Improving Network Understanding

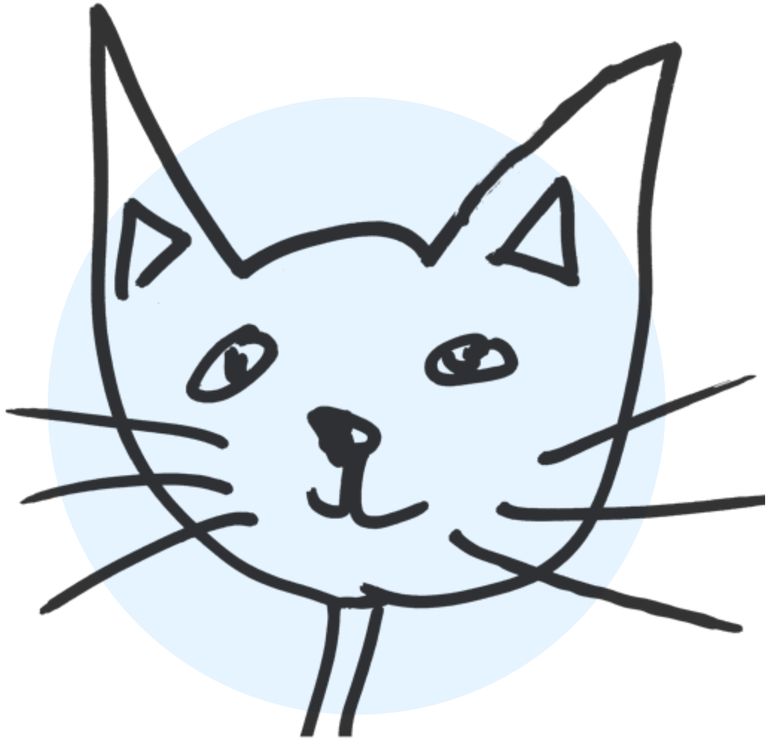

Rüdiger Birkner

PhD Defense
September, 27 2021

ETH zürich

1969

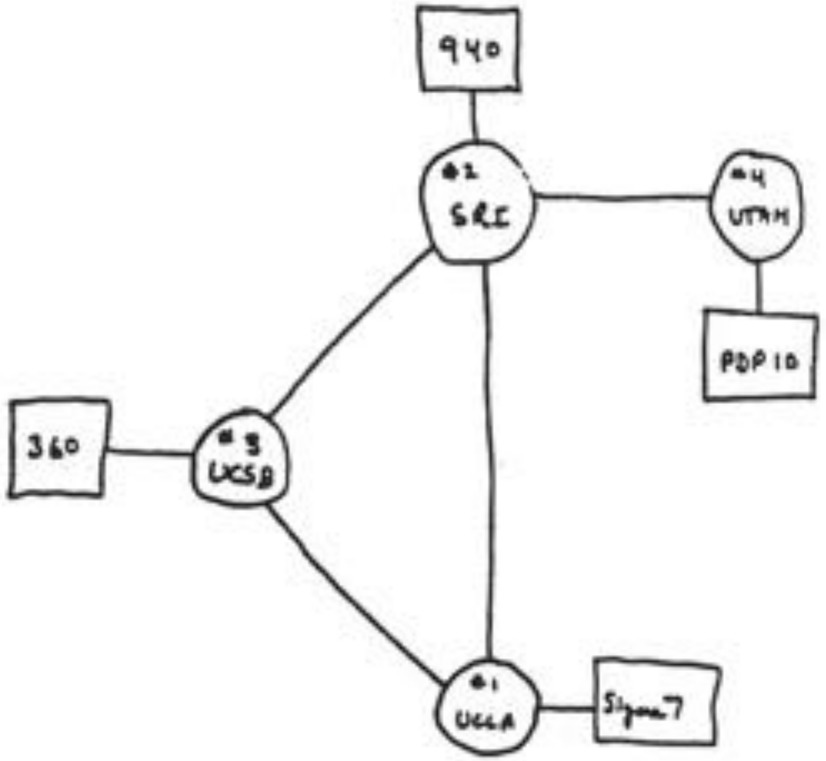

THE ARPA NETWORK

DEC 1969

4 NODES

2021

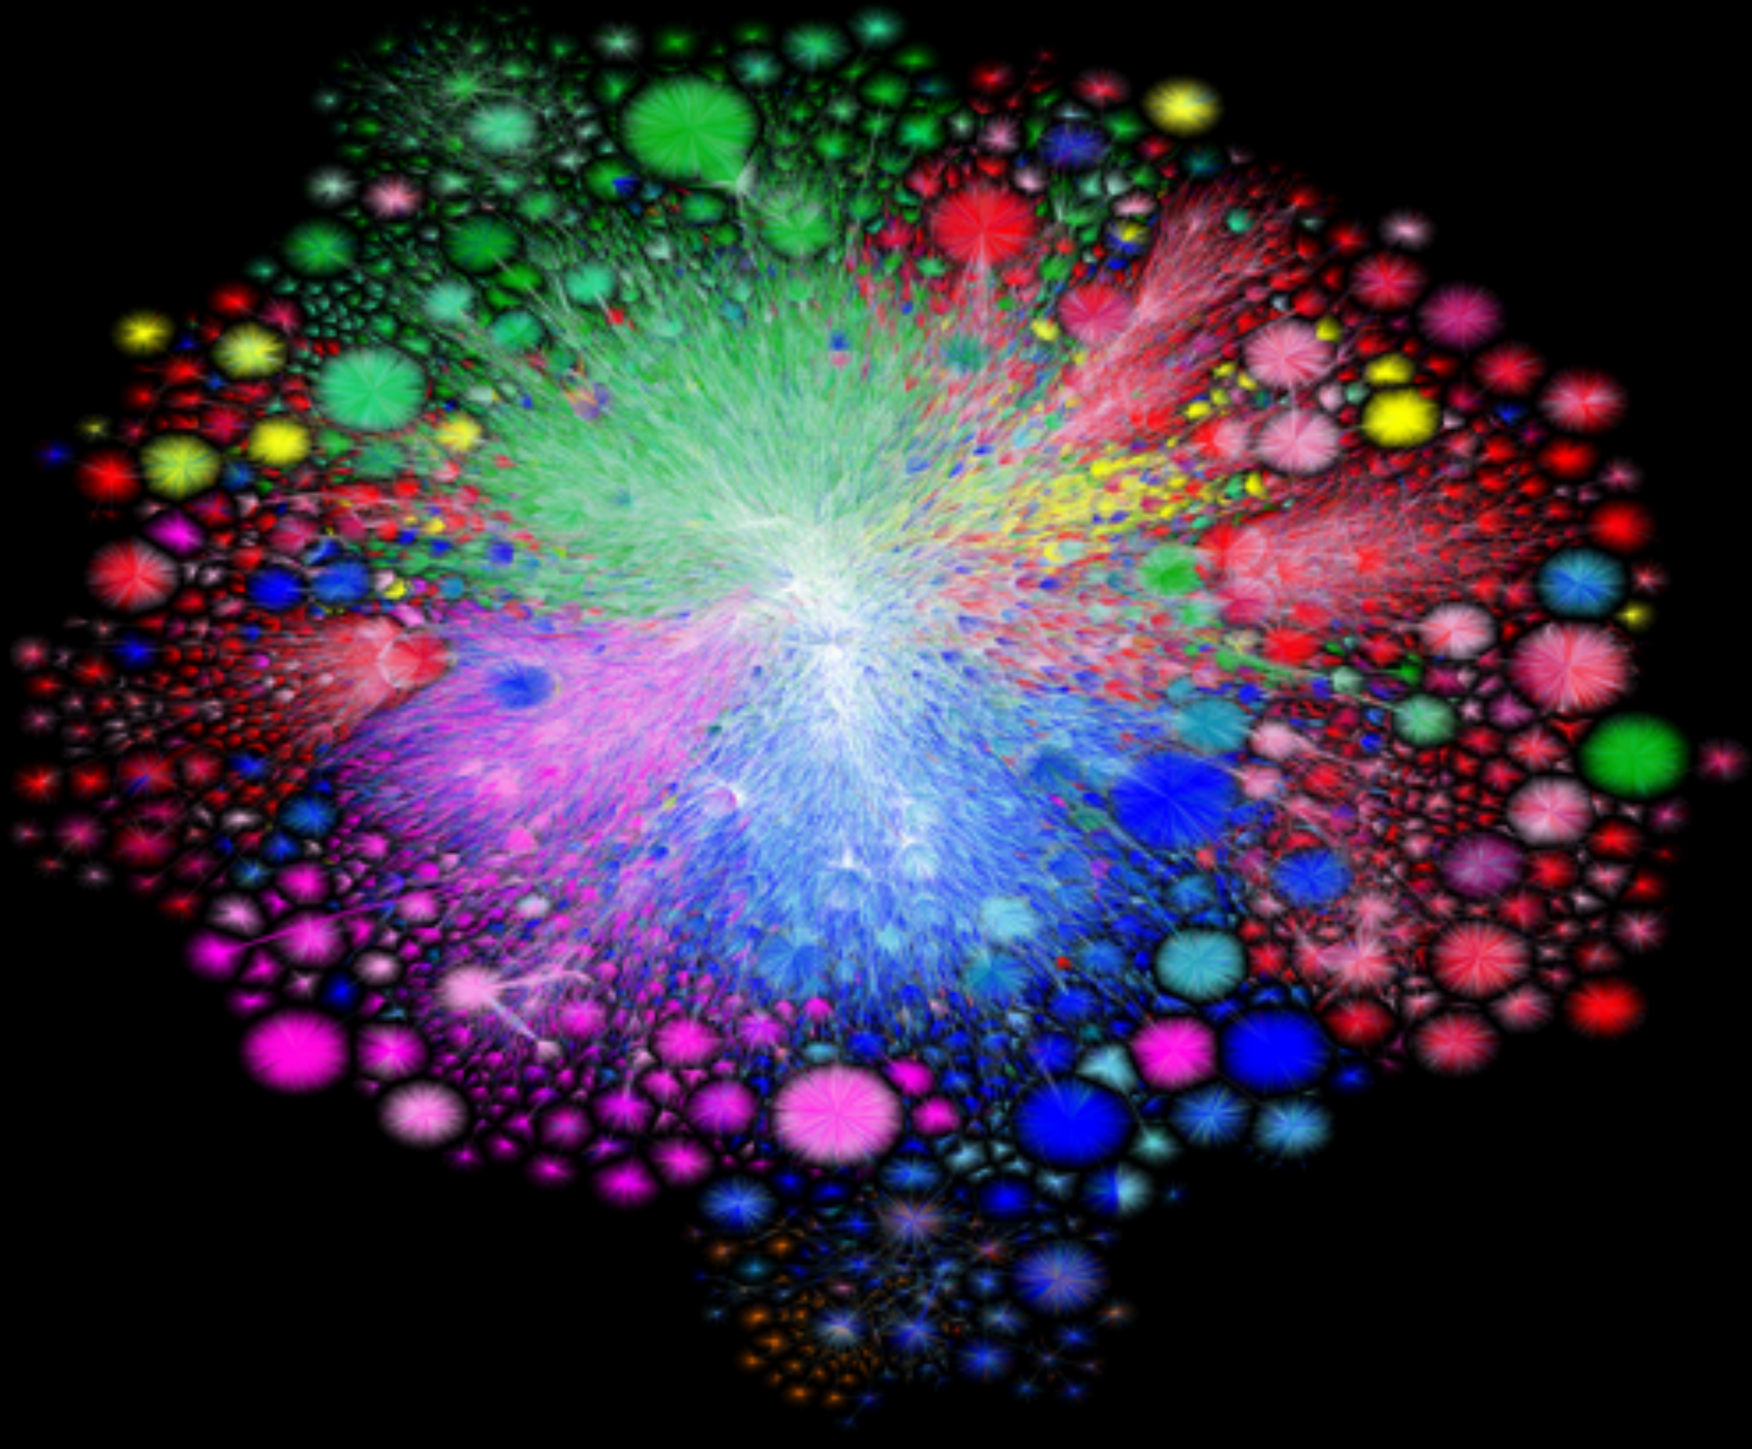

Over the years,
the Internet has seen tremendous growth

	ARPANET (1969)	Internet (2021)
size	4 nodes	70 000 networks
traffic	kbps	Tbps
use case	remote access	collaboration entertainment shopping …

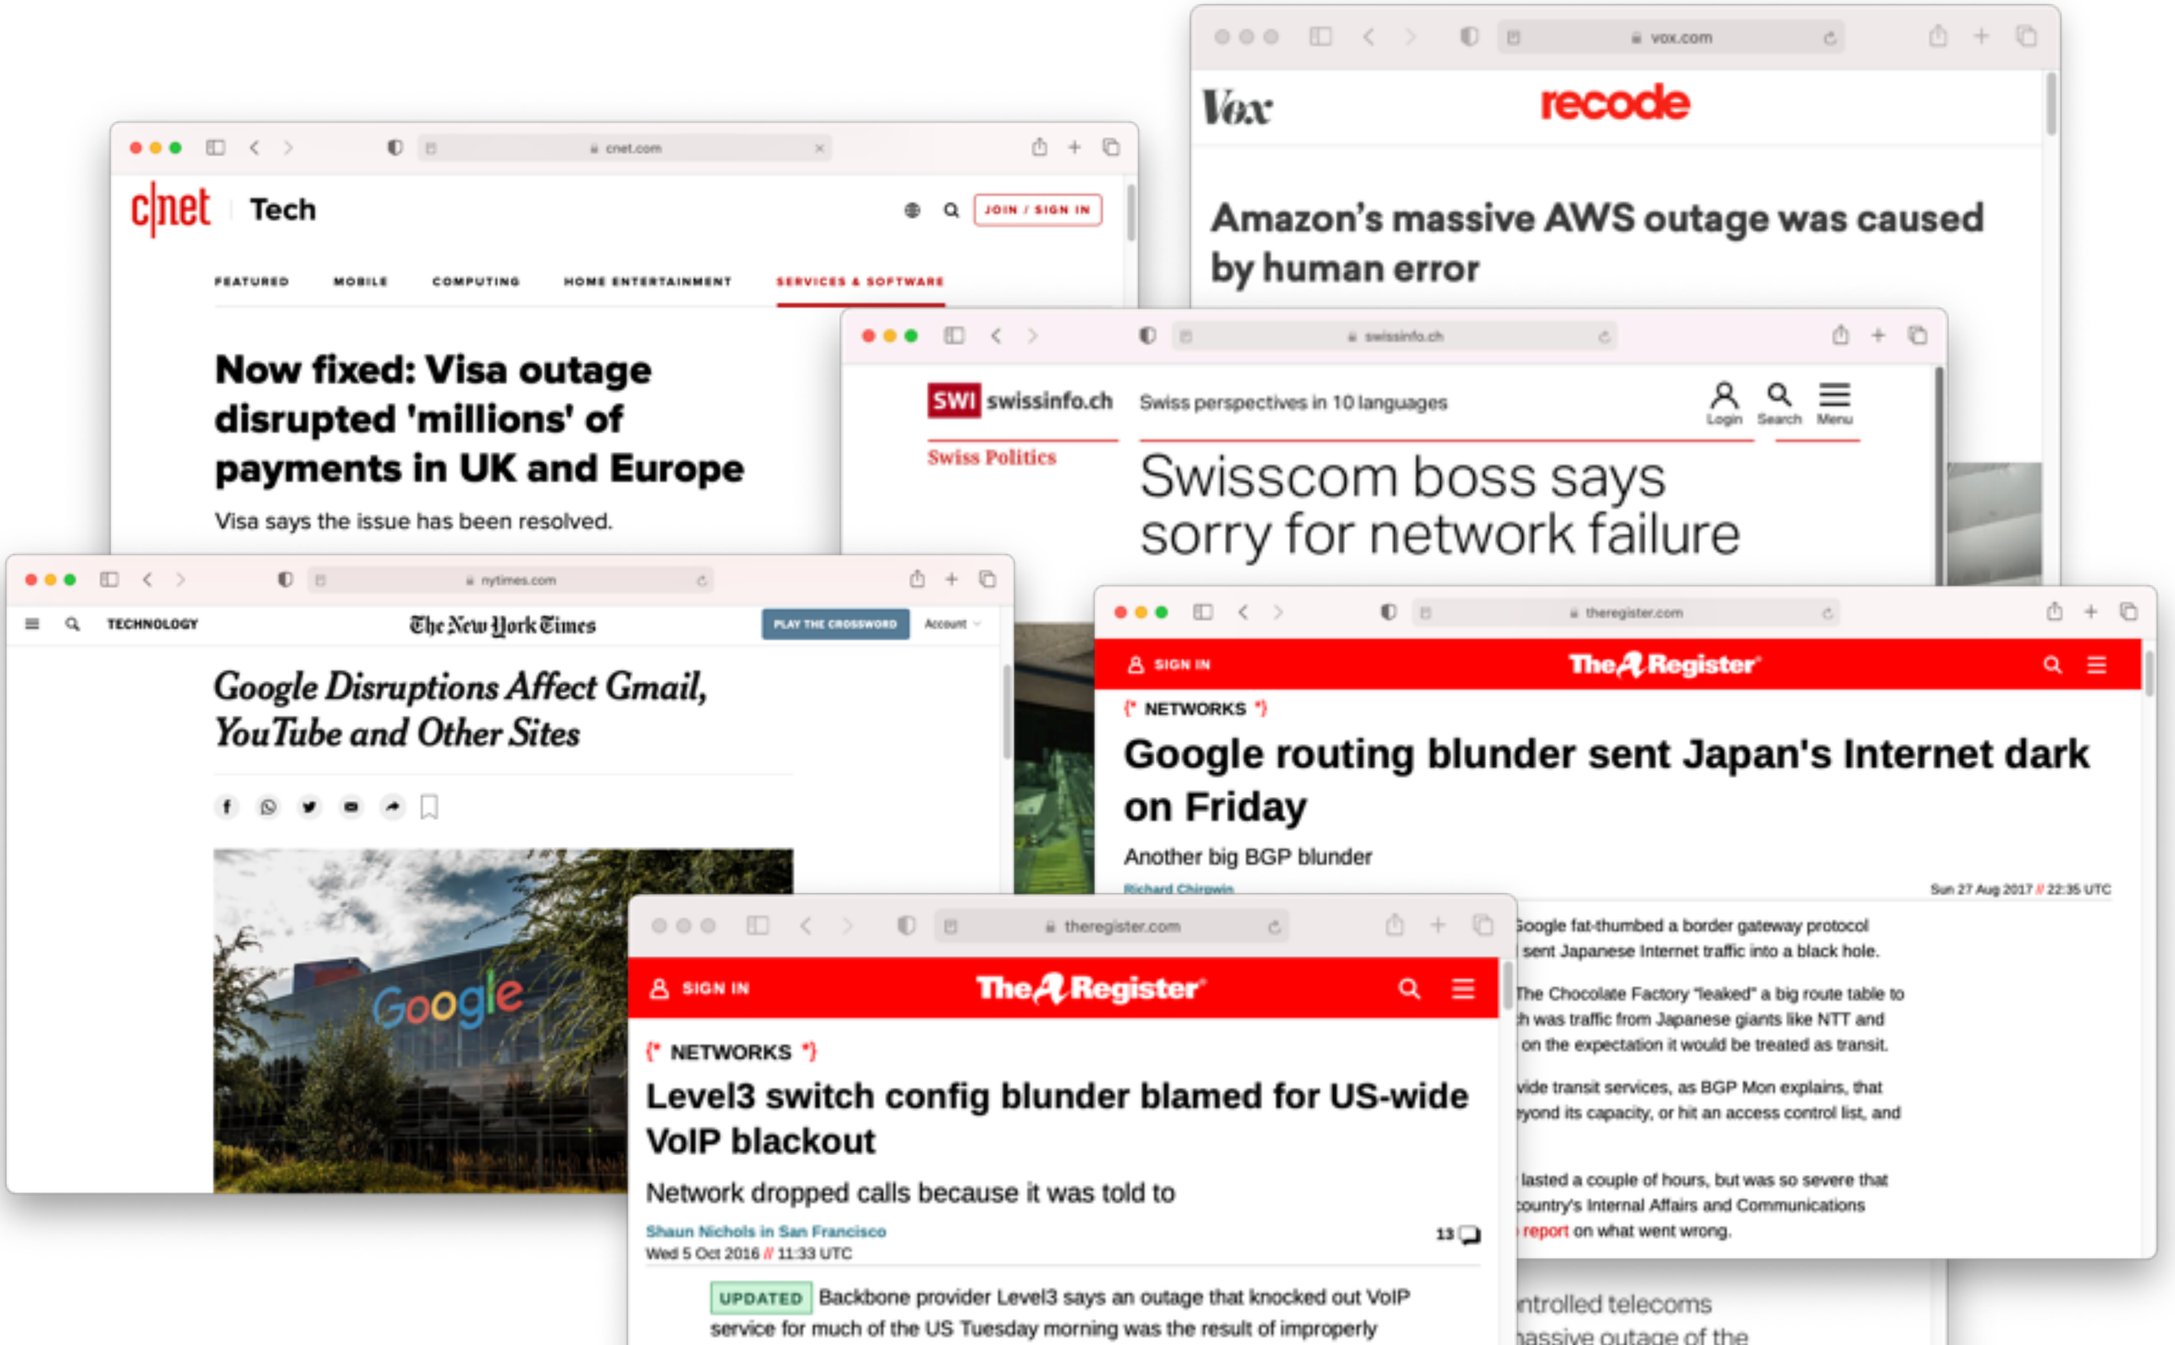
vox.com
Vox
recode
Amazon's massive AWS outage was caused by human error
cnet.com
c|net
Tech
JOIN / SIGN IN
FEATURED
MOBILE
COMPUTING
HOME ENTERTAINMENT
SERVICES & SOFTWARE
Now fixed: Visa outage disrupted 'millions' of payments in UK and Europe
Visa says the issue has been resolved.
swissinfo.ch
SWI swissinfo.ch
Swiss perspectives in 10 languages
Login
Search
Menu
Swiss Politics
Swisscom boss says sorry for network failure
nytimes.com
TECHNOLOGY
The New York Times
PLAY THE CROSSWORD
Account
Google Disruptions Affect Gmail, YouTube and Other Sites
theregister.com
SIGN IN
The Register
{* NETWORKS *}
Google routing blunder sent Japan's Internet dark on Friday
Another big BGP blunder
Richard Chirgwin
Sun 27 Aug 2017 // 22:35 UTC
Google fat-thumbed a border gateway protocol
sent Japanese Internet traffic into a black hole.
The Chocolate Factory "leaked" a big route table to
h was traffic from Japanese giants like NTT and
on the expectation it would be treated as transit.
vide transit services, as BGP Mon explains, that
yond its capacity, or hit an access control list, and
lasted a couple of hours, but was so severe that
country's Internal Affairs and Communications
report on what went wrong.
theregister.com
SIGN IN
The Register
{* NETWORKS *}
Level3 switch config blunder blamed for US-wide VoIP blackout
Network dropped calls because it was told to
Shaun Nichols in San Francisco
Wed 5 Oct 2016 // 11:33 UTC
13
UPDATED Backbone provider Level3 says an outage that knocked out VoIP service for much of the US Tuesday morning was the result of improperly
ntrolled telecoms
assive outage of the

