

AFLNet Five Years Later: On Coverage- Guided Protocol Fuzzing

Ruijie Meng

PhD student at NUS

Incoming tenure-track faculty at CISPA

Co-authors: Van-Thuan Pham, Marcel Böhme and Abhik Roychoudhury

What is AFLNet?

First **code**- and **state**-coverage guided *protocol* fuzzer (ICST'20, Tool)

March 2020




Released on GitHub

Contributions to Community

First **code**- and **state**-coverage guided *protocol* fuzzer (ICST'20, Tool)

March 2020

June 2025



Released on GitHub

Research Contribution

Practice Contribution

Education Contribution

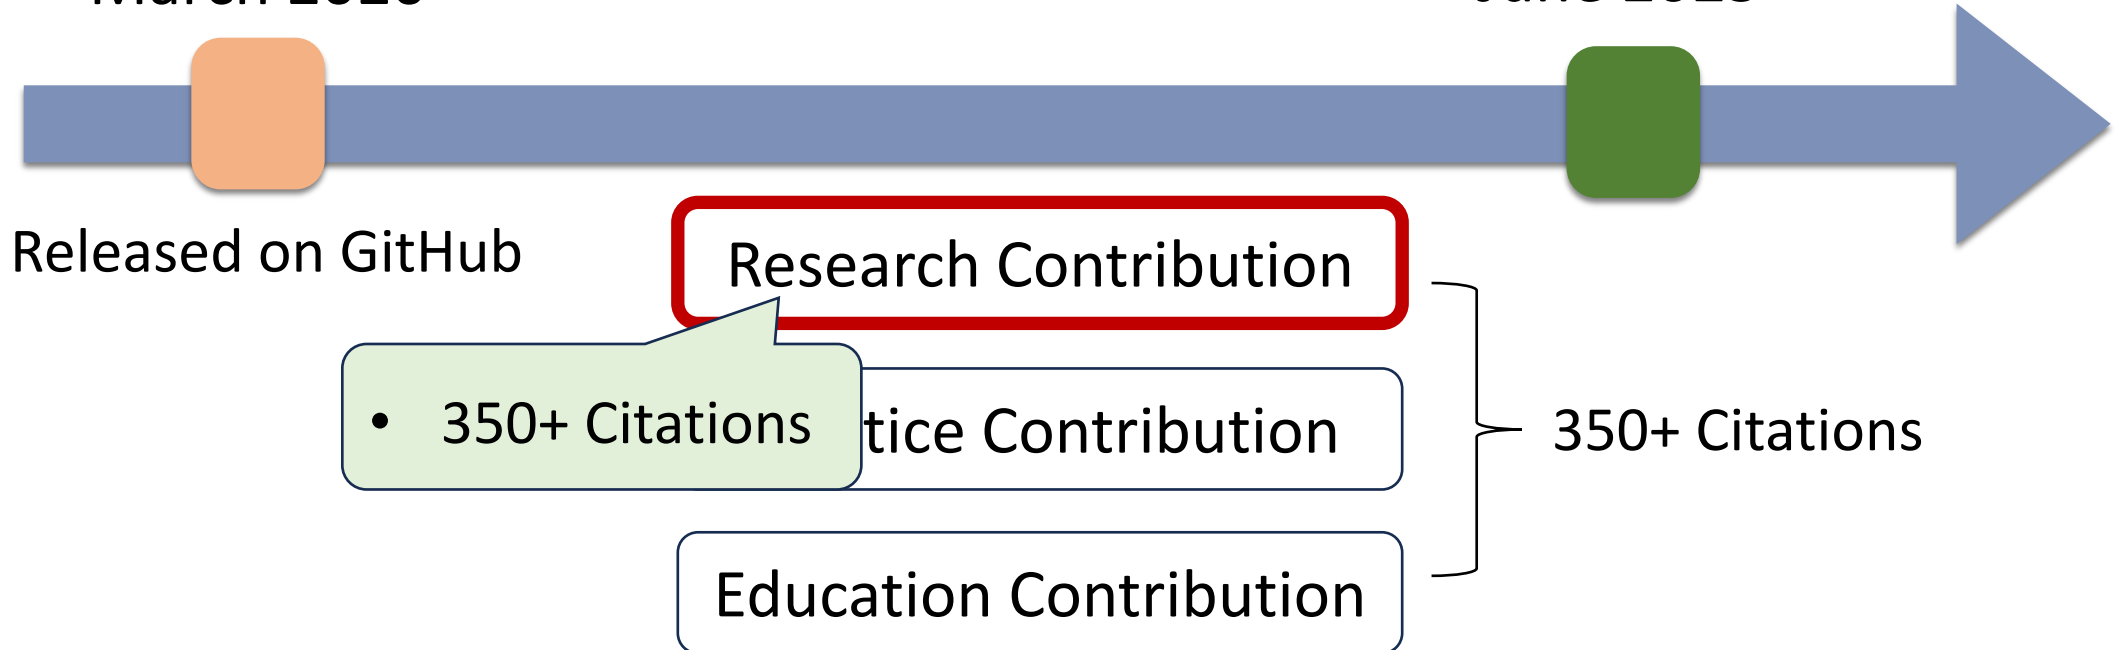


Contributions to Community

First **code**- and **state**-coverage guided *protocol* fuzzer (ICST'20, Tool)

