

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

11-2020

### CoinWatch: A clone-based approach for detecting vulnerabilities in cryptocurrencies

Qingze HUM

*Singapore University of Technology and Design*

Wei Jin TAN

*Singapore University of Technology and Design*

Shi Ying TEY

*Singapore University of Technology and Design*

Latasha LENUS

*Singapore University of Technology and Design*

Ivan HOMOLIAK

*Technical University of Brno*

See next page for additional authors  
Show this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Information Security Commons](#), [Software Engineering Commons](#), and the [Technology and Innovation Commons](#)

---

#### Citation

HUM, Qingze; TAN, Wei Jin; TEY, Shi Ying; LENUS, Latasha; HOMOLIAK, Ivan; LIN, Yun; and SUN, Jun. CoinWatch: A clone-based approach for detecting vulnerabilities in cryptocurrencies. (2020). *2020 3rd IEEE International Conference on Blockchain: November 2-6, Rhodes, Greece: Proceedings*. 17-25. Available at: [https://ink.library.smu.edu.sg/sis\\_research/6031](https://ink.library.smu.edu.sg/sis_research/6031)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylids@smu.edu.sg](mailto:cherylids@smu.edu.sg).

---

**Author**

Qingze HUM, Wei Jin TAN, Shi Ying TEY, Latasha LENUS, Ivan HOMOLIAK, Yun LIN, and Jun SUN

# CoinWatch: A Clone-Based Approach For Detecting Vulnerabilities in Cryptocurrencies

Qingze Hum,<sup>†</sup> Wei Jin Tan, Shi Ying Tey, Latasha Lenus  
Singapore University of Technology and Design,  
<sup>†</sup>qzhum1996@gmail.com

Yun Lin  
National University of Singapore,  
dcsliny@nus.edu.sg

Ivan Homoliak  
FIT, Brno University of Technology,  
ihomoliak@fit.vutbr.cz

Jun Sun  
Singapore Management University,  
sunjunhqq@gmail.com

**Abstract**—Cryptocurrencies have become very popular in recent years. Thousands of new cryptocurrencies have emerged, proposing new and novel techniques that improve on Bitcoin’s core innovation of the blockchain data structure and consensus mechanism. However, cryptocurrencies are a major target for cyber-attacks, as they can be sold on exchanges anonymously and most cryptocurrencies have their codebases publicly available. One particular issue is the prevalence of code clones in cryptocurrencies, which may amplify security threats. If a vulnerability is found in one cryptocurrency, it might be propagated into other cloned cryptocurrencies. In this work, we propose a systematic remedy to this problem, called COINWATCH (*CW*). Given a reported vulnerability at the input, *CW* uses the code evolution analysis and a clone detection technique for the indication of cryptocurrencies that might be vulnerable. We applied *CW* on 1094 cryptocurrencies using 4 CVEs and obtained 786 true vulnerabilities present in 384 projects, which were confirmed with developers and successfully reported as CVE extensions.

## I. INTRODUCTION

Bitcoin is designed to be a peer-to-peer electronic cash system, introducing an innovative consensus mechanism and incentives for participants. All participants in the consensus protocol must run a full node. Bitcoin has inspired hundreds of cryptocurrencies utilizing blockchain technology to serve as a decentralized means of transferring crypto-assets. There were 98 cryptocurrencies introduced in 2016, 412 introduced in 2017, and 726 introduced in 2018. Given that Bitcoin’s codebase [1] is publicly available, developers may easily fork its repository from GitHub and make custom modifications, which they might later publish as their own cryptocurrency. In this way, many cryptocurrencies might become popular and adopted by the public. However, many cryptocurrencies are also “popular” due to various high profile security incidents [2], [3], [4], [5].

Detecting and patching the vulnerabilities manually is prohibitively costly. Due to the decentralized nature of cryptocur-




Figure 1: A histogram of forked projects from Bitcoin and their clone ratios as compared to Bitcoin v0.17.0.

rencies, whenever a patch is created, individual miners have to patch their own nodes independently. In the case of the inflation bug in Bitcoin, the vulnerability was fixed by only 31.65% of all nodes [6] even after a full year from the disclosure.<sup>1</sup> The inflation bug [7] was discovered in September 2018 and allowed the inflation of the total Bitcoin supply. Upon analysis, it was found that this vulnerability had been present since 2017 but has been patched only in 2019. Nevertheless, this vulnerability is shared by 31 cryptocurrencies that emerged as modifications of Bitcoin clones. Patching the vulnerability across all such cryptocurrencies is of paramount importance for their financial stability.

Our empirical study shows that 786 cryptocurrencies were directly or indirectly forked from a version of Bitcoin. Figure 1 depicts the histogram of code similarity (i.e., clone ratio) of cryptocurrencies against the Bitcoin codebase. In most of the cases, the clone ratio is below 30%; however, some of them have this ratio even higher than 50%. Therefore, once a vulnerability is discovered in one cryptocurrency, the vulnerability is very likely to be propagated among its clones. Editing cloned codebases is a common phenomenon

<sup>1</sup>Note that as of 30<sup>th</sup> September 2020, the ratio of updated nodes is even worse and is equal to 14.42%. See <https://luke.dashjr.org/programs/bitcoin/files/charts/security.html?201817144>.

in software engineering and maintenance. Nevertheless, the popularity of forking codebases of cryptocurrencies makes the problem unique in that: (1) the clones are spread across a vast number of projects, which makes overall clone management difficult, (2) the code among cryptocurrency projects is of high similarity, and (3) the financial impact of neglecting a propagated vulnerability might be extremely costly.

In this work, we propose COINWATCH (*CW*), a clone-based approach that indicates propagated vulnerabilities across cryptocurrencies. At the input, *CW* takes the code fragment relevant to a known vulnerability in a *target* cryptocurrency and the code repositories of *monitored* cryptocurrencies. First, *CW* analyzes when the vulnerability of the target cryptocurrency was introduced within its incremental code history (e.g., Git). Next, *CW* goes through the suspicious cryptocurrencies and infers whether they were forked before the vulnerability was introduced. In the positive case, *CW* conducts clone detection to find out whether the given vulnerability has been already fixed and if not it reports a warning.

