

Fuzzing for Security and Resilience

Gilles Coremans

`gilles.coremans@vub.be`



SOFTWARE
LANGUAGES
LAB

Test cases are great!



Quick to find bugs or regressions



Writing tests improves understanding of code

... but limited



Production deployments always have more cases!

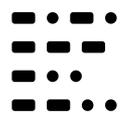


Writing good tests takes a lot of time

A different approach: Fuzzing



What if we just tried every input?



Fuzzing tests programs by
sending random inputs



Unlike other automatic testing
techniques: no false positives!



Testing Oracles



Fuzzing for security has “easy” oracles:

- Segmentation faults
- AddressSanitizer
- etc



For other bug types we have to get creative

- Check violations of expected properties
- Check for unexpected errors
- Check execution time/latency

Speeding up fuzzing



Grammar-based Fuzzing

Generate structured inputs that conform to what the program expects



Feedback-driven Fuzzing

Determine which inputs are interesting, then mutate those

Feedback-driven fuzzing



Coverage tells us whether an input is interesting



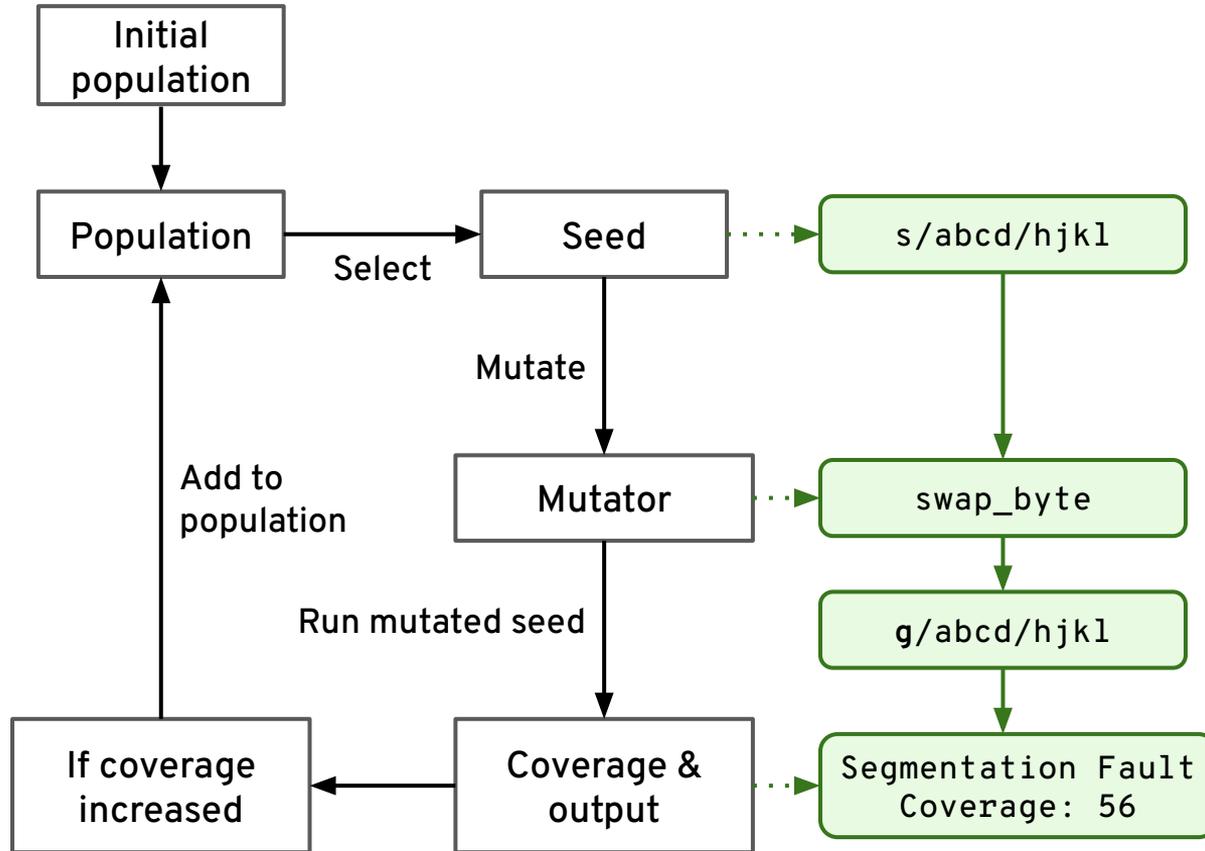
Mutating interesting inputs yields more interesting inputs



