

Internet Research Task Force
Internet-Draft
Intended status: Informational
Expires: August 20, 2013

M.H. Behringer
Cisco
G. Huston
Asia Pacific Network Information Centre
February 18, 2013

A Framework for Defining Network Complexity
draft-irtf-ncrg-complexity-framework-00.txt

Abstract

Complexity is a widely used parameter in network design, yet there is no generally accepted definition of the term. Complexity metrics exist in a wide range of research papers, but most of these address only a particular aspect of a network, for example the complexity of a graph or software. There is a desire to define the complexity of a network as a whole, as deployed today to provide Internet services. This document provides a framework to guide research on the topic of network complexity.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 20, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

| | |
|--|---|
| 1. Introduction | 2 |
| 2. Current Understanding of Network Complexity | 2 |

| | | |
|------|---|---|
| 2.1. | The Behavior of a Complex Network | 2 |
| 2.2. | Robust Yet Fragile | 3 |
| 2.3. | The Complexity Cube | 3 |
| 3. | Towards Defining Network Complexity | 3 |
| 3.1. | General Observations | 3 |
| 3.2. | The Problem Space | 3 |
| 3.3. | Technical Debt | 4 |
| 3.4. | Layering considerations | 5 |
| 4. | Possible Directions of Research | 5 |
| 4.1. | Definitions and Metrics | 5 |
| 4.2. | Comparative Analysis | 6 |
| 4.3. | Containment, Control or Reduction of Complexity | 6 |
| 4.4. | Use Cases | 6 |
| 5. | Security Considerations | 7 |
| 6. | Acknowledgements | 7 |
| 7. | References | 7 |
| | Authors' Addresses | 8 |

1. Introduction

During the design phase of a network complexity plays a key role. Network designers generally seek to find the simplest design that fulfils a set of requirements. As no objective definition of network complexity exists, subjective measures are used to come to a conclusion. The resulting diverging views on what constitutes complexity subsequently lead to conflicts in design teams. While most people would agree that complexity is an important factor in network design, today's design decisions are made based on a rough estimation of the network's complexity, rather than a solid understanding.

The goal of this document is to define a framework for network complexity research. This framework describes related research and current understanding of the topic, as well as outlining some ways research could be taken forward. Specifically, contributions are invited in all of the areas mentioned.

Many references to existing research in the area of network complexity are listed on the Network Complexity Wiki [wiki]. That wiki also contains background information on previous meetings on the subject, previous research, etc.

2. Current Understanding of Network Complexity

2.1. The Behavior of a Complex Network

While there is no generally accepted definition of network complexity, there is some understanding of the behavior of a complex network. It has some or all of the following properties:

- o Self-Organization: A network runs some protocols and processes without external control; for example a routing process, failover mechanisms, etc. The interaction of those mechanisms can lead to

a complex behaviour.

- o Un-predictability: In a complex network, the effect of a local change on the behaviour of the global network may be unpredictable.
- o Emergence: A network has an emergent property if a small local change produces a large scale, seemingly unrelated state or result.
- o Non-linearity: An input into the network produces a non-linear result.
- o Fragility: A small local input can break the entire system.

2.2. Robust Yet Fragile

Networks typically follow the "robust yet fragile" paradigm: They are designed to be robust against a set of failures, yet they are very vulnerable to other failures. Doyle [Doyle] explains the concept with an example: The Internet is robust against single component failure, but fragile to targeted attacks. The "robust yet fragile" property also touches on the fact that all network designs are necessarily making trade-offs between different design goals. The simplest one is articulated in "The Twelve Networking Truths" RFC1925 [RFC1925]: "Good, Fast, Cheap: Pick any two (you can't have all three)." In real network design, trade-offs between many aspects have to be made, including, for example, issues of scope, time and cost in the network cycle of planning, design, implementation and management of a network platform.