We have built a proof-of-concept tool for *CW*, which to the date scans 1094 cryptocurrencies and compares them to Bitcoin, serving as the target cryptocurrency. For evaluation purposes, we selected the top four cryptocurrency-oriented CVEs in terms of severity. In sum, *CW* reported 786 propagated vulnerabilities across 384 cryptocurrencies, achieving a true positive rate of 89.7%. We notified the developers of affected cryptocurrencies and reported vulnerabilities found to NVD NIST. In turn, we obtained 4 CVE extensions: **CVE-2018-17144**, **CVE-2016-10724**, **CVE-2016-10725**, and **CVE-2019-7167**, which are described at <https://github.com/JinBean/CVE-Extension/>.

## II. AN ILLUSTRATING EXAMPLE

In this section, we illustrate how *CW* works while taking CVE-2018-17144 as an example.

### A. CVE-2018-17144 Explained

This vulnerability is a result of *reachable assertion* weakness CWE-617 [8], and it enables the crashing of any node in the network by crafting a transaction that double spends (see Figure 2), which causes a DoS of all nodes that receive a mined block containing such a transaction. The vulnerability was fixed in Bitcoin release 0.16.3 on Sep 18, 2018. We depict the vulnerable period of this vulnerability across various Bitcoin clones in Table I.

### B. Using COINWATCH

Given the CVE description and details about a vulnerability, the codebase of the target cryptocurrency (i.e., Bitcoin) is analyzed for the specific code segments causing the vulnerability:

- 1) the bug-fixing commit (b8f8019) is identified in GitHub,
- 2) the bug-introducing commits (eecffe5, 5083079, 3533fb4) are identified in GitHub,
- 3) the first cut of potentially vulnerable cryptocurrencies is generated based on the fork dates of input candidates.

Cryptocurrency	Timeframe (MM/DD)	Months Vulnerable
Argentum	12/16-11/18	23
Chaincoin	12/16-10/18	22
CreativeCoin	12/16-11/18	23
DigiByte	12/16-9/18	21
EliCoin	2/18-11/18	9
Irlecoin	6/18-present	not fixed
Litecoin	12/16-9/18	21
Lynx	12/6-10/18	22
Machinecoin	4/17-10/18	18
MatrixCoin	12/16-10/18	22
Methuselah	4/18-10/18	6
MinexCoin	7/17-4/19	21
MktCoin	10/17-present	not fixed
PlatinCoin	7/18-2/20	19
Quasarcoin	11/17-present	not fixed
United Bitcoin	12/16-9/18	21
Unitus	4/17-11/18	19

Table I: Vulnerable period of CVE-2018-17144 across Bitcoin clones.

Hence, if a cryptocurrency was created before the bug-introducing commits in its “parent,” it is unlikely that it will contain the bug.

In the next step of *CW*, the code segment that patches the vulnerability is identified. Given this information, *CW* performs a test using clone detection to check for the existence of the vulnerable code (and the fix) within all monitored cryptocurrencies. The results mark vulnerable cryptocurrencies: they contain the code with the vulnerability but not the fix.

### C. Applying *CW* on CVE-2018-17144

The full detection process consists of scanning the codebase of all monitored cryptocurrencies (1094) from our dataset while searching for the code that causes the vulnerability. On completion, *CW* reports a candidate set of *suspicious* cryptocurrencies that might be vulnerable with a high likelihood.

**PigeonCoin.** On 26 September 2018, PigeonCoin was attacked [9] due to CVE-2018-17144, resulting in a loss of 235 million PGN (i.e., \$15,000 USD). The fix [10] of the vulnerability was only uploaded to their GitHub repository on 27 September 2018: 9 days after it was added to Bitcoin. If *CW* had monitored the PigeonCoin, the vulnerability would be discovered and fixed before its exploitation.

```

Double-Spend Vulnerability (Bitcoin 0.15.3, validation.cpp):

void UpdateCoins(const CTransaction& tx, CCoinsViewCache& inputs,
    CTxUndo &txundo, int nHeight) {
    // mark inputs spent
    if (!tx.IsCoinBase()) {
        txundo.vprevout.reserve(tx.vin.size());
        for (const CTxIn &txin : tx.vin) {
            txundo.vprevout.emplace_back();
            bool is_spent = inputs.SpendCoin(txin.prevout,
                &txundo.vprevout.back());
            assert(is_spent);
        }
    }
    // add outputs
    AddCoins(inputs, tx, nHeight);
}

```

Figure 2: Propagated code of DoS vulnerability (CVE-2018-17144).

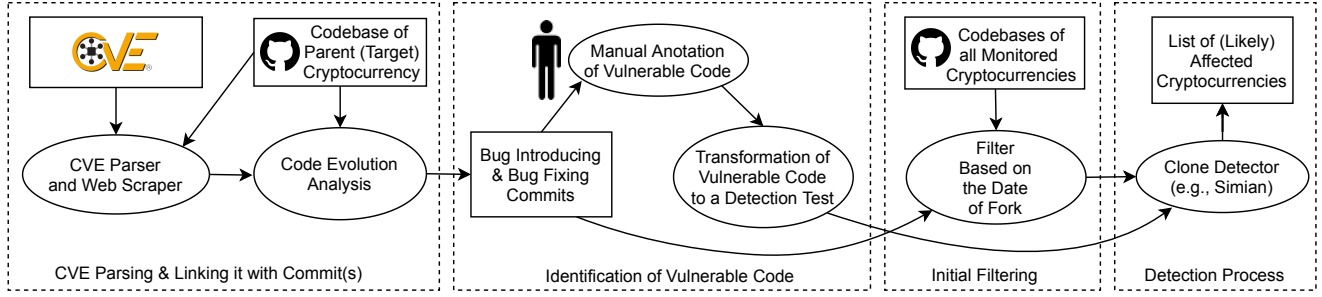


Figure 3: Overview of COINWATCH.

### III. COINWATCH

#### A. Overview

The overall workflow of  $CW$  is shown in Figure 3. First, a target CVE [11] is provided at input together with data publicly obtainable from its structured details. After picking a target CVE,  $CW$  performs a code evolution analysis of the parent project to obtain bug fixing and bug introducing commits. The bug introducing and fixing commits are then manually annotated to minimize the code responsible for the vulnerability and the fix, respectively. Then,  $CW$  generates a detection test for a clone detector. Before running the clone detector, the list of monitored projects is filtered based on the date of the fork to narrow down the search space by pruning unrelated and unaffected projects. Then, at the heart of  $CW$  is the clone detector, which, given the filtered codebases of monitored cryptocurrencies, identifies and reports cloned projects that are likely to be affected by the vulnerability.