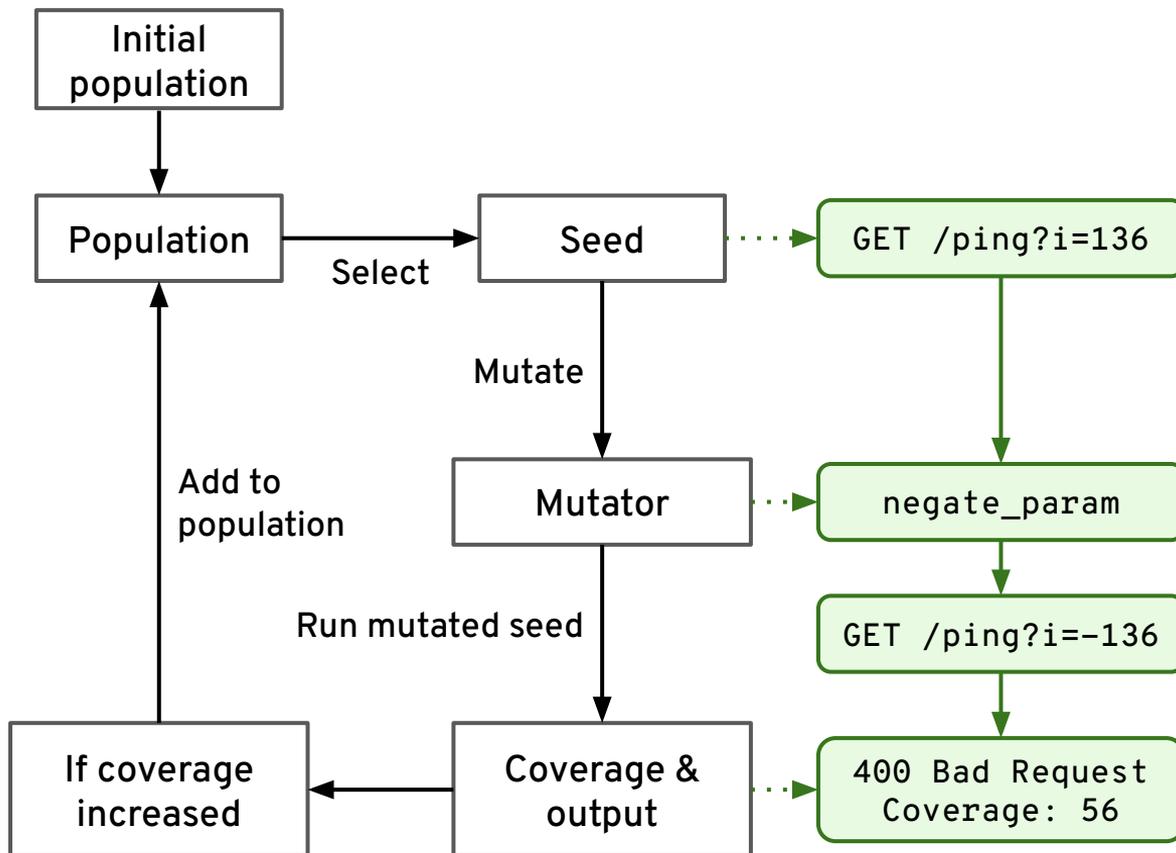
Many vulnerabilities found this way

IJG jpeg ¹	libjpeg-turbo ^{1 2}	libpng ¹
libtiff ^{1 2 3 4 5}	mozjpeg ¹	PHP ^{1 2 3 4 5 6 7 8}
Mozilla Firefox ^{1 2 3 4}	Internet Explorer ^{1 2 3 4}	Apple Safari ¹
Adobe Flash / PCRE ^{1 2 3 4 5 6 7}	sqlite ^{1 2 3 4...}	OpenSSL ^{1 2 3 4 5 6 7}
LibreOffice ^{1 2 3 4}	poppler ^{1 2...}	freetype ^{1 2}
GnuTLS ¹	GnuPG ^{1 2 3 4}	OpenSSH ^{1 2 3 4 5}
PuTTY ^{1 2}	ntpd ^{1 2}	nginx ^{1 2 3}
bash (post-Shellshock) ^{1 2}	tcpdump ^{1 2 3 4 5 6 7 8 9}	JavaScriptCore ^{1 2 3 4}
pdfium ^{1 2}	ffmpeg ^{1 2 3 4 5}	libmatroska ¹
libarchive ^{1 2 3 4 5 6 ...}	wireshark ^{1 2 3}	ImageMagick ^{1 2 3 4 5 6 7 8 9 ...}
BIND ^{1 2 3 ...}	QEMU ^{1 2}	lcms ¹
Oracle BerkeleyDB ^{1 2}	Android / libstagefright ^{1 2}	iOS / ImageIO ¹
FLAC audio library ^{1 2}	libsndfile ^{1 2 3 4}	less / lesspipe ^{1 2 3}
strings (+ related tools) ^{1 2 3 4 5 6 7}	file ^{1 2 3 4}	dpkg ^{1 2}
rcs ¹	systemd-resolved ^{1 2}	libyaml ¹
Info-Zip unzip ^{1 2}	libtasn1 ^{1 2 ...}	OpenBSD pfctl ¹
