Chair of Communication Networks
Department of Electrical and Computer Engineering
Technical University of Munich

Adversarial Network Benchmarking: A Data-Driven Approach

Andreas Blenk (TU Munich, University of Vienna)°*

Joint work with:

Johannes Zerwas*, Patrick Kalmbach*, Laurenz Henkel*, Sebastian Lettner, Gábor Rétvári^, Wolfgang Kellerer*, Stefan Schmid°

**Technical University of Munich, Germany*

^Budapest University of Technology and Economics, Hungary

°Faculty of Computer Science, University of Vienna, Austria

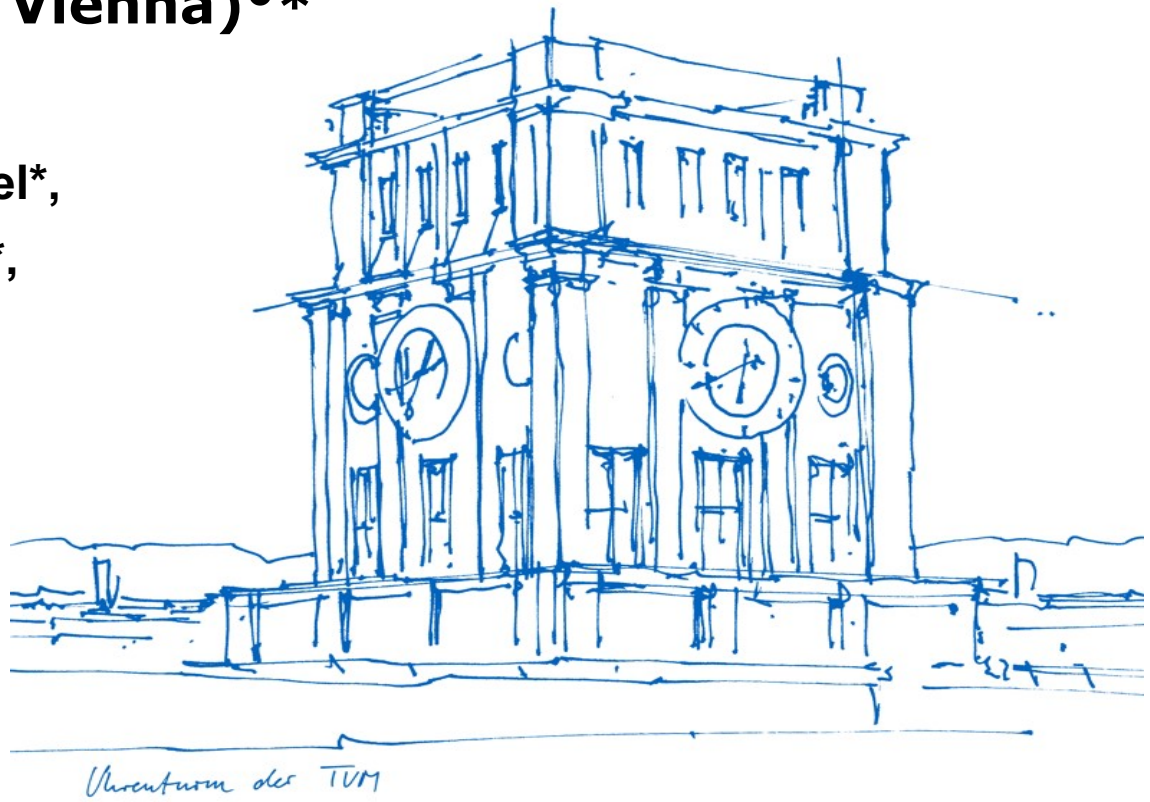

MIR^3 – September 2020 in Garching

Self-***Driving*** and Data ***Driven*** Networks?

It is not about cars!

It is about net… me… co… …olves!

