linked lists. In CSL, every node has a predecessor and a successor; however, in a Skip Graph, this condition is not always met. Therefore, there need not exist a circle encompassing all nodes in a Skip Graph.

6.2 Variants of Graph Powers

The r^{th} power of a graph $\mathcal{G}$ is the graph that results after adding edges between nodes that are separated by a path length of up to r in $\mathcal{G}$ while maintaining the same number of nodes. Computing powers of a given graph has important relevance to network design. A graph power reduces the diameter of the original graph by adding long distance edges across nodes. The extra edges that are added increase the resilience of a network against failures. However, the extra edges add to the infrastructure and bookkeeping costs and may affect the symmetry property of a network.

Observation 1. *Given a non-complete graph $\mathcal{G}$ (that is, a graph that is not a clique) of diameter d, the r^{th} power $\mathcal{G}^r$ of $\mathcal{G}$ has diameter $\lceil \frac{d}{r} \rceil$.*

This follows from the fact that every pair of nodes on the diameter of $\mathcal{G}$ which were separated by path length r are separated by pathlength 1 in $\mathcal{G}^r$. Thus the diameter of $\mathcal{G}^r$ is $\lceil \frac{d}{r} \rceil$. Specifically, raising a graph $\mathcal{G}$ to its d^{th} power results in a complete graph whose diameter is 1.

While a graph power operation reduces the diameter of a graph drastically, it adds a lot of new edges. We consider a few variants of the power operation that effect a significant reduction in diameter by introducing fewer new edges.

A variant of $\mathcal{G}^r$ is a graph, $\mathcal{G}+E_r$, that is constructed by adding edges between nodes of $\mathcal{G}$ are that separated by a path length of *exactly* r. This variant is interesting because, the number of new edges added to the original graph is much lesser compared to that in the case of graph powers. Another variant is a graph in which we add edges between only a subset of pairs of nodes ("important" nodes) that are separated by a path length of exactly r in $\mathcal{G}$. We are interested in understanding the effects of such operations on properties such as diameter, connectivity and symmetry.

Specifically, if the initial graph, $\mathcal{G}$, is the minimal CSL, i.e., a cycle over n nodes, we derive a class of CSL construction techniques called *polygon insertion*. The underlying ideas were developed in [Patil et al., 2009]. Below we describe some example polygon insertions and discuss the results. We consider undirected graphs unless otherwise mentioned.

6.2.1 $\mathcal{G}+E_r$–Graphs as Circular Skip Lists

When the initial graph, $\mathcal{G}$, is a cycle over n nodes, the following hold:

Observation 2. *If n is even, for every node in $\mathcal{G}$, (1) there exist exactly two nodes that are at a distance r, where $1 \leq r \leq \frac{n}{2}-1$, and (2) exactly one node at distance $\frac{n}{2}$.*

Observation 3.

If n is odd, for every node in $\mathcal{G}$, there exist exactly two nodes that are at a distance r, where $1 \leq r \leq \left\lfloor \frac{n}{2} \right\rfloor$.

This has interesting implications for $\mathcal{G} + E_r$ graphs: depending on the values of n and r, four types of CSL structures result as we show below. We make observations that are relevant to optimal topology design concerning symmetry, diameter, hamiltonicity and connectivity. The proofs for these observations are beyond the scope of the current work.

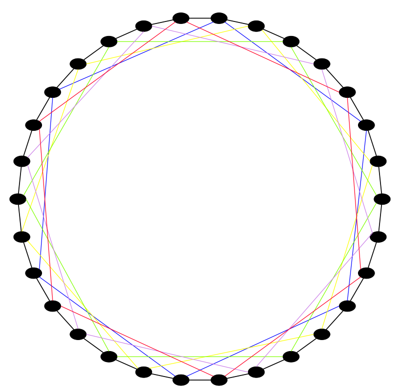

Figure 6.1: $n = 30$, $r = 5$

Case–1: When $1 \leq r < \frac{n}{2}$ **and** $0 \equiv n \; mod \, (r)$ (Figure 6.1)

1. Adding edges between pairs of nodes separated by a distance r results in a CSL in which every node is a vertex in one of the r "regular polygons" of size (number of edges), $\frac{n}{r}$.

2. There are r polygons of size $\frac{n}{r}$ inserted. Hence, the number of edges in the resulting CSL is $2n$.

3. With every node having a degree of 4, the CSL is 4–regular.

4. Being a CSL, it has a connectivity of at least 2.

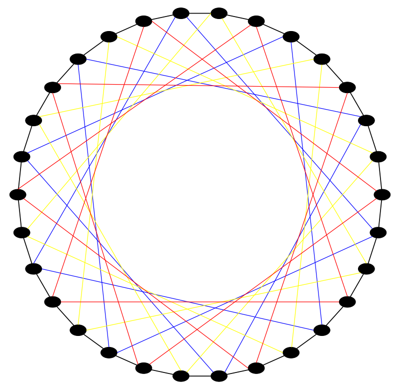

Figure 6.2: $n = 30$, $r = 9$

Case–2: When $1 \leq r \leq \left\lfloor \frac{n}{2} \right\rfloor$ **and** $0 \not\equiv n \; mod\,(r)$ (Figure 6.2)

1. Adding edges between pairs of nodes separated by a distance r results in a CSL that has $gcd\,(n, r)$ number of regular polygons each of size $\frac{n}{gcd(n,r)}$.

2. Hence, the number of edges in the resulting CSL is $2n$.

3. With every node having a degree of 4, the CSL is 4–regular.

4. Being a CSL, it has a connectivity of at least 2.

Case–3: When $1 \leq r \leq \left\lfloor \frac{n}{2} \right\rfloor$**,** $0 \not\equiv n \; mod\,(r)$ **and** n **and** r **are *coprime*** (Figure 6.3)

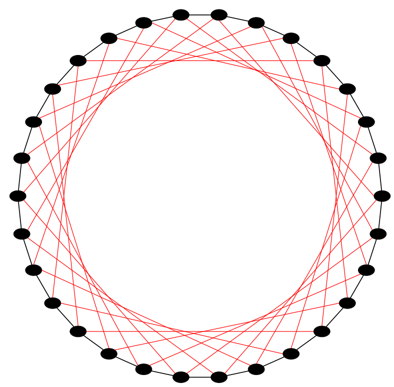

Figure 6.3: $n = 30, r = 7$

1. Adding edges between pairs of nodes separated by a distance r results in a CSL that has a hamiltonian inside the circle. Again, the number of edges in the resulting CSL is $2n$.

2. With every node having a degree of 4, the CSL is 4–regular.

3. This CSL structure has a *hamiltonian decomposition*: the edges in the CSL can be divided such that two edge-disjoint hamiltonian circuits result. Here, the outer circle is one hamiltonian circuit. The inserted polygon is another. They do not share an edge.

4. As a corollary, the CSL is 4–connected.

Case–4: When n is even and $r = \frac{n}{2}$ (Figure 6.4)

1. Adding edges between pairs of nodes that are separated by a distance $\frac{n}{2}$ results in a “cartwheel” like structure.

2. The number of edges in the resulting graph is $\frac{3n}{2}$.

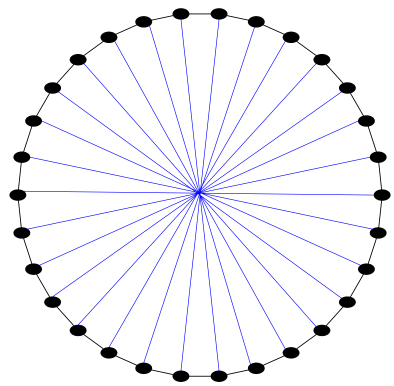

Figure 6.4: $n = 30$, $r = 15$

3. The resulting CSL is 3–regular.

4. Being a CSL, it has a connectivity of at least 2.

Diameter Reduction due to the $\mathcal{G} + E_r$ Operation

Case–1: When $1 \leq r < \frac{n}{2}$ **and** $0 \equiv n \, mod \, (r)$

1. if $r = 2$, $d = \left\lceil \frac{n}{4} \right\rceil$

2. else if $0 \equiv n \, mod \, (2r)$, $d = \left\lfloor \frac{n}{2r} \right\rfloor + \left\lfloor \frac{r-1}{2} \right\rfloor$

3. else, $d = \left\lfloor \frac{n}{2r} \right\rfloor + \left\lfloor \frac{r}{2} \right\rfloor$

Observation 4. *If* $r \cdot s = n$*, then,* $d\left(\mathcal{G} + E_r\right) = d\left(\mathcal{G} + E_s\right)$

The diameter of the initial graph $\mathcal{G}$ is $\left\lfloor \frac{n}{2} \right\rfloor$. The diameter of a $\mathcal{G} + E_r$ graph is a function of r, when $r \neq 2$. Thus, in order to achieve the greatest reduction in diameter from $\left\lfloor \frac{n}{2} \right\rfloor$, we need to construct a $\mathcal{G} + E_r$ such that r is optimal. We do this by finding the minima of d:

$$d = f(r) = \left\lfloor \frac{n}{2r} \right\rfloor + \left\lfloor \frac{r}{2} \right\rfloor$$

We find that the value of r that minimizes the diameter is, $r_{min} = \sqrt{n}$. And the minimum diameter of the CSL upon inserting the set of edges $E_{\sqrt{n}}$ is, $d_{min} = \sqrt{n}$.

The result holds if we consider the expression, $d = \left\lfloor \frac{n}{2r} \right\rfloor + \left\lfloor \frac{r-1}{2} \right\rfloor$, as well.

Case–2: When $1 \leq r \leq \left\lfloor \frac{n}{2} \right\rfloor$ **and** $0 \not\equiv n \, mod \, (r)$: We have not been able to find an exact expression for the diameter in this case. However, $d \leq \left\lfloor \frac{n}{2 \cdot r} + \frac{r}{2} \right\rfloor$. Therefore, $d_{min} = r_{min} \leq \sqrt{n}$.

Case–3: When $1 \leq r \leq \left\lfloor \frac{n}{2} \right\rfloor$**,** $0 \not\equiv n \, mod \, (r)$ **and** n **and** r **are *coprime***: We have not been able to find an exact expression for the diameter in this case as well. The same lower bound as above applies: $d \leq \left\lfloor \frac{n}{2 \cdot r} + \frac{r}{2} \right\rfloor$. Therefore, $d_{min} = r_{min} \leq \sqrt{n}$.

Case–4: When n **is even and** $r = \frac{n}{2}$: In this case, $d = \left\lceil \frac{n}{4} \right\rceil$. This corresponds with Observation 4.

Variable Polygon Insertion

We show two variants of the above CSLs: (1) single polygon insertion, wherein we add only a subset of E_r edges to achieve significant diameter reduction; and (2) we construct a CSL by inserting multiple polygons of different sizes.

The class of CSLs we discussed above, $\mathcal{G} + E_r$, are highly symmetric and efficient in terms of diameter. However, the construction of the CSLs involves addition of n edges ($\frac{n}{2}$ in **Case–4**). It might be possible to achieve significant diameter reduction by adding less number of strategically placed long range edges. Also, it might be necessary in some applications to fortify only a subset of nodes by adding extra edges amongst them.

For example, Figure 6.5 shows a CSL wherein $n = 30$ and $r = 5$. A single polygon of size $\frac{n}{r}$ is added to a circle.