March 2020

June 2025



Contributions to Community

First **code**- and **state**-coverage guided *protocol* fuzzer (ICST'20, Tool)

March 2020

June 2025



Released on GitHub
RTSP, FTP, SSH, TLS,
SMTP

Research Contribution

Practice Contribution

350+ Citations, 870+ Stars,
17 protocols, more targets


- 870+ stars on GitHub and 17 protocols supported
- Used for challenging targets, e.g., 5G protocols, smart home ecosystem, medical imaging applications and automotive systems

Contributions to Community

First **code**- and **state**-coverage guided *protocol* fuzzer (ICST'20, Tool)

March 2020

June 2025



Released on GitHub

Research Contribution

Practice Contribution

Education Contribution

- Introduced in graduate courses, e.g., at University of Melbourne and Carnegie Mellon University

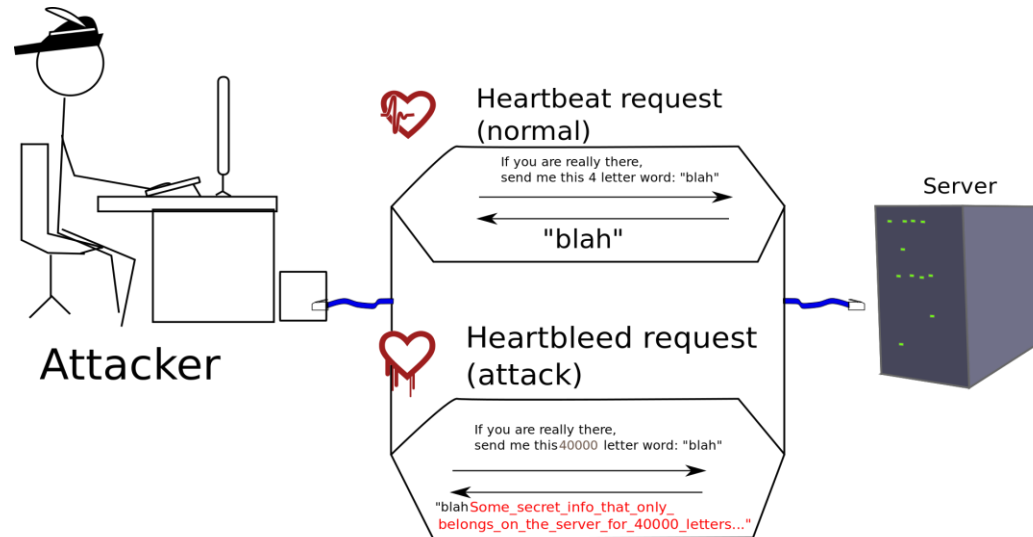
350+ Citations, 870+ Stars,
17 protocols, more targets,
course materials

**Why has AFLNet generated such impact
in a short period?**

Why do we need stateful protocol fuzzing?

Testing Protocols is Important

- Network protocols are backbone of critical infrastructure
- Bugs in network protocols damage a lot



*Heartbleed Vulnerability in
OpenSSL in 2014*

Network protocols must be automatically and continuously tested
for security vulnerabilities

But...Testing Protocols is Challenging

- **Statefulness** and **sequences of inputs** pose challenges

```
220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed.
MKD demo
257 Directory created.
CWD demo
250 Requested file action okay, completed.
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
```

*FTP Example of message exchange between client (**red**) and server (black)*

But...Testing Protocols is Challenging

- **Statefulness** and **sequences of inputs** pose challenges

```
220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed
MKD demo
257 Directory created.
CWD demo
250 Requested file action
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
```

Statefulness:

sending the *same* input message twice might yield a *different* response every time based on *states*

FTP Example of message exchange between client (red) and server (black)

But...Testing Protocols is Challenging

- **Statefulness** and **sequences of inputs** pose challenges


```
220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed.
MKD demo
257 Directory created.
CWD demo
250 Requested file action okay, completed
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
```

Sequence Input:
Reaching a state
requires a sequence
of messages

FTP Example of message exchange between client (red) and server (black)

Existing Technique – Model Checking

- A verification technique, but in practice for “bug finding”



Existing Technique – Model Checking

- A verification technique, but in practice for “bug finding”
- $\mathcal{M} \models \boxed{\varphi}$
- Effective but with Limitations:
 - A temporal logic property needs to be provided

Existing Technique – Model Checking

- A verification technique, but in practice for “bug finding”
- $\boxed{\mathcal{M}}| = \varphi$
- Effective but with Limitations:
 - A temporal logic property needs to be provided
 - System modeling is not trivial

inconsistently with each other [12, 18]. Most approaches abstract away the environment behind a model [2, 11], but writing abstract models is labor-intensive (taking in some cases multiple person-years [2]), models are rarely 100% accurate, and they tend to lose

“S2E: A Platform for In-Vivo Multi-Path Analysis of Software Systems”, ASPLOS’11

and the CompCert compiler [120]. Yet, there are, unfortunately, two main inter-related limitations of formally-verified systems: (1) they require years of PhD-level expertise, with the specification often larger than the verified code itself, and (2) the resulting systems lack many of the features and/or performance of their non-verified counterparts. In