This dissertation:

How can we assist network operators
in managing their network safely and reliably?

Network understanding
is a manual and time-consuming task

data access	rudimentary tools
	low-level data
	distributed across the network
data overload	hundreds of devices
	Tbps of traffic
	more than 900*k* destinations

Assisting network operators through automated network understanding

Network Configurations

Forwarding Behavior

Network Validators

Network Configurations

Forwarding Behavior

Network Validators

Config2Spec

[NSDI'20]

Configuring a network is an indirect process

network specification

```
loadbalancing(BAS,LAU)
waypoint(LUG,GEN,ZUR)
...
```

configurations

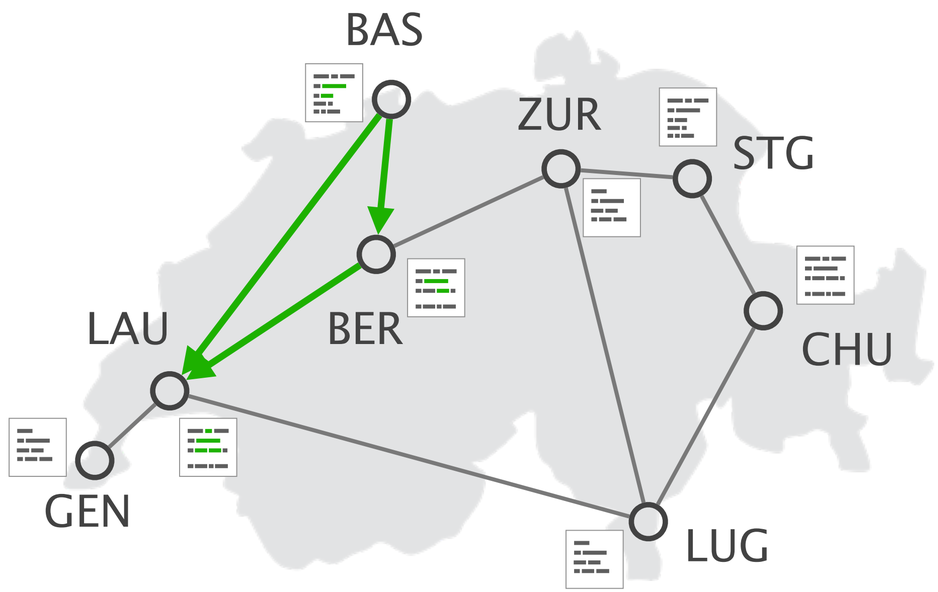

Configuring a network is an indirect process

network specification

```
loadbalancing(BAS,LAU)
waypoint(LUG,GEN,ZUR)
...
```

configurations

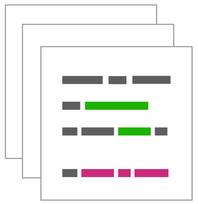

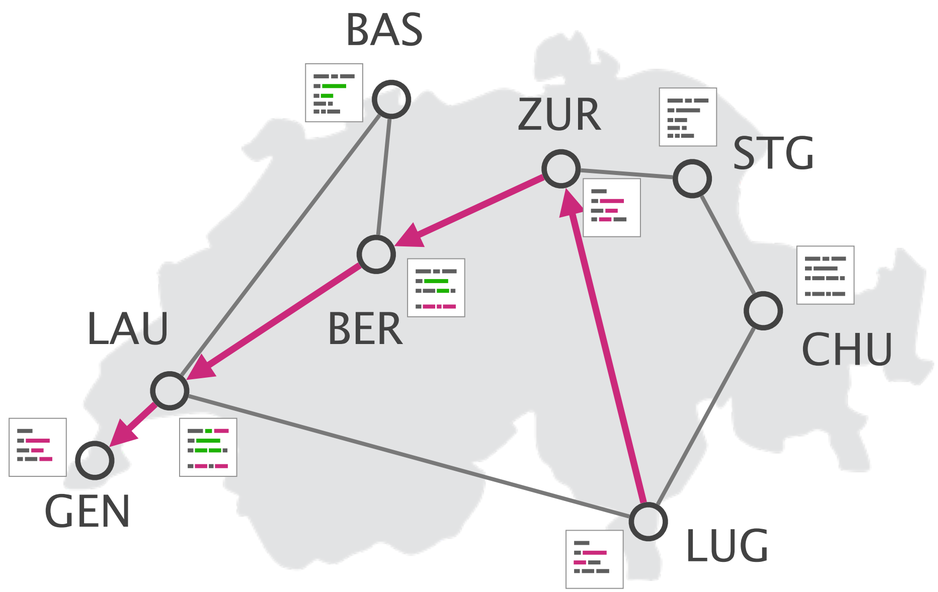

Understanding all the policies a network enforces, is an extremely cumbersome and difficult process