#### B. Details

**CVE Parsing.** After a new vulnerability in the parent project is reported, the  $CW$  triggers its execution. First,  $CW$  parses the input CVE and extracts details about CVE: (1) the date of publishing, (2) keywords from the description using a text mining for short texts [12], (3) external links pointing to resources of the version control system (e.g., issues, pull requests, release notes), and (4) the list of cryptocurrencies affected by the CVE along with the programming language they were written in. Then,  $CW$  performs a check whether the vulnerability is dependent on the code specific to the

cryptocurrency since some vulnerabilities may originate from used underlying protocols (e.g., TLS). Finally, the software release versions corresponding to the bug and the fixes are extracted using web scrapping techniques and utilized for a code evolution analysis that takes into account the incremental code history of a project.

**Code Evolution Analysis.**  $CW$  uses the code evolution of a target project (e.g., Bitcoin) to discover the commit IDs containing both the fix and the bug.  $CW$  sets the entry point as the release date of the vulnerability and proceeds back in time while iterating over the issues, until it links the fixing commit(s) with the input CVE. To track down the commit ID,  $CW$  utilizes the SZZ algorithm proposed by Sliwerski, Zimmermann, and Zeller [13]. SZZ algorithm was developed as an approach to identify bug-introducing commits in a software repository. By leveraging on the issue-tracking system of GitHub, we improve the accuracy using the *git-blame* command for finding the bug-introducing commits. An open implementation of the SZZ algorithm called *SZZ Unleashed* [14] provides insight on how to apply it in practice.

In this work, we built on the SZZ algorithm and extend it for our purpose to identify bug-fixing and bug-introducing commits. Using regular expression matching (lines 1 to 5 of Algorithm 1, we search through all issues that have been fixed, resolved, closed, or labeled as a "bug".  $CW$  uses keywords such as "CVE", "CVE-ID" and extracted keywords from the CVE details as the string pattern for regular expression matching. This allows  $CW$  to locate the bug-fixing commits, which are stored at  $G_f$ . To find the bug-introducing commit  $G_b$ ,  $CW$  uses git-blame functionality, which annotates the line of code with information from the revision that modified the line as the last one (i.e., *getPrevCommits()*). We emphasize

**Input :**  $I$ : Issues that match criteria  
**Input :**  $S$ : Regex string  
**Output:**  $G_b := \emptyset$ : Git commits with bugs  
**Output:**  $G_f := \emptyset$ : Git commits with fixes

```

1 foreach  $i$  in  $I$  do
2   | foreach  $c$  in  $i.commits$  do
3     |  $G_f := G_f \cup RE.match(S, c)$ 
4   | end
5 end
6 foreach  $cf$  in  $G_f$  do
7   |  $G_b := G_b \cup getPrevCommits(cf)$ ;
8 end
9 return  $G_f, G_b$ 

```

Algorithm 1: An extended SZZ algorithm.

**Input :**  $V_c$ : set of monitored projects  
**Input :**  $G_b, G_f$ : bug-introducing and fixing commits  
**Output:**  $V_{vuln}$ : set of potentially vulnerable projects

```

1 foreach  $p$  in  $V_c$  do
2   | if  $p.repoForkDate \geq minDate(G_b) \wedge$ 
3     |  $p.repoForkDate \leq maxDate(G_f)$  then
4     |  $V_{vuln} := V_{vuln} \cup p$ ;
5   | end
6 return  $V_{vuln}$ 

```

Algorithm 2: Initial filtering of monitored projects.

that  $\mathcal{CW}$  takes a conservative approach to pick the oldest date for the bug-introducing commit and the newest date for the bug-fixing commit. This method allows us to track down both bug-introducing and bug-fixing commits, with which we can estimate the time-frame of the presence of the bug in any codebase. We use this time-frame as a window representing its propagation within a project that was created by forking the target project.

**Identification of Vulnerable Code.** After obtaining bug fixing and bug introducing commits,  $\mathcal{CW}$  requires manual annotation of the code in these commits (i.e., the code responsible for the vulnerability and the fix). This requires understanding the cause of the vulnerability; however, this manual effort is only a one-time action (per vulnerability), while all monitored projects are checked automatically later. Next,  $\mathcal{CW}$  transforms the annotated code to a detection test suitable for a clone detector.

**Filtering of Monitored Projects.** We provide a simple method for the initial filtering of monitored cryptocurrencies in Algorithm 2, which is based on the timestamp of forking these projects from the parent project. Letting  $V_c$  be the set of monitored cryptocurrency projects, if  $p \in V_c$  is created after  $G_b.date$  and before  $G_f.date$ , we consider  $V_c$  for further analysis. Otherwise, if the repository of  $p$  was created before the bug was introduced or after the bug was fixed, we ignore this case since the bug had not existed at the time of forking the project. The list of candidate projects forms our first cut of potentially vulnerable projects. Algorithm 2 uses the output of Algorithm 1 to identify the dates when the vulnerability was created and presented until by checking the timestamps associated with the particular commit IDs.

**Detection Process.** Finally, using the generated detection tests,  $\mathcal{CW}$  applies the clone detector to find clones of the parent project in the monitored list of cryptocurrencies, which are then analyzed using  $G_f$  and  $G_b$  to indicate such cloned projects that are potentially vulnerable.

#### IV. IMPLEMENTATION

**Experiment Setup.** All the experiments we performed ran on the machine with a quad-core i7-7500U CPU clocked at 2.70GHz and equipped with 16 GB of memory. We conducted a simulation experiment where we applied  $\mathcal{CW}$  on the versions of projects dating before the version where the vulnerability was disclosed. The simulation experiment evaluated whether some cryptocurrencies might have been fixed earlier if our approach has been adopted. To instantiate a clone detector, we chose Simian [15], a similarity analyzer that identifies duplication in C++ code as well as other tools that use text-based string comparison to identify type I clones. The Type I clones represent the situation when code fragments in a particular code-base are the same, omitting the comments and indentations [16], [17]. Note that Type II clones are syntactically similar code fragments, which are usually created from Type I clones by renaming variable names or modifying data types. Finally, Type III clones are created due to additions,

deletions, or modifications of code fragments in Type I or Type II clones.