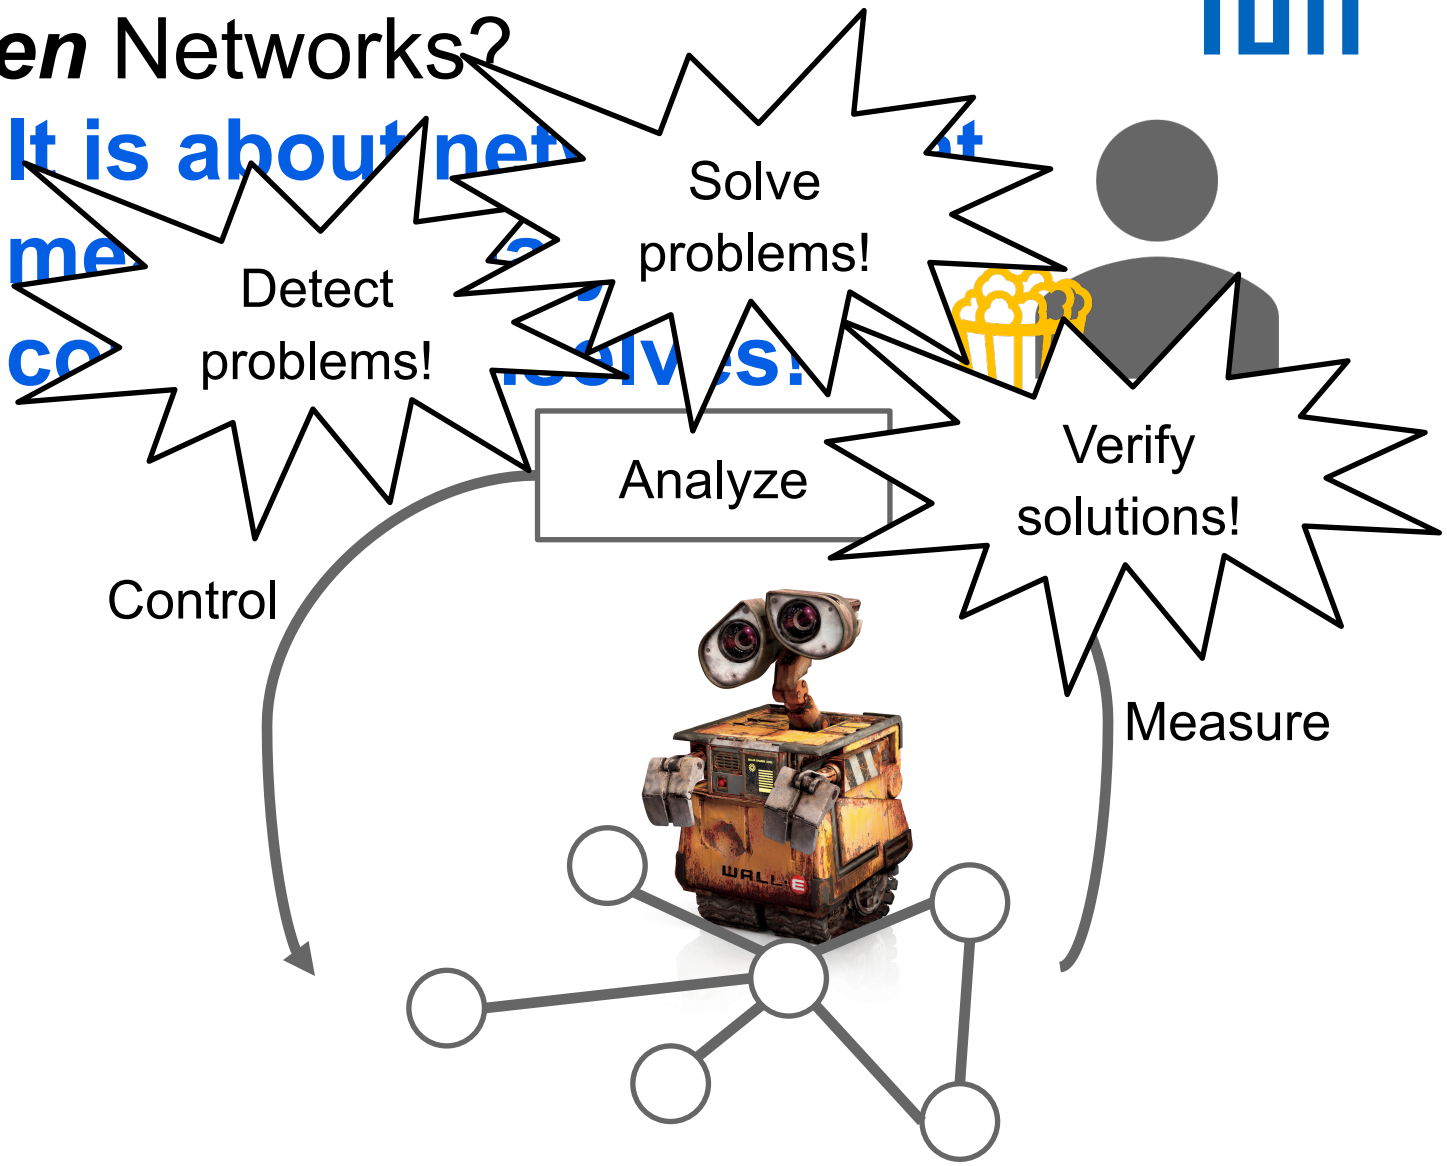

N. Feamster and J. Rexford, “Why (and How) Networks Should Run Themselves,” CoRR, vol. abs/1710.11583, 2017.

But Data-Driven Systems Can be Tricked

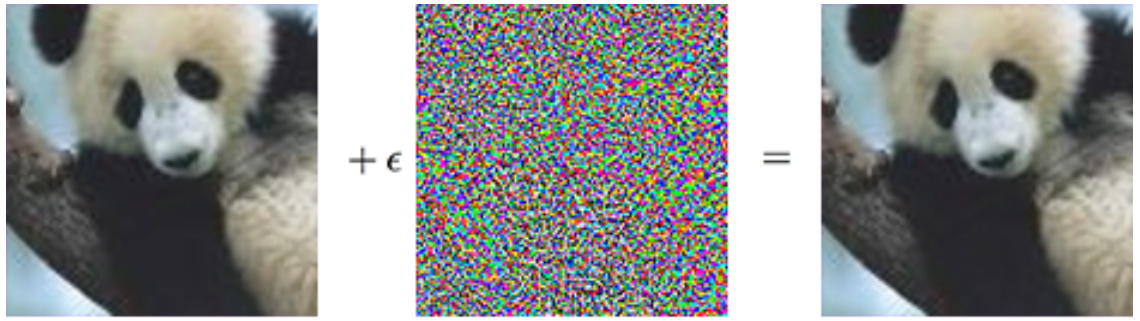

In typical ML applications

(Self) Driving Under the Influence:
Intoxicating Adversarial Network Inputs

Roland Meier[1], Thomas Holterbach[1], Stephan Keck[1], Matthias Stähli[1], Vincent Lenders[2], Ankit Singla[1], Laurent Vanbever[1]

ACM HotNets 2019

[1] ETH zürich
[2] Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
armasuisse

… and in networking

Source: https://openai.com/blog/adversarial-example-research/

Adversarial Input Only Critical for Machine Learning?

Adversarial Input

Machine Learning-based solution

Adversarial Input

Solution designed by human

Why?

… but this is also true for existing solutions by human!

Adversarial input is not only critical for self-driving networks …
It's already a problem!

Benchmarking Network Algorithms, Architectures etc…

The Traditional Way …

Not always available

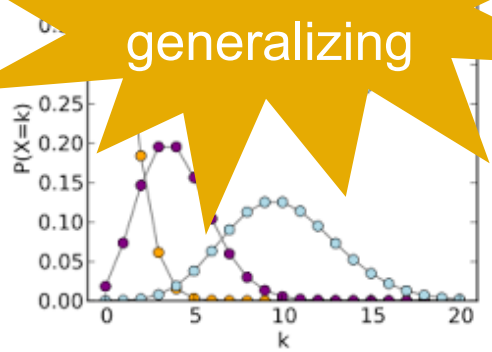

Not generalizing

Hmm… Biased?

Alternative opponent?

Traces | **Models** | **Human‘s Best Guesses** | **Data-Driven**

Our idea: Use ML to automatically find adversarial input to benchmark legacy and self-driving networks