NetBSD bpf ¹	man & mandoc ^{1 2 3 4 5 ...}	IDA Pro [reported by authors]
clamav ^{1 2 3 4 5 6}	libxml2 ^{1 2 4 5 6 7 8 9 ...}	glibc ¹
clang / llvm ^{1 2 3 4 5 6 7 8 ...}	nasm ^{1 2}	ctags ¹
mutt ¹	procmail ¹	fontconfig ¹
pdksh ^{1 2}	Qt ^{1 2...}	wavpack ^{1 2 3 4}
redis / lua-cmsgpack ¹	taglib ^{1 2 3}	privoxy ^{1 2 3}
perl ^{1 2 3 4 5 6 7...}	libxmp	radare2 ^{1 2}
SleuthKit ¹	fwknop [reported by author]	X.Org ^{1 2}
exifprobe ¹	jhead [?]	capnproto ¹
Xerces-C ^{1 2 3}	metacam ¹	djvulibre ¹
exiv ^{1 2}	Linux btrfs ^{1 2 3 4 6 7 8}	Knot DNS ¹
curl ^{1 2 3}	wpa_supplicant ¹	libde265 [reported by author]
dnsmasq ¹	libbgp (1)	lame ^{1 2 3 4 5 6}
libwmf ¹	uudecode ¹	MuPDF ^{1 2 3 4}
imlib2 ^{1 2 3 4}	libraw ¹	libbson ¹
libsass ¹	yara ^{1 2 3 4}	W3C tidy-html5 ¹
VLC ^{1 2}	FreeBSD syscons ^{1 2 3}	John the Ripper ^{1 2}
screen ^{1 2 3}	tmux ^{1 2}	mosh ¹
UPX ¹	indent ¹	openjpeg ^{1 2}
MMIX ¹	OpenMPT ^{1 2}	rxvt ^{1 2}
dhcpcd ¹	Mozilla NSS ¹	Nettle ¹
mbed TLS ¹	Linux netlink ¹	Linux ext4 ¹
Linux xfs ¹	botan ¹	expat ^{1 2}
Adobe Reader ¹	libav ¹	libical ¹