“Software Security Analysis in 2030 and Beyond: A Research Roadmap”, TOSEM’25

Existing Technique – Model Checking

- A verification technique, but in practice for “bug finding”
- $\mathcal{M} \models \varphi$
- Effective but with Limitations:
 - A temporal logic property needs to be provided
 - System modeling is not trivial
 - Bugs are reported in a counter-example trace not inputs



Existing Technique – Fuzzing

- Stateful blackbox fuzzing (e.g., Peach)
 - Writing models still involves much manual effort and expertise
 - Learn nothing from past execution
- Stateless greybox fuzzing (e.g., AFL)
 - Almost no manual effort and expertise required
 - Neither know states or message sequences

“One of the things that I struggle with is the limitation AFL seems to have, in that it only performs fuzzing with one input (a file). For many systems, such as network protocols, it would be useful if fuzzing could be done on a sequence of inputs. This sequence of inputs might be for example messages necessary to complete a handshake in TLS/TCP.”

- Paul (a member of the AFL's user group) [5]

“I'm interested in doing something fairly non-traditional and definitely not currently supported by AFL. I would like to perform fuzzing of a large and complex external server that cannot easily be stripped down into small test cases.”

- Tim Newsham (a member of the AFL's user group) [5]

Why has AFLNet generated such impact in a short period?

- Providing a practical solution for this long-standing problem

AFLNet Approach

```


Input : Server program  $\mathcal{P}$ , Sniffer traces  $T$ , IPSM  $S$ 
Output: Crashes  $C_{\times}$ , Corpus  $C$ , and IPSM  $S$ 
1 Corpus  $C \leftarrow \emptyset$ ; Crashes  $C_{\times} \leftarrow \emptyset$ ; Bitmap  $B \leftarrow \emptyset$ 
2 for each trace  $t \in T$  do                                ▷ Pre-processing Phase
3   Sequence  $M \leftarrow \text{parse}(t)$ 
4   Corpus  $C \leftarrow C \cup \{M\}$ 
5   Response  $R \leftarrow \text{send}(\mathcal{P}, M, B)$ 
6   IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
7 LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
8 repeat                                                    ▷ Fuzzing Phase
9   if ( $\text{cur\_time} - lpt > \text{MaxTimeGap}$ ) then
10    State  $s \leftarrow \text{choose\_state}(S)$ 
11    Sequence  $M \leftarrow \text{choose\_sequence\_to\_state}(C, s)$ 
12     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
        (i.e., split  $M$  in subsequences such
        that  $M_1$  is the message sequence to
        drive  $\mathcal{P}$  to arrive at state  $s$ , and mes-
        sage sequence  $M_2$  is selected to be
        mutated)
13  else                                                    ▷ Interleaving Seed Selection
14    Sequence  $M \leftarrow \text{choose\_sequence\_from\_queue}(C)$ 
15     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
        (i.e., randomly select subsequence
         $M_2$  to be mutated)
16  for  $i$  from 1 to  $\text{energy}(M)$  do
17    Sequence  $M' \leftarrow \langle M_1, \text{mutate}(M_2), M_3 \rangle$ 
18    Response  $R \leftarrow \text{send}(\mathcal{P}, M', B)$ 
19    if  $\mathcal{P}$  has crashed then
20      Crashes  $C_{\times} \leftarrow C_{\times} \cup \{M'\}$ 
21      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
22    else if  $\text{is\_interesting}(M', B)$  then
23      Corpus  $C \leftarrow C \cup \{M'\}$ 
24      IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
25      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
26 until timeout reached or abort
  
```

Recording and replay
for fuzzing

AFLNet Approach

```

220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed.
MKD demo
257 Directory created.
CWD demo
250 Requested file action okay, completed.
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
    
```



Input : Server program \mathcal{P} , Sniffer traces T , IPSM S
Output: Crashes C_x , Corpus C , and IPSM S


```

1 Corpus  $C \leftarrow \emptyset$ ; Crashes  $C_x \leftarrow \emptyset$ ; Bitmap  $B \leftarrow \emptyset$ 
2 for each trace  $t \in T$  do                                ▷ Pre-processing Phase
3   Sequence  $M \leftarrow \text{parse}(t)$ 
4   Corpus  $C \leftarrow C \cup \{M\}$ 
5   Response  $R \leftarrow \text{send}(\mathcal{P}, M, B)$ 
6   IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
7 LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
8 repeat                                                    ▷ Fuzzing Phase
9   if ( $\text{cur\_time} - lpt$ ) >  $\text{MaxTimeGap}$  then
10    State  $s \leftarrow \text{choose\_state}(S)$ 
11    Sequence  $M \leftarrow \text{choose\_sequence\_to\_state}(C, s)$ 
12     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
    (i.e., split  $M$  in subsequences such
    that  $M_1$  is the message sequence to
    drive  $\mathcal{P}$  to arrive at state  $s$ , and mes-
    sage sequence  $M_2$  is selected to be
    mutated)
13  else                                                    ▷ Interleaving Seed Selection
14    Sequence  $M \leftarrow \text{choose\_sequence\_from\_queue}(C)$ 
15     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
    (i.e., randomly select subsequence
     $M_2$  to be mutated)
16  for  $i$  from 1 to  $\text{energy}(M)$  do
17    Sequence  $M' \leftarrow \langle M_1, \text{mutate}(M_2), M_3 \rangle$ 
18    Response  $R \leftarrow \text{send}(\mathcal{P}, M', B)$ 
19    if  $\mathcal{P}$  has crashed then
20      Crashes  $C_x \leftarrow C_x \cup \{M'\}$ 
21      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
22    else if  $\text{is\_interesting}(M', B)$  then
23      Corpus  $C \leftarrow C \cup \{M'\}$ 
24      IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
25      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
26 until timeout reached or abort
    
```

