

PhD Forum: A Clean-Slate Framework for Enabling Centralized Control and Management in Data Networks

Hammad Iqbal
University of Pittsburgh, USA
hiqbal@cs.pitt.edu

Abstract—Our research provides a framework for the design of clean-slate network control and management plane using the abstraction of 4D architecture. We utilize and extend 4D’s concept of a logically centralized Decision plane that is responsible for managing network-wide resources, and present mechanisms for network operators to centrally plan, control, and manage their networks while ensuring network robustness and scalability. Our work is especially relevant in the context of ISPs and large geographically dispersed enterprise networks where stringent demands on network availability and reliability, coupled with the prevalent problems in large-scale distributed network management, make the task of network management a difficult and costly challenge.

I. INTRODUCTION AND MOTIVATION

Internet was not designed for the widespread global deployment that it is used for today. Internet which started with a few connected nodes in the 1970s currently serves over a billion users around the globe, with millions of network devices used to access and run the network. This exponential growth has outpaced the network control and management tools available to the Internet designers and operators, leaving the present day Internet difficult to maintain.

The root cause of the difficulty in network management is the fact that the data networks are controlled by a variety of distributed routing algorithms (e.g. OSPF, IS-IS, BGP, etc.), each working independently to achieve some network-wide objective, while operating collectively on diverse physical network devices. This has created a situation where each network functionality (e.g. inter-domain and intra-domain routing) maintains a distinct state across many different physical devices and is governed by its own set of configuration rules and protocol logic, making it extremely difficult to control their interactions. Consequently, the management of typical data networks requires extensive manual configuration of individual protocol parameters, leaving the networks fragile [1], [2] and insecure [3].

Effective control and management of networks is especially a challenge for large and geographically dispersed networks such as large ISP and enterprise networks where it is important to efficiently manage the network resources across a large number of geographically dispersed and heterogeneous network devices, while meeting strict constraints on network availability and reliability.

Incremental solutions to overcome this limitation, including the use of better management tools, has been ineffective as they try to match the pace of changes in various device operations and technical advances. Additionally, newer services and objectives beyond best-effort routing places new demands on network control algorithms that are difficult to meet in the presence of intricate inter-dependencies and distributed state and logic. This often necessitates the error prone and laborious process of indirectly inducing desired behavior in dynamic protocol operation through static configurations, e.g. traffic engineering [4].

Centralization of network control and logic is an alternative approach in tackling the challenge of management complexity. The main motivation behind this approach is the reduction in complexity from decoupling control functionality from data forwarding devices and using centralized algorithms for network control instead of distributed implementations of the same. Centralization of network decision-making naturally allows for simpler implementations and provides a single point of network interface for management and policy specification. This can be a significant improvement in large network management where individual control of thousands of network devices is a very difficult and costly task. This approach of network control centralization is the basic tenant of the 4D architecture [1]. The 4D architecture advocates a new layering design of the IP networks that separates the task of packet forwarding, a data layer function, from the task of network control, an operation and management function. This separation of data and control layers is in contrast with the current practice where the data forwarding mechanism and control logic are intertwined inside monolithic network devices, such as network routers or switches.

We propose a framework for reducing the management complexity in data networks, using a centralized design for network management. By extending 4D architecture’s vision of centralized Decision and management to a *logically centralized* Decision plane, our framework improves the reliability of the architecture and provides a more robust and scalable design. This work is especially relevant to large enterprise and ISP networks where reliable operation and effective control over network resources is currently a significant challenge. Successful implementation and deployment of proposed de-

sign holds the promise of drastically improved support for network management, lower network operation and management (O&M) costs, increased network reliability, and lower cost of network devices.

II. FRAMEWORK FOR SCALABLE AND ROBUST DECISION PLANE DESIGN

Our research focuses on the design of logically centralized clean-slate Decision plane using the 4D clean-slate network paradigm as the basis for developing an efficient, robust, and reliable network control architecture. The following objectives shape our framework and form the basis of this research:

- **Scalability:** The control and management tools currently used by the network managers become unwieldy as the size of the managed network increases. This suggests that a new control architecture must be scalable in network size, both geographically and numerically (in terms of the number of managed devices)
- **Robustness:** One of the foremost concerns about a centralized control architecture is its ability to recover from failures in contrast with the distributed alternative. Therefore, a new centralized architecture must be carefully designed to achieve similar overall robustness as found in the distributed architecture of the current Internet. Specifically, it should be dynamically adaptable to failures in the managed devices as well as within the control devices themselves.
- **Optimality:** Any acceptance of the new design, and its success, would depend on its comparative performance with the present Internet. This implies that the planned architecture would need to support operations at similar or better rate than what is currently achievable. Moreover, desirable features and operations that are not offered by the current control and management tools would increase its utility and help in its acceptance.