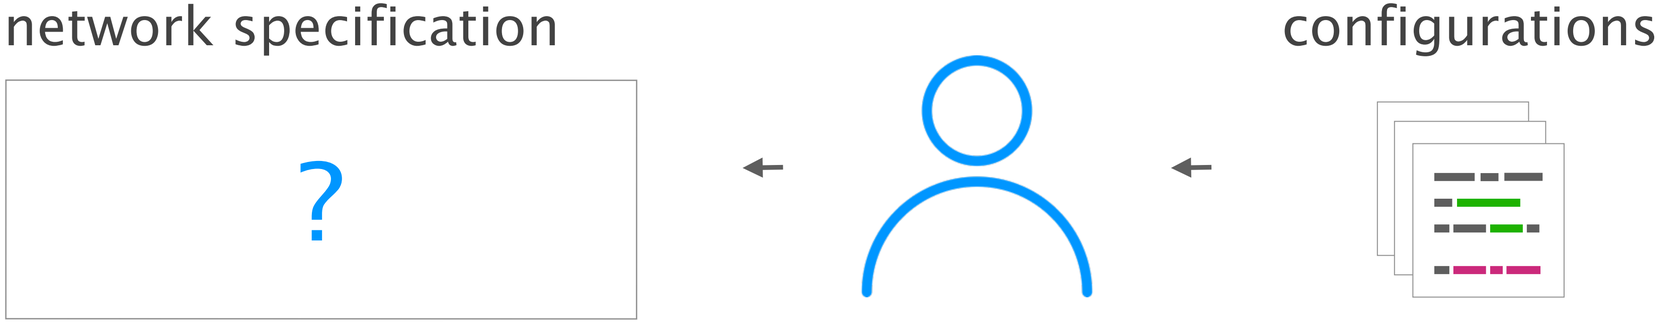

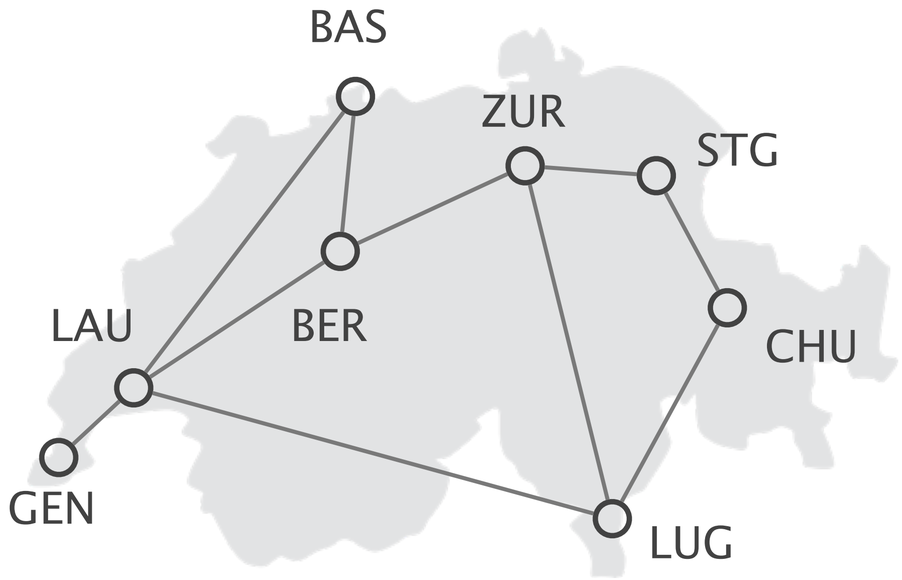

Config2Spec helps operators understand the policies their network configuration enforces

Input

Network Configuration

Failure Model

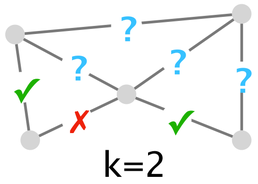

Config2Spec

Output

Network Specification

```
loadbalancing(BAS, LAU)
reachability(BAS, GEN)
reachability(BAS, LAU)

…

reachability(GEN, ZUR)
reachability(LUG, STG)
```

Config2Spec relies on a combination of data plane analysis and control plane verification

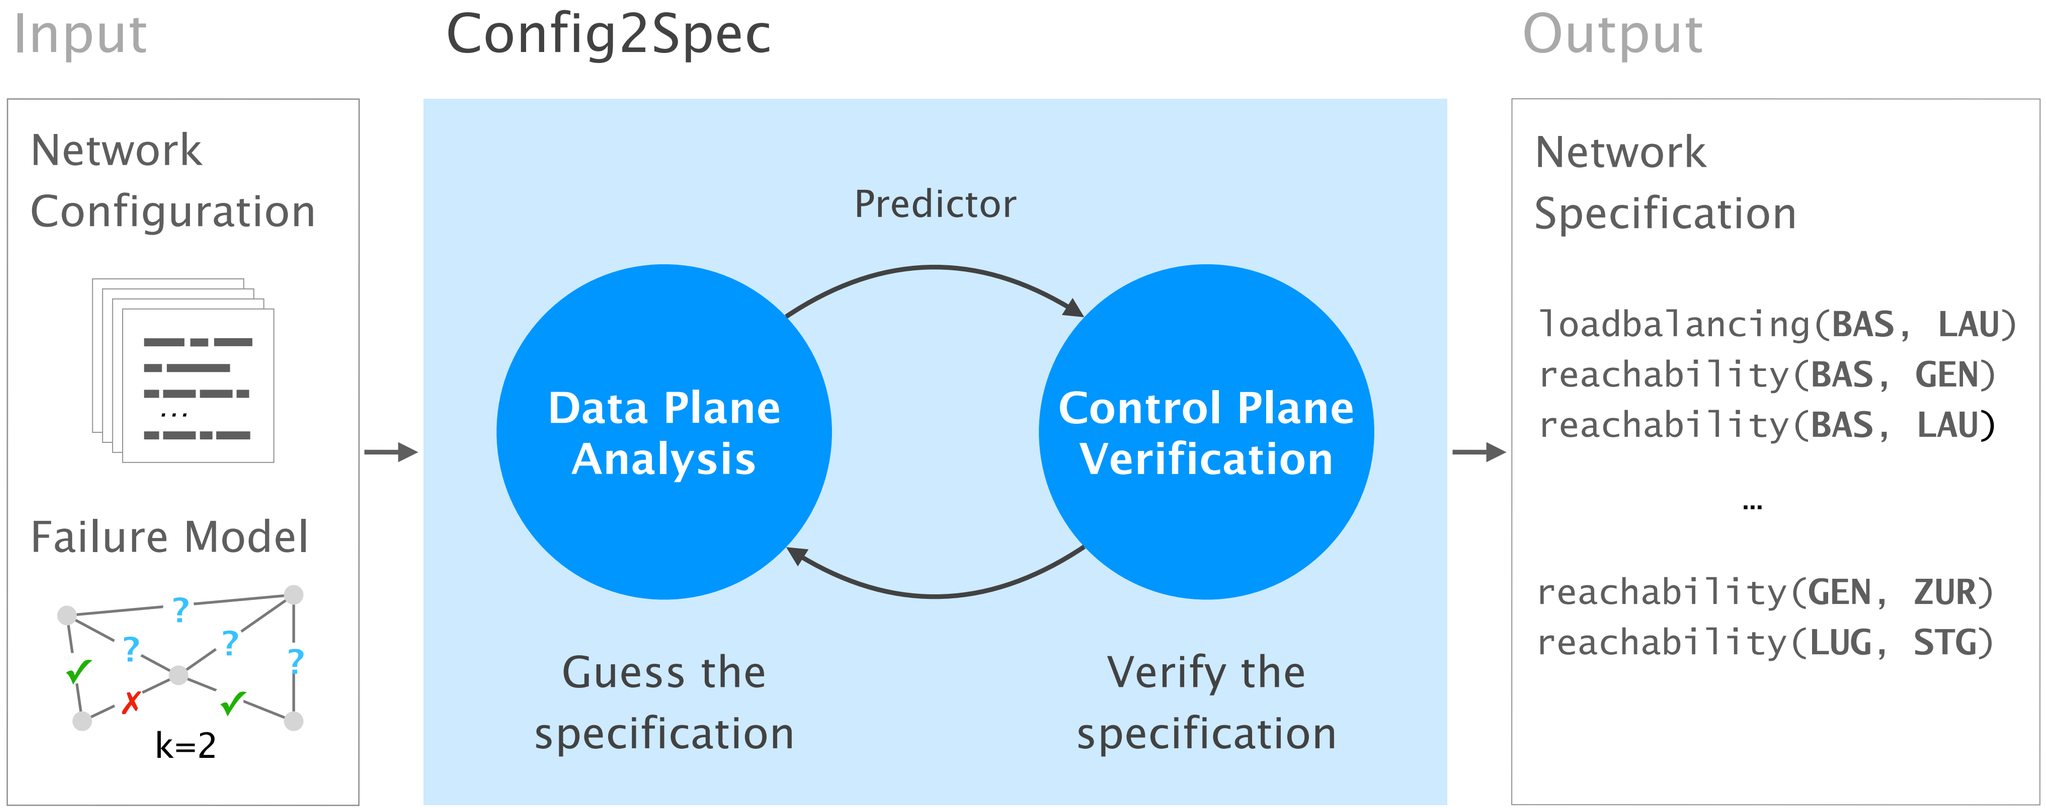

data plane analysis

control plane verification

data plane analysis

control plane verification

The network specification is the intersection of the policies that hold for every concrete environment