AFLNet Approach

```

220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed.
MKD demo
257 Directory created.
CWD demo
250 Requested file action okay, completed.
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
    
```



Input : Server program \mathcal{P} , Sniffer traces T , IPSM S

Output: Crashes C_x , Corpus C , and IPSM S


```

1 Corpus  $C \leftarrow \emptyset$ ; Crashes  $C_x \leftarrow \emptyset$ ; Bitmap  $B \leftarrow \emptyset$ 
2 for each trace  $t \in T$  do                                ▷ Pre-processing Phase
3   Sequence  $M \leftarrow \text{parse}(t)$ 
4   Corpus  $C \leftarrow C \cup \{M\}$ 
5   Response  $R \leftarrow \text{send}(\mathcal{P}, M, B)$ 
6   IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
7 LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
8 repeat                                                    ▷ Fuzzing Phase
9   if  $(\text{cur\_time} - lpt) > \text{MaxTimeGap}$  then
10    State  $s \leftarrow \text{choose\_state}(S)$ 
11    Sequence  $M \leftarrow \text{choose\_sequence\_to\_state}(C, s)$ 
12     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
13    (i.e., split  $M$  in subsequences such
14    that  $M_1$  is the message sequence to
15    drive  $\mathcal{P}$  to arrive at state  $s$ , and mes-
16    sage sequence  $M_2$  is selected to be
17    mutated)
18  else                                                    ▷ Interleaving Seed Selection
19    Sequence  $M \leftarrow \text{choose\_sequence\_from\_queue}(C)$ 
20     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
21    (i.e., randomly select subsequence
22     $M_2$  to be mutated)
23  for  $i$  from 1 to  $\text{energy}(M)$  do
24    Sequence  $M' \leftarrow \langle M_1, \text{mutate}(M_2), M_3 \rangle$ 
25    Response  $R \leftarrow \text{send}(\mathcal{P}, M', B)$ 
26    if  $\mathcal{P}$  has crashed then
27      Crashes  $C_x \leftarrow C_x \cup \{M'\}$ 
28      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
29    else if  $\text{is\_interesting}(M', B)$  then
30      Corpus  $C \leftarrow C \cup \{M'\}$ 
31      IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
32      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
33 until timeout reached or abort
    
```

AFLNet Approach

```

220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed.
MKD demo
257 Directory created.
CWD demo
250 Requested file action okay, completed.
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
    
```



Input : Server program \mathcal{P} , Sniffer traces T , IPSM S

Output: Crashes C_x , Corpus C , and IPSM S


```

1 Corpus  $C \leftarrow \emptyset$ ; Crashes  $C_x \leftarrow \emptyset$ ; Bitmap  $B \leftarrow \emptyset$ 
2 for each trace  $t \in T$  do                                ▷ Pre-processing Phase
3   Sequence  $M \leftarrow \text{parse}(t)$ 
4   Corpus  $C \leftarrow C \cup \{M\}$ 
5   Response  $R \leftarrow \text{send}(\mathcal{P}, M, B)$ 
6   IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
7 LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
8 repeat                                                    ▷ Fuzzing Phase
9   if ( $\text{cur\_time} - lpt$ ) >  $\text{MaxTimeGap}$  then
10    State  $s \leftarrow \text{choose\_state}(S)$ 
11    Sequence  $M \leftarrow \text{choose\_sequence\_to\_state}(C, s)$ 
12     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
    (i.e., split  $M$  in subsequences such
    that  $M_1$  is the message sequence to
    drive  $\mathcal{P}$  to arrive at state  $s$ , and mes-
    sage sequence  $M_2$  is selected to be
    mutated)
13  else                                                    ▷ Interleaving Seed Selection
14    Sequence  $M \leftarrow \text{choose\_sequence\_from\_queue}(C)$ 
15     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
    (i.e., randomly select subsequence
     $M_2$  to be mutated)
16  for  $i$  from 1 to  $\text{energy}(M)$  do
17    Sequence  $M' \leftarrow \langle M_1, \text{mutate}(M_2), M_3 \rangle$ 
18    Response  $R \leftarrow \text{send}(\mathcal{P}, M', B)$ 
19    if  $\mathcal{P}$  has crashed then
20      Crashes  $C_x \leftarrow C_x \cup \{M'\}$ 
21      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
22    else if  $\text{is\_interesting}(M', B)$  then
23      Corpus  $C \leftarrow C \cup \{M'\}$ 
24      IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
25      LastPathTime  $lpt \leftarrow \text{cur\_time}$ 
26 until timeout reached or abort
    
```