The fuzzing loop for a command line tool



The fuzzing loop for a REST API



Our research: fuzzing for micro-service backends



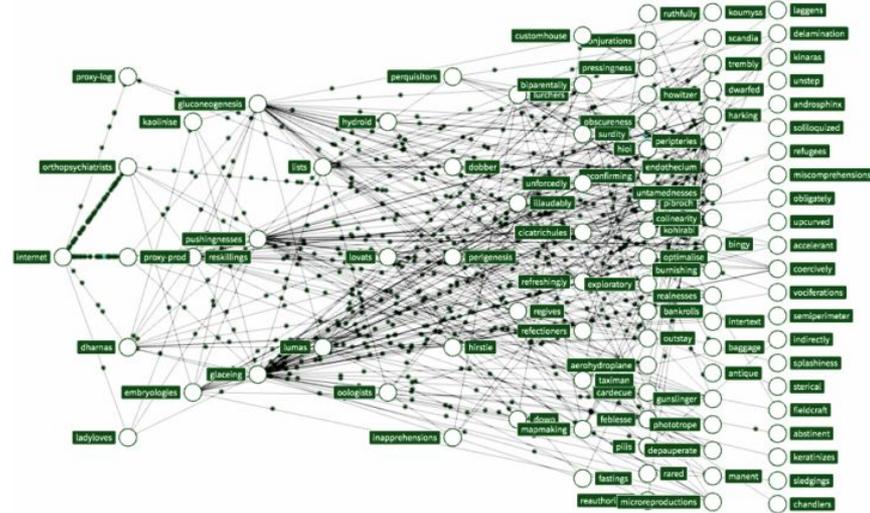
Backend split into many small services which communicate



Used for large, scalable, constantly evolving backends



Developers need to consider fault tolerance



Fuzzing for resilience defects



Resilience defects aren't triggered on developer machines or in testing



With fuzzing, we can find these bugs automatically in testing environments

Resilience defects

```
1 @PostMapping("/pong")
2 public ResponseEntity<Pong> pongPost(RestTemplate rest, @RequestBody Increment i) {
3     long inc = i.increment();
4     if(inc >= 0 && inc < 10000) {
5         long cur = counter.addAndGet(inc);
6         return ResponseEntity.ok().body(new Pong(cur));
7     } else {
8         return ResponseEntity.badRequest().build();
9     }
10 }
```

200 OK if in range

400 Bad Request if out of range

```
1 @PostMapping("/ping")
2 public ResponseEntity<Ping> pingPost(RestTemplate rest, @RequestBody Increment i) {
3     long cur;
4     ResponseEntity<Pong> ponge;
5     try {
6         ponge = rest.postForEntity("http://pong/pong", i, Pong.class);
7         cur = counter.addAndGet(i.increment());
8     } catch (HttpClientErrorException e) {
9         logger.warn(e);
10        return ResponseEntity.badRequest().build();
11    } catch (RestClientException e) {
12        // Retry
13        logger.error(e);
14        ponge = rest.postForEntity("http://pong/pong", i, Pong.class);
15        cur = counter.addAndGet(i.increment());
16    }
17    Pong pong = ponge.getBody();
18    return ResponseEntity.ok(new Ping(cur, pong.id()));
19 }
```

Check and handle 400 response

Check for connection error and retry

Fails on two errors
Does not handle 400 response

Can't be found without fault injection!