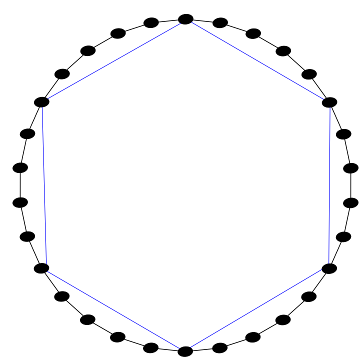

Figure 6.5: A single undirected polygon insertion

It can be shown that:

1. When $0 \equiv n\, mod\, (r)$, $d = \left\lfloor \frac{n}{2r} \right\rfloor - r - 1$
2. When $0 \not\equiv n\, mod\, (r)$, $d = \left\lceil \frac{n}{2r} \right\rceil - r - 1$

We can also get optimal structures by inserting multiple polygons, either of the same size or of different sizes. Below, we describe one such technique called *polygon halving insertion*.

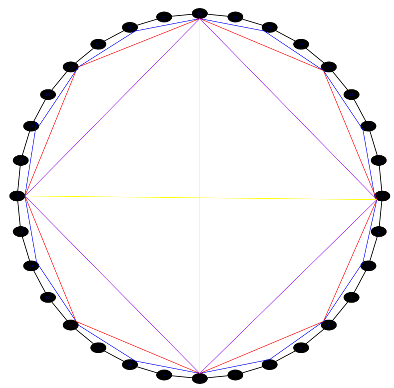

Figure 6.6: Polygon-Halving Insertion

Figure 6.6 shows an insertion, in which multiple polygons are inserted by successively halving their sizes. A node is chosen as the "root" arbitrarily. This is the starting point of the insertion. In the figure, the first inserted polygon has a size of $n/2$. The next polygon, which has $n/4$ sides, is inserted starting from the root node. Successive polygons are inserted by halving until we reach the smallest polygon, a triangle or two diameter lines as shown in the figure.

A polygon-halving insertion has a similar routing complexity as a tree rooted at the "root" node that we have selected. If the outermost inserted polygon has l sides, then the diameter of the topology can be computed from the relation,

$$d = \frac{n}{l} + \log l$$

When $l = \frac{n}{2}$, as in the above example,

$$d = \log n + 1$$

In the polygon insertion approach, as can be seen in the above figures, some of the nodes have a higher degree than the others. They can be likened to "super peers". Polygon insertion strategies, although efficient have a low load distribution value. In environments like grids, where all nodes are not necessarily symmetric, these serve to be useful topologies. Moreover, this is not a major shortcoming as a symmetric topology can be easily achieved in a polygon insertion by replicating the structure of Figure 6.6 at every node. Such a topology has a diameter of $\log n$ and is symmetric, load balanced and robust. However, the better load distribution and robustness come at much higher infrastructure and book-keeping costs. As a matter of fact, the topology thus obtained is equivalent to the topology of the well known Chord [Stoica et al., 2003] DHT.

Polygon halving is also closely related to a class of networks called base–b networks that we discuss in the next section.

6.3 (Balanced) Base–b Networks

We present a class of circular skip lists whose construction technique is derived from representations of numbers in different bases. An exemplar of this kind of an approach is the Chord skip list, which is used for lookup in p2p networks. Chord forms a circular skip list over $n = 2^k$ nodes (i.e., node identifiers in the range $\left[0, 2^{k-1}\right]$), where $k \in \mathbb{N}$. Each node has $log\ n$ edges formed using the following rule: connect to a node whose identifier is 2^i skips ahead, where $0 \leq i < k$. Thus, every node has a distinct "successor" node (due to the skip–2^0) thus forming a logical directed circle; and a set of directed chords to other nodes along the logical circle. It can be shown that the Chord skip list has a diameter of $log\ n$. A routing algorithm can be devised on the skip list in the manner of binary search whose complexity is $O\left(log\ n\right)$.

The approach that we present in this section can be considered as a generalization of Chord. Chord uses binary or base–2 representation for its skip list. We propose a technique to construct skip lists for any base–b, $b \geq 2$. Chord forms a directed graph topology. Our technique can be used to construct both directed and undirected (bidirected) skip list topologies.

We first define the terminology used in this section and then proceed to discuss the construction technique.

Base–b Numeral System: In a base–b numeral system, numbers are represented using an alphabet, l, of length b. The symbols in the alphabet are non-negative: $l = \{0, \ldots, b-1\}$. For example, the alphabet of base–6 numeral system is, $l = \{0, 1, 2, 3, 4, 5\}$.

The number of decimal numbers that can be uniquely represented using a base–b–string of length k is, b^k, viz; $\left[0, b^k - 1\right]$.

Balanced Base–b Numeral System: Unlike a base–b numeral system, in a balanced base–b numeral system, an alphabet can have negative symbols. There are several ways of forming a balanced base–b alphabet of length b. However, we choose an alphabet, l, such that the number of negative symbols is roughly equal to the number of positive symbols ("balanced").

If b is odd: $l = \left\{-\left\lfloor \frac{b}{2} \right\rfloor, \ldots, \left\lfloor \frac{b}{2} \right\rfloor\right\}$. For example, the alphabet of balanced base–5 numeral system is, $l = \{-2, -1, 0, 1, 2\}$.

If b is even, the alphabet is either, $l = \left\{-\frac{b}{2} - 1, \ldots, \frac{b}{2}\right\}$ or $l = \left\{-\frac{b}{2}, \ldots, \frac{b}{2} - 1\right\}$. The alphabet of balanced base–6 numeral system is, $l = \{-2, -1, 0, 1, 2, 3\}$ or $l = \{-3, -2, -1, 0, 1, 2\}$.

The number of decimal numbers that can be uniquely represented using a balanced base–b string of length k is, b^k. If b is odd, the range of numbers that can be represented is: $\left[-\frac{b^k-1}{2}, \frac{b^k-1}{2}\right]$. If b is even, the range of numbers that can be represented is either, $\left[-(b^k - \frac{b^{k+1}-b}{2(b-1)}), \frac{b^{k+1}-b}{2(b-1)}\right]$ or $\left[-\frac{b^{k+1}-b}{2(b-1)}, b^k - \frac{b^{k+1}-b}{2(b-1)}\right]$.

The above discussion relates to the construction of CSLs, which we call base–b and balanced base–b networks. A base–b network corresponds to a directed graph topology, whereas a balanced base–b network corresponds to an undirected (or bidirected) graph topology.

We illustrate balanced base–b networks by constructing a *balanced ternary (base–3) network*.

Balanced Ternary Network: Consider base $b = 3$, with alphabet, $l_3^b = \{-1, 0, 1\}$. If $k = 3$, the number of decimal numbers that can be uniquely represented using l is, $b^k = 27$, viz; $[-13, 13]$. Further, the first k positions in the ternary system form the set, $p = 3^i$, where $0 \leq i < k$. Thus, $p = \{3^0, 3^1, 3^2\} = \{1, 3, 9\}$.

Consider an undirected empty graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, featured with the following:

1. The number of nodes in $\mathcal{G}$, $n = |\mathcal{V}| = b^k = 27$. Each node, $v \in \mathcal{V}$, takes a unique identifier, v_i, in the range $[-13, 13]$. For the sake of convenience, the 27 identifiers can also be represented alternatively using the range $[0, 26]$.

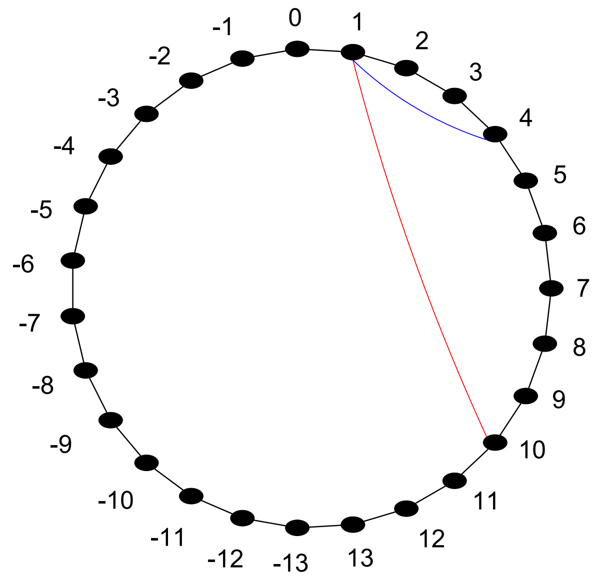

Figure 6.7: Constructing a balanced ternary network: node 1 has edges to nodes at skips 1, 3 and 9.

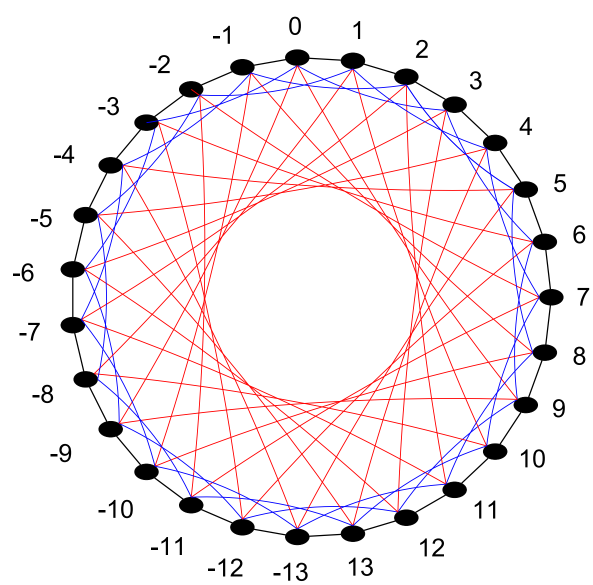

Figure 6.8: A balanced ternary network with $n = 27$

Here, the negative numbers $[-13, -1]$ are mapped to $[14, 26]$ through congruence relations over 27. For instance, $14 \equiv -13 \; mod \; (27)$.

2. The first k positions in the ternary system, $p = \{1, 3, 9\}$, form the *skip list*. Each node, $v \in \mathcal{V}$, has an edge to nodes at skips 1, 3 and 9. For example, node 3 has an edge to nodes 4, 6 and 12. Node -13 has edges to nodes -12, -10 and -4. Node 6 has edges to nodes 7, 9 and -12 (because $15 \equiv -12 \, mod \, (27)$).

The edges resulting from connections to nodes at a skip of 1 result in a circle covering all the nodes $[-13, 13]$. The edges due to skips 3 and 9 form chords within the circle. Please refer to Figures 6.7 and 6.8.

The balanced ternary alphabet, $l = \{-1, 0, 1\}$, implies that at any node the following operations are allowed: (1) one skip can be taken backwards (anti-clockwise) on the above logical circle, (2) no skip is taken, and (3) one skip can be taken forwards (clockwise) on the logical circle. This implies that all pairwise distances and routes can be represented in terms of the sums and differences of the skips.

In other words, the balanced ternary representation of the distance between two identifiers on the logical circle gives the routing scheme. For example, the distance between 4 and -4 on the circle is -8. The number -8 is represented in balanced ternary as, (-1)01. In terms of routing over the balanced ternary network this is equivalent to taking the first skip (9 steps) backward from 4 to -5; then take the last skip (1 step) forward to reach -4.

The diameter is the highest distance on the wrapped number line between a pair of numbers, which is 13 (or -13). The number 13 in balanced ternary representation is 111. This implies a skip of 9, 3 and 1 each. Similarly, -13 is represented as (-1)(-1)(-1), implying 3 backward skips. Thus, the diameter of the above graph is $|p| = k = log_3\, 27 = 3$.