AFLNet Approach

```

220 LightFTP server v2.0a ready
USER foo
331 User foo OK. Password required
PASS foo
230 User logged in, proceed.
MKD demo
257 Directory created.
CWD demo
250 Requested file action okay, completed.
STOR test.txt
150 File status okay
226 Transfer complete
LIST
150 File status okay
226 Transfer complete
QUIT
221 Goodbye!
    
```



Input : Server program \mathcal{P} , Sniffer traces T , IPSM S

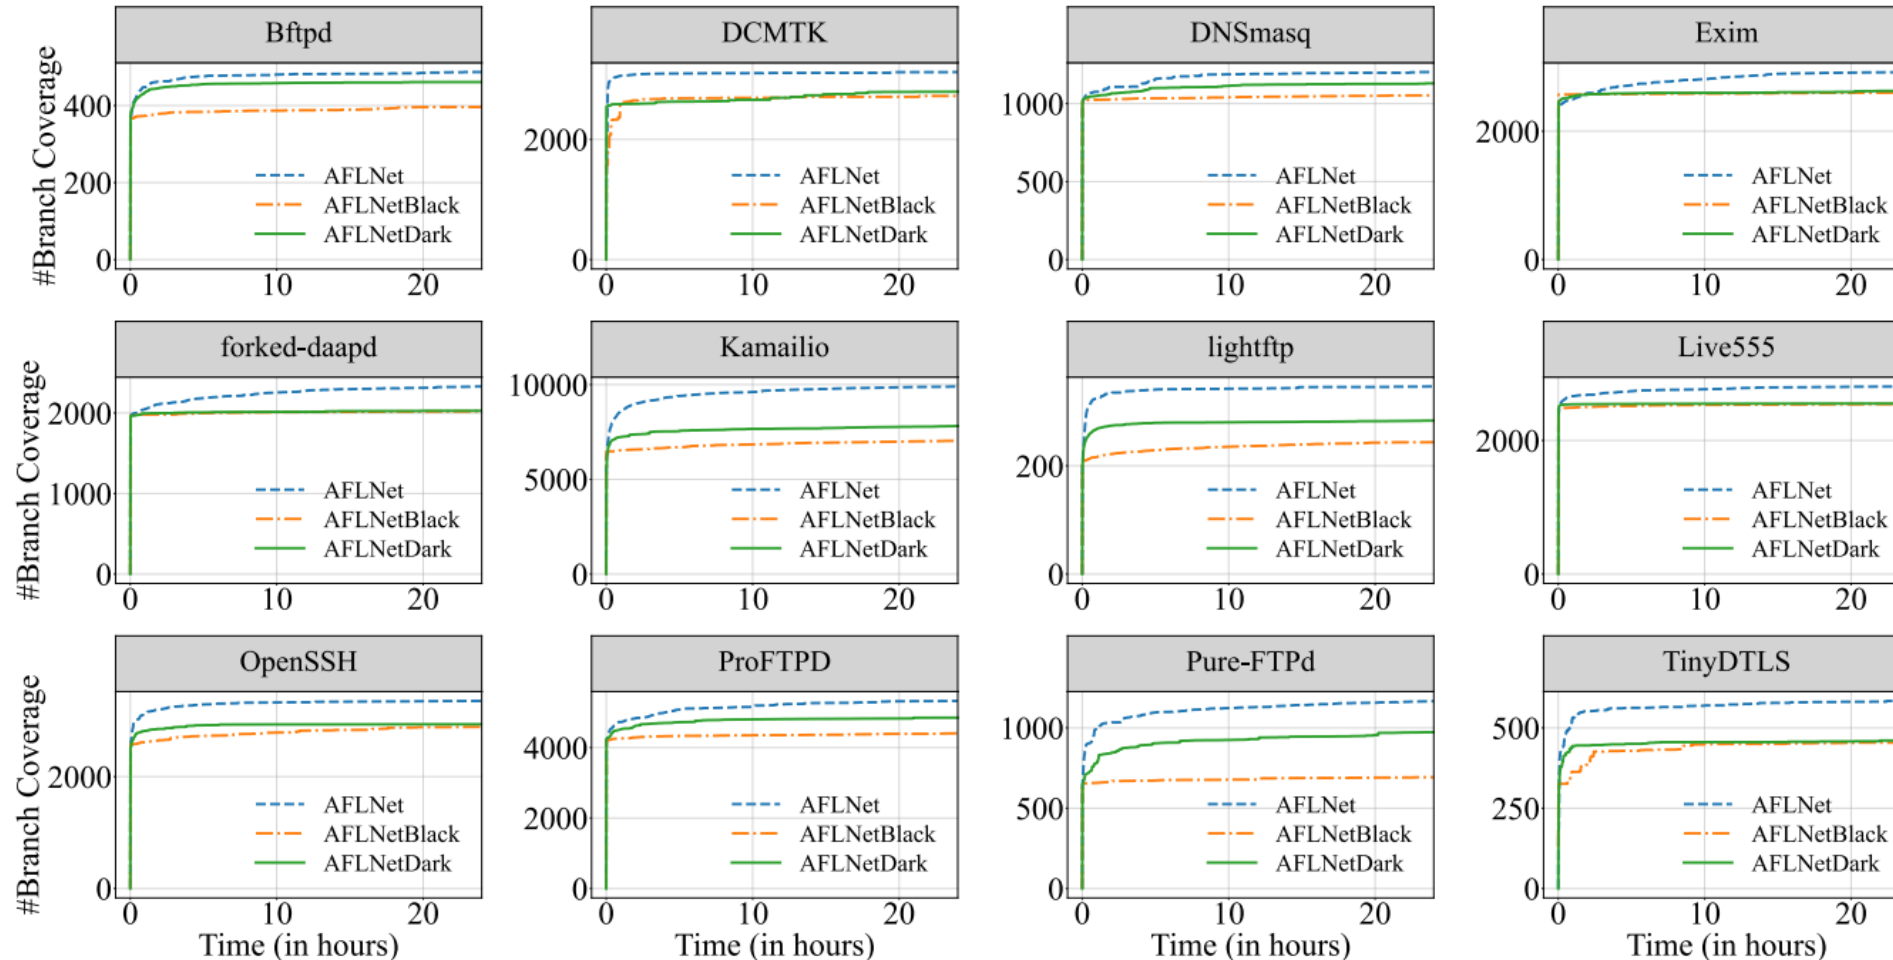
Output: Crashes C_x , Corpus C , and IPSM S

```

1 Corpus  $C \leftarrow \emptyset$ ; Crashes  $C_x \leftarrow \emptyset$ ; Bitmap  $B \leftarrow \emptyset$ 
2 for each trace  $t \in T$  do                                ▷ Pre-processing Phase
3   Sequence  $M \leftarrow \text{parse}(t)$ 
4   Corpus  $C \leftarrow C \cup \{M\}$ 
5   Response  $R \leftarrow \text{send}(\mathcal{P}, M, B)$ 
6   IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
7 LastPathTime  $\text{lpt} \leftarrow \text{cur\_time}$ 
8 repeat                                                    ▷ Fuzzing Phase
9   if ( $\text{cur\_time} - \text{lpt} > \text{MaxTimeGap}$ ) then
10    State  $s \leftarrow \text{choose\_state}(S)$ 
11    Sequence  $M \leftarrow \text{choose\_sequence\_to\_state}(C, s)$ 
12     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
    (i.e., split  $M$  in subsequences such
    that  $M_1$  is the message sequence to
    drive  $\mathcal{P}$  to arrive at state  $s$ , and mes-
    sage sequence  $M_2$  is selected to be
    mutated)
13  else                                                    ▷ Interleaving Seed Selection
14    Sequence  $M \leftarrow \text{choose\_sequence\_from\_queue}(C)$ 
15     $\langle M_1, M_2, M_3 \rangle \leftarrow M$ 
    (i.e., randomly select subsequence
     $M_2$  to be mutated)
16  for  $i$  from 1 to  $\text{energy}(M)$  do
17    Sequence  $M' \leftarrow \langle M_1, \text{mutate}(M_2), M_3 \rangle$ 
18    Response  $R \leftarrow \text{send}(\mathcal{P}, M', B)$ 
19    if  $\mathcal{P}$  has crashed then
20      Crashes  $C_x \leftarrow C_x \cup \{M'\}$ 
21      LastPathTime  $\text{lpt} \leftarrow \text{cur\_time}$ 
22    else if  $\text{is\_interesting}(M', B)$  then
23      Corpus  $C \leftarrow C \cup \{M'\}$ 
24      IPSM  $S \leftarrow \text{updateIPSM}(S, R)$ 
25      LastPathTime  $\text{lpt} \leftarrow \text{cur\_time}$ 
26 until timeout reached or abort
    
```