Resilience defects

```
1 @PostMapping("/pong")
2 public ResponseEntity<Pong> pongPost(RestTemplate rest, @RequestBody Increment i) {
3     long inc = i.increment();
4     if(inc >= 0 && inc < 10000) {
5         long cur = counter.addAndGet(inc);
6         return ResponseEntity.ok().body(new Pong(cur));
7     } else {
8         return ResponseEntity.badRequest().build();
9     }
10 }
```

200 OK if in range

400 Bad Request if out of range

1. These bugs won't be found without fault injection
2. Some bugs require multiple faults
3. Some bugs require specific inputs and faults

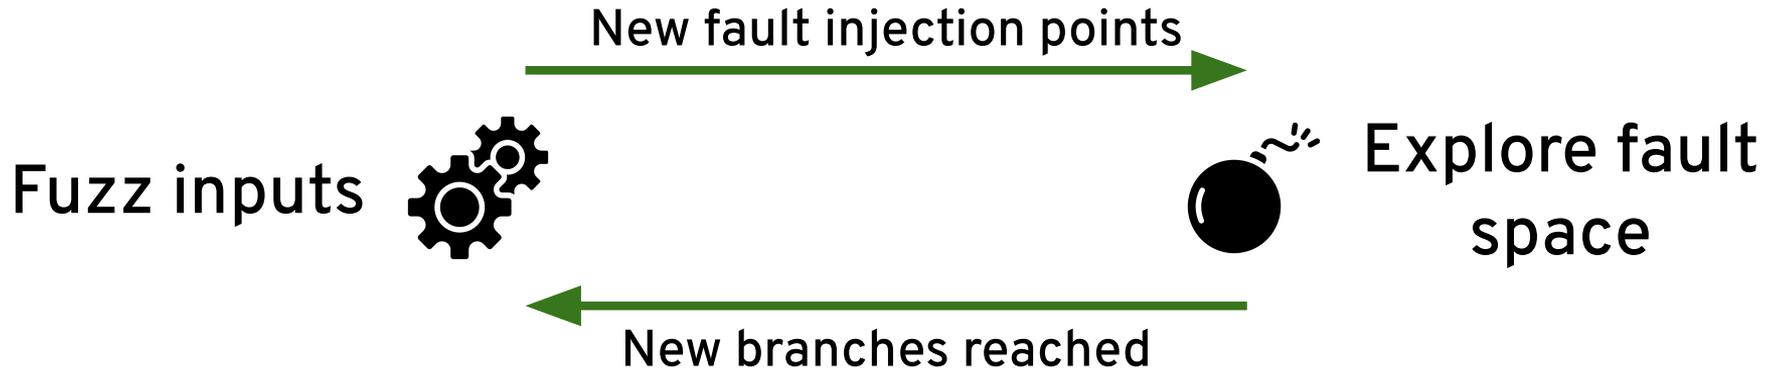
```
1 @PostMapping
2 public Respo
3     long cur
4     Response
5     try {
6         pong
7         cur
8     } catch(
9         logg
10        retu
11    } catch(RestClientException e) {
12        // Retry
13        logger.error(e);
14        ponge = rest.postForEntity("http://pong/pong", i, Pong.class);
15        cur = counter.addAndGet(i.increment());
16    }
17    Pong pong = ponge.getBody();
18    return ResponseEntity.ok(new Ping(cur, pong.id()));
19 }
```

Check for connection error and retry

Fails on two errors
Does not handle 400 response

Can't be found without fault injection!

Faults and inputs



We can use fuzzing to explore faults as well!