Towards Automated Network Optimization and Design

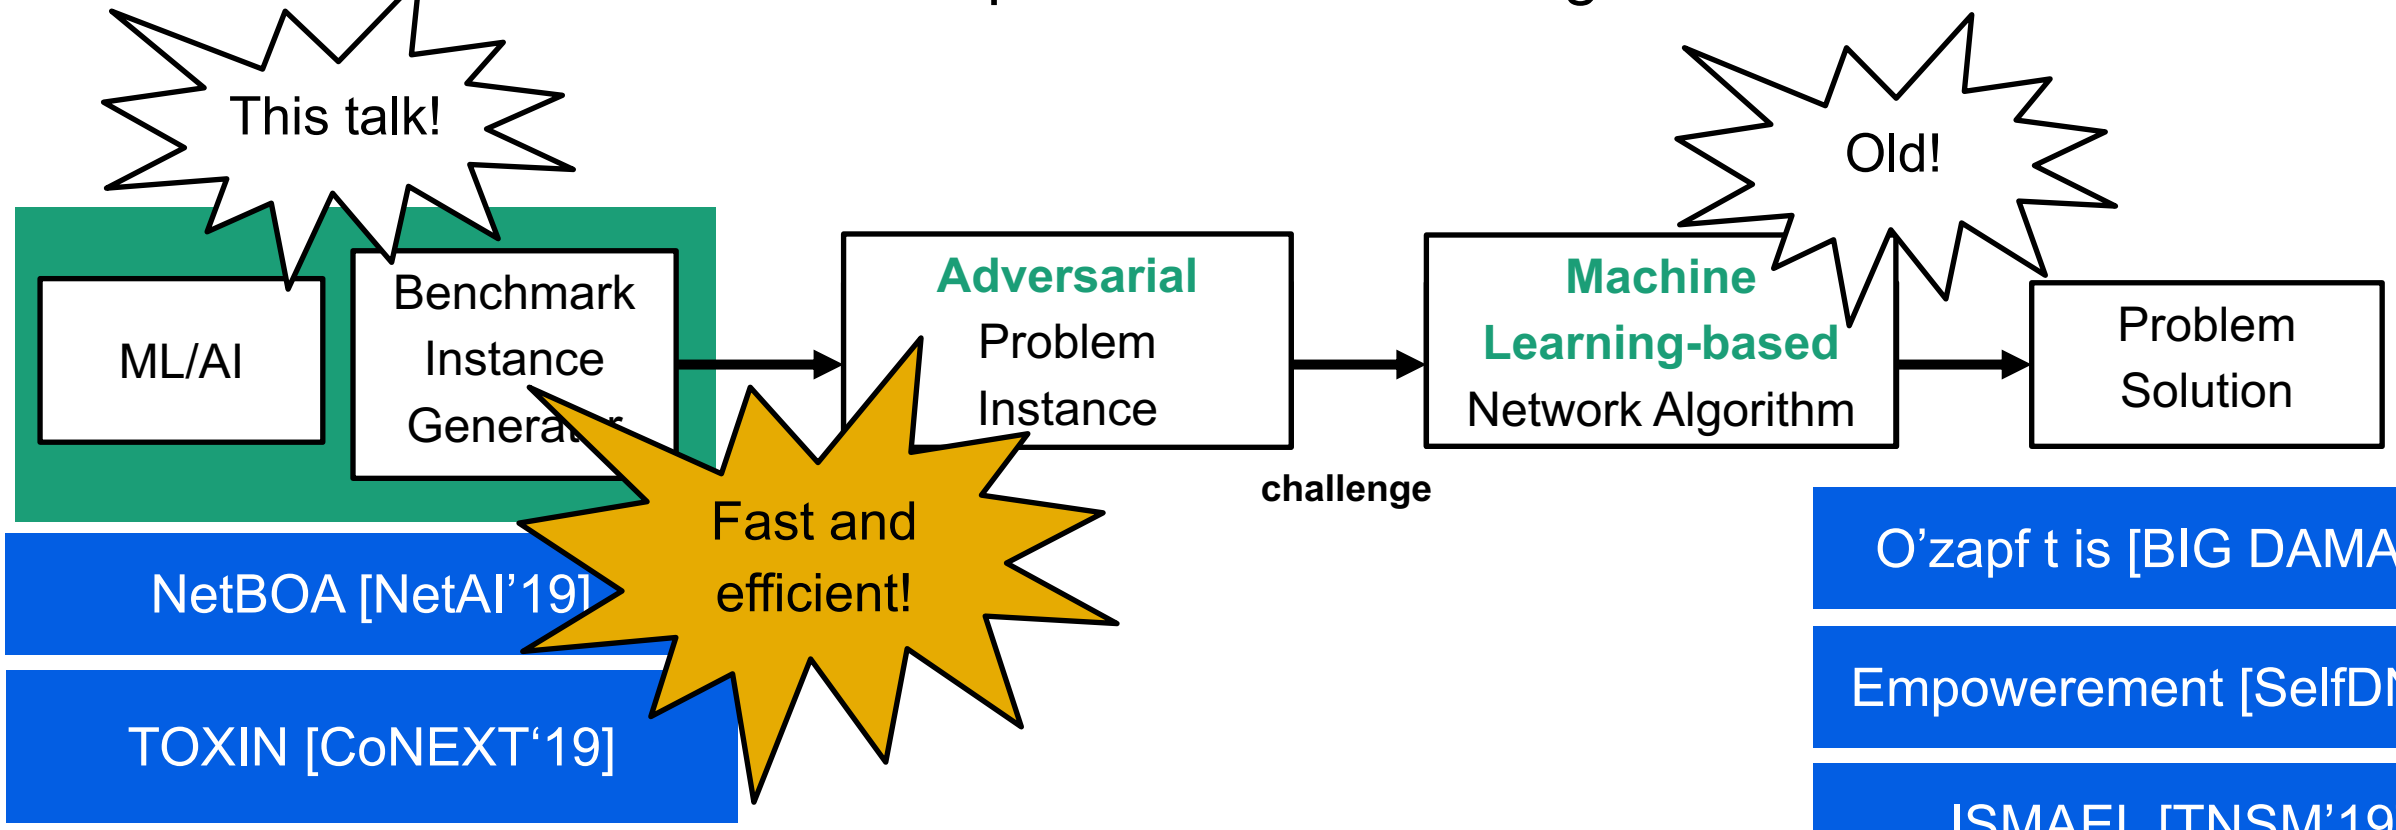

O'zapf t is [BIG DAMA'17]

Empowerement [SelfDN'18]

ISMAEL [TNSM'19]

The Traditional Way!

ML/AI vs ML/AI and Human

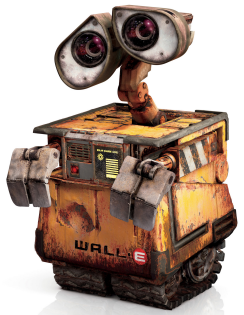

Data-Driven Adversarial Network Benchmarking in Data Centers: NetBOA

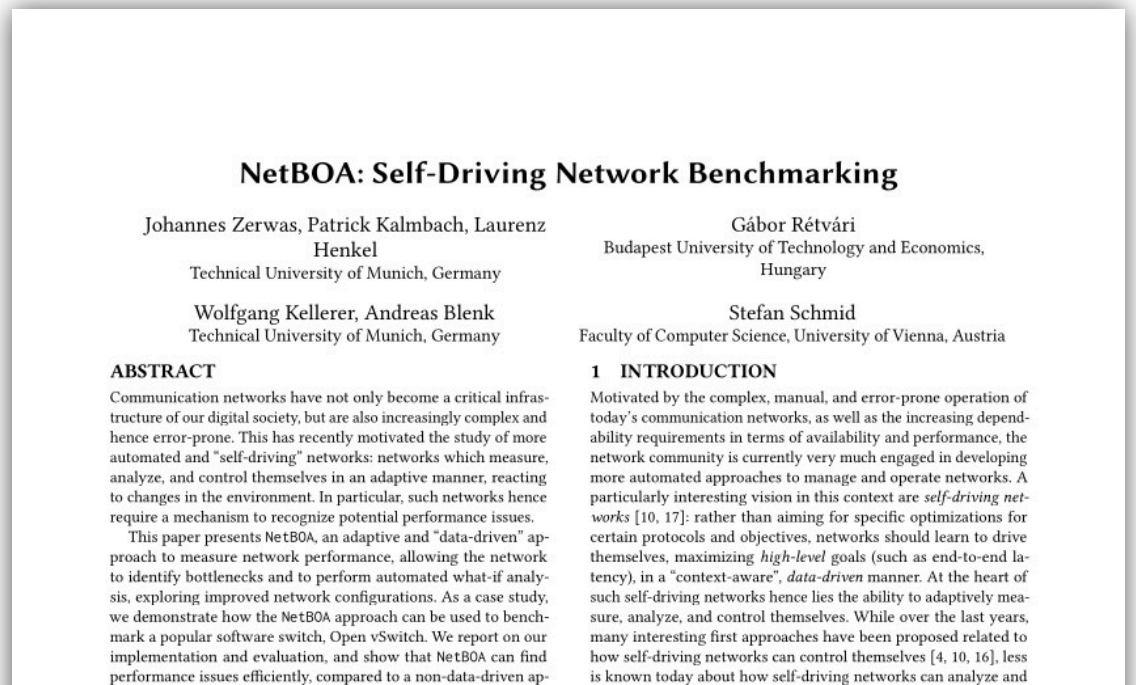

NetBOA: Self-Driving Network Benchmarking

Johannes Zerwas, Patrick Kalmbach, Laurenz Henkel
Technical University of Munich, Germany

Gábor Rétvári
Budapest University of Technology and Economics, Hungary

Wolfgang Kellerer, Andreas Blenk
Technical University of Munich, Germany

Stefan Schmid
Faculty of Computer Science, University of Vienna, Austria

ABSTRACT

Communication networks have not only become a critical infrastructure of our digital society, but are also increasingly complex and hence error-prone. This has recently motivated the study of more automated and "self-driving" networks: networks which measure, analyze, and control themselves in an adaptive manner, reacting to changes in the environment. In particular, such networks hence require a mechanism to recognize potential performance issues.