A physically centralized Decision plane design was investigated in [5], [6] where replication of Decision Elements (DE) was used to ensure Decision plane robustness to DE failures. However complete centralization of control logic invokes questions about the robustness of design in the face of failures; even more so since the distributed nature of the Internet was a design choice to prevent its failure due to any localized event. An alternative design approach is where the logical Decision plane is distributed over physically independent DEs. In this design, each DE controls a subset of the whole network, and works collaboratively with other DEs to achieve overall network control. We believe this addition of distributed structure is necessary to make the 4D architecture scalable with the size of the network, as well as in making it more robust to DE failures.

Realizing our architectural vision involves solving several research problems in optimal architecture design, protocol development, and performance evaluation. Specifically, we need to investigate the physical design of logically centralized Decision plane, fault-tolerance mechanisms that allow

rapid and smooth *fail-over*, protocols that will govern system operation, and evaluate the performance against what is achievable in the current Internet. Moreover, routing strategies beyond the currently available shortest-path routing need to be investigated in context of the centralized control structure of 4D architecture. We explore the motivation of the identified problems in the following sections.

A. Physical Design of the Distributed 4D Architecture

An efficient physical design of Decision plane is necessary for the efficient operation of logically centralized 4D architecture. The physical design includes the number, placement, and connectivity of DEs within an Autonomous System (AS), as well as the connectivity between switches and DEs. Given the crucial task of the Decision plane, an optimal physical design is necessary to avoid problems associated with higher response time to failures and convergence delays. We presented a design of Decision plane connectivity in [7], where we used optimization techniques to determine a placement of DE that minimizes Decision plane response time.

B. Design of Fault-tolerance Mechanisms

The protocols and algorithms that govern today's Internet operation were designed to operate in a distributed fashion to increase the resilience of the network to any localized failure, even keeping the network operable in case of nuclear attacks as it's early specifications suggested. This emphasis on network resilience and survivability has been a main feature of the Internet design and needs to be preserved in any new architecture. In the case of logical centralization of control logic, it thus becomes important to ensure that the Decision plane efficiently handles failures of data plane devices (e.g. switches), physical links between various network components, and failures of DEs. Since a number of data plane devices will be connected to each DE, it will be necessary to reassign the devices in DE failure scenario. This reassignment of devices needs to be done fast enough so that the network operation is not interrupted, implying a dynamic fail-over design where the Decision plane automatically reconfigures itself in case of failures. In [8], we present an adaptive algorithm that dynamically assigns data plane devices to DEs in the network, while ensuring minimum convergence delay and optimal load balancing at the Decision plane.

C. Design of the Decision Plane Protocol

The Decision plane operation will be governed by a protocol running distributively over DEs, where this protocol will be responsible for the management of DEs to provide a uniform network-wide Decision plane. We note that in order to effectively provide network-wide control and meet the design goals, protocol design needs to conform to the basic functional requirements of robustness to multiple failures, minimum pre-configuration, extensibility, and scalability. Our design of DPP protocol in [8] provides a uniform interface to the managed devices while implementing control logic distributively over the Decision plane.

D. Optimal Routing in Distributed 4D Architecture

Routing algorithms in the current Internet utilize shortest-paths for sending traffic in intra-domain routing. Shortest-path routing strategy, while simple enough for distributed computation, provides non-optimal delays and is considered a handicap in situations where finer control of traffic is required. This has an effect of making it difficult for network managers to effectively engineer their networks. Routing strategy that minimizes the total traffic delay optimizes the network for its intended purpose but is difficult to deploy distributively as non-stationary traffic flows make it very hard to solve the resource-intensive computations needed for online network optimization. The centralization of control logic holds promise of enabling such network optimization as it provides a platform for monitoring the network for slowly changing traffic flows and using that information to optimize the network accordingly. Moreover, planned maintenance events also provides such an opportunity where a centralized Decision plane can move traffic flows from paths and devices that will be affected by the anticipated event. We propose to develop routing techniques that will utilize the network-wide information available at the Decision layer to optimize the total delay of the network.