In general, the following results can be proved for balanced base–b networks:

1. Given a base–b, and a k, $k \in \mathbb{Z}^+$, the number of nodes that can be packed in a balanced base–b network is, $n = b^k$.

2. The diameter of such a network with an alphabet l is, $d = \max\{|i| \in l\} \cdot k$, where $\max\{|i| \in l\} \approx \frac{b}{2}$. That is to say, $d \approx \frac{b}{2}\, log_b\, n$.

3. A balanced base–b network is $2k$–regular.

Ternary Network: The ternary alphabet is, $l_3 = \{0, 1, 2\}$. If $k = 4$, the number of decimal numbers that can be uniquely represented using l is $b^k = 81$ viz; $[0, 80]$. Further, the first k positions in the ternary system form the set, $p = 3^i$, where $0 \leq i < k$. Thus, $\{3^0, 3^1, 3^2, 3^3\} = \{1, 3, 9, 27\}$.

Let $\mathcal{G}\,(\mathcal{V}, \mathcal{E})$ be a directed graph featured with the following properties:

1. The number of nodes in $\mathcal{G}$, $n = |\mathcal{V}| = b^k = 81$. Each node $v \in \mathcal{V}$ takes a unique identifier v_i, in the range $[0, 80]$.

2. The first k positions in the ternary system, $p = \{1, 3, 9, 27\}$, form the *skip list*. Each node, $v \in \mathcal{V}$ has an edge to nodes at skips 1, 3, 9 and 27 $(mod\,(81))$. For example, node 3 has an edge to nodes 4, 6, 12 and 30.

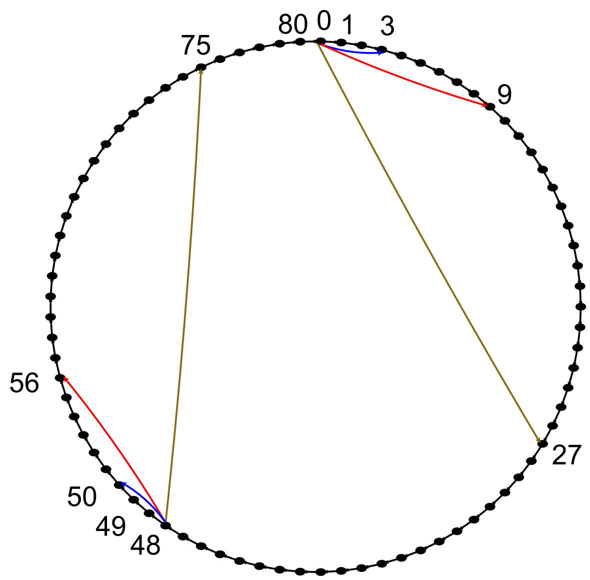

Figure 6.9: Constructing a ternary network: nodes 0 and 48 have connected to edges at skips of 1, 3, 9 and 27

The ternary alphabet $l = \{0, 1, 2\}$ implies that, from a node 0, 1 or 2 skips can be taken forwards (clockwise). Since the pairwise distance between the numbers on the directed circle above is in the range $[0, 80]$, all distances is the above CSL can be represented in ternary. In other words all pairwise distances and routes can be represented in terms of the sums of the skips. Thus, the ternary representation of the distance between a pair of nodes on the directed circle gives the routing scheme for that pair.

The diameter of the above graph is the highest number that can represented using the ternary skip list, which is 80, represented as 2222. This implies, take 2 skips each of 27, 9, 3 and 1 steps. Thus, the diameter of this network is 8.

In general, the following results can be proved for a base–b network:

1. Given a base–b, and a k, $k \in \mathbb{Z}^{+}$, the number of nodes that can be packed in a network is, $n = b^{k}$.

2. The diameter of such a network is, $d = \max\{|i| \in l\} \cdot k$, where $\max\{|i| \in l\} = b - 1$. That is to say, $d = (b-1)\ log_b\, n$.

3. A base–b network is k–regular.

Faultily faultless,

icily regular

splendidly null,

dead perfection

no more.

Lord Alfred Tennyson, *Maud; A Monodrama, 1855. Part I, section ii.*

7

Regular Graphs and Optimal Network Design

Regular graphs have elicited a lot of interest in both theory and applications. As we discussed in Section 2.1, all graphs pertinent to the degree-diameter problem are regular. Large random regular graphs have been observed to have several optimal features such as high connectivity, low diameter and hamiltonicity [Robinson and Wormald, 1994; Wormald, 1999].

The *symmetry* requirement is very common in many network design problems. In Section 1.1, we discussed several applications where symmetry—in one form or the other—is an important design consideration.

In case of distributed indices, symmetry translates to a topology where each node is expected to take the same amount of bookkeeping cost in managing the distributed index. Bookkeeping cost arises from keeping track of other nodes and their addresses in DHT *finger tables*. Symmetry is addressed by modelling the index in the form of a graph in which every node has the same degree (number of edges incident on a node), viz; a *regular graph*.

In case of NCW, symmetry translates to building networks that are ideally *vertex transitive* (and/or *edge transitive*). Again, all vertex and edge transitive graphs are necessarily regular graphs. In case of CDNs, symmetry translates to designing networks such that the nodes and/or edges have uniform (or nearly uniform) communication load, which is to say nodes and/or edges have similar values of betweenness centrality.

In sections 2.2.2 and 2.2.3, we discussed topologies useful in DHTs, complex networks and communication networks, such as, rings, hypercubes, homogeneous isotropic networks and entangled networks, all of which are symmetric. Recently, there are several efforts in search of universal optimal structures for topology design. Structural and spectral analyses of regular graph families such as Cayley graphs and Ramanujan graphs show them to have high efficiency, high resilience to failures and low susceptibility to congestion. Regular graphs emerge as the fundamental structures for designing optimal networks.

Thus, the underlying universal structure of a symmetric topology is the regular graph. The family of regular graph topologies forms an interesting subset of optimal topologies having its own unique properties that deserves to be studied separately. Further, almost all regular graphs are hamiltonian as shown by [Robinson and Wormald, 1994]. This implies almost all regular graphs are circular skip lists. In the previous chapters we observed that CSL are an important class of topologies useful across many application domains.

In this regard, we study regular graphs and bring out several properties and their relationships to the optimality in networks. The objective of this chapter is to explore theoretical underpinnings of regular graphs that are pertinent to network design. The important results from this work can be found in [Patil and Srinivasa, 2010].

7.1 Connected Regular Graphs: Basics

A graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is called r–regular if $\forall v \in \mathcal{V}, degree(v) = r$. The term r is called the *regularity* of the graph $\mathcal{G}$. We first examine properties of *undirected* regular graphs relevant to optimal network design. We will also briefly address regular directed graphs in Section 7.4.

In the following sections, whenever we refer to regular graphs, we shall be referring to **undirected regular connected graphs**, unless specified otherwise.

Further, we shall not be considering multiedges, as well as self loops.

We begin with the following observations about regular connected graphs.

Postulate 1. *If a connected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where the number of nodes, $|\mathcal{V}| = n$, $n > 2$, is r–regular, then $r \geq 2$.*

Postulate 2. *It is always possible to build a 2–regular connected graph for any given set of nodes $n > 2$. Such a 2–regular graph is in the form of a single hamiltonian circuit encompassing all n nodes.*

Postulate 3. *The smallest number of nodes required to build an r–regular connected graph is $(r + 1)$, which is the clique with $(r + 1)$ nodes.*

Postulate 4. *It is not possible to build an r–regular graph with n nodes, where both r and n are odd numbers, since the total degree of a graph is always even.*

During the course of our work, we encountered some conjectures like the following: (1) every r–regular graph is also $(r - 1)$–regular. That is, it is possible to obtain an $(r - 1)$–regular graph from an r–regular graph by removing one or more edges; (2) since a connected 2–regular graph is a hamiltonian circuit, every connected r–regular graph with $r > 2$ has a hamiltonian circuit.

If the first conjecture were to be true, it has applications in handling failures in distributed systems. If one or more connections fail, and it is not possible to construct an r–regular graph with the set of nodes, we should always be able to build an $r - 1$ regular graph and retain the symmetric degree centrality property.

If the second conjecture were to be true, we can guarantee that an r–regular graph with n nodes with $r > 2$ has a diameter less than or equal to $\frac{n}{2}$. Also, a

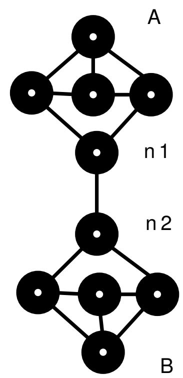

Figure 7.1: A 3–regular graph that has no hamiltonian circuit nor does it have 3 edge independent paths between nodes.

hamiltonian undirected graph is at least 2–connected. That is, there are at least two edge-independent (as well as, node-independent) paths between any pair of nodes in the graph.

However, we can refute both these conjectures using the above postulates. Further, Figure 7.1 clearly refutes the second conjecture.

7.2 Extending Regular Graphs

In this section, we describe a simple algorithm to construct regular graphs of arbitrary number of nodes and arbitrary regularity by a procedure called *extension*. We also address the question of how to extend a given regular graph to accommodate arrival of new nodes without affecting the regularity.

Theorem 1. *If r is even (i.e. $r = 2q$, where $q \in \mathbb{N}$), it is always possible to build an r–regular graph over n nodes, where $n \geq r + 1$.*

Proof. Let $\mathcal{G}_k = (\mathcal{V}, \mathcal{E})$ be a $2q$–regular graph over k nodes where $k \geq 2q + 1$. In order to add a new node v' into $\mathcal{V}$, first find a *matching*[1] m of q edges. For every edge, $e_{ij} \in m$, do the following: (1) disconnect it from one of its end nodes, say j, and connect the free end to v', thus effecting the edge $e_{iv'}$ (2) add an edge between v' and j, $e_{v'j}$. Note that at the end of step (2), the degrees of every node pair (i, j) remain $2q$, whereas the degree of v' increases by 2. Thus, after q such operations, the degree of v' is $2q$, which is the same as all other nodes in the graph. Thus, we get $\mathcal{G}_{k+1}$, a $2q$–regular graph with $k + 1$ nodes.

The crucial step in the above construction is to find a matching of size q. The existence of a matching is a sufficient condition for extension. The existence of a matching of size q is easily proved based on the following known results: (1) a graph, $\mathcal{G}$, with minimum degree, $\delta \geq 2$, has a cycle, C, of length, $|C| \geq \delta + 1$, and (2) a graph $\mathcal{G}$ with $|C| \geq l,\ l \in \mathbb{N}$, has a matching of length at least $\left\lfloor \frac{l}{2} \right\rfloor$.

From the above results, Since a $2q$–regular graph has a matching of length at least, q, we can extend it indefinitely by 1 node. Hence the theorem. $\square$

Figure 7.2 illustrates Theorem 1 being applied on a 4–regular graph with 5 nodes. It is extended by one node by rearranging two edges and adding two more edges. Another way of looking at the above algorithm is this: for each edge in matching m, "break" the edge in the middle, and connect the new node between the dangling ends of the edge. At the end of this process, we get a regular graph with a new node.

The extension algorithm essentially needs to find a cycle C, whose length is

[1] A matching on a graph $\mathcal{G}$ is a subset, m, of $\mathcal{E}$ such that no two edges in m share a vertex.