**Testing Simian.** Before we conducted any clone detection experiment, we first ran a test to verify that Simian is capable of accurate detection of cloned code within the full code structure. Therefore, we performed a positive test and compared a cloned repository of Bitcoin-0.17 with a portion of its own code. This identified the correct file and line numbers of the entire code snippet. We then performed a negative test with a code that did not exist within Bitcoin-0.17. We found that to minimize false positives, we had to set the code line threshold to the exact length of each snippet of vulnerable code. This ensures that only cryptocurrencies containing the full code causing the vulnerability are detected. However, doing so increased the chance of false negatives; due to the time-sensitive nature of cryptocurrencies and a large attack surface, we adopted this setting since it is more important to detect truly vulnerable projects than to miss potentially vulnerable ones.

**Identifying Vulnerable Code.** Using the CVE database specific to Bitcoin [18] (i.e., the parent project), we reviewed Bitcoin’s CVEs to select vulnerabilities that are contained in its codebase versions in the range  $\langle 2016, 2018 \rangle$ , a time frame where many alternative cryptocurrencies were created. Once a CVE was selected for analysis,  $\mathcal{CW}$  searched for the bug fixing and bug introducing commits. Then, we manually annotated the code segments of these commits, which are related to the vulnerability and the fix, based on which,  $\mathcal{CW}$  generated detection tests for Simian. Next,  $\mathcal{CW}$  parsed information from the Bitcoin version where the vulnerability was introduced in and the version it was eventually fixed at.

**Clone Detection in Monitored Projects.** We wrote Python scripts that scraped *coingecko.com* to collect a dataset of the names of all existing cryptocurrencies, the link of their GitHub repositories, and the language that the cryptocurrency was written in. Out of the 1970 cryptocurrencies that we had, 43 were listed with a broken GitHub link or were just links to websites. We then used Git to clone all valid repositories. Since we wanted to identify vulnerabilities from Bitcoin clones, we removed all cryptocurrencies that had not used C++. We were then left with 1094 cryptocurrencies. The folder structure for the analysis is comprised of one main folder, which we will call the *comparison folder*. Inside the *comparison folder* were two more folders; the first containing the vulnerable code and the second containing the code of the cryptocurrency we wanted to compare against. Python scripts were used to automate the moving of the code from our database into the comparison folder, running the Simian clone detection, and then moving the code back. This was done iteratively for each cryptocurrency in the database. Simian produced an XML file after every iteration containing the results of the clone detection. The XML file contains the file path of each detected clone, the number of lines detected, and the processing time, as seen in Figure 4. Another script then ran to parse each of the Simian outputs, which gave us the list of suspicious vulnerable cryptocurrencies. To support the validity of our methodology,

```

<simian version="2.5.10">
  <check ignoreCharacterCase="true" ignoreCurlyBraces="true" ignoreIdentifierCase="true" ignoreModifiers="true" ignoreStringCase="true"
    threshold="6">
    <set lineCount="6" fingerprint="5eb7165c145ded7b7469ac9ac53e677b">
      <block sourceFile="\simian\bin\coin\bitcoin-0.16.3\src\test\skiplist_tests.cpp" startLineNumber="24" endLineNumber="31"/>
      <block sourceFile="\simian\bin\coin\Dogecoin\src\test\skiplist_tests.cpp" startLineNumber="25" endLineNumber="32"/>
    </set>
    <set lineCount="6" fingerprint="4ceb8650e14442de03bb7eb1b8f9b9b4">
      <block sourceFile="\simian\bin\coin\bitcoin-0.16.3\src\key.cpp" startLineNumber="203" endLineNumber="211"/>
      <block sourceFile="\simian\bin\coin\Dogecoin\src\key.cpp" startLineNumber="72" endLineNumber="80"/>
    </set>
    <set lineCount="6" fingerprint="7b80d7542d174e7bdb2e532d6a96e1bb">
      <block sourceFile="\simian\bin\coin\bitcoin-0.16.3\src\test\script_tests.cpp" startLineNumber="1282" endLineNumber="1289"/>
      <block sourceFile="\simian\bin\coin\Dogecoin\src\test\script_tests.cpp" startLineNumber="964" endLineNumber="971"/>
    </set>
    ...
    <summary duplicateFileCount="348" duplicateLineCount="34092" duplicateBlockCount="2369" totalFileCount="455" totalRawLineCount="160822"
      totalSignificantLineCount="94838" processingTime="1624"/>
  </check>
</simian>

```

Figure 4: An example of XML output from Simian.

we also experimented with Zcash (and its CVEs) serving as a parent project (see vulnerability CVE-2019-7167).

**Responsible Disclosure.** Finally, before disclosing vulnerabilities in the cloned projects, we filed 4 CVE extensions with MITRE. Once the CVEs were accepted, we informed the developers of these projects about the vulnerability and the fix recommended in the CVE description.

## V. EVALUATION

We evaluated  $\mathcal{CW}$  through various experiments. The experiments were designed to answer the following research questions (RQs):

- **RQ1:** Are clones prevalent in cryptocurrencies?
- **RQ2:** Can we accurately detect propagated vulnerabilities through clones?
- **RQ3:** Are the propagated vulnerabilities real problems?

RQ1 examines our hypothesis, i.e., clones are prevalent in cryptocurrencies, and therefore identifying vulnerabilities by tracking clones is an important problem. RQ2 evaluates the accuracy of our approach. RQ3 is designed to question the usefulness of  $\mathcal{CW}$  in practice and reports on true positive and false positive rates.

### A. RQ1: Clone Prevalence

To answer RQ1, we systematically applied  $\mathcal{CW}$  to the latest versions of the cryptocurrencies within our dataset scraped from Coingecko as of October 2018.<sup>2</sup> In total, there were initially 2079 cryptocurrencies as of Oct 2018, and 4389 cryptocurrencies as of Apr 2019. We note that a large number of new cryptocurrencies were created in this short period of time. There are also different types of cryptocurrencies; e.g, 2431 of them are ERC-20 tokens [19], which we consider as out-of-scope.<sup>3</sup> Out of the remaining 2079 cryptocurrencies that

we scraped, 1094 cryptocurrencies are written in C++, the same programming language as Bitcoin is written in. The breakdown of the remaining cryptocurrencies can be found in Table II.