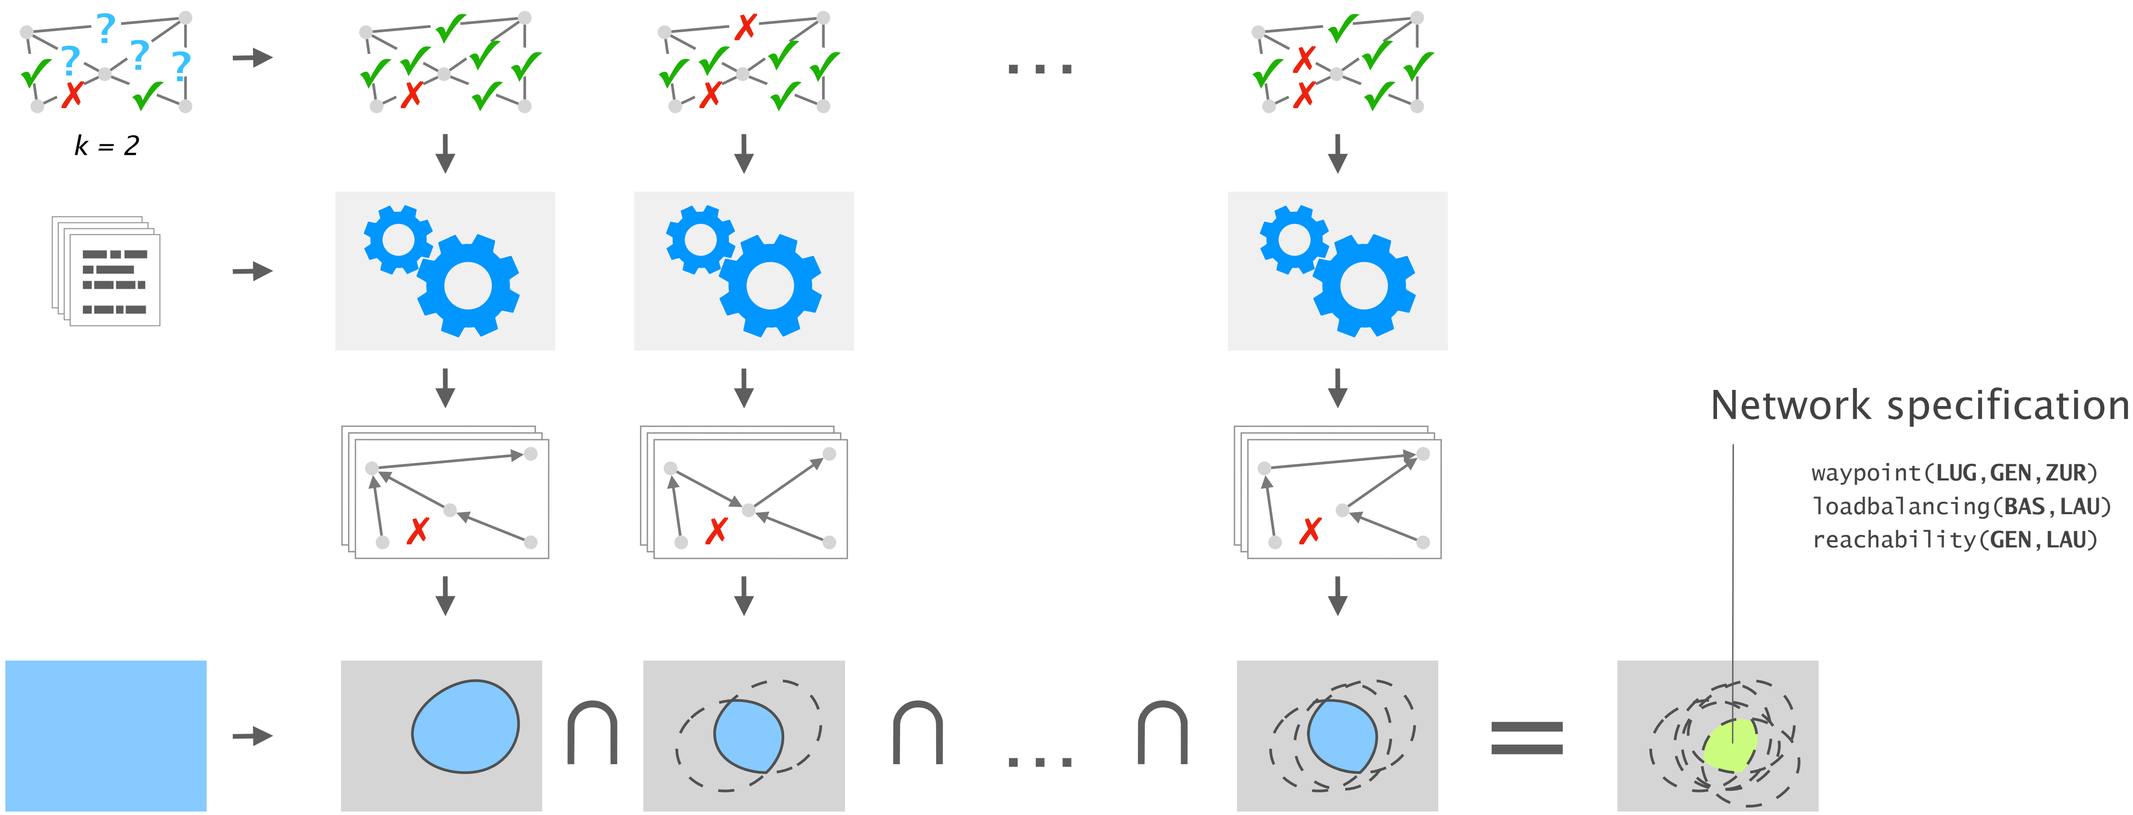

data plane analysis

control plane verification

The network specification is the set of policies that the verifier determined to hold for the failure model

