

# Verification of Programmable Networks



**Dragoș Dumitrescu**

Advisor: Prof. Dr. Ing. Dragoș Niculescu

Facultatea de Automatică și Calculatoare  
Universitatea Politehnica București

This dissertation is submitted for the degree of  
*Doctor of Philosophy*

Bucharest

May 2022



## Summary

**Keywords:** networks, model-checking, verification, P4, SDN, BGP, network verification, network correctness, cloud networking.

Computer networks are at the center of modern day lives, they are the building blocks of all compute systems. As such, they need to be highly reliable and bug-free to ensure the correct functioning of safety and mission critical applications. The advent of network data- and control-plane programmability as key enablers of multi-tenant cloud infrastructures has made networks more and more complex; it is thus increasingly difficult to manually assess their correctness. In order to tackle this problem, a new area of research – network verification – has seen significant success in the past decade.

In this dissertation, I will report on my research in making network verification algorithms practical in modern programmable networks. Practicality is explored in two flavors: ease of use and performance. In what ease of use is concerned, I explore implicit correctness properties which must hold for all networks under consideration, thus requiring a *zero-specification* effort from administrators. In what performance is concerned, the main requirement is to have verification tools scale reasonably well to large public cloud-provider networks, which may be comprised of thousands of routers, servers and network appliances.

The contributions of this thesis are positioned at the layers of programmability involved in a typical cloud network environment and show how my work tackles some of the hurdles encountered at each layer. Dataplane programming (e.g., P4) deals with programming forwarding devices; software defined networking (e.g. OpenFlow) handles the orchestration of multiple network devices while control plane configuration (e.g., routing policies) handles dissemination of control-plane information across the entire provider network. My research answers questions at all these layers. First of all, my contributions to Vera (Stoenescu et.al., *Debugging P4 programs with Vera* in SIGCOMM'18), Vera2 (Dumitrescu et.al. Dataplane verification for P4 presentation in NetPL'19) and bf4 (Dumitrescu et.al., *bf4: Towards bug-free P4 programs* in SIGCOMM'20) show how detecting bugs in P4 programs can be made practical and easy to use. Secondly, netdiff (Dumitrescu et.al. *Equivalence and its applications to network dataplanes* in NSDI'19) is an algorithm which employs network dataplane equivalence to find

bugs in SDN controllers. My third contributions makes BGP simulation and model-checking efficient for production-scale cloud networks.

**Contributions summary.** The positioning of my work in the landscape of network verification hints at the prominent research endeavor tackled within my thesis: "finding efficient and scalable decision procedures in the context of complex and programmable network policies".

The most important principle my work is driven by consists of *ease of specification* or even *zero-specification*. My efforts in both programmable network dataplane and SDN controller verification use implicit notions of correctness and tailor decision algorithms based on their specifics.

My work in control plane verification builds upon tools which are innately customized checks meant to increase scalability. Thus, they obey the same previously advertised *zero-specification* principle. While this design decision is not part of my contribution, the approach is nonetheless worth mentioning as an alternative to existing forms of *zero-specification*. Given the properties under verification, my contributions focus on scaling BGP simulation algorithms to production networks. This is the second principle behind my work: building verification solutions which *scale to large networks*.

The main research question of this thesis can be summarized as follows: "Can one efficiently verify the correctness of programmable networks with little (preferably no) specification effort from network users?".

**Note on previously published material.** Some of the chapters in this thesis contain excerpts of already published scientific papers. The contents of sections 4.1 and 2.6.2.8 will undergo submission for publication in the near future.

**Thesis outline.** Starting from the lowest layer in the OSI model, in Chapter 2, I report on my research in P4 programmable dataplane verification. The focus lies on zero-specification verification of undefined behaviors in P4 dataplanes. Section 2.5 describes a new algorithm to infer runtime filters or produce code fixes in such way as to guarantee no bugs are present at runtime. Section 2.6 offers a detailed design and implementation report of P4 verification tools.

In chapter 3, I dive into network dataplane equivalence, another implicit form of specifying network correctness. The observation is that virtual networks are often regarded from distinct perspectives – what the tenant had intended and what the provider has actually implemented. The gap between these perspectives is bridged by a highly complex piece of distributed middleware which is prone to bugs and inconsistencies. The aim of this paper is to show correctness by checking equivalence between the two perspectives. Using this method, we have come across rare bugs in Neutron – OpenStack’s networking driver.

Chapter 4 presents my contributions to network control-plane simulation in large cloud networks. In section 4.1, I focus on the methods to make BGP simulation in DC networks scale to large production DCs.

Chapter 5 draws the conclusions of this dissertation and relates them to the research questions initially formulated. It also provides insights into research avenues which I find interesting to explore in the future.



# Table of Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Programmability in cloud networks . . . . .	2
1.2	Finding network bugs . . . . .	7
1.3	Contributions summary . . . . .	10
1.4	Network Verification Background. . . . .	11
<b>2</b>	<b>P4 Verification</b>	<b>15</b>
2.1	A P4 primer . . . . .	15
2.2	Bugs in P4 programs . . . . .	18
2.3	Debugging P4 programs with Vera . . . . .	19
2.3.1	Translating P4 to SEFL . . . . .	19
2.3.2	Evaluation . . . . .	23
2.4	Vera2 – Verification Condition Generation for P4 . . . . .	26
2.5	bf4 – towards bug-free P4 programs . . . . .	28
2.5.1	Introduction . . . . .	28
2.5.2	Motivation . . . . .	30
2.5.3	Solution space . . . . .	32
2.5.4	bf4 overview . . . . .	34
2.5.5	Evaluation . . . . .	48
2.5.6	Related work . . . . .	54
2.5.7	Conclusions . . . . .	54
2.6	Anatomy of a P4 Verification Engine . . . . .	55
2.6.1	Introduction . . . . .	55
2.6.2	The P4 verification ingredients . . . . .	56
2.6.3	Reachability with packet processing . . . . .	72
2.6.4	Takeaways and next steps . . . . .	78

---

<b>3</b>	<b>Equivalence and its applications to network dataplane verification</b>	<b>81</b>
3.1	Introduction . . . . .	81
3.2	Goals . . . . .	83
3.3	Approaches to checking equivalence . . . . .	85
3.3.1	Existing solutions fall short . . . . .	86
3.4	Dataplane equivalence with netdiff . . . . .	87
3.4.1	Dataplane symbolic execution . . . . .	89
3.4.2	Equivalence between pathsets . . . . .	90
3.4.3	Correctness and complexity . . . . .	92
3.5	Correctness of netdiff . . . . .	93
3.6	Implementation . . . . .	95
3.6.1	OpenStack Neutron Integration . . . . .	96
3.7	Evaluation . . . . .	98
3.7.1	Neutron bugs . . . . .	98
3.7.2	Checking a large Neutron deployment . . . . .	103
3.7.3	P4 equivalence . . . . .	104
3.8	Related work . . . . .	105
3.9	Conclusions . . . . .	106
<b>4</b>	<b>Scaling BGP verification</b>	<b>109</b>
4.1	Making datacenter network simulation fast . . . . .	110
4.1.1	Introduction . . . . .	110
4.1.2	Illustration . . . . .	112
4.1.3	Algorithms . . . . .	120
4.1.4	Evaluation . . . . .	132
4.1.5	Related Work . . . . .	133
4.1.6	Conclusion . . . . .	134
<b>5</b>	<b>Conclusions</b>	<b>137</b>
5.1	Future work . . . . .	139