This paper presents NetBOA, an adaptive and "data-driven" approach to measure network performance, allowing the network to identify bottlenecks and to perform automated what-if analysis, exploring improved network configurations. As a case study, we demonstrate how the NetBOA approach can be used to benchmark a popular software switch, Open vSwitch. We report on our implementation and evaluation, and show that NetBOA can find performance issues efficiently, compared to a non-data-driven ap-

1 INTRODUCTION

Motivated by the complex, manual, and error-prone operation of today's communication networks, as well as the increasing dependability requirements in terms of availability and performance, the network community is currently very much engaged in developing more automated approaches to manage and operate networks. A particularly interesting vision in this context are *self-driving networks* [10, 17]: rather than aiming for specific optimizations for certain protocols and objectives, networks should learn to drive themselves, maximizing *high-level* goals (such as end-to-end latency), in a "context-aware", *data-driven* manner. At the heart of such self-driving networks hence lies the ability to adaptively measure, analyze, and control themselves. While over the last years, many interesting first approaches have been proposed related to how self-driving networks can control themselves [4, 10, 16], less is known today about how self-driving networks can analyze and

Use Cases of This Talk: Data Center Network Benchmarking

(1) Benchmarking Open vSwitch: NetBOA

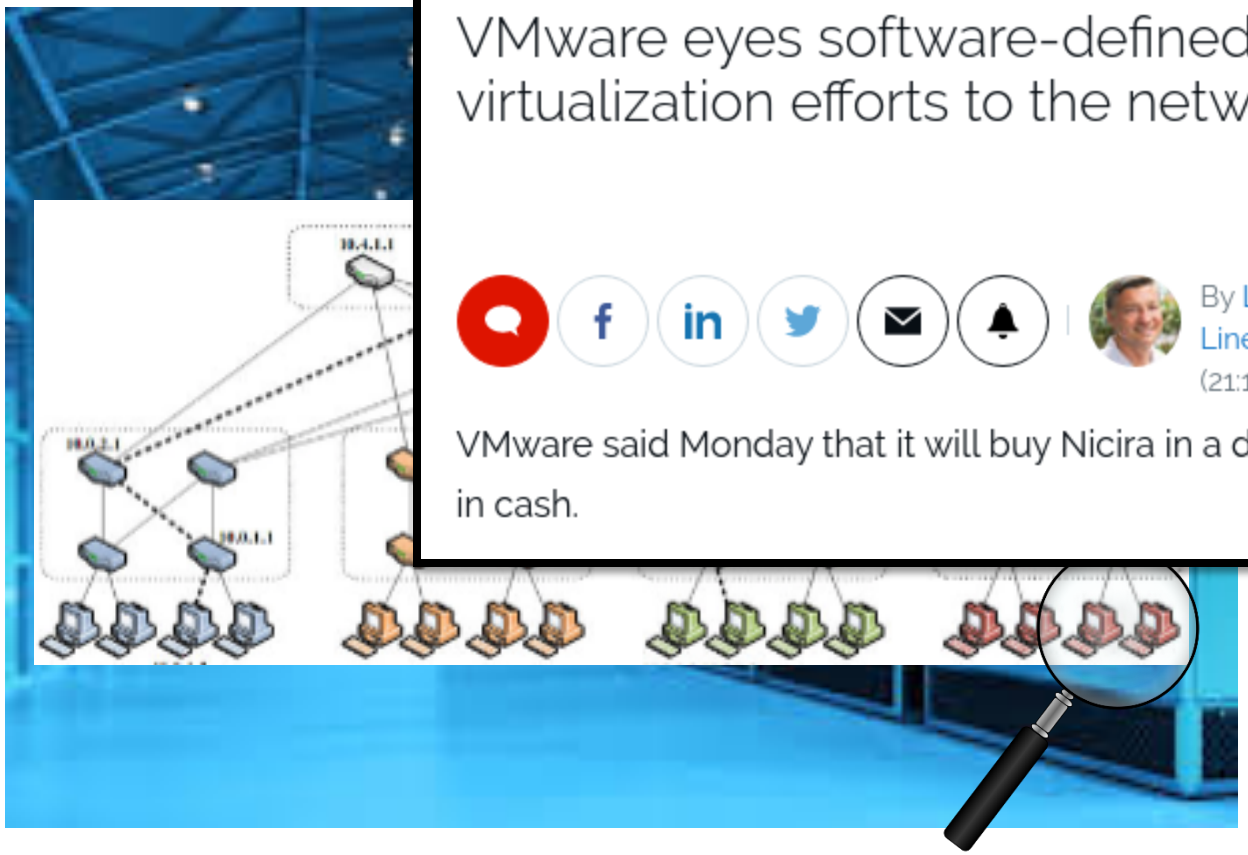

VMware buys Nicira for $1.05 billion

VMware eyes software-defined networking as it aims to take its virtualization efforts to the network.

By Larry Dignan for Between the Lines | July 23, 2012 -- 20:11 GMT (21:11 BST) | Topic: Cloud

VMware said Monday that it will buy Nicira in a deal valued at $1.05 billion in cash.