2.3. The Complexity Cube

Complex tasks on a network can be done in different components of the network. For example, routing can be controlled by central algorithms, and the result distributed (e.g., OpenFlow model); the routing algorithm can also run completely distributed (e.g., routing protocols such as OSPF or ISIS), or a human operator could calculate routing tables and statically configure routing. Behringer [Behringer] defines these three axes of complexity as a "complexity cube" with three axes: Network elements, central systems, and human operators. While different functions can be shifted between these axes of the network, the overall complexity may change.

3. Towards Defining Network Complexity

3.1. General Observations

Any analysis of practical network complexity must take a wide range of parameters into account, also parameters which are hard to measure, for example the human element. Human error constitutes in most cases of critical outages the trigger condition; therefore any analysis ignoring the human factor cannot address the full picture. [insert a reference that 70%(?) of critical outages have a human origin]

3.2. The Problem Space

When discussing network complexity, a large number of influencing factors have to be taken into account to arrive at a full picture, for example:

- o State in the network: Contains the network elements, such as routers, switches (with their OS, including protocols), lines, central systems, etc. The number and algorithmical complexity of the protocols on network devices for example.
- o Human operators: Complexity manifests itself often by a network that is not completely understood by human operators. Human error is a primary source for catastrophic failures, and therefore must be taken into account.
- o Classes / templates: Rather than counting the number of lines in a configuration, or the number of hardware elements, more important is the number of classes from which those can be derived. In other words, it is probably less complex to have 1000 interfaces which are identically configured than 5 that are completely different configured.
- o Dependencies and interactions: The number of dependencies between elements, as well as the interactions between them has influence on the complexity of the network.
- o TCO (Total cost of ownership): TCO could be a good metric for network complexity, if the TCO calculation takes into account all influencing factors, for example training time for staff to be able to maintain a network.
- o Benchmark Unit Cost is a related metric that indicates the cost of operating a certain component. If calculated well, it reflects at least parts of the complexity of this component. Therefore, the way TCO or BUC are calculated can help to derive a complexity metric.
- o Churn / rate of change: The change rate in a network itself can contribute to complexity, especially if a number of components of the overall network interact.

Networks differ in terms of their intended purpose (such as is found in differences between enterprise and public carriage network platforms, and in their intended role (such as is found in the differences between so-called "access" networks and "core" transit networks). The differences in terms of role and purpose can often lead to differences in the tolerance for, and even the metrics of, complexity within such different network scenarios. This is not necessarily a space where a single methodology for measuring complexity, and defining a single threshold value of acceptability of complexity, is appropriate.

3.3. Technical Debt

Many changes in a network are made with a dependency on the existing network. Often, a suboptimal decision is made because the optimal decision is hard or impossible to realise at the time. Over time, the number of suboptimal changes in themselves cause significant complexity, which would not have been there had the optimal solution been implemented.

The term "technical debt" refers to the accumulated complexity of sub-optimal changes over time. As with financial debt, the idea is that also technical debt must be repaid one day by cleaning up the network or software.

3.4. Layering considerations

In considering the larger space of applications, transport services, network services and media services, it is feasible to engineer responses for certain types of desired applications responses in many different ways, and involving different layers of the so-called network protocol stack. For example, quality of Service could be engineered at any of these layers, or even in a number of combinations of different layers.

Considerations of complexity arise when mutually incompatible measures are used in combination (such as error detection and retransmission at the media layer in conjunction with the use TCP transport protocol), or when assumptions used in one layer are violated by another layer. This results in surprising outcomes that may result in complex interactions. This has lead to the perspective that increased layering frequently increases complexity [RFC3439].

While this research work is focussed network complexity, the interactions of the network with the end-to-end transport protocols, application layer protocols and media properties are relevant considerations here.

4. Possible Directions of Research

The problem space of network complexity is very large, as many influencing factors contribute to the overall complexity of a network. The following sections outline areas for research.

4.1. Definitions and Metrics

In the context of general network operations, as well as in the context of standardisation of protocols a common definition of the term "network complexity" would be useful. It would also be useful to have a metric for the complexity of a protocol or network design, such that two candidate proposals can be objectively compared. This could happen in a bottom-up approach, where metrics for parts of a network are combined to an overall metric; or in a top-down approach where a global metric or vector of metrics is broken down into the components of a network.