As we conjecture that these cryptocurrencies are clones of various versions of Bitcoin, using  $\mathcal{CW}$ , we systematically apply Simian to compare each of these cryptocurrencies with different versions of Bitcoin. Part of the results are summarized in Table III, where the second column shows the overall number of lines of code (LOC) in the project and the following columns show the clone ratio of each cryptocurrency in respect to different versions of Bitcoin. The clone ratio is defined as follows:

$$ratio_{clone} = \frac{K}{T}, \quad (1)$$

where  $K$  is the number of LOC cloned from Bitcoin and  $T$  is the total number of LOC. Note that Table III shows the clone ratio of a subset of 14 cryptocurrencies that were among the top 50 by the market capitalization as of the time of making the experiment. The average clone ratio was found to be 30.7%. However, if we compare the highest clone ratio among all versions of the same set of cryptocurrencies written in C++, the average clone ratio increases to 46.6%. Furthermore, multiple cryptocurrencies have a clone ratio of more than 90%. Hence, we conclude that code clones are indeed prevalent in cryptocurrencies.

Language	The Number of Projects
C++	1094
Javascript	334
C	65
Go	65
Python	36
Java	30
Others	455

Table II: Distribution of programming languages across cryptocurrencies considered in our work.

<sup>2</sup><https://www.coingecko.com/en>

<sup>3</sup>Note that ERC-20 tokens are built on the Ethereum platform and contain a standardized smart contract code that is parametrized by particular tokens.

Cryptocurrency	LOC	Bitcoin 0.17.0	Bitcoin 0.16.3	Bitcoin 0.15.2	Bitcoin 0.14.2	Bitcoin 0.13.2	Bitcoin 0.12.0	Bitcoin 0.11.2
Bitcoin Cash	68,545	26.0%	32.2%	<b>34.2%</b>	27.3%	20.3%	17.2%	13.1%
Bitcoin Diamond	51,102	55.6%	70.2%	77.2%	<b>98.1%</b>	73.6%	56.2%	42.3%
DogeCoin	38,073	33.2%	40.8%	43.9%	55.2%	67.1%	77.5%	<b>91.6%</b>
Monero	57,232	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Zcash	51,509	21.1%	24.4%	25.1%	33.8%	<b>42.2%</b>	46.0%	40.7%
Dash	65,104	40.24%	47.23%	52.52%	<b>63.93%</b>	46.25%	38.62%	28.64%
DigiByte	63,769	<b>90.90%</b>	69.32%	60.81%	42.02%	30.94%	24.14%	18.85%
Electroneum	37,040	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
EOS	96,066	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Litecoin	57,014	76.73%	<b>98.17%</b>	86.11%	58.55%	43.75%	33.27%	25.83%
Nano	26,987	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Qtum	112,867	37.52%	<b>48.18%</b>	42.68%	28.82%	21.48%	16.30%	12.68%
Ravencoin	62,200	63.33%	81.09%	<b>82.41%</b>	55.93%	41.22%	31.45%	24.27%
Steem	108,561	0.01%	0.02%	0.02%	0.02%	0.02%	0.02%	<b>0.07%</b>

Table III: Percentage of cloned code as compared to various Bitcoin versions.

Vulnerability	Parent Project	# Investigated Projects	TPR	FPR
<b>CVE-2018-17144</b>	Bitcoin	31	100.0%	0.00%
<b>CVE-2016-10724</b>	Bitcoin	422	89.3%	10.7%
<b>CVE-2016-10725</b>	Bitcoin	422	89.3%	10.7%
<b>CVE-2019-7167</b>	Zcash	1	100.0%	0.00%

Table IV: A summary of the results obtained by CW.

### B. RQ2: Accuracy of COINWATCH

To answer this question, we systematically apply CW to multiple vulnerabilities discovered in recent years and see whether they are propagated to other cryptocurrencies. We focus on such vulnerabilities due to the following reasons. First, these are high-profile vulnerabilities. Second, these vulnerabilities are reported between the years 2016 to 2018, in which we evinced a boom of a vast number of cryptocurrencies. A vulnerability propagates itself across clones since it is not yet discovered at the time when the parent project is cloned. The propagation of a vulnerability is highly dependent on when the source code was cloned and how long the vulnerability has existed before it was fixed. If a cryptocurrency forked the code during the time frame when the bug was presented in the source code, there is a high chance of the clone containing vulnerable code. As a result, we found out that CW is capable of accurate detection of vulnerabilities in cloned projects as proven by filing 4 CVE extensions.

**Case Study.** In the following, we focus on CVE-2018-17144. Among 985 pre-filtered cryptocurrencies monitored using CW, we found that as of October 2018, 31 cryptocurrencies had Type I clones of a vulnerable code found in Bitcoin and hence were exposed to this vulnerability. We monitored the situation after we alerted vulnerable projects and our CVE extension was published at NVD. In February 2019, the results from CW suggest that 22 of the 31 cryptocurrencies have patched the vulnerability, while in June 2020 there were still 4 projects unfixed. This implies that attackers had at least a month’s time frame from the time the CVE was announced in September 2018 to attack the 31 vulnerable cryptocurrencies, exposing up to \$35,585,594 market capital (as of 07 October 2018).

Moreover, in February 2019 and June 2020, attackers could still attack the remaining 9 and 4 unfixed cryptocurrencies, respectively.

### C. RQ3: Is a clone equal to a vulnerability?

To answer this question, we investigate how many of the candidate cryptocurrencies that are discovered using CW are indeed vulnerable. This is a challenging problem since we must be able to determine whether the affected code is reached or not, and whether its execution would result in the same reported vulnerability as in the target project. Therefore, we attempt to answer the question in two ways. First, we manually examined all the patches of 4 CVEs (that we extended) and cross-checked them with cryptocurrencies reported by CW. For each patch, we checked if the vulnerable cloned code was already modified (by maintainers) in the most recent version of the project, and if so, whether it was modified in a way similar to the fix of the CVE in the parent project. In the positive case, we considered a discovery made by CW on a particular cryptocurrency as a true vulnerability. The second way we answer the question was manual checking with the developers/maintainers of the project.