Fault injection



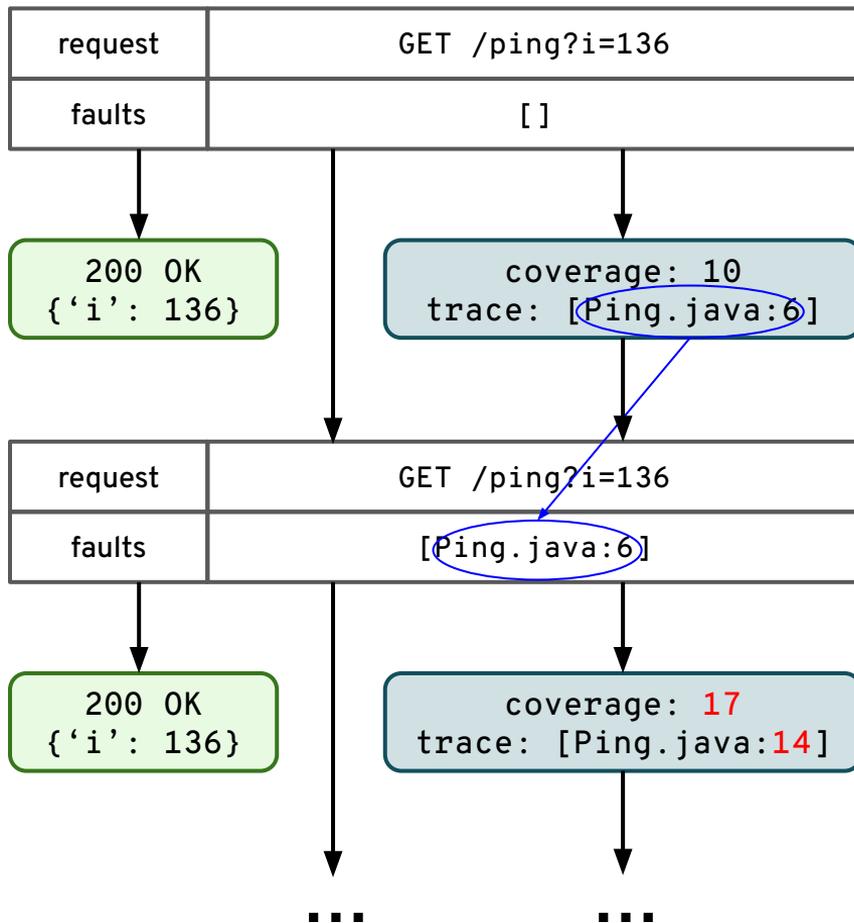
Faults are part of the seed along with request



Fault injection points are found during execution



Faults added to seed using mutators



Onweer



Automated resilience tester
based on these concepts



Faults are collected and injected by
an instrumentation agent



Faults are connection error exceptions
thrown during REST requests



Evaluation

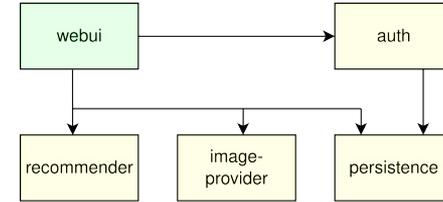
? RQ1: Does fault injection increase coverage?

? RQ2: Does fault injection cover code fuzzing alone can't reach?

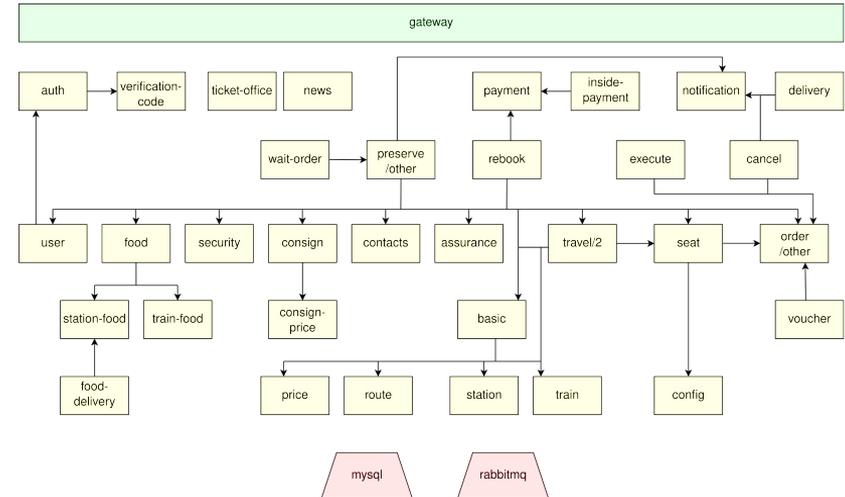
? RQ3: Does adding fault injection find resilience defects?

