Distinguishing Legitimate and Fake/Crude Antivirus Software

Masaki Kasuya†, Kenji Kono†‡
†Dept. of Information and Computer Science
Keio University
Yokohama, Japan
‡CREST, Japan Science and Technology Agency
Email: kasuya@sslab.ics.keio.ac.jp, kono@ics.keio.ac.jp

***Abstract*—Fake antivirus (AV) software, a kind of malware, pretends to be a legitimate AV product and frightens computer users by showing fake security alerts, as if their computers were infected with malware. In addition, fake AV urges users to purchase a "commercial" version of the fake AV. In this paper, we search for an *indicator* that captures behavioral differences in legitimate AV and fake AV. The key insight behind our approach is that legitimate AV behaves *differently* in clean and infected environments, whereas fake AV behaves *similarly* in both environments, because it does not analyze malware in the infected environments. We have investigated three potential indicators, file access pattern, CPU usage, and memory usage, and found that memory usage is an effective indicator to distinguish legitimate AV from fake AV. In an experiment, this indicator identifies all fake AV samples (39 out of 39) as fake and all legitimate AV products (8 out of 8) as legitimate. It is impractical for fake AV to evade this indicator because to do so it would require it to detect malware infections, just as legitimate AV does.**

***Keywords*—Antivirus Software; Fake Antivirus Software; Behavior Analysis; Malware**

I. INTRODUCTION

Fake antivirus (AV) software is a severe threat to our computer systems. It pretends to be actual AV software and shows false security warnings to users as if their computer systems were infected with malicious software (malware). Fake AV persuades victim users to purchase a useless, "commercial" version of the fake AV to eliminate bogus threats. "Crude AV" poses a similar threat to fake AV. It differs from fake AV in that it detects malware, but its detection quality is too low to be practical.

The threat of fake AV is real. Symantec detected 43 million installation attempts of fake AV from July 2008 to Jun 2009 [6]. According to Rajab et al. [15], fake AV accounts for 15% of all malware detected by Google's malware detection infrastructure [14]. Fake AV business earns tremendous revenue. Stone-Gross et al. revealed that three kinds of fake AV have earned more than $130 million dollars [16]. McAfee disclosed that the annual revenue of one vendor of fake AV exceeded $180 million dollars [13]. To distribute fake AV, software download sites are sometimes exploited. In fact, one famous download site, CNET, distributed a fake AV sample called RegGenie in 2012 [5].

Fake/crude AV is similar to a social engineering attack; the victim users are deceived and never suspect that fake/crude AV is not legitimate because it is carefully designed to look like legitimate AV. One approach for defending against fake/crude AV is to use signature-based approaches. However, signature-based approaches are exploit-specific, and a signature must be prepared for each instance of fake/crude AV. Therefore, these approaches cannot detect previously unseen instances of fake/crude AV.

In this paper, we reveal behavioral differences between legitimate and fake/crude AV, and propose an *indicator* that captures the behavioral differences in AV software. The key insight behind our approach is that legitimate AV behaves differently in 1) clean environments and 2) infected environments, while fake/crude AV would not show such differences. A clean environment is the one in which no malware has been installed, whereas an infected environment is the one in which malware has been installed. Fake/crude AV is not expected to show behavioral differences in clean and infected environments because it does not analyze malware samples in the infected environment. On the other hand, legitimate AV instances are expected to show the differences because they deeply analyze suspicious instances in the infected environments.

Our indicator can be used in software download sites such as CNET [1] and PCMAG [2]. When AV samples are uploaded to download sites with the tag indicating AV, they can be checked as to whether they are legitimate or fake/crude. Since our indicator works by capturing behavioral differences instead of using signature, it can detect the latest fake/crude AV instances.

This paper shows that "memory usage" is an effective indicator of fake/crude AV software. Surprisingly, crude AV does not show differences in memory usage between clean and infected environments. In our approach, the memory usages of AV-like software in clean and infected environments are compared statistically. The Levene Test [10], an inferential statistic, is used to assess the equality of variances in memory usage. If a software sample that is suspected to be fake/crude AV has statistically the same distribution of memory usage in both environments, it is considered fake. Otherwise, it is considered legitimate.

Since our indicator is based on behavioral differences, fake/crude AV has to mimic the behaviors of legitimate AV in order to evade it. It is not easy to the mimic behaviors of legitimate AV. If fake/crude AV samples change their memory consumption at random, our approach can detect fake/crude AV correctly because it compares memory consumption under several settings, i.e., clean/clean, infected/infected, and clean/infected.

To demonstrate the usefulness of our approach, we have conducted experiments on 39 "real" fake/crude AV samples and 8 legitimate AV products. The results show that our indicator can identify all 39 fake/crude samples, which means

 ISBN: 978-1-61208-298-1

there are no false negatives, and all 8 legitimate products, which means there are no false positives.