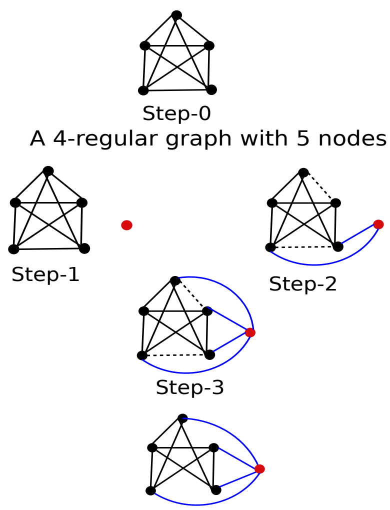

Figure 7.2: Extending a 5-node 4-regular graph by one node.

$|C| \geq 2q$ (such a cycle has two matchings each of size at least q). This can be done by a modified version of the standard depth first search (DFS) algorithm. Thus the complexity of the above algorithm is, $O(|V| + |E|)$.

Extension of regular graphs has implications in network design. For example, in DHT design it is quite common to find questions of being able to extend a DHT in the face of node arrivals. We can see that it is possible to extend DHT topologies while keeping it regular.

For odd regularities, the technique used in Theorem 1 does not work. We need to extend an odd regular graph by two nodes at a time. For this, we can extend Theorem 1 for odd regularities as shown below in Theorem 2. Figure 7.3 illustrates the extension procedure for odd regular graphs.

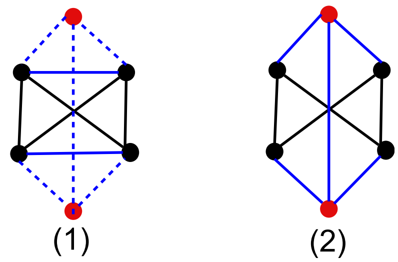

Figure 7.3: Extending a 4-node 3-regular graph by two nodes.

Theorem 2. *Given an r–regular graph with $n \geq r+1$ nodes where r is odd (i.e. $r = 2q+1$, where $q \in \mathbb{N}$) and n is even, there exists an r–regular graph with $n+2$ nodes.*

Proof. Given a $(2q+1)$–regular graph $\mathcal{G}_k = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{V}| = k$, where k is even, we show how to extend it by adding two nodes v' and v'' respectively to $\mathcal{V}$.

Find two size q matchings m' and m'' in the graph, such that $m' \cap m'' = \emptyset$. Use m' to connect to v' and m'' to v'', as described previously. Now, both v' and v'' have a degree of $2q$. Finally, add an edge between v' and v''. Thus, we have a $(2q+1)$–regular graph with $k+2$ nodes.

Let us note that the union of the two matchings, $m' \cup m''$, need not be a matching. Further, we can prove that it is always possible to find matching m' and m'', using arguments similar to the ones in Theorem 1. □

Note that Theorem 2 represents the *minimal extension* we can make to a graph with odd regularity. If an r–regular graph with n nodes has to exist where r is odd, then n should be even and can only be extended in steps of 2. Again, the above algorithm needs to find a cycle C of length, $|C| \geq 2q$. Thus, the complexity is, $O(|V| + |E|)$.

Extending regular graphs by "merging": While the minimal extension theorems show theoretical underpinnings for extending regular graphs, in practice it is often desirable to build regular graphs much faster than extending them by the minimal extension possible. For example, if there are a large number of nodes over which a DHT has to be constructed, it is often desirable to start the DHT construction process parallely involving several subsets of nodes. These subsets can then merge into one another to finally result in a single DHT over the entire network. An example of this is GHT [Ratnasamy et al., 2003], a geographic hash table implemented over wireless sensor networks, where locality plays a major role in the DHT performance. It is easier for local nodes to form networks amongst themselves and then merge into a global network. Similarly, merging is also necessary when a network has to repair itself from a partitioning. Here, we address merging of regular graphs without losing the regularity property.

Theorem 3. *Given any r–regular graph with n nodes and $r \geq 2$, there exists an r–regular graph of $n + r - 1$ nodes.*

Proof. We prove this by merging an $(r-2)$–regular clique with an r–regular graph. Note that an $(r-2)$–regular clique has $r-1$ nodes and is fully connected. So the combined graph would have $n + r - 1$ nodes.

Let $\mathcal{G}^r$ be the given graph and $\mathcal{C}^{r-2}$ be the $(r-2)$–regular clique to be merged with $\mathcal{G}^r$. In order to perform the merging, take each node v_i of $\mathcal{C}^{r-2}$ in sequence.

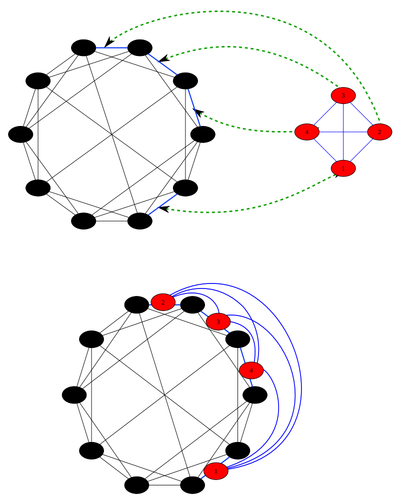

Figure 7.4: Merging a K_3 into a 5–regular 10 node graph

For each such v_i pick an arbitrary edge from $\mathcal{G}^r$ and *insert* v_i in the "middle" of the edge. We now note that v_i, which already had degree $r-2$, now has degree r after inserting it in the middle of an existing edge. Once all nodes from $\mathcal{C}^{r-2}$ are inserted, we get one combined graph with regularity r.

However, we need to address a caveat here. Suppose there is an edge $(u, v) \in \mathcal{G}^r$. This edge can take insertions from nodes $v_i, v_j \in \mathcal{C}^{r-2}$ iff (v_i, v_j) is not an edge in $\mathcal{C}^{r-2}$. Failing which, we would have multiple edges between v_i and v_j in the combined graph.

In order to ensure that the formation of multi-graphs can always be prevented, $\mathcal{G}^r$ needs to have at least as many edges as the number of nodes in $\mathcal{C}^{r-2}$. Since $\mathcal{C}^{r-2}$ is an $(r-2)$–regular clique, it has $r-1$ nodes. If the number of nodes in $\mathcal{G}^r$ is n, then it has $\frac{nr}{2}$ edges. Since n has a lower bound of $(r+1)$ (Postulate 3), the inequality $\frac{nr}{2} > r - 1$ is satisfied trivially. □

We can extend Theorem 3 by allowing the merging of any $(r-2)$–regular graph (not necessarily a clique) with an r–regular graph as long as the number of nodes in the former is no more than the number of edges in the latter. It is possible to merge graphs with a larger number of nodes than that specified above, but the maximum node constraint is a safe upper bound below which, merging is always possible.

7.3 Regularity, Connectivity and Hamiltonicity

The *edge-connectivity* of a graph is the smallest number of edges (or the edge cutset) whose removal would partition the graph. Edge connectivity has direct correspondence with the resilience of a network in the face of connection failures. In the following sections, when we refer to the connectivity of a graph, we are referring to edge-connectivity only. We shall not be addressing the counterpart of node-connectivity.

Menger's theorem states that the connectivity of a graph is equal to the maximum number of pairwise independent paths between any two nodes in the graph. Also, as we had noted earlier, the presence of a hamiltonian circuit guarantees connectivity of at least 2. In this section, we discuss some observations relating regularity, connectivity and hamiltonicity of graphs.

7.3.1 Regularity and Connectivity

We observed earlier that a regularity of r need not mean r–connectivity (Figure 7.1). We can find graphs that are r–regular but k–connected where $k < r$.

Postulate 5. *The maximal connectivity of an r–regular graph is r. This follows from the following two facts: (1) the maximal connectivity of a graph is equal to the minimal degree in the graph [Harary, 1962], and (2) in a regular graph all nodes have the same degree equal to the regularity; hence the minimal degree is the regularity.*

Postulate 6. *A connected 2–regular graph is exactly 2–connected. This follows from the fact that all connected 2–regular graphs are circles, which are 2–connected.*

Theorem 4. *Given an r–regular, k–connected graph ($k \leq r$) over n nodes with the following constraints: r is even (i.e. $r = 2q$, where $q \in \mathbb{N}$) and $n \geq 2r$ (i.e. $n \geq 4q$); it is always possible to extend this graph by any number of nodes without altering the regularity r or edge connectivity k.*

Proof. A k–connected, $2q$–regular graph $\mathcal{G}^{k,2q}$ is one that is $2q$–regular and there exists at least one set of k edges, whose removal partitions the graph into two disconnected components.

Let the graph be represented as two components C_1 and C_2 connected with k bridges across them. From the extension theorem (Theorem 1), we can extend either C_1 or C_2 arbitrarily without adding any edges across C_1 and C_2. This keeps the graph $2q$–regular and k–connected after extension.

However, for the extension theorem to work for any component C_i, we see that C_i should have at least $2q$ nodes. This brings us to the constraint that the number of nodes $n \geq 4q$. Given this, at least one of C_1 or C_2 is guaranteed to have at least $2q$ nodes and can be arbitrarily extended. □

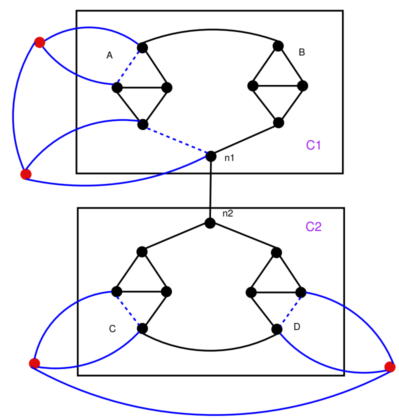

Figure 7.5: Extension without changing the connectivity

Since the smallest number of nodes, n, to create a $2q$–regular graph is only $2q+1$, it is interesting to explore whether we can extend k–connected, $2q$–regular graphs having $n < 4q$ nodes. As of now, we do not have any proof or refutation for this conjecture.

The significance of the above theorem for network design is as follows. Often, it is cheaper to build r–regular graphs with a connectivity less than r. For example, in NCW, while it is ideal to have the maximum possible connectivity r, it might not be cost effective; hence, building/extending a network with connectivity $k < r$ might be an useful alternative. This is also useful when a distributed system is modelled as a metric space and there is a distance function or a cost associated with each edge that is proportional to the (physical or logical) distance between the incident nodes. In such situations, it is easier to maintain r–regularity using short-range connections, rather than connecting nodes in far away corners of the network. The connectivity, k, which even though is less than r, may be the minimal connectivity guarantee that the design offers for such a network.

The extension theorem for k–connected regular graphs is significant in this respect, where connectivity can be maintained (or at least is not compromised) when extending the graph. The extension theorem for k–connected graphs can be extended to odd regularity in an analogous fashion as follows.

Theorem 5. *Given an r–regular, k–connected graph, where r is odd, $n \geq 2r$ and n is even, there exists an r–regular, k–connected graph with $n + 2$ nodes.*

Proof. The proof is similar to the one above, except that in this case we consider the extension theorem for odd regularity (Theorem 2). □

Both the above extension theorems require the number of nodes, n, to be at least twice the regularity, r. This value is much higher than the minimum required value of $r + 1$ for forming an r–regular graph. However, this is not a limiting constraint since in most real world networks $n \gg r$.

We see that the smallest connected r–regular graph, which is a clique of $r + 1$ nodes is exactly r–connected. If the connectivity has to decrease, it means more nodes need to be added to this graph increasing the value of n.