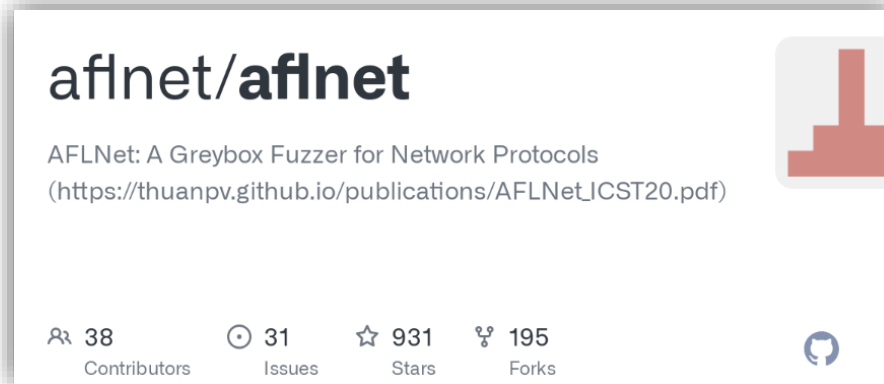
Multiple Feedback Modes

- Coverage in black-, dark-, and grey-modes



Open-Sourced Tool

- AFLNet acts as the client
 - Follow real-world architectures of network protocols
 - Reduce manual effort in understanding protocols or modifying source code



however, a better option is AFLnet (<https://github.com/aflnet/aflnet>) which allows you to define network state with different type of data packets.

best practice in fuzzing network services from the AFL++ document

Why has AFLNet generated such impact in a short period?