MORE FROM LARRY

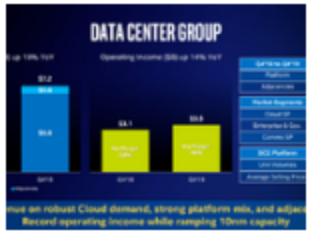

ring: Netflow, PAN, RSPAN

ated Control: low, OVSDB protocol

Network Traffic Generation in a Testbed

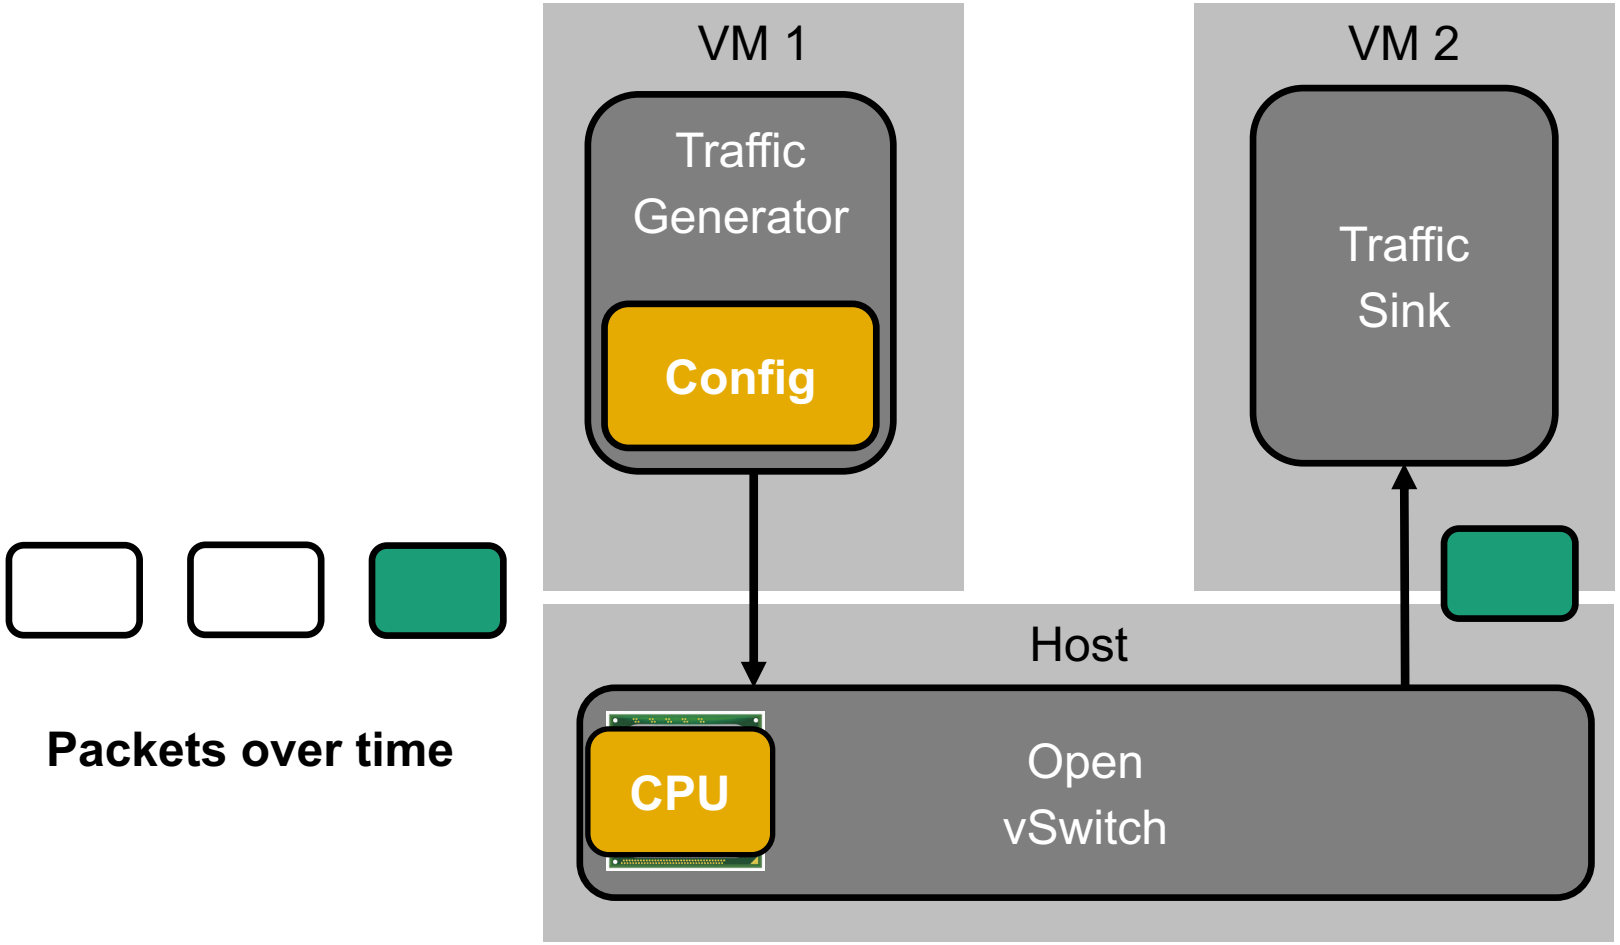

Match	Rule
	Forward
*	DROP

Goal: Find network traffic configuration that maximizes CPU load

Network Benchmarking is Challenging: Complex and Huge Configuration Space

How many packets to send? How should headers look like? What protocol to use? When to send packets? Etc.

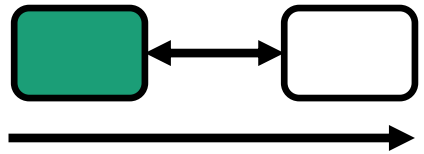

Packet Inter Arrival Time [1ms – 13ms]

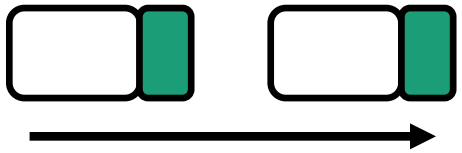

VLANs [1-5]

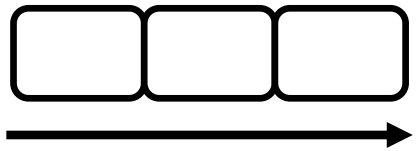

Number of Network Packets [1000 – 5000]

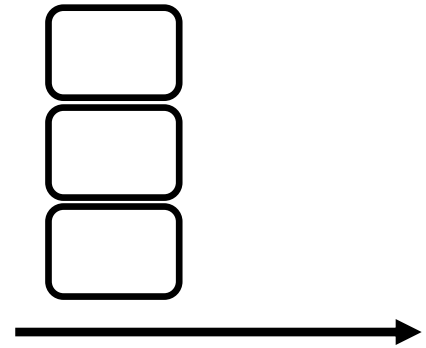

Batch Size [1-5]