**True Positives VS False Positives.** The summary of our results is presented in Table IV. In the case of CVE-2018-17144, due to an extremely severe consequence of this double-spending vulnerability, many popular cloned cryptocurrencies quickly applied the fix after the CVE was released, e.g. Dash, Ravencoin, Bitcoin Cash, etc. This is not the case for low-profile vulnerabilities or for smaller cryptocurrencies, which are less diligent in keeping their codebases updated. Pigeoncoin is an example of a cryptocurrency that suffered from this vulnerability, and which was also exploited. We analyzed different releases of Pigeoncoin and found that it contained the double-spending vulnerability from its official launch on 07 April 2018 until 27 September 2018. The reason for the bug fix on 27 September was the attack made on it. If CW were to monitor Pigeoncoin at the time the Bitcoin patch was released, we could have identified the vulnerability and prevented the attack. In the case of CVE-2018-17144 and CVE-2019-7167, there were



```

DoS Vulnerability (Bitcoin Diamond 0.14, validation.cpp)

for (const auto& tx : block.vtx)
    if (!CheckTransaction(*tx, state, block.nVersion &
        VERSIONBITS_FORK_BCD && isBCDBlock, false))
        return state.Invalid(false, state.GetRejectCode(), state.
            GetRejectReason(),
            sprintf("Transaction check failed (tx hash %s) %s", tx->
                GetHash().ToString(), state.GetDebugMessage()));

```

Figure 5: Propagated code for DoS vulnerability in Bitcoin Diamond.

```

DoS Vulnerability (Bitcoin Stash, validation.cpp)

if ( !CheckRegularTransaction(*tx, state, false) ) {
    return state.Invalid(
        false, state.GetRejectCode(), state.GetRejectReason(),
        sprintf("Transaction check failed (txid %s) %s",
            tx->GetId().ToString(), state.GetDebugMessage()));
}

```

Figure 6: Propagated code for DoS vulnerability in Bitcoin Stash.

no false positives since all candidate cryptocurrencies reported were confirmed to contain these vulnerabilities.

On the other hand, *CW* reported a few false positives as well. For example, although we initially found out that 422 cryptocurrencies had enabled the alert system (causing CVE-2016-10724 and CVE-2016-10725 in the parent project), 45 cryptocurrencies reported at the output of *CW* did not reference it during compilation and thus are not vulnerable. Due to this, the true positive rate of *CW* was deteriorated to 89.3% for these particular vulnerabilities.

**Type II and III Clones and Possible False Negatives.** Type II and III clones come into play when the codebase is edited by changing the way certain functions work or are named. One example of a Type III clone that we identified was presented in Bitcoin Diamond Version 0.14.0 and 0.13.0. The code can be seen in Figure 5. Bitcoin Diamond modified its function *CheckTransaction()* to include new arguments that are specific to Bitcoin Diamond itself. An example of Type II clone was identified in Bitcoin Stash as seen in Figure 6, where the *CheckTransaction()* function was renamed to *CheckRegularTransaction()*. These clones cannot be detected using Simian as the clone detector, and they require more sophisticated tools which we will consider in our future work. We contacted the development team of these two projects regarding the possible vulnerabilities but did not get the response; hence, we cannot confirm whether these two cases represent false negatives.

## VI. DISCUSSION

**Scope of Clones.** The results presented in this paper are related mainly to Bitcoin, as it is currently the dominant cryptocurrency and has the most number of forks from its GitHub repository. However, there are other cryptocurrencies written in different programming languages (or are close-sourced) that are similar to Bitcoin but would not be detected as Type I clones. There is a trade-off between false positives and false negatives if we use more complex types of code clone detection techniques such as the ones searching for Type II or Type III clones. In some cases we would have to conduct

a decompilation of binaries and do the analysis on abstract syntax trees of assembly rather than a verification of source code for match. This is the direction that we will investigate in our future work.

**Manual Annotation.** A semi-automated nature of *CW* is another limitation, which stems from the manual annotation of the vulnerable code and fix within the bug fixing and bug introducing commits. However, this is a one-time action per vulnerability and can be performed in a relatively short time, while later *CW* continues automatically in checking the full list of monitored projects. In our future work, we plan to design simple heuristic methods that can automate this step.

**Non-Standard Assumptions.** *CW* assumes that if a cryptocurrency project was forked after the bug was introduced in the parent project and before the bug was fixed, it is likely to suffer from the vulnerability. However, we note that sometimes forked projects may cherry-pick commits from their parents or apply patches mitigating some security issues, which we currently do not distinguish in *CW*. However, we emphasize that such selected projects represent only the first cut of potentially vulnerable candidate projects and if they contain a fix, the clone detector will mark them as safe. In the worst case, if the fix is presented in a modified form that is not matched by the clone detector, they are reported as false positives – this causes a notification to the maintainers, who just ignore it.

## VII. RELATED WORK

**Security of Cryptocurrencies.** The security of cryptocurrencies has been undermined by several high profile incidents. Bitcoin [20] itself has been described to be vulnerable to several attacks such as double-spending [21], transaction malleability [22], networking attacks [23] [24], attacks targeting mining [25] [26] [27], and others [2] [3]. As a research topic, Bitcoin’s security is attractive due to a broad number of use cases proposed for its underlying technology, blockchain. Bitcoin has inspired several cryptocurrencies to build and improve different aspects of its code to increase its transaction throughput (eg. Bitcoin Cash) or provide greater anonymity (eg. Bitcoin Private). The intense scrutiny of Bitcoin’s security has led to a gap in research on the effects of vulnerabilities from Bitcoin propagating to other cryptocurrencies that have copied parts of Bitcoin’s code. These cryptocurrencies also store value but are not as valuable as Bitcoin and thus do not have as many developers and resources to build a body of research around their security. Therefore, we developed *CW* to address the security problem of propagated vulnerabilities from Bitcoin while focusing on real-world attacks that have been identified and fixed within Bitcoin’s code.