- Providing a practical solution for this long-standing problem
- **Open Science Approach**

More Research is In Progress

Recent Progress in Stateful Fuzzing

What is a **state**?

- SGFuzz (Usenix Sec'22)
- StateAFL (EMSE'22)
- ChatAFL (NDSS'24)
- ...

How to maximize **syntactic validity** of each message?

- ChatAFL (NDSS'24)
- ...

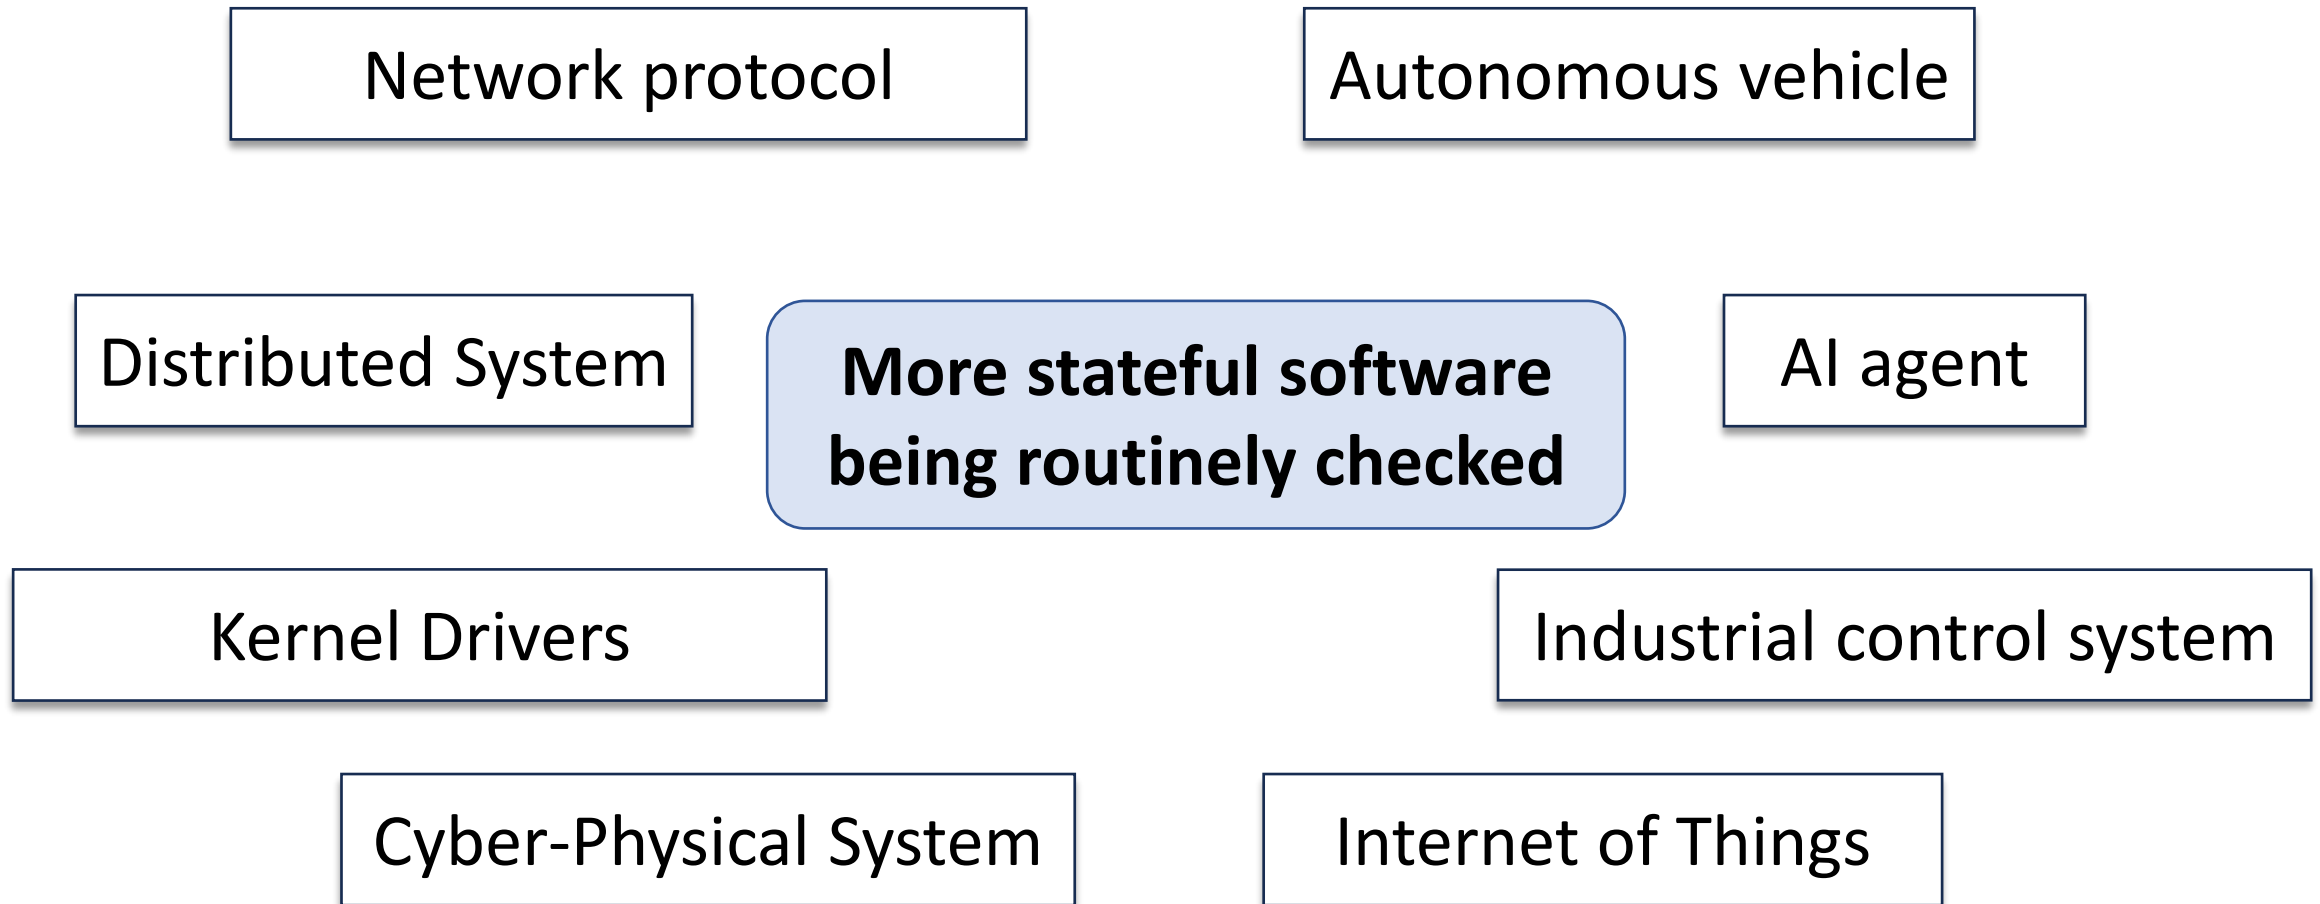
Protocol **Environment** Fuzzing

- ChaosAFL (ArXiv'23)
- EnvFuzz (CCS'23)
- ...

How to maximize **fuzzing throughput**?

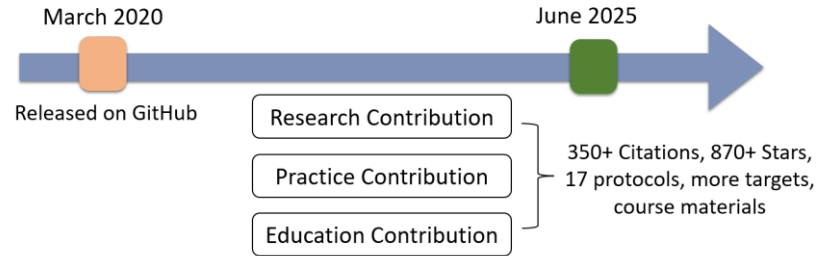
- Nyx-Net (EuroSys'22)
- SnapFuzz (ISSTA'22)
- ...

Path Forward

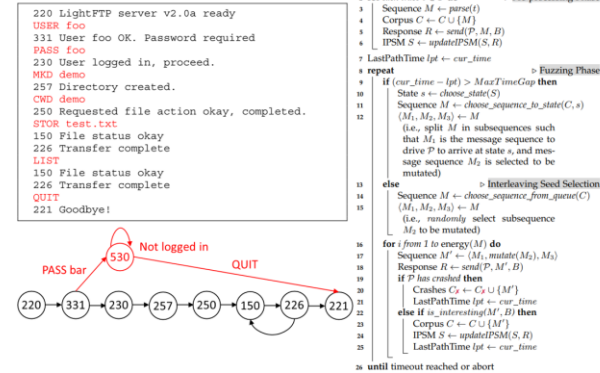


What is AFLNet?

First **code**- and **state**-coverage guided **protocol** fuzzer (ICST'20, Tool)



AFLNet Approach



Why has AFLNet generated such impact in a short period?

- Providing a practical solution for this long-standing problem
- Open Science Approach

Recent Progress in Stateful Fuzzing

What is a **state**?

- SGFuzz (Usenix Sec'22)
- StateAFL (EMSE'22)
- ChatAFL (NDSS'24)
- ...

How to maximize **syntactic validity** of each message?

- ChatAFL (NDSS'24)
- ...

Protocol **Environment** Fuzzing

- ChaosAFL (ArXiv'23)
- EnvFuzz (CCS'23)
- ...

How to maximize **fuzzing throughput**?

- Nyx-Net (EuroSys'22)
- SnapFuzz (ISSTA'22)
- ...

Thanks!