NetBOA: The Bayesian Optimization Measurement Loop

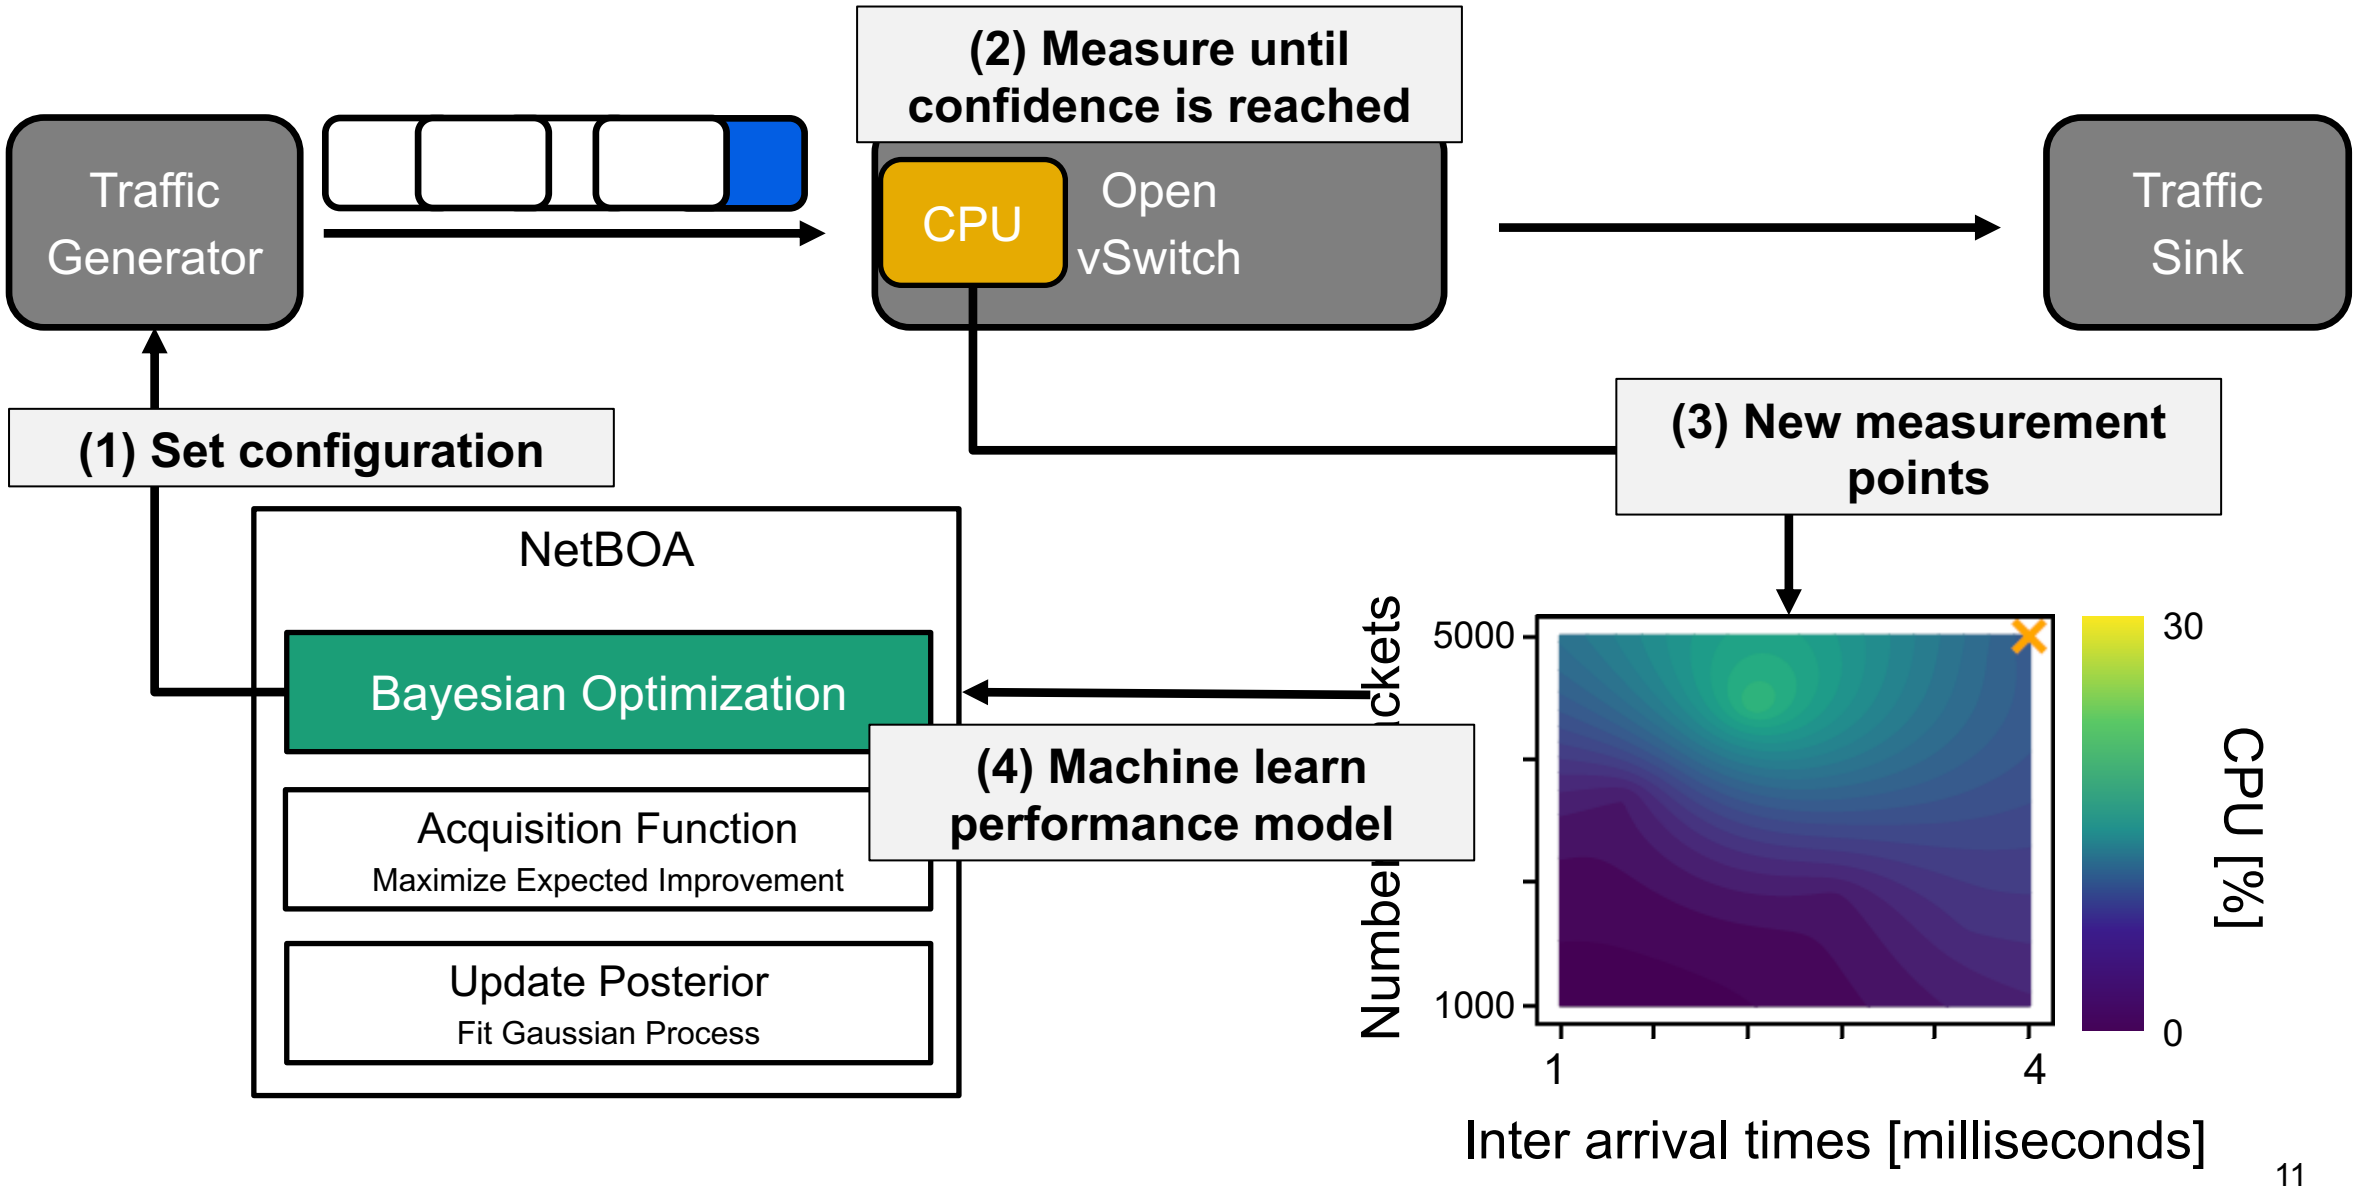

OVS Performance for Number of Packets and Inter-arrival Times

NetBOA finds this implementation weak spot! 30% CPU increase!