This brings us to the conjecture that given a value of r there is a maximum value of n below which a regular graph of r–regularity has to be at least k–connected. It would be interesting to find this threshold. As of now we do not have any proof or refutation for this conjecture. [Imase et al., 1985] have developed some interesting results in this connection by studying regular *directed graphs*. They derive relations linking the number of nodes, n, regularity, r, diameter, d, and edge connectivity, λ. They show, for example, that the edge connectivity, λ is equal to r in an r–regular graph if $n \geq r^{d-1} + d^2 - 2$. They also develop similar results for connectivity less than r given a diameter d.

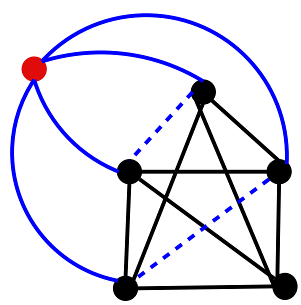

Figure 7.6: Extension while maintaining connectivity

A side effect of Theorems 4 and 5 is also a guarantee that we can always build and extend r–regular graphs that are r–connected, whenever r is even. This is explained by the following theorem.

Theorem 6. *Starting from an r–regular, r–connected graph, it is always possible to build an r–regular, r–connected graphs over any n, when r is even (i.e., $r = 2q$, where $q \in \mathbb{N}$).*

Proof. Note that this theorem does not trivially entail from Theorem 4 for k–

connected graphs due to the constraint there on the number of nodes, $n \geq 4q$. This is because, the smallest $2q$–regular, $2q$–connected graph is a clique with $2q+1$ nodes. Therefore, here, $n \geq 2q+1$.

We prove the theorem using the extension theorem for even regular graphs as follows. Consider the way extension is done. When a new node needs to be added to the graph, first, q edges in the existing graph are "broken". To guarantee extensibility, this is done after ensuring that a matching of q edges is available.

Therefore, each of the q pairs of nodes now have one path less between them. As a result, the connectivity of the graph reduces to $2q-1$ (or $r-1$). Next, the $2q$ dangling ends of the broken edges are connected to the new node to obtain the extended graph of regularity $2q$. Now, each of the q pairs of nodes have a new path between them through the newly added node, thus restoring their pairwise $2q$ (or r) edge independence. Also, the newly added node is connected to $2q$ distinct nodes in the graph, each of which has $2q$ edge independent paths to all nodes. Thus, the connectivity of the graph remains $2q$ (or r). Hence the result follows. □

Note that in this case of r–connectivity, there is no constraint on the number of nodes n to be at least $2r$. We can extend this result to odd regular graphs analogously.

Theorem 7. *Starting from an r–regular, r–connected graph, it is always possible to build r–regular, r–connected graphs over any $n+2$ using the extension theorem.*

Proof. The proof is similar to the above proof, except that we consider the extension theorem with odd regularity (Theorem 2). □

Hence, when connectivity is r, extension without compromising on the connectivity is always possible. However, for regular graphs where the connectivity is less than r, there is a constraint on the minimum number of nodes in order to guarantee extensibility.

Connectivity of regular graphs also has important applications in consensus problems in distributed systems, where nodes are allowed to fail in a Byzantine fashion (failure due to arbitrariness in behavior) [Lamport et al., 1982].

7.3.2 Regularity and Hamiltonicity

We have similar observations with respect to regularity and hamiltonicity; and the absence of it.

In an undirected graph of n nodes, the presence of a hamiltonian circuit bounds the diameter of the graph to $\lceil \frac{n}{2} \rceil$. It also guarantees 2–connectivity. For directed graphs, a hamiltonian circuit is the smallest regular graph (with one directed edge per node) that keeps the graph strongly connected. Hamiltonicity hence plays an important role in network design.

[Bermond et al., 1998] study hamiltonian decompositions in wrapped butterfly networks (WBNs). WBNs comprise of several edge-disjoint hamiltonian decompositions. Hamiltonian networks are advantageous for algorithms that make use of the ring structure (many DHT algorithms do this). A hamiltonian decomposition allows the load to be equally distributed, also making the network robust.

Similarly, topologies based on hypercubes of dimension $2d$ can be decomposed into d hamiltonians.

[Gould, 1991] provides a detailed survey of the hamiltonian cycle problem and reports several sufficiency conditions for hamiltonicity. Given a set of n nodes and a regularity number r, it is quite natural to expect the question whether it is possible to construct a hamiltonian r–regular graph over these n nodes. To answer this, we start with the following observation.

Postulate 7. *The smallest r–regular graph, i.e. a clique of $n = r + 1$ nodes is hamiltonian.*

Theorem 8. *Starting from an r–regular hamiltonian graph, it is always possible to build an r–regular graph over n nodes, while preserving hamiltonicity, when r is even and $n \geq r + 1$.*

Proof. Let $\mathcal{G}$ be an r–regular hamiltonian, and let $h = v_1, v_2, \ldots v_n, v_1$ be a hamiltonian cycle in $\mathcal{G}$.

Now extend $\mathcal{G}$ using the extension theorem for even graphs (Theorem 1) as follows. While extending $\mathcal{G}$, using a matching m, let at least one edge (v_i, v_j), which is in h, be in m. After extension, since v' will be incident on two edges in h, it will now be a part of the extended hamiltonian cycle, $h + v'$. We have thus successfully inserted an extra node into the hamiltonian cycle and retained the hamiltonicity property in the new graph. □

We can analogously prove the corresponding result for odd regular graphs using the extension theorem for odd regularity (Theorem 2).

Theorem 9. *It is always possible to build a hamiltonian r-regular graph over n nodes, where r is odd, n is even and $n \geq r + 1$.*

While hamiltonicity is a desirable property in applications such as DHT and NCW, there are applications where *not* having a hamiltonian circuit is desirable. An example is of hybrid networks that have a mix of fixed and mobile components. For such networks it is desirable to not have a single overarching data structure over the entire network. It may be desirable to model a DHT for hybrid networks as a set of distinct components connected by bridges, with an appropriate key distribution mechanism. Changes to a DHT structure in the face of churn is usually expensive. In hybrid networks, it is desirable to isolate the mobile part of the network from the fixed part. All these may require to be performed in a way that does not affect the symmetry property of the DHT. In such cases the absence of hamiltonicity becomes an important property.

For addressing such issues, we begin with the following observation.

Postulate 8. *It is not possible to have a connected 2–regular graph that is non-hamiltonian. In a 2–regular graph, all nodes have a degree 2. Unless there are an infinite number of nodes, it is not possible to have anything other than a logical cycle encompassing all the nodes, to build a connected 2–regular graph.*

Corollary 1. *Given an r–regular, non-hamiltonian graph with n nodes, where r is even and $n \geq 2r$, it is always possible to extend this graph over any number of nodes and retain the regularity and non-hamiltonian properties.*

Proof. The proof for this can be seen as a corollary to Theorem 4. In a non-hamiltonian graph, there exists at least one edge, e', that needs to be traversed twice before completing a circuit that encompasses all the nodes in the graph.

Even though the removal of this edge does not partition the graph, we can consider this edge to divide the graph into two parts C_1 and C_2, where C_1 is the part that was traversed before the first traversal of e' combined with the part traversed after the second traversal; and C_2 is the part traversed in between the first traversal and the second.

We can readily see that C_2 is a connected component and we can also show that C_1 is a connected component since it has at least one node of the bridge in common between the first and third traversals. Once we show that e' logically divides the graph into two, with at least one part containing at least r nodes, we can use Theorem 4 to extend the graph. □

We can analogously extend this to graphs with odd regularity.

Corollary 2. *Given an r–regular, non-hamiltonian graph over n nodes, where r is odd, n is even, and $n \geq 2r$, there exists an r–regular, non-hamiltonian graph with $n+2$ nodes.*

7.4 Regular Directed Graphs

Until now, we have only addressed regular graphs that are undirected. Directed graphs involve several issues like differences between indegree and outdegree and strong connectivity that make them harder to analyze than undirected graphs.

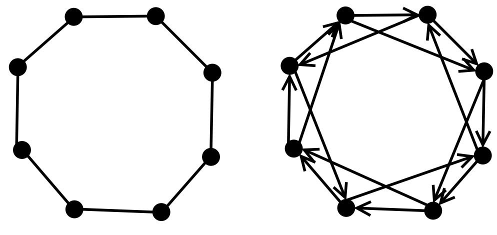

Figure 7.7: Diameter can be reduced in directed regular graphs, without affecting regularity

Since many networks are modelled naturally as directed graphs, we briefly explore the symmetry property for directed graphs in this section.

A directed graph is r–inregular if all nodes have an indegree of r; similarly, a directed graph is r–outregular if all nodes have an outdegree of r. Finally, if a directed graph is r–regular if it is both r–inregular and r–outregular.

An undirected graph can be rewritten as a directed graph by representing each undirected edge as a pair of directed edges in opposite directions. Given this, we face the question of whether designing a diameter optimal r–regular directed graph necessarily entails representing directed edges in pairs, making the graph into an undirected graph.

Figure 7.7 refutes such an assertion. Given 8 nodes and 8 undirected edges (or 16 pairs of directed edges), the best we can do for regularity-preserving diameter reduction, is to form a 2–regular hamiltonian cycle having a diameter 4. This is shown in Figure 7.7(a). However, when the pairwise coupling of the directed edges are broken, we can obtain a directed 2–regular graph with diameter 3, as shown in Figure 7.7(b).

Further, we can make the following observations over directed r–regular graphs.

Postulate 9. *The smallest number of nodes required for a strongly connected, directed r–regular graph is $r+1$. Such a graph is in the form of a clique, where every node has r outgoing edges connecting every other node in the graph, directly.*

Theorem 10. *An r–regular directed graph over n nodes, where $n \geq r + 1$, can be extended by a single node, regardless of whether the regularity is odd or even.*

Proof. The extension theorem for directed graphs is similar to Theorem 2 concerning odd regular undirected graphs. Consider an r–regular directed graph G and an incoming node v. We can think of v as having r holes for incoming connections and r dangling outgoing edges looking for connections. We now break r edges from G and connect them to v, such that they create holes on r distinct nodes in G. The r outgoing edges from v are then connected to these holes.

Note that any node in G can host both an outgoing connection, as well as an incoming connection from v. This requires a maximum of only r nodes in G, which is satisfied trivially. □

In directed regular graphs, we do not face incompatibility between regularity and number of nodes. If the graph is r-regular in both indegree and outdegree, the total degree of every node is always even. This makes it possible to have r–regular directed graphs over any number of nodes, as long as the number of nodes are at least $r + 1$.

The connectivity of a directed graph is the size of the smallest edge cutset, whose removal will no longer make the graph strongly connected. Given this, we can prove results analogous to Theorem 4 for directed graphs as well. Similarly, hamiltonian circuits in directed graphs can be extended in the same spirit as of Theorem 8.

Do not worry if you have built your castles in the air. They are where they should be. Now put the foundations under them.

Henry David Thoreau

8

Concluding Remarks

Networked systems are ubiquitous and form an important part of our everyday lives. There is an ever increasing need to exchange material, energy and information. As a result of which the problems related to optimal network design provide new and interesting research challenges to computer science.

In this thesis, we focus on certain aspects of optimal network design to present a combination of (1) performance modelling and experiments, (2) empirical observations and analysis, and (3) development of algorithms and theoretical results.