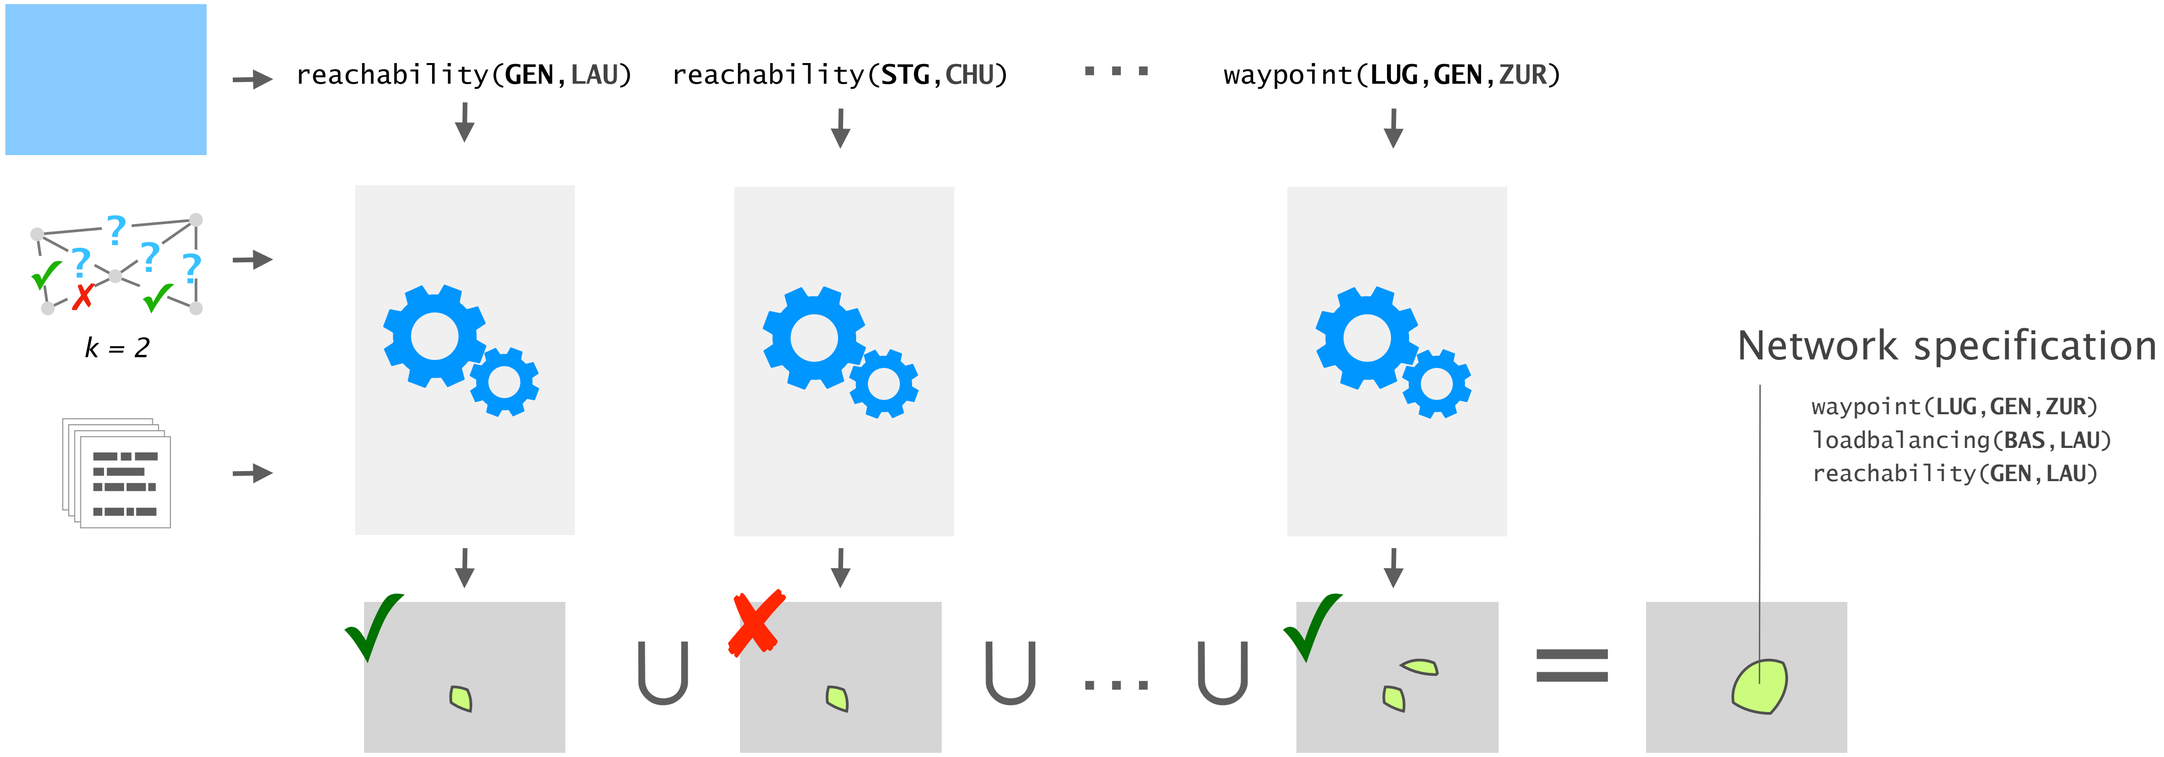

Config2Spec leverages their individual strengths

approach	data plane analysis	control plane verification
	all policies for **one** concrete env.	**one** policy for the **entire** failure model
good at	quickly pruning the candidate set	verifying a small candidate set

Config2Spec mines the network's full specification from its configuration and the required failure tolerance

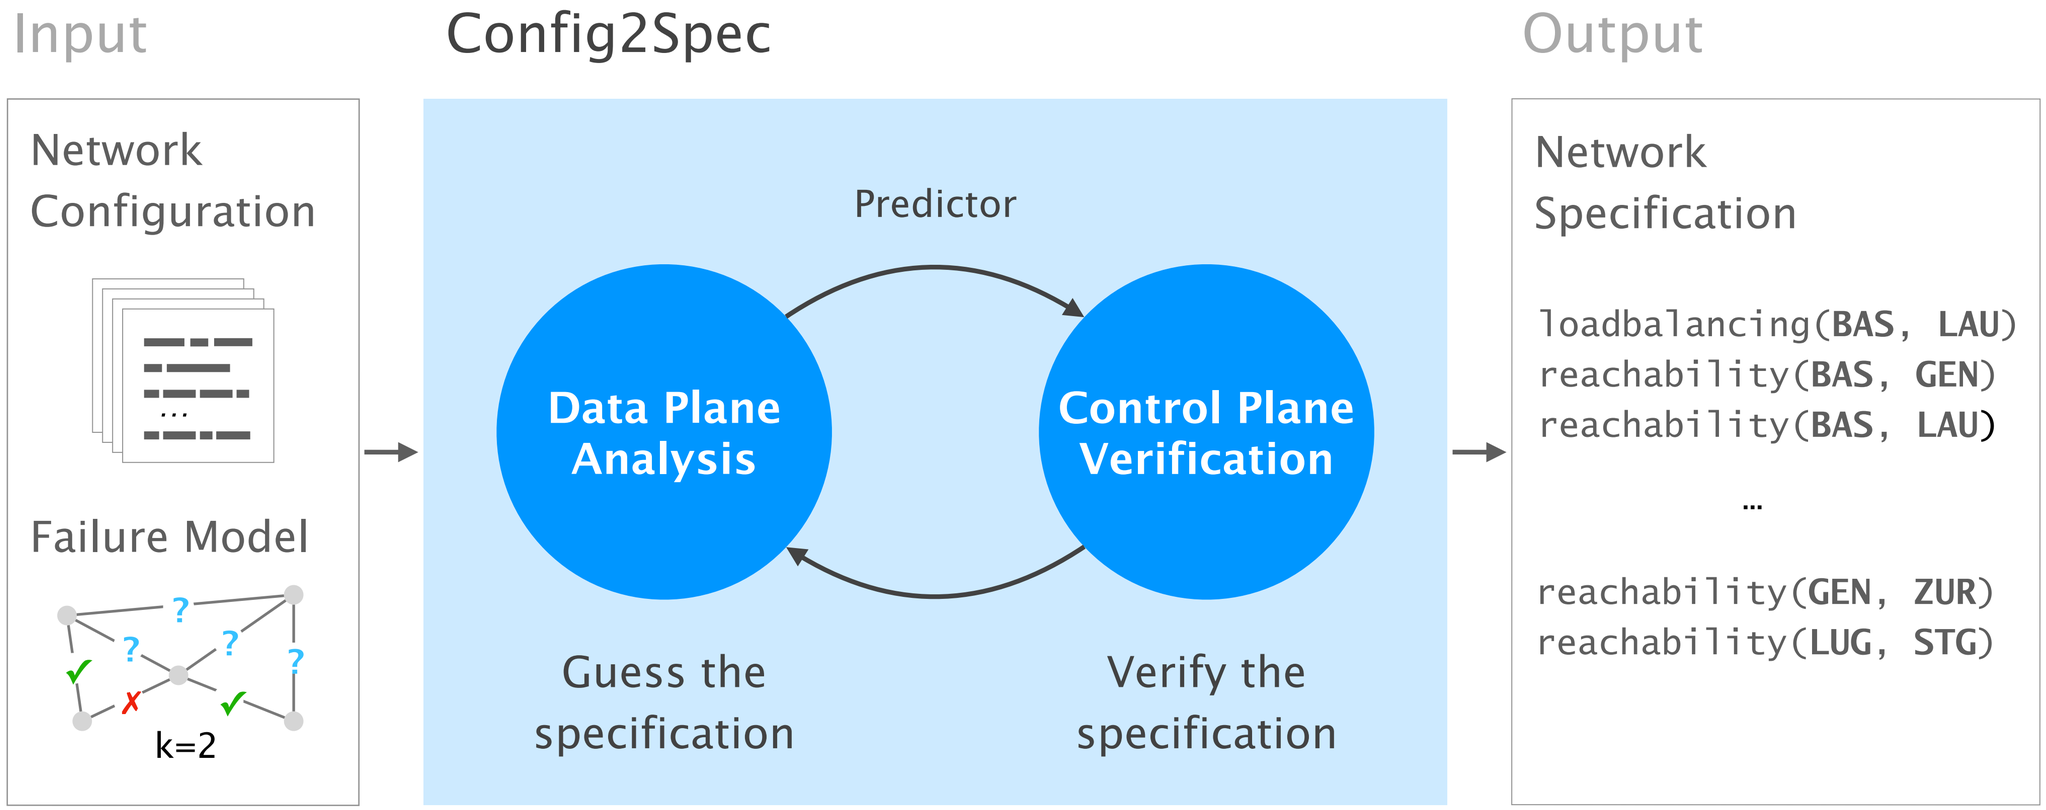

Config2Spec can be improved further using three domain-specific techniques

policy trimming

policy-aware selection

policy grouping

We fully implemented Config2Spec and show its practicality

Implementation	5k lines of Python and Java using Batfish and Minesweeper
Methodology	generated configs using NetComplete employing OSPF, BGP
	for a small, medium, and large network with 33, 70, and 158 routers

For failure models with few concrete environments, data plane analysis on its own provides fastest progress