5000 packets, IAT 1 ms

4000 packets, IAT 3 ms

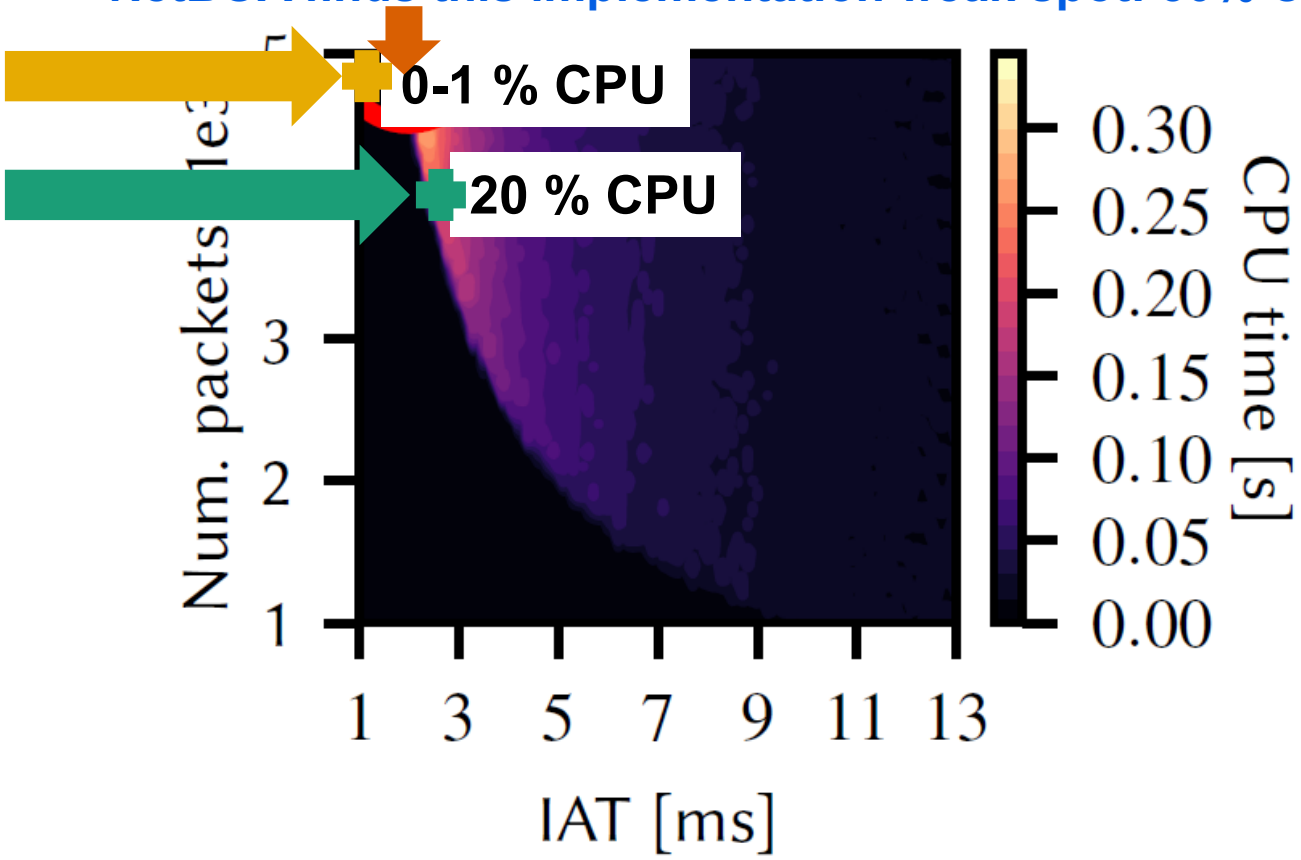

- Performance models are non-trivial
- **Surprising:** Sending less network packets over time can lead to significantly higher CPU

Why? Let Us Look At OvS Behavior!

Match	
1	Drop

Match	
1	Drop
2	Drop

Match	
1	Drop
2	Drop
3	Drop
…	…
N	Drop

Match	
2	Drop
3	Drop
…	…
N	Drop
N	Drop

Match	
1	Drop
2	Drop
3	Drop
…	…
N	Drop

Every packet triggers 2 times a costly array resizing operation!

OvS rule timeout 10 seconds

- We are using the OvS switch with the **Megaflow Cache enabled**
- For instance for 5000 packets: We trigger roughly every >2 ms a flow insertion + removal

→ **Forcing OvS to continuously run through the array + resizing it**

Summary

Adversarial input can harm your systems!

This talk: Data-Driven approach to automatically generate adversarial input to find weak spots, security holes … to make your systems bullet-proof!

Information missing in this talk: measurement details, simulation details, details on the used machine learning and artificial intelligence algorithms, … anything else :D?

Use concepts like NetBOA to receive continuous feedback about your solutions/implementations

References

[BIG DAMA'17] Blenk, Andreas; Kalmbach, Patrick; Schmid, Stefan; Kellerer, Wolfgang: o'zapft is: Tap Your Network Algorithm's Big Data! ACM SIGCOMM 2017 Workshop on Big Data Analytics and Machine Learning for Data Communication Networks (Big-DAMA), 2017

[SelfDN'18] Kalmbach, Patrick; Zerwas, Johannes; Babarczi, Péter; Blenk, Andreas; Kellerer, Wolfgang; Schmid, Stefan: Empowering Self-Driving Networks. Proceedings of the Afternoon Workshop on Self-Driving Networks - SelfDN 2018, ACM Press, 2018

[NetAI'19] Zerwas, Johannes; Kalmbach, Patrick; Henkel, Laurenz; Retvari, Gabor; Kellerer, Wolfgang; Blenk, Andreas; Schmid, Stefan: NetBOA: Self-Driving Network Benchmarking. ACM SIGCOMM 2019 Workshop on Network Meets AI & ML (NetAI '19), 2019

[CoNEXT'19] Lettner, Sebastian; Blenk, Andreas: Adversarial Network Algorithm Benchmarking. The 15th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '19 Companion), ACM, 2019

[TNSM'19] Zerwas, Johannes; Kalmbach, Patrick; Schmid, Stefan; Blenk, Andreas: Ismael: Using Machine Learning To Predict Acceptance of Virtual Clusters in Data Centers. IEEE Transactions on Network and Service Management, 2019

Thank you!

Questions?

What Could be Seen as Related

- Algorithmic complexity attacks (software domain):
  - SlowFuzz
  - PerfFuzz
- *Automated Synthesis of Adversarial Workloads for Network Functions*, ACM Sigcomm 2018
- **Policy Injection: A Cloud Dataplane DoS Attack**, ACM Sigcomm DEMO 2018

Why Important?

Implementation aspects can harm performance

Could even be used to attack your systems!