III. RESEARCH CHALLENGES

There are several open research challenges in our proposed work. The use of routing strategies beyond shortest-path would require the use of a mix of theoretical and experimental techniques. Although there is vast existing literature in this area, mechanisms need to be devised that would allow a distributed set of DEs to solve multi-objective optimization problems that necessarily arise in minimum-delay routing. Planned maintenance and long-term network-wide traffic patterns also present opportunities for network optimization that need further refinement and analysis.

Similarly, there are several obstacles in experimental evaluation of our framework. The first one being the absence of any simulation tool, in our knowledge, that can be directly used to model and simulate the Decision plane operation. We utilized the ns-2 platform in our initial evaluation by building new modules that enable us to evaluate the performance of our Decision plane design. Our evaluation has been restricted by the limited topological data that publicly exists about large real-world ISP and enterprise networks. We plan to investigate the use of other tools, especially large-scale networking testbeds e.g. GENI. This would particularly be useful in our design of minimum-delay routing as real topological and traffic data would allow accurate modeling and realistic investigations.

IV. CONCLUSION

We presented the vision of a robust and scalable control and management framework, within the context of 4D architecture, that holds promise of mitigating the problems associated with the management of large networks. The various building blocks of our design were identified along with the initial results. Some of hurdles that we anticipate to cross in reaching our objectives were also noted. It is our hope that this

framework would enable better control and lower operational costs for the network operators, and would open new avenues of research and innovation in network management.

V. RELATED WORK

Several recent studies have embraced centralization of network logic as a way of overcoming management complexity or providing new services that are presently difficult to implement. Greenberg et al. [1] provide a comprehensive survey of the issues in network control and management, and propose the architectural vision embraced and extended in this paper.

Centralized control has been explored in BGP design where RCP [9], [10] was proposed as a logically centralized point for computing BGP routes and improving the scalability of large networks. However, RCP is limited to BGP route computation and does not extend to Interior Gateway Protocol (IGP) routes.

Recently, CONMAN [11] utilized the concept of management plane and centralization in the design and operation of “network managers” that are used to manage the protocols running on individual routers.

Several efforts in open router design [12], [13] have also advocated migration of control functions away from routers to reduce their complexity, where they utilize “control elements” for the implementation of distributed network algorithms, and design protocols to enable communication between different network elements. In contrast, our work uses 4D’s approach of network-wide decision making and presents a robust and scalable design for the Decision plane that implements router assignment, routing table computation, and failure recovery tasks.

REFERENCES

- [1] A. Greenberg *et al.*, “A clean slate 4D approach to network control and management,” *SIGCOMM CCR.*, vol. 35, no. 5, 2005.
- [2] D. A. Maltz, G. Xie, J. Zhan, and H. Zhang, “Routing design in operational networks: A look from the inside,” in *Proc. ACM SIGCOMM.* ACM Press, 2004.
- [3] A. Wool, “A quantitative study of firewall configuration errors,” *IEEE Computer*, vol. 37, no. 6, 2004.
- [4] B. Fortz, J. Rexford, and M. Thorup, “Traffic engineering with traditional IP routing protocols,” *IEEE Communications Magazine*, vol. 40, 2002.
- [5] A. Greenberg *et al.*, “Refactoring network control and management: A case for the 4D architecture,” Carnegie Mellon University, Tech. Rep. CMU-CS-05-117, Sept 2005.
- [6] H. Yan *et al.*, “Tesseract: A 4D network control plane.” in *USENIX NSDI*, 2007.
- [7] H. Iqbal and T. Znati, “Distributed control plane for 4D architecture,” in *Globecom 2007.* IEEE, 2007.
- [8] —, “Design of clean-slate network control and management plane,” University of Pittsburgh, Tech. Rep. PITT-CS-TR-09-168, August 2009.
- [9] N. Feamster *et al.*, “The case for separating routing from routers,” in *ACM SIGCOMM FDNA Workshop*, 2004.
- [10] M. Caesar *et al.*, “Design and implementation of a routing control platform,” in *USENIX NSDI*, 2005.
- [11] H. Ballani and P. Francis, “CONMan: a step towards network manageability,” *SIGCOMM CCR*, vol. 37, no. 4, 2007.
- [12] T. V. Lakshman *et al.*, “The SoftRouter architecture,” in *HotNets-III*, 2004.
- [13] L. Yang, R. D. T. Anderson, and R. Gopal, “RFC 3746 Forwarding and Control Element Separation Framework,” 2004.