! Method: Compare Onweer with and without fault injection and RESTler

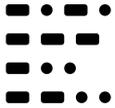
TeaStore [von Kistowski et al. 2018]



TrainTicket [Zhou et al. 2018]



RQ1: Code coverage



Coverage improves for TeaStore!



But not for TrainTicket?

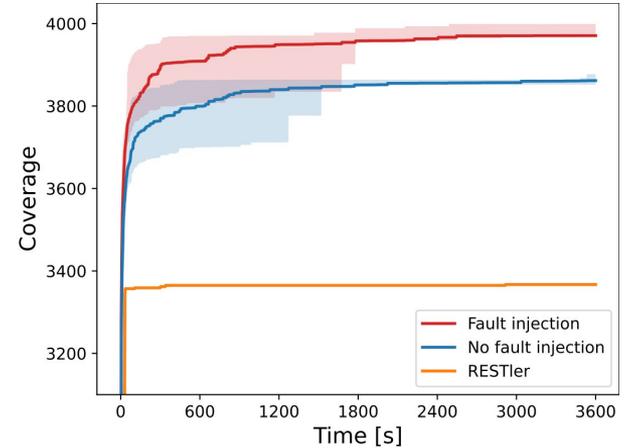


RQ2: Because all coverage increase is in fault handlers

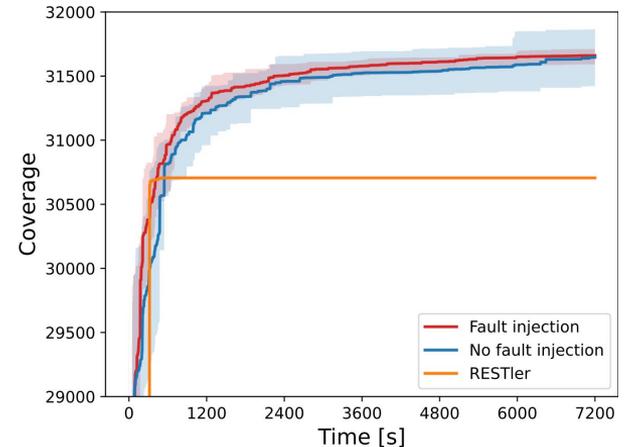


TrainTicket has no fault handling

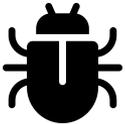
TeaStore



TrainTicket



RQ3: Finding resilience defects



We find several resilience defects in both projects!

TrainTicket

request	GET /api/v1/foodservice/foods/y%C2%87%C2%80j/%C3%83%C2%A8/%24%C3%B2%F2%95%92%85/%F1%8C%9F%8E%F3%A4%BC%9E%1E%C2%8F
faults	[food-FoodServiceImpl.java:242]

500 Internal Server Error
Intercession at food-FoodS...

RQ3: Finding resilience defects



We find several resilience defects in both projects!

TeaStore

request	GET /products
---------	---------------

200 OK
[{'id': 54, 'name': "Black Tea"}, ...]

request	POST /cart/add {'id': ???}
---------	----------------------------

links	index: 0 from: [0]['id'] to: ['id']
-------	---

faults	[webui-CartRest.java:67, webui-CartRest.java:67, webui-CartRest.java:67]
--------	--

500 Internal Server Error
ServiceLoadBalancerResult was null!

Conclusion

- Are you applying fuzzing, chaos engineering, resilience testing?
 - We would love to hear about your experiences and challenges
 - We're eager to put our tools to the test on your code
- We can help:
 - Test and secure your devices using fuzzing
 - Test and secure your backends using fuzzing
 - Tailor fault injection to the specifics of your application
 - Develop bespoke fuzzers
- Within individual and collaborative research and development projects (O&O, ICON, ...)