NetBOA vs Random Search

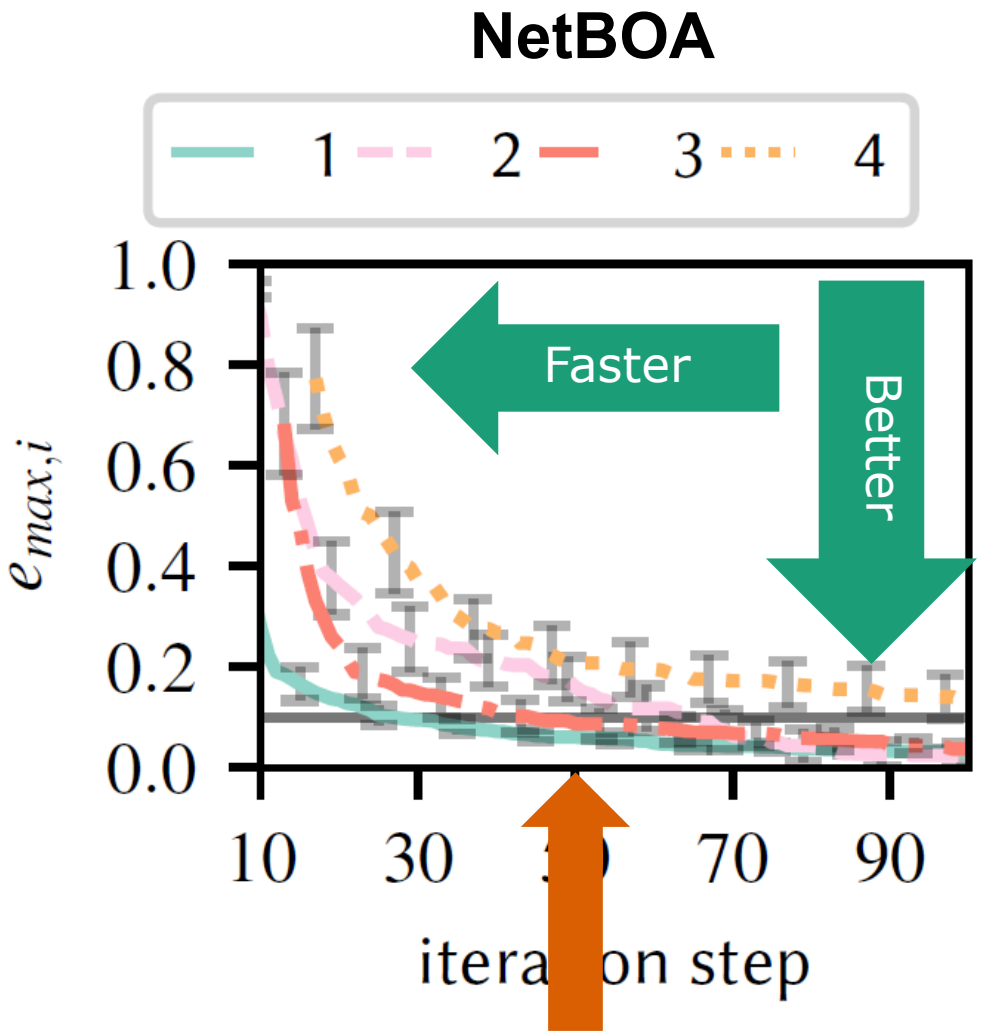

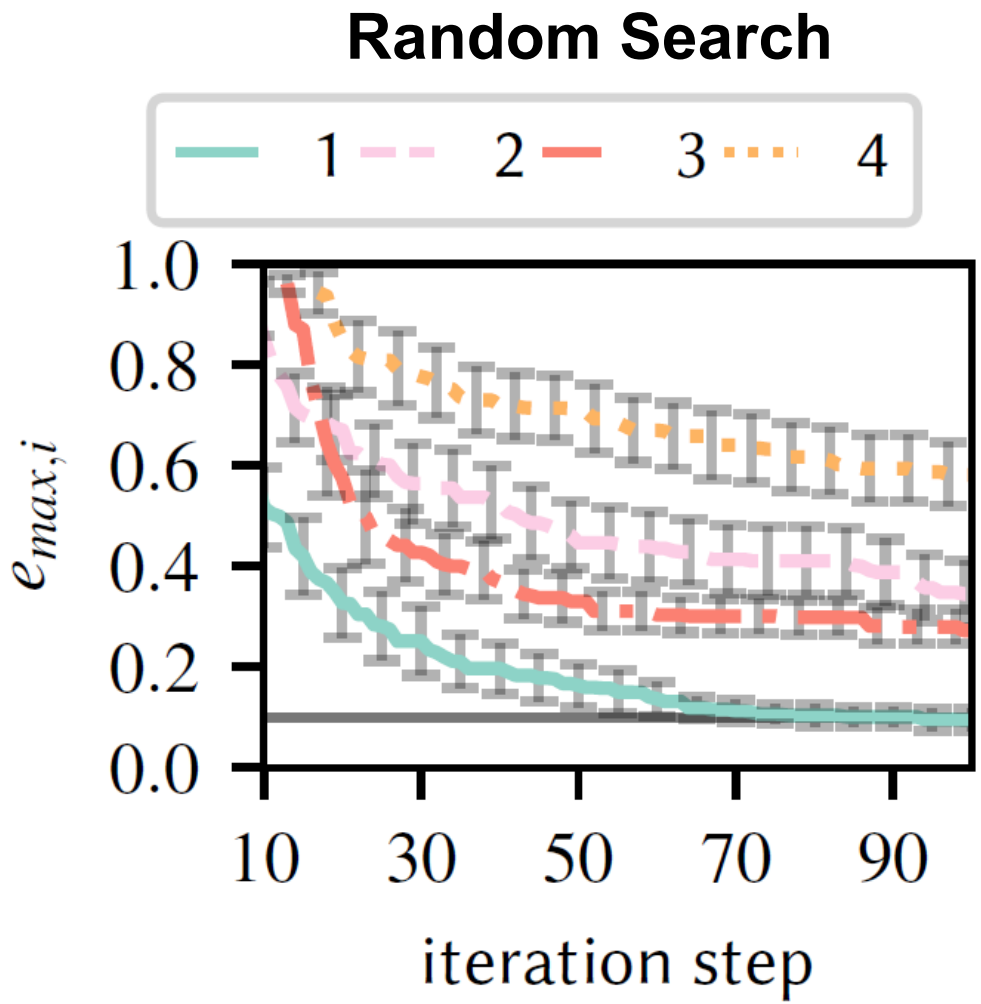

24 % higher CPU utilization

Bayesian Optimization: NetBOA for Inter Arrival Time (IAT) Parameter

Update Gaussian Process at runtime

Sampling from Gaussian Process gives confidence

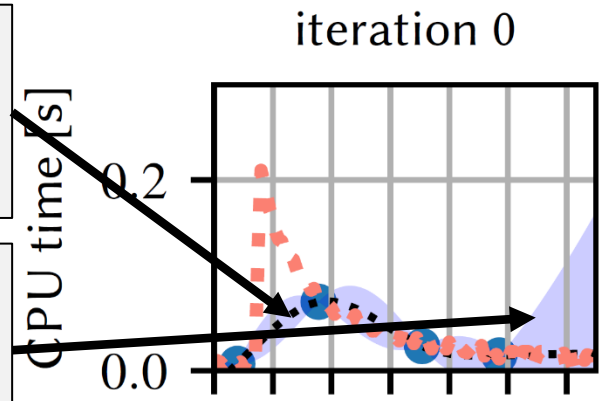

ping criteria
rts search

Expe
Improv
guides search