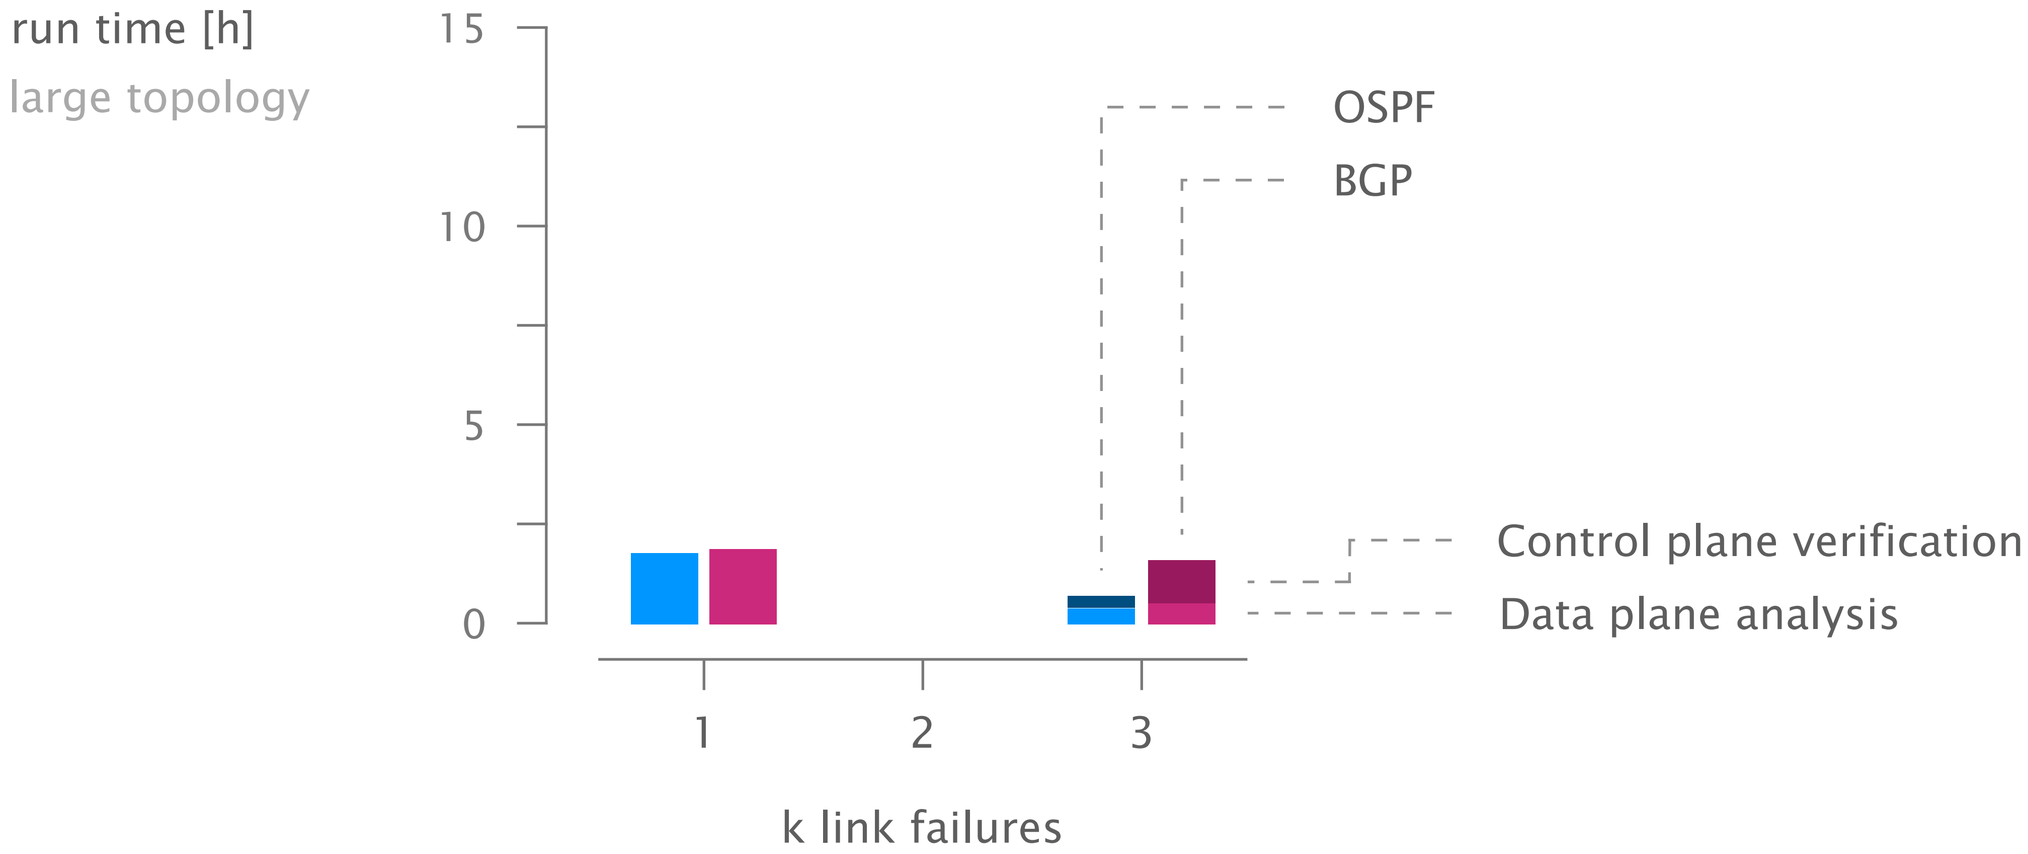

For failure models with a high failure bound, policy trimming reduces the candidate space significantly