We start by modelling the performance requirements of a network in terms of its structural properties. Using a genetic algorithm framework, we evolve networks that are optimal under different trade-offs. We observe patterns of recurring structures or classes of topology which provide useful design guidelines for a network designer. The following are the prominent classes that we observe: (1) hubs–and–spokes, (2) circular hubs, (3) symmetric multihubs, and (4) circular skip lists.

We also observe that these classes correspond to specific subspaces in an optimal topology space. We characterize the space of optimal topologies to show the "phase transitions" between classes of topologies. In particular, we show the sharp emergence of circular skip lists once the severe constraints on cost are relaxed. In fact, circular skip lists are the most prominent class of topologies as we show in our analysis.

Next, we study circular skip lists in detail to understand their properties. Circular skip lists exhibit a number of structural properties that are considered optimal in network design: symmetric centrality measures, hamiltonian decompositions, low diameters, symmetric distance measures, and high connectivity. Based on our observations and analysis, we conjecture that circular skip lists form potential underpinnings to designing optimal network topologies in multiple application domains.

Further, there are several theoretical results concerning special families of graphs such as Cayley graphs, Ramanujan graphs, random regular graphs and expanders that have optimal properties similar to the ones exhibited by circular skip lists. Please refer to Section 2.1 for a detailed discussion in this regard. Theoretical results in the domain of DHTs also show that ring–like structures are best suited for DHT design (Section 2.2.2). These findings are in agreement with our conjecture.

In the first part of the thesis we evolve optimal topologies to observe different classes of topologies. In the second part of this thesis, we focus on the most prominent class—circular skip lists. In Chapter 6, we develop simple algorithms for constructing circular skip lists that have optimal properties such as low diameter, symmetry, hamiltonicity and high connectivity. In Chapter 7, we address the problem of preserving the aforementioned optimal properties in the face of perturbations such as addition of new nodes to a network. In particular, we prove theoretical results about regular graphs, most of which are circular skip lists.

This work is part of a larger vision, which is to develop a deeper theoretical understanding of network design in terms of general design principles. In this light, the present work has lead to several interesting future directions. In the next section, we briefly discuss two research problems which we want to work on in the future.

8.1 Future Directions

Distributed systems work in inherently uncertain and dynamic environments. One source of uncertainty is internal: failure of machines and communication links; frequent joining or leaving of machines; and, mobility of machines. The other source of uncertainty is external: performance requirements of a network might change over time; it might be necessary for a network to operate under different levels of optimality at different times.

A limitation of the present work is that topology design does not consider dynamics such as the above. A network's performance requirements are modelled *a priori*. The optimization is one shot and static, while a network's performance requirements might change over time. Similarly, the network topology itself might change due to perturbations. The changes in topology might render the network suboptimal.

Thus, another critical parameter that decides a network's optimality in addition to efficiency, robustness and cost is *adaptability*. This leads to the following two research problems concerning: (1) adaptability of network topologies in the face of structural changes, and (2) adaptability of network topologies in the face of changing performance requirements.

8.1.1 Structural Signatures and Snapping

In the face of perturbations, it is difficult to guarantee the optimality of a network topology. For example, when a communication link fails or a new machine joins a network, a topology which was symmetric before this event might cease to remain symmetric. The objective of this work is to develop ways of *reconstructing* an optimal topology in the face of disruptions. They can also be seen as ways of reaching an optimal topology under a given set of constraints starting from an arbitrary sub-optimal topology.

This is the motivation for exploring *snapping models*. A snapping model has embedded in it the knowledge of optimal topologies. We assume that through the process of topology breeding, we have a set of known optimal topologies under various constraints. The basis for the reconstruction is the conjecture that optimal topologies have a structural "signature" that we can compute. The challenge then is to use such a signature to make suitable topology changes so that a network that has been rendered sub-optimal reaches a known optimality. We call this process snapping.

Presently, we are exploring this problem (1) to develop combinations of graph theoretical properties that can form a signature, and (2) to develop snapping algorithms such that nodes in a network can use these signatures and work towards optimality in a distributed manner.

8.1.2 Adaptable Networks under Varying Performance Requirements

We would like to extend our topology breeding work to discover networks that are not only efficient, robust, inexpensive, but also *adaptable*. By an adaptable network we imply a network that can be (re)configured to one of several desired configurations with low transition costs. Here, each configuration corresponds to a performance requirement. For instance, one desired configuration may want a minimum guaranteed efficiency, and another may want guaranteed robustness.

We believe that such adaptive networks will be useful in applications like network centric warfare, adaptive supply networks, labour supply networks in information technology, disaster management and simulating terrorist cell networks. Please refer to [Malik et al., 2010] for some preliminary work with respect to providing scalable data services through adaptable networks.

Presently, we are working on the following aspects of the problem: (1) defining metrics for transition cost, and (2) developing methods to establish equivalence of two topologies in terms of performance.

A

List of Related Publications

Journal and Conference Papers

1. S. Patil, S. Srinivasa, S. Mukherjee, A.R. Rachakonda and V. Venkatasubramanian. Breeding Diameter-Optimal Topologies for Distributed Indexes. *Complex Systems*, 18(2(c)):175–194, Complex Systems Publications, Inc., 2009.

2. S. Patil, S. Srinivasa and V. Venkatasubramanian. Classes of Optimal Topologies under Multiple Efficiency and Robustness Constraints. *In Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 4940–4945, IEEE, 2009.

3. S. Patil. Designing Optimal Network Topologies under Multiple Efficiency and Robustness Constraints. *In PhD Forum Publication of the International Conference on Distributed Computing and Networking (ICDCN)*, pages 25–28, 2010.

4. S. Patil and S. Srinivasa. Theoretical Notes on Regular Graphs as Applied to Optimal Network Design. *In Proceedings of the International Conference on Distributed Computing and Internet Technology (ICDCIT)*, pages 236–242, Springer, 2010.

5. T. Malik, R. Prasad, S. Patil, A. Chaudhary and V. Venkatasubramanian. Providing Scalable Data Services in Ubiquitous Networks. *In Proceedings of the International Conference on Database Systems for Advanced Applications (DASFAA)*, pages 445–457. Springer, 2010.

Talks and Posters

1. Designing Optimal Network Topologies under Multiple Efficiency and Robustness Constraints. *Research Poster Competition, IBM IRL Collaborative Academia Research Exchange (I-CARE)*, New Delhi, October 2009.

2. Designing Optimal Network Topologies under Multiple Efficiency and Robustness Constraints. *Research Poster Session, PhD forum at International Conference on Distributed Computing and Networking (ICDCN)*, Kolkata, January 2010.

Bibliography

E. Abrahamson and L. Rosenkopf. Social network effects on the extent of innovation diffusion: A computer simulation. *Organization Science*, 8(3):289–309, 1997.

L. Adamic, B. Huberman, A. Barabási, R. Albert, H. Jeong, and G. Bianconi. Power-law distribution of the world wide web. *Science*, 287(5461):2115a, 2000.

N. Adler and K. Smilowitz. Hub-and-spoke network alliances and mergers: Price-location competition in the airline industry. *Transportation Research Part B*, 41 (4):394–409, 2007.

A. Agrawal, P. Klein, and R. Ravi. When trees collide: an approximation algorithm for the generalized steiner problem on networks. *SIAM Journal on Computing*, 24(3):440–456, 1995.

S. Akers and B. Krishnamurthy. A group-theoretic model for symmetric interconnection networks. *IEEE Transactions on Computers*, 38(4):555–566, 1989.

R. Albert and A. Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47–97, 2002.

R. Albert, H. Jeong, and A. Barabási. Diameter of the world-wide web. *Nature*, 401(6749):130–131, 1999.

D. Alberts, J. Garstka, and F. Stein, editors. *Network centric warfare: Developing and leveraging information superiority*. Dept. of Defense. Center for Advanced Concepts and Technology (ACT), 1999.

M. Alderighi, A. Cento, P. Nijkamp, and P. Rietveld. Network competition–the coexistence of hub-and-spoke and point-to-point systems. *Journal of Air Transport Management*, 11(5):328–334, 2005.

D. Alderson, L. Li, W. Willinger, and J. Doyle. Understanding Internet topology: principles, models, and validation. *IEEE/ACM Transactions on Networking*, 13 (6):1205–1218, 2005.

N. Alon, S. Hoory, and N. Linial. The Moore bound for irregular graphs. *Graphs and Combinatorics*, 18(1):53–57, 2002.

J. Aspnes and G. Shah. Skip graphs. *ACM Transactions on Algorithms (TALG)*, 3 (4):37–es, 2007.

J. Aspnes and Y. Yin. Distributed Algorithms for Maintaining Dynamic Expander Graphs. *Relation*, 10(1.74):8616, 2008.

B. Awerbuch and C. Scheideler. Towards a scalable and robust DHT. *Theory of Computing Systems*, 45(2):234–260, 2009.

T. Bäck. *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, USA, 1st edition, 1996.

G. Bagler. Analysis of the Airport Network of India as a complex weighted network. *Physica A: Statistical Mechanics and its Applications*, 387(12):2972–2980, 2008.

P. Bak. *How nature works: the science of self-organized criticality*, volume 212. Copernicus New York, 1996.

E. Bannai and T. Ito. On finite Moore graphs. *Journal of the Faculty of Science, Sect. 1A*, 20:191–208, 1973.

A. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

A. Barabási and E. Bonabeau. Scale-free Networks. *Scientific American*, 288(5): 50–59, 2003.

A. Barabási and R. Crandall. *Linked: The new science of networks*. Basic Books, 2002.

F. Barahona and G. Reinelt. An application of combinatorial optimization to statistical physics and circuit layout design. *Operations Research*, 36(3):493–513, 1988.

A. Barrat, M. Barthélemy, R. Pastor-Satorras, and A. Vespignani. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences*, 101(11):3747–3752, 2004.

M. Barthélemy, A. Barrat, R. Pastor-Satorras, and A. Vespignani. Characterization and modeling of weighted networks. *Physica A: Statistical Mechanics and Its Applications*, 346(1-2):34–43, 2005.

A. Bento, M. Cropper, A. Mobarak, and K. Vinha. The effects of urban spatial structure on travel demand in the United States. *Review of Economics and Statistics*, 87(3):466–478, 2005.

J. Bermond, E. Darrot, O. Delmas, and S. Perennes. Hamilton circuits in the directed wrapped Butterfly network. *Discrete Applied Mathematics*, 84(1-3): 21–42, 1998.

T. Berners-Lee, W. Hall, J. Hendler, K. O'Hara, N. Shadbolt, and D. Weitzner. A framework for web science. *Foundations and Trends in Web Science*, 1(1): 1–130, 2006.

N. Biggs. *Algebraic graph theory*, volume 67, pages 180–190. Cambridge Univ Press, 2 edition, 1993.

B. Bollobás. *Random graphs*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2nd edition, 2001.

B. Bollobás. *Extremal graph theory*. Dover Publications, 2004.

J. Bondy and U. Murty. *Graph Theory with Applications*, volume 290, page 235. MacMillan London, 1976.

K. Börner, S. Sanyal, and A. Vespignani. Network science. *Annual review of information science and technology*, 41(1):537–607, 2007.

W. G. Bridges and S. Toueg. Impossibility of Directed Moore Graphs. *Journal of Combinatorial Theory, Series B*, 29(3):339–341, 1980.

J. Brueckner and Y. Zhang. A model of scheduling in airline networks: how a hub-and-spoke system affects flight frequency, fares and welfare. *Journal of transport Economics and Policy*, 35(2):195–222, 2001.