**Clones among Cryptocurrencies.** Some cryptocurrencies fork existing repositories and launch their own cryptocurrencies with a specific focus. As these cryptocurrencies are changing only a portion of the forked codebase before launching, a large portion of their code is cloned from the specific version it was forked from. Prior work suggests that code clones increase maintenance effort [28] [29], might cause bugs (Li et. al

analyzed Linux, FreeBSD, Apache [30]), and signal errors in other parts of a codebase. It is known that clone genealogy information can be used to identify clones that may benefit new approaches and clone management [31]. Based on the findings of the authors of [31], aggressive code refactoring might not be necessary for code clones that are volatile and that the technique of refactoring needs to be complemented by other code maintenance approaches. Prior research analyzing package dependencies [32] utilizes a systematic approach of measuring the attack surface exposed by individual vulnerabilities through component level dependency analysis. In this work, cloned cryptocurrencies do not utilize Bitcoin as a package but the concept that the attack surface is dependent on code utilized from elsewhere. The authors of [33] analyzed code diversity in the cryptocurrency projects basing on the source code similarity. The intention of the authors was to examine the extent to which new cryptocurrencies provide innovations.

**Code Clone Detection.** Basic code clone detection techniques are performed using several tools and generally follow a 6-step process [34] to generate clone pairs. Code cloning can be used for the detection of code refactoring, greater efficiency of development, etc. However, this methodology has been used for other purposes due to the large volume of open-source code that can be analyzed, such as detecting vulnerabilities. For example, Viertel et al. [35] utilized this technique for detecting vulnerabilities in codebases of Java-written projects, in a similar fashion than *CW* does. A similar environment is represented by smart contracts in Ethereum or other smart contract platforms. Early work focused on finding semantic clones in smart contracts that contained vulnerabilities in the smart contract logic is presented in [36]. A tool detecting previously discovered smart contract vulnerabilities through clones was proposed, and can quickly identify other smart contracts that are semantically equivalent (and thus might contain the same set of vulnerabilities). The automatic classification of Ponzi smart contracts based on code clones was proposed in [37], and it bases on the fact that related smart contracts have fundamentally similar programming logic. Clone detection was also performed after transferring the code into an abstract form such as strings [38], tokens [39], and abstract syntax trees (AST) [40]. An efficient heuristic technique basing on edit distance and AST for pair-wise comparison of smart contract byte codes was proposed in [41]. Techniques such as line comparison based on dot plots [42], comparing whole files [43], pattern recognition of code characteristics through clustering [44], finding clones within a syntactic unit through post-processing [39], and pre-processing [45] are also used for clone detection.

## VIII. CONCLUSION

In this work, we proposed COINWATCH, an approach that monitors a specified set of cryptocurrency projects for the presence of vulnerabilities disclosed in their parent projects. We applied COINWATCH on 1094 monitored cryptocurrencies

using 4 CVEs at the input, and we identified 786 true vulnerable projects, which were confirmed with developers and successfully reported as CVE extensions. Our results show that many attacks could have been prevented should COINWATCH has been monitoring the attacked projects.

In future work, we plan to extend code clone detection to Type II and III clones to increase the detection performance as well as eliminate a need for a manual step of annotating the code within bug fixing and introducing commits.

## REFERENCES

- [1] B. Core. (2019) Bitcoin codebase at github version control system. [Online]. Available: <https://github.com/bitcoin/bitcoin>
- [2] I. Homoliak, S. Venugopalan, Q. Hum, and P. Szalachowski, "A security reference architecture for blockchains," in *2019 IEEE International Conference on Blockchain (Blockchain)*. IEEE, 2019, pp. 390–397.
- [3] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, "Sok: Research perspectives and challenges for bitcoin and cryptocurrencies," in *2015 IEEE Symposium on Security and Privacy*, 2015, pp. 104–121.
- [4] W. Wang, D. T. Hoang, Z. Xiong, D. Niyato, P. Wang, P. Hu, and Y. Wen, "A survey on consensus mechanisms and mining management in blockchain networks," *arXiv preprint arXiv:1805.02707*, pp. 1–33, 2018.
- [5] M. Conti, E. S. Kumar, C. Lal, and S. Ruj, "A survey on security and privacy issues of bitcoin," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, 2018.
- [6] A. Zmudzinski. (2019) Research: 60% of all bitcoin full-nodes are still vulnerable to inflation bug. [Online]. Available: <https://cointelegraph.com/news/research-60-of-all-bitcoin-full-nodes-are-still-vulnerable-to-inflation-bug>
- [7] B. Core. (2018) Cve-2018-17144 full disclosure. [Online]. Available: <https://bitcoincore.org/en/2018/09/20/notice/>
- [8] Mitre. (2019) Cwe-617: Reachable assertion. [Online]. Available: <https://cwe.mitre.org/data/definitions/617.html>
- [9] R. Esteves. (2018) Pigeoncoin (pgn) hacked due to bitcoin protocol bug, copycat coins in danger? [Online]. Available: <https://www.newsbtc.com/2018/10/06/pigeoncoin-pgn-hacked-due-to-bitcoin-protocol-bug-copycat-coins-in-danger>
- [10] Pigeoncoin. (2018) Update validation.cpp (#21). [Online]. Available: <https://github.com/Pigeoncoin/pigeoncoin/commit/cc4885cc8e176e4e68a6ef966638b6728191c853>
- [11] MITRE Corporation, "Common Vulnerabilities and Exposures (CVE)," 2020. [Online]. Available: <https://cve.mitre.org/>
- [12] C. Carretero-Campos, P. Bernaola-Galván, A. Coronado, and P. Carpena, "Improving statistical keyword detection in short texts: Entropic and clustering approaches," *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 6, pp. 1481–1492, 2013.
- [13] J. Śliwerski, T. Zimmermann, and A. Zeller, "When do changes induce fixes?" in *ACM sigsoft software engineering notes*, vol. 30, no. 4. ACM, 2005, pp. 1–5.
- [14] M. Borg, O. Svensson, K. Berg, and D. Hansson, "Szz unleashed: An open implementation of the szz algorithm-featuring example usage in a study of just-in-time bug prediction for the jenkins project," *arXiv preprint arXiv:1903.01742*, 2019.
- [15] S. Harris. (2003) Simian – similarity analyser. [Online]. Available: <https://www.harukizaemon.com/simian>
- [16] C. K. Roy, M. F. Zibran, and R. Koschke, "The vision of software clone management: Past, present, and future (keynote paper)," in *2014 Software Evolution Week-IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering (CSMR-WCRE)*. IEEE, 2014, pp. 18–33.
- [17] C. K. Roy, "Detection and analysis of near-miss software clones," in *2009 IEEE International Conference on Software Maintenance*. IEEE, 2009, pp. 447–450.
- [18] Bitcoin.it. (2019) Bitcoin common vulnerabilities and exposures. [Online]. Available: [https://en.bitcoin.it/wiki/Common\\_Vulnerabilities\\_and\\_Exposures](https://en.bitcoin.it/wiki/Common_Vulnerabilities_and_Exposures)
- [19] V. B. Fabian Vogelsteller. (2015) Eip 20: Erc-20 token standard. [Online]. Available: <https://eips.ethereum.org/EIPS/eip-20>