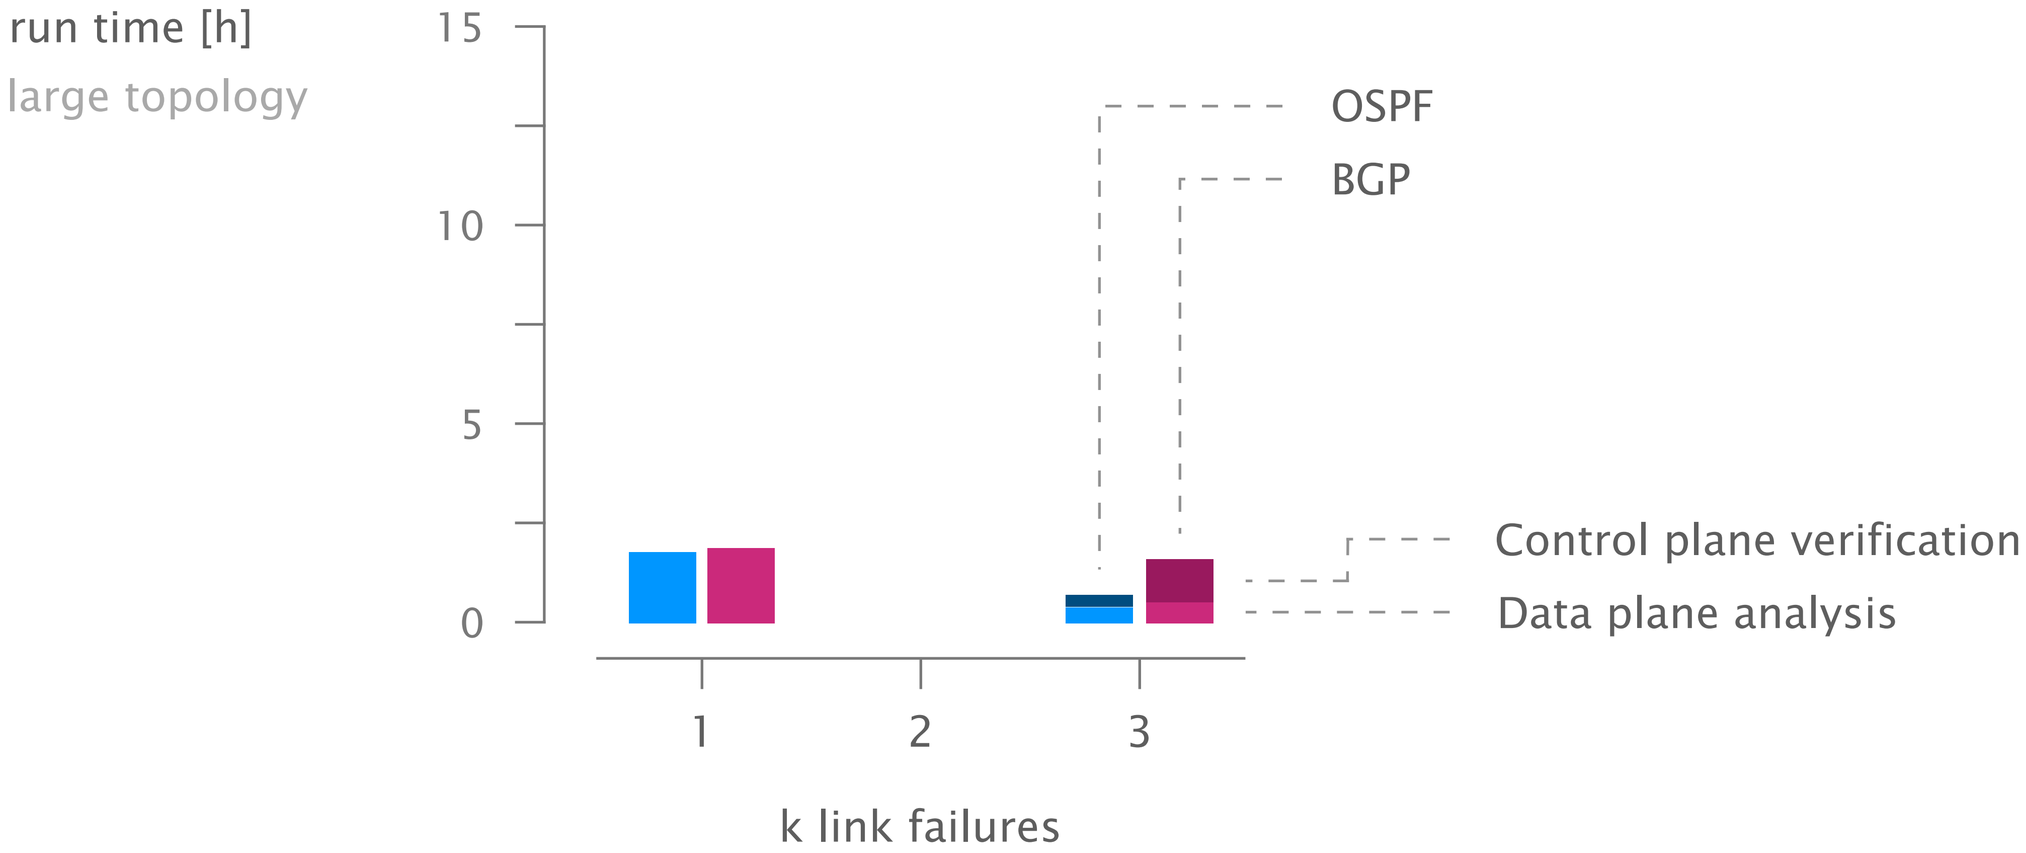

Config2Spec mines the specification for realistic networks in few hours

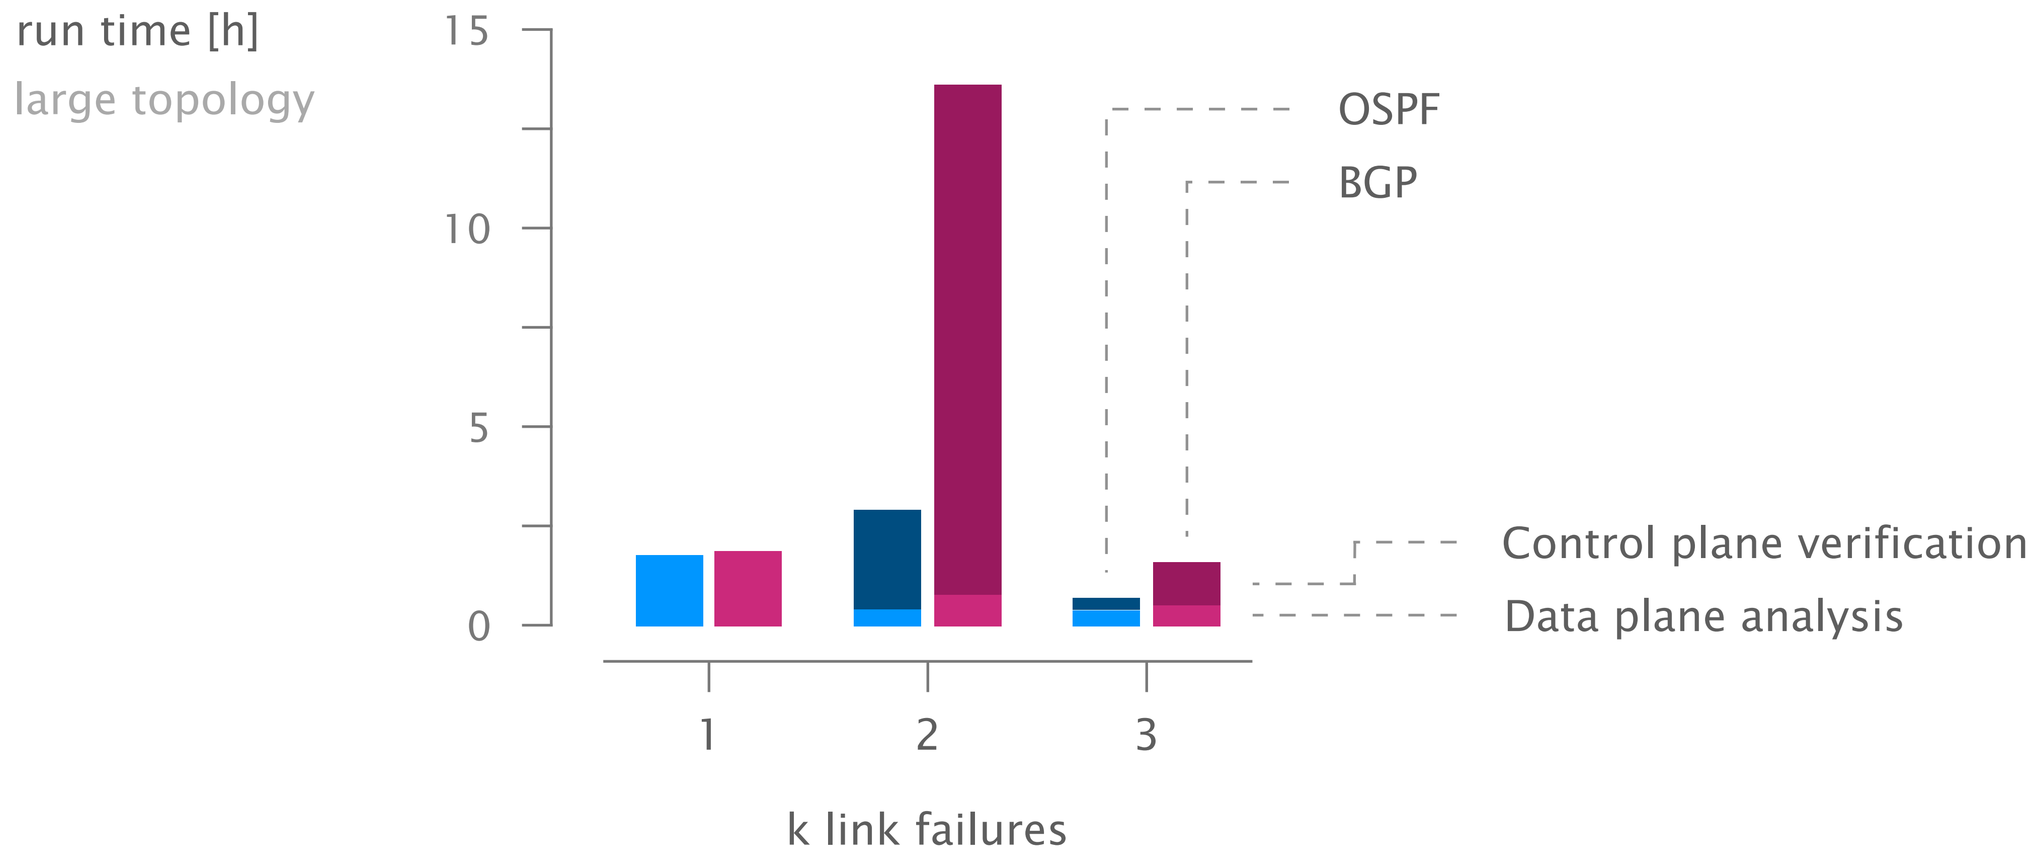

Config2Spec is useful beyond network understanding

adoption of validation tools

checking the correctness of the configs

configuration streamlining

synthesising semantically-equivalent configs

what-if analysis

analysing the impact of a config change

How can we assist network operators
in managing their network safely and reliably?

by improving network understanding!

Config2Spec
[NSDI'20]

Net2Text
[NSDI'18]

Metha
[NSDI'21]

Tomorrow's network understanding

Noisy data	What if the data is incomplete or wrong? detect anomalies/bugs, clean the data
Rich specifications	What characteristics does a specification need? dynamic specifications, control-plane policies
More input data	Can we find additional insights by combining data? new data sources, network provenance

Network Configurations

Config2Spec

[NSDI'20]

Forwarding Behavior

Net2Text

[NSDI'18]

Network Validators

Metha

[NSDI'21]

Thesis Publications

[NSDI'18]

Rüdiger Birkner, Dana Drachsler Cohen, Laurent Vanbever, and Martin Vechev

Net2Text: Query-Guided Summarization of Network Forwarding Behaviors

USENIX NSDI 2018. Renton, WA, USA

[NSDI'20]

Rüdiger Birkner, Dana Drachsler Cohen, Laurent Vanbever, and Martin Vechev

Config2Spec: Mining Network Specifications from Network Configurations

USENIX NSDI 2020. Santa Clara, CA, USA

[NSDI'21]

Rüdiger Birkner*, Tobias Brodmann*,
Petar Tsankov, Laurent Vanbever, and Martin Vechev

Metha: Network Verifiers Need To Be Correct Too!

USENIX NSDI 2021. Online

*These authors contributed equally to this work

Supplemental Publications

[SOSR'17a]

Rüdiger Birkner, Arpit Gupta, Nick Feamster, and Laurent Vanbever

SDX-Based Flexibility or Internet Correctness? Pick Two!

ACM SOSR 2017. Santa Clara, CA, USA

[SOSR'17b]

Robert MacDavid, Rüdiger Birkner, Ori Rottenstreich,
Arpit Gupta, Nick Feamster, and Jennifer Rexford

Concise Encoding of Flow Attributes in SDN Switches

ACM SOSR 2017. Santa Clara, CA, USA

[SIGCOMM'21]

Tibor Schneider, Rüdiger Birkner, and Laurent Vanbever

Snowcap: Synthesizing Network-Wide Configuration Updates

ACM SIGCOMM 2021. Online