J. Camacho, R. Guimera, and L. Amaral. Analytical solution of a model for complex food webs. *Physical Review E*, 65(3):30901–30904, 2002a.

J. Camacho, R. Guimerà, and L. Amaral. Robust patterns in food web structure. *Physical review letters*, 88(22):228102–228106, 2002b.

S. Cantono and G. Silverberg. A percolation model of eco-innovation diffusion: the relationship between diffusion, learning economies and subsidies. *Technological Forecasting & Social Change*, 76(4):487–496, 2009.

J. Carlson and J. Doyle. Highly optimized tolerance: A mechanism for power laws in designed systems. *Physical Review E*, 60(2):1412–1427, 1999.

S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of Internet topology using k-shell decomposition. *Proceedings of the National Academy of Sciences*, 104(27):11150–11154, 2007.

S. Çamtepe, B. Yener, and M. Yung. Expander graph based key distribution mechanisms in wireless sensor networks. In *Proceedings of the IEEE International Conference on Communications*, volume 5, pages 2262–2267. IEEE, 2006.

A. Cebrowski. The Implementation of Network-Centric Warfare. *Office of Force Transformation, Department of Defense*, 2005.

M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed Steiner problems. *Journal of Algorithms*, 33(1):73–91, 1999.

Q. Chen, H. Chang, R. Govindan, and S. Jamin. The origin of power laws in Internet topologies revisited. In *Proceedings of the IEEE International Conference on Computer Communications*, volume 2, pages 608–617. IEEE, 2002.

M. Chlebík and J. Chlebíková. Approximation hardness of the steiner tree problem on graphs. *Algorithm Theory*, pages 95–99, 2002.

N. Christakis and J. Fowler. The collective dynamics of smoking in a large social network. *New England Journal of Medicine*, 358(21):2249–2258, 2008a.

N. Christakis and J. Fowler. Estimating Peer Effects on Health in Social Networks. *Journal of Health Economics*, 27:1400–1405, 2008b.

N. Christin, A. Weigend, and J. Chuang. Content availability, pollution and poisoning in file sharing peer-to-peer networks. In *Proceedings of the ACM conference on Electronic commerce*, pages 68–77. ACM, 2005.

F. Chung. Diameters of graphs: old problems and new results. *Congressus Numerantium*, 60:295–317, 1987.

V. Colizza, A. Barrat, M. Barthélemy, and A. Vespignani. The role of the airline transportation network in the prediction and predictability of global epidemics. *Proceedings of the National Academy of Sciences*, 103(7):2015–2020, 2006.

F. Comellas and J. Gómez. New large graphs with given degree and diameter six. *Networks*, 34:154–161, 1999.

L. Costa, F. Rodrigues, G. Travieso, and P. Boas. Characterization of complex networks: A survey of measurements. *Advances in Physics*, 56:167–242, 2007.

R. Cowan. Network models of innovation and knowledge diffusion. *Clusters, networks and innovation*, pages 29–53, 2005.

N. de Bruijn. A combinatorial problem. *Koninklijke Netherlands: Academe Van Wetenschappen*, 49:758–764, 1946.

T. Decker, B. Monien, and R. Preis. Towards optimal load balancing topologies. In *Proceedings of Euro-Par*, pages 277–287. Springer, 2000.

A. Dekker and B. Colbert. Network robustness and graph topology. In *Proceedings of the Australasian Conference on Computer Science*, volume 26, pages 359–368. ACM, 2004.

S. Delre, W. Jager, T. Bijmolt, and M. Janssen. Will it spread or not? The effects of social influences and network topology on innovation diffusion. *Journal of Product Innovation and Management*, 27(2):267–282, 2007.

N. Dennis. Industry consolidation and future airline network structures in Europe. *Journal of Air Transport Management*, 11(3):175–183, 2005.

M. J. Dinneen and P. R. Hafner. New results for the degree/diameter problem. *Networks*, 24(7):359–367, 1994.

L. Donetti, P. Hurtado, and M. Munoz. Entangled networks, synchronization, and optimal network topology. *Physical Review Letters*, 95(18):188701–188705, 2005.

L. Donetti, F. Neri, and M. Muñoz. Optimal network topologies: Expanders, Cages, Ramanujan graphs, Entangled networks and all that. *Journal of Statistical Mechanics: Theory and Experiment*, 2006:P08007–P08024, 2006.

R. Dougherty and V. Faber. The degree-diameter problem for several varieties of Cayley graphs. I: The Abelian case. *SIAM journal on discrete mathematics*, 17 (3):478–519, 2004.

Z. Drezner. *Facility location: a survey of applications and methods*. Springer

Series in Operations Research and Financial Engineering. Springer Verlag, 1st edition, 1995.

D. Du, B. Lu, H. Ngo, and P. Pardalos. Steiner tree problems. *Encyclopedia of optimization*, 5:227–290, 2001.

P. Erdös, A. Rényi, and V. Sós. On a problem in the theory of graphs. *Publ. Math. Inst. Hungar. Acad. Sci*, 7:215–235, 1962.

S. Eubank, H. Guclu, V. Anil Kumar, M. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184, 2004.

G. Exoo and R. Jajcay. Dynamic cage survey. *The Electronic Journal of Combinatorics*, 15:48, 2008.

M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of the ACM International Conference on Computer Communication*, pages 251–262. ACM, 1999.

M. Fiol, I. Alegre, and J. Yebra. Line digraph iterations and the (d, k) problem for directed Graphs. In *Proceedings of the ACM International Symposium on Computer Architecture*, pages 174–177. ACM, 1983.

G. Fletcher, H. Sheth, and K. Börner. Unstructured peer-to-peer networks: Topological properties and search performance. *Agents and Peer-to-Peer Computing*, 3601:14–27, 2005.

S. Forbes. The effect of air traffic delays on airline prices. *International Journal of Industrial Organization*, 26(5):1218–1232, 2008.

J. Fowler and N. Christakis. The Spread of Obesity in a Large Social Network Over 32 Years. *The New England Journal of Medicine*, 357(4):370–379, 2007.

J. Fowler and N. Christakis. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *British Medical Journal*, 337(dec04 2):a2338, 2008.

P. Fraigniaud and P. Gauron. D2B: A de Bruijn based content-addressable network. *Theoretical Computer Science*, 355(1):65–79, 2006.

M. Freedman and R. Vingralek. Efficient peer-to-peer lookup based on a distributed trie. In *Proceedings of the International Workshop on Peer-to-Peer Systems*, pages 66–75. Springer, 2002.

H. Friedman. On the impossibility of certain Moore graphs. *Journal of Combinatorial Theory, Series B*, 10(3):245–252, 1971.

M. Garey, R. Graham, and D. Johnson. The complexity of computing Steiner minimal trees. *SIAM journal on applied mathematics*, 32(4):835–859, 1977.

N. Garg, G. Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group Steiner tree problem. *Journal of Algorithms*, 37(1):66–84, 2000.

C. Gkantsidis, M. Mihail, and A. Saberi. Random walks in peer-to-peer networks: algorithms and evaluation. *Performance Evaluation*, 63(3):241–263, 2006.

C. Godsil, G. Royle, and C. Godsil. *Algebraic graph theory*, volume 8, page 207. Springer New York, 2001.

M. González, C. Hidalgo, and A. Barabási. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.

R. Gould. Updating the hamiltonian problem-a survey. *Journal of Graph Theory*, 15(2):121–157, 1991.

R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. In *Proceedings of the IEEE International Conference on Computer Communications*, volume 3, pages 1371–1380. IEEE, 2002.

C. Gropl, S. Hougardy, T. Nierhoff, and H. Promel. Approximation algorithms for the Steiner tree problem in graphs. *Steiner trees in industry*, pages 235–279, 2001.

M. Guida and F. Maria. Topology of the Italian airport network: A scale-free small-world network with a fractal structure? *Chaos, Solitons & Fractals*, 31 (3):527–536, 2007. ISSN 0960-0779.

R. Guimerá and L. A. N. Amaral. Modeling the world-wide airport network. *The European Physical Journal B*, 38(2):381–385, 2004.

R. Guimerá, A. Diaz-Guilera, F. Vega-Redondo, A. Cabrales, and A. Arenas. Optimal network topologies for local search with congestion. *Physical Review Letters*, 89(24):248701–248705, 2002.

R. Guimerá, S. Mossa, A. Turtschi, and L. Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences*, 102(22):7794–7799, 2005.

R. Guimerá, M. Sales-Pardo, and L. Amaral. A network-based method for target selection in metabolic networks. *Bioinformatics*, 23(13):1616–1622, 2007.

K. Gummadi, R. Gummadi, S. Gribble, S. Ratnasamy, S. Shenker, and I. Stoica. The impact of DHT routing geometry on resilience and proximity. In *Proceedings of the ACM International Conference on Computer Communications*, pages 381–394. ACM, 2003.

P. Hafner. Large Cayley graphs and digraphs with small degree and diameter. *Computational Algebra and Number Theory*, pages 291–302, 1992.

F. Harary. The maximum connectivity of a graph. *Proceedings of the National Academy of Sciences of the United States of America*, 48(7):1142–1146, 1962.

F. Harary. *Graph Theory*, pages 89, 112 and 175. Addison-Wesley, Reading, MA, 1994.

F. Harary and E. Palmer. *Graphical enumeration*, volume 123. Academic Press, New York, 1st edition, 1973.

A. Hoffman and R. Singleton. On Moore graphs with diameters 2 and 3. *IBM Journal of Research and Development*, 4(5):497–504, 1960.

S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

S. Hsieh and T. Hsiao. The k-degree Cayley graph and its topological properties. *Networks*, 47(1):26–36, 2006.

K. Hwang. *Advanced computer architecture: parallelism, scalability, programmability*, volume 348. McGraw-Hill, New York, 1993.

M. Imase, T. Soneoka, and K. Okada. Connectivity of regular directed graphs with small diameters. *IEEE Transactions on Computers*, 100(34):267–273, 1985.

H. Jeong, S. Mason, A. Barabási, and Z. Oltvai. Lethality and centrality in protein networks. *Nature*, 411(6833):41–42, 2001.

T. Jiang and D. Mubayi. Connectivity and separating sets of cages. *Journal of Graph Theory*, 29(1):35–44, 1998.

M. Kaashoek and D. Karger. Koorde: A simple degree-optimal distributed hash table. *Peer-to-Peer Systems II*, 2735/2003:98–107, 2003.

J. Kleinberg. Navigation in a small world. *Nature*, 406(6798):845–845, 2000.

A. Kumar, S. Merugu, J. Xu, E. Zegura, and X. Yu. Ulysses: A robust, low-diameter, low-latency peer-to-peer network. *European Transactions on Telecommunications*, 15(6):571–587, 2004.

R. Kumar, J. Novak, and A. Tomkins. Structure and evolution of online social networks. *Link Mining: Models, Algorithms, and Applications*, 4:337–357, 2010.

L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.

V. Latora and M. Marchiori. Is the Boston subway a small-world network? *Physica A: Statistical Mechanics and its Applications*, 314(1-4):109–113, 2002.

C. Law and K. Siu. Distributed construction of random expander graphs. In *Proceedings of the IEEE International Conference on Computer Communications*, pages 2133–2143. IEEE, 2003.