- [20] S. Nakamoto *et al.*, "Bitcoin: A peer-to-peer electronic cash system," *Communications of the ACM*, 2008.
- [21] G. O. Karame, E. Androulaki, A. Capkun, Srdjan, J. A. Kroll, and E. W. Felten, "Double-spending fast payments in bitcoin," in *Proceedings of the 2012 ACM conference on Computer and communications security*, 2012, pp. 906–917.
- [22] C. Decker and R. Wattenhofer, "Bitcoin transaction malleability and mtgox," in *European Symposium on Research in Computer Security*, 2014, pp. 313–326.
- [23] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on bitcoin's peer-to-peer network," in *24th {USENIX} Security Symposium ({USENIX} Security 15)*, 2015, pp. 129–144.
- [24] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking bitcoin: Routing attacks on cryptocurrencies," in *2017 IEEE Symposium on Security and Privacy (SP)*, 2017, pp. 375–392.
- [25] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," *Communications of the ACM*, vol. 61, no. 7, pp. 95–102, 2018.
- [26] A. Sapirshtein, Y. Sompolinsky, and A. Zohar, "Optimal selfish mining strategies in bitcoin," in *International Conference on Financial Cryptography and Data Security*, 2016, pp. 515–532.
- [27] I. Eyal, "The miner's dilemma," in *2015 IEEE Symposium on Security and Privacy*, 2015, pp. 89–103.
- [28] A. Monden, D. Nakae, T. Kamiya, S.-i. Sato, and K.-i. Matsumoto, "Software quality analysis by code clones in industrial legacy software," in *Proceedings Eighth IEEE Symposium on Software Metrics*, 2002, pp. 87–94.
- [29] M. Mondal, C. K. Roy, M. S. Rahman, R. K. Saha, J. Krinke, and K. A. Schneider, "Comparative stability of cloned and non-cloned code: An empirical study," in *Proceedings of the 27th Annual ACM Symposium on Applied Computing security*, 2012, pp. 1227–1234.
- [30] Z. Li, S. Lu, S. Myagmar, and Y. Zhou, "Cp-miner: Finding copy-paste and related bugs in large-scale software code," *IEEE Transactions on software Engineering*, vol. 32, no. 3, pp. 176–192, 2006.
- [31] M. Kim, V. Sazawal, D. Notkin, and G. Murphy, "An empirical study of code clone genealogies," in *ACM SIGSOFT Software Engineering Notes*, vol. 30, no. 5, 2005, pp. 187–196.
- [32] S. Zhang, X. Zhang, X. Ou, L. Chen, N. Edwards, and J. Jin, "Assessing attack surface with component-based package dependency," in *International Conference on Network and System Security*, 2015, pp. 405–417.
- [33] P. Reibel, H. Yousaf, and S. Meiklejohn, "Short paper: An exploration of code diversity in the cryptocurrency landscape," in *International Conference on Financial Cryptography and Data Security*. Springer, 2019, pp. 73–83.
- [34] C. K. Roy, J. R. Cordy, and R. Koschke, "Comparison and evaluation of code clone detection techniques and tools: A qualitative approach," *Science of computer programming*, vol. 74, no. 7, pp. 470–495, 2009.
- [35] F. P. Viertel, W. Brunotte, D. Strüber, and K. Schneider, "Detecting security vulnerabilities using clone detection and community knowledge," in *SEKE*, 2019, pp. 245–324.
- [36] H. Liu, Z. Yang, C. Liu, Y. Jiang, W. Zhao, and J. Sun, "Eclone: Detect semantic clones in ethereum via symbolic transaction sketch," in *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2018, pp. 900–903.
- [37] W. Chen, Z. Zheng, J. Cui, E. Ngai, P. Zheng, and Y. Zhou, "Detecting ponzi schemes on ethereum: Towards healthier blockchain technology," in *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, 2018, pp. 1409–1418.
- [38] J. H. Johnson, "Substring matching for clone detection and change tracking," in *ICSM*, vol. 94, 1994, pp. 120–126.
- [39] Y. Higo, Y. Ueda, T. Kamiya, S. Kusumoto, and K. Inoue, "On software maintenance process improvement based on code clone analysis," in *International Conference on Product Focused Software Process Improvement*, 2002, pp. 185–197.
- [40] R. Koschke, R. Falke, and P. Frenzel, "Clone detection using abstract syntax suffix trees," in *2006 13th Working Conference on Reverse Engineering*, 2006, pp. 253–262.
- [41] P. Hartel, I. Homoliak, and D. Reijtsbergen, "An Empirical Study into the Success of Listed Smart Contracts in Ethereum," *IEEE Access*, vol. 7, pp. 177 539–177 555, 2019.
- [42] M. Rieger, "Effective clone detection without language barriers," Ph.D. dissertation, Verlag nicht ermittelbar, 2005.
- [43] U. Manber *et al.*, "Finding similar files in a large file system," in *Usenix Winter*, vol. 94, no. 5, 1994, pp. 1–10.
- [44] E. Merlo, G. Antoniol, M. Di Penta, and V. F. Rollo, "Linear complexity object-oriented similarity for clone detection and software evolution analyses," in *20th IEEE International Conference on Software Maintenance, 2004. Proceedings.*, 2004, pp. 412–416.
- [45] J. R. Cordy, T. R. Dean, and N. Synnyskyy, "Practical language-independent detection of near-miss clones," in *Proceedings of the 2004 conference of the Centre for Advanced Studies on Collaborative research*, 2004, pp. 1–12.