For example, one such approach to a complexity metric is described in [Chun]:

"A complexity metric could make these arguments objective. A good metric would be based on quantifiable, concrete measurements of the system properties that induce implementation difficulties, complex interactions and failures, and so forth. Many metrics are possible. A perfect metric would be intuitive and easy to calculate, and would correlate with other, more subjective metrics, such as lines of code or system designers' experience. "We build on the observation that much of system design centers on issues of state - the required state must be defined and operations for constructing and using it must be developed - but in distributed systems, one state can derive from states stored on other nodes. To calculate its state, a node must hear from the remote nodes that store the dependencies. This adds additional dependencies on the network and intermediate node states required to relay input states to the node in question. Thus, not only are a given piece of state's dependencies distributed, there are also more of them. "We conjecture that the complexity particular to networked systems arises from the need to ensure state is kept in sync with its distributed dependencies."

4.2. Comparative Analysis

In the foreseeable future it is unlikely to define a single, objective metric that includes all the relevant aspects of complexity. In the absence of such a global metric, a comparative approach could be easier.

For example, if two network architectures are compared against each other, it may be possible to ignore the network layout and device hardware if those are the same in both cases. In such specific comparisons it should be considerably easier to find valid metrics, and to compare the approaches objectively.

4.3. Containment, Control or Reduction of Complexity

In some disciplines such as software engineering, complexity is relatively well understood, as well as metrics and methods to reduce it. Such approaches can be applied in the networking industry to achieve the same result.

4.4. Use Cases

While it is hard to define a universal set of metrics for network complexity, special use cases should be documented to serve as examples, and to stimulate discussion. Such use cases could come out of different areas:

- o Documented examples of "catastrophic failure": While the cause of complexity is hard to understand, the result may be a catastrophic outage, which can be reverse-engineered to understand the root causes. The knowledge from this process may give insight into root causes of complexity.
- o A detailed complexity analysis of a particular network or protocol. Even if this analysis may not be complete or fully objective, it would be useful to learn about different approaches.
- o Analysis of existing networks, protocols or components from an insider point of view, discussing in detail where the perceived complexity in the set-up is, and how this could be changed to reduce complexity.
- o Work in related areas, for example a detailed analysis of the total cost of ownership, and how this could be mapped into a complexity metric.

5. Security Considerations

This document does not discuss any specific security considerations.

6. Acknowledgements

The motivations and framework of this overview of studies into network complexity is the result of many meetings and discussions, with too many people to provide a full list here. However, key contributions have been made by: John Doyle, Jon Crowcroft, Mark Handley, Fred Baker, Paul Vixie, Lars Eggert, Bob Briscoe, Keith Jones, Bruno Klauser, Steve Youell, Joel Obstfeld.

The authors would like to acknowledge the contributions of Rana Sircar, Ken Carlberg and Luca Caviglione in the preparation of this Research Group document.

7. References

[Behringer]

Behringer, M.H., "Classifying Network Complexity",
Proceedings of the ACM Re-Arch'09, December 2009.

[Chun]

Chun, B-G., Ratnasamy, S. and E. Eddie, "NetComplex: A Complexity Metric for Networked System Design", 5th Usenix Symposium on Networked Systems Design and Implementation NSDI 2008, April 2008, <<http://berkeley.intel-research.net/sylvia/netcomp.pdf>>.

[Doyle]

Doyle, J.C., "The 'robust yet fragile' nature of the Internet", PNAS vol. 102 no. 41 14497-14502, October 2005.

[RFC1925]

Callon, R., "The Twelve Networking Truths", RFC 1925, April 1996.

[RFC3439] Bush, R. and D. Meyer, "Some Internet Architectural Guidelines and Philosophy", RFC 3439, December 2002.

[wiki] "Network Complexity Wiki", , <<http://networkcomplexity.org/>>.

Authors' Addresses

Michael H. Behringer
Cisco

Email: mbehring@cisco.com

Geoff Huston
Asia Pacific Network Information Centre

Email: gih@apnic.net