T. Lengauer. *Combinatorial algorithms for integrated circuit layout*. John Wiley & Sons, Inc. New York, USA, 1990.

W. Li and X. Cai. Statistical analysis of airport network of China. *Physical Review E*, 69(4):46106–46112, 2004.

D. Loguinov, J. Casas, and X. Wang. Graph-theoretic analysis of structured peer-to-peer systems: routing distances and fault resilience. *IEEE/ACM Transactions on Networking*, 13(5):1107–1120, 2005.

E. Loz and J. Širáň. New record graphs in the degree-diameter problem. *Australasian Journal of Combinatorics*, 41:63–80, 2008.

E. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys and Tutorials*, 7(2):72–93, 2005.

W. Magnus, A. Karrass, and D. Solitar. *Combinatorial group theory: Presentations of groups in terms of generators and relations*. Dover publications, 1976.

D. Magoni and J. Pansiot. Analysis of the autonomous system network topology. *ACM SIGCOMM Computer Communication Review*, 31(3):26–37, 2001.

P. Mahadevan, D. Krioukov, X. Dimitropoulos, B. Huffaker, M. Fomenkov, K. Claffy, and A. Vahdat. The Internet AS-level topology: three data sources and one definitive metric. *ACM SIGCOMM Computer Communication Review*, 36(1):17–26, 2006.

T. Malik, R. Prasad, S. Patil, A. Chaudhary, and V. Venkatasubramanian. Providing Scalable Data Services in Ubiquitous Networks. In *Proceedings of the Database Systems for Advanced Applications (DASFAA)*, pages 445–457. Springer, 2010.

D. Malkhi, M. Naor, and D. Ratajczak. Viceroy: A scalable and dynamic emulation of the butterfly. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 183–192. ACM, 2002.

G. Manku, M. Bawa, and P. Raghavan. Symphony: Distributed hashing in a small world. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems*, volume 4, pages 127–140. USENIX Association, 2003.

P. McAfee and V. Te Velde. Dynamic pricing in the airline industry. *Handbook on Economics and Information Systems, Ed: T.J. Hendershott*, 2006.

B. McKay and C. Praeger. Vertex-transitive graphs which are not Cayley graphs. *Journal of the Australian Mathematical Society-Series A*, 56(1):53–63, 1994.

B. McKay, M. Miller, and J. Širáň. A note on large graphs of diameter two and given maximum degree. *Journal of Combinatorial Theory, Series B*, 74(1): 110–118, 1998.

N. Meade and T. Islam. Modelling and forecasting the diffusion of innovation–A 25-year review. *International Journal of Forecasting*, 22(3):519–545, 2006.

A. Medina, I. Matta, and J. Byers. On the origin of power laws in Internet topologies. *ACM SIGCOMM Computer Communication Review*, 30(2):18–28, 2000.

M. Meringer. Fast generation of regular graphs and construction of cages. *Journal of Graph Theory*, 30(2):137–146, 1999.

M. Miller and J. Širán. Moore graphs and beyond: A survey of the degree/diameter problem. *Electronic Journal of Combinatorics*, 61:1–63, 2005.

D. Molle, S. Richter, and P. Rossmanith. A faster algorithm for the Steiner tree problem. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, volume 3884, pages 561–570. Springer, 2006.

R. J. Mondragón C. Optimal networks, congestion and Braess' paradox. In *Proceedings from the Workshop on Interdisciplinary systems approach in perfor-*

mance evaluation and design of computer & communications systems. ACM, 2006.

A. Motter and Z. Toroczkai. Introduction: Optimization in networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 17:026101–026104, 2007.

M. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.

M. Newman and M. Girvan. Mixing patterns and community structure in networks. *Statistical Mechanics of Complex Networks*, pages 66–87, 2003.

G. Pandurangan, P. Raghavan, and E. Upfal. Building low-diameter P2P networks. In *Proceedings of the IEEE Symposium On Foundations Of Computer Science*, volume 42, pages 492–499. IEEE, 2002.

B. Parhami and M. Rakov. Perfect difference networks and related interconnection structures for parallel and distributed systems. *IEEE Transactions on Parallel and Distributed Systems*, 16(8):714–724, 2005.

S. Patil and S. Srinivasa. Theoretical Notes on Regular Graphs as applied to Optimal Network Design. In *Proceedings of the International Conference on Distributed Computing and Internet Technology*, pages 236–242. Springer, 2010.

S. Patil, S. Srinivasa, S. Mukherjee, A. Rachakonda, and V. Venkatasubramanian. Breeding Diameter-Optimal Topologies for Distributed Indexes. *Complex Systems*, 18(2(c)):175–194, 2009.

C. Plaxton, R. Rajaraman, and A. Richa. Accessing nearby copies of replicated objects in a distributed environment. *Theory of Computing Systems*, 32(3):241–280, 1999.

W. Pugh. Skip lists: a probabilistic alternative to balanced trees. *Communications of the ACM*, 33(6):668–676, 1990.

C. Qu, W. Nejdl, and M. Kriesell. Cayley DHTs–A Group-Theoretic Framework for Analyzing DHTs Based on Cayley Graphs. *Semantic Web and Peer-to-Peer Part II*, pages 89–105, 2006.

H. Rahmandad and J. Sterman. Heterogeneity and Network Structure in the Dynamics of Diffusion: Comparing Agent-Based and Differential Equation Models. *Management Science*, 54(5):998–1014, 2008.

S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *Proceedings of the ACM International Conference on Computer Communications*, pages 161–172. ACM, 2001.

S. Ratnasamy, I. Stoica, and S. Shenker. Routing algorithms for DHTs: Some open questions. *Peer-to-Peer Systems*, 2429:45–52, 2002.

S. Ratnasamy, B. Karp, S. Shenker, D. Estrin, R. Govindan, L. Yin, and F. Yu. Data-centric storage in sensornets with GHT, a geographic hash table. *Mobile networks and applications*, 8(4):427–442, 2003.

R. Robinson and N. Wormald. Almost all regular graphs are Hamiltonian. *Random Structures and Algorithms*, 5(2):363–374, 1994.

E. Rogers. *Diffusion of innovations*. Free Press, 1995. ISBN 0029266718.

M. Rosvall, A. Trusina, P. Minnhagen, and K. Sneppen. Networks and cities: An information perspective. *Physical Review Letters*, 94(2):28701–28705, 2005.

A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, volume 11, pages 329–350. Springer, 2001.

T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro. A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, 167(1):96–115, 2005.

P. Sarnak. What is an expander? *Notices of the AMS*, 51(7):762–763, 2004.

M. Schlosser, M. Sintek, S. Decker, and W. Nejdl. HyperCuP–hypercubes, ontologies, and efficient search on peer-to-peer networks. In *Proceedings of Agents and Peer-to-Peer Computing, First International Workshop, AP2PC 2002, Revised and Invited Papers*, pages 112–124. Springer, 2003.

P. Sen, S. Dasgupta, A. Chatterjee, P. Sreeram, G. Mukherjee, and S. Manna. Small-world properties of the Indian railway network. *Physical Review E*, 67 (3):36106–36111, 2003.

J. Šiagiová and J. Širáň. A note on large Cayley graphs of diameter two and given degree. *Discrete Mathematics*, 305(1-3):379–382, 2005.

J. Sienkiewicz and J. Hołyst. Statistical analysis of 22 public transport networks in Poland. *Physical Review E*, 72(4):46127–46138, 2005.

R. Singleton. There is no irregular Moore graph. *The American Mathematical Monthly*, 75(1):42–43, 1968.

T. Small, B. Li, and B. Liang. Topology Affects the Efficiency of Network Coding in Peer-to-Peer Networks. In *Proceedings of IEEE International Conference on Communications*, pages 5591–5597. IEEE, 2008.

A. Smart, L. Amaral, and J. Ottino. Cascading failure and robustness in metabolic networks. *Proceedings of the National Academy of Sciences*, 105(36):13223–13228, 2008.

A. Steger and N. Wormald. Generating random regular graphs quickly. *Combinatorics, Probability and Computing*, 8(04):377–396, 1999.

I. Stoica, R. Morris, D. Liben-Nowell, D. Karger, M. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Transactions on Networking*, 11(1):17–32, 2003.

D. Stouffer, J. Camacho, W. Jiang, and L. Nunes Amaral. Evidence for the existence of a robust pattern of prey selection in food webs. *Proceedings of the Royal Society B*, 274(1621):1931–1940, 2007.

S. Strogatz. Exploring complex networks. *Nature*, 410(6825):268–276, 2001.

D. Stutzbach and R. Rejaie. Understanding churn in peer-to-peer networks. In *Proceedings of the ACM SIGCOMM Conference on Internet Measurement*, pages 189–202. ACM, 2006.

A. Tatem, S. Hay, and D. Rogers. Global traffic and disease vector dispersal. *Proceedings of the National Academy of Sciences of the United States of America*, 103(16):6242–6247, 2006.

E. Tri Baskoro, M. Miller, and J. Plesník. On the structure of digraphs with order close to the Moore bound. *Graphs and Combinatorics*, 14(2):109–119, 1998.

A. Valente, A. Sarkar, and H. Stone. Two-peak and three-peak optimal complex networks. *Physical Review Letters*, 92(11):118702–118706, 2004.

T. Valente. Network models and methods for studying the diffusion of innovations. *Models and Methods in Social Network Analysis*, pages 98–116, 2005.

T. Valente and R. Davis. Accelerating the diffusion of innovations using opinion leaders. *The Annals of the American Academy of Political and Social Science*, 566(1):55–67, 1999.

V. Venkatasubramanian. A theory of design of complex teleological systems: Unifying the Darwinian and Boltzmannian perspectives. *Complexity*, 12(3): 14–21, 2007.

V. Venkatasubramanian, S. Katare, P. Patkar, and F. Mu. Spontaneous emergence of complex optimal networks through evolutionary adaptation. *Computers & chemical engineering*, 28(9):1789–1798, 2004.

P. Wang, B. Xu, and J. Wang. A note on the edge-connectivity of cages. *the electronic journal of combinatorics*, 10(2):1–4, 2003.

D. Watts and S. Strogatz. Collective dynamics of "small-world" networks. *Nature*, 393(6684):440–442, 1998.

R. Williams and N. Martinez. Simple rules yield complex food webs. *Nature*, 404 (6774):180–183, 2000.

W. Willinger, D. Alderson, and J. Doyle. Mathematics and the internet: A source of enormous confusion and great potential. *Notices of the American Mathematical Society*, 56(5):586–599, 2009.

N. Wormald. Models of random regular graphs. *Surveys in combinatorics*, 276: 239–298, 1999.

W. Xiao and B. Parhami. Some mathematical properties of Cayley digraphs with applications to interconnection network design. *International Journal of Computer Mathematics*, 82(5):521–528, 2005.

W. Xiao and B. Parhami. Cayley graphs as models of deterministic small-world networks. *Information Processing Letters*, 97(3):115–117, 2006.

W. Xiao, W. Chen, M. He, W. Wei, and B. Parhami. Biswapped Networks and Their Topological Properties. In *Proceedings of the ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing*, pages 193–198, 2007.

J. Xu, A. Kumar, and X. Yu. On the fundamental tradeoffs between routing table size and network diameter in peer-to-peer networks. *IEEE Journal on Selected Areas in Communications*, 22(1):151–163, 2004.

L. Zhao, Y. Lai, K. Park, and N. Ye. Onset of traffic congestion in complex networks. *Physical Review E*, 71(2):26125–26133, 2005.

S. Zhou and R. Mondragón. Accurately modeling the Internet topology. *Physical Review E*, 70(6):66108–66116, 2004.