The remainder of this paper is organized as follows. Section II describes the differences between fake AV and crude AV, and the current criteria to distinguish them from legitimate AV. Section III explains our basic approach and shows how to distinguish fake/crude AV from legitimate AV by using the Levene Test. Section IV presents our experimental results. Discussion and related work are presented in Sections V and VI. Section VII concludes the paper.

II. Fake AV and Crude AV

There are two types of malicious AV software: fake AV and crude AV. To understand the difficulties of distinguishing between fake/crude AV and legitimate AV, this section describes the behaviors of fake and crude AV, and briefly introduces recent guidelines to distinguish fake/crude AV and legitimate AV.

A. Fake AV

Fake AV mimics the behavior of legitimate AV and shows bogus security warnings without scanning for malware infections in the victims' file systems. To make the behavior resemble that of legitimate AV, fake AV searches the file system to obtain file and/or directory names to be displayed in warning messages. For example, Security Antivirus [8], a fake AV, displays the following message:

Virus name:
Virus.Win32.Faker.a
Infected file:
C:\Documents and Settings\Kasuya\Recent\snl2w.dll
Description:
These programs steal MSN Messenger passwords...

Pathname, `C:\Documents and ... snl2w.dll`, is the real one in the victim's file system. By showing real pathnames in the victim's file system, the fake AV deceives victims into believing the machine is infected with malware and encourages them to the purchase a product version of the fake AV.

The directory traverse of fake AV makes it difficult to distinguish it from legitimate AV. When observed from the outside, fake AV traverses directories just as legitimate AV does. Security Antivirus traverses most directories that all legitimate AV products commonly access. According to our investigation, the access coverage of Security Antivirus is over 99.7% (= 2393 / 2400). This result means that Security Antivirus carefully takes access patterns of legitimate AV into account. In other words, we cannot use directory traversal as an indicator to distinguish fake AV from legitimate AV.

B. Crude AV

Crude AV is low-quality AV software whose detection accuracy is too low to be useful. Crude AV differs from fake AV in that it scans file systems for malware and detects the infection. At the same time, crude AV differs from legitimate AV in that it cannot detect a large portion of widely deployed malware. To confirm that the detection rate of crude AV is very low, we measured the detection rate of Anti-Virus Elite [9], a well-known crude AV. We installed 905 unique instances of malware in Windows XP SP3. Anti-Virus Elite detected only 74 samples. The detection rate was 8.2%. Kaspersky, an example of legitimate AV, detected all 905 samples, i.e., its detection rate was 100.0%.

Crude AV is usually classified into malware. According to VirusTotal [3], an online antivirus scan service, Anti-Virus Elite is classified as malware in 65% (28 out of 43) of commercial AV products. Crude AV is malware because there are sites that urge the visitors to buy a "product" version, which in most cases is just as poor as the crude AV.

Crude AV blurs the boundary between fake and legitimate AV, and makes it more difficult to distinguish fake/crude AV from legitimate AV. Crude AV traverses file systems and inspects suspicious files that may contain malware. Aside from the quality of detection, crude AV behaves very similarly to legitimate AV.

C. Current Criteria to Distinguish Legitimate and Fake/Crude AV

Recently, a security industry has published a white paper [7] to help end-users identify fake security products such as fake AV. The document provides a helpful checklist to judge whether the users' computers are infected with fake AV. The checklist says, for instance, that fake AV reports an unreasonably high number of infections, shows a popup window frequently that warns your machine is infected with malware, and suggests the purchase of a commercial version.

This checklist is useful for manual inspection for discovering fake AV. However, these criteria do not suit automated distinction of fake/crude and legitimate AV. Suppose that we attempt to automate the process of counting the number of reported infections. Since the reports are shown in natural language, it is not easy for computers to understand the reports. Even if we could interpret the reports, there are samples of fake AV that do not show up in a lot of reports. For example, Anti Spyware Expert has only 18 reports of infection.

Furthermore, this checklist does not address how to distinguish crude AV from legitimate AV. Since crude AV behaves similarly to legitimate AV aside from the quality of detection, it is almost impossible to draw up a guideline for crude AV.

III. Searching for indicators

In this section, we describe our search for fake/crude AV indicators. As mentioned above, the key insight behind our approach is that legitimate AV behaves differently in clean and infected environments while fake/crude AV behaves similarly in both environments. A good indicator captures behavioral differences only in legitimate AV.

A. What is a good indicator?

A fake/crude AV indicator should satisfy at least three requirements. First, it should be applicable to as many instances of fake/crude AV as possible. In particular, it should be applicable to previously-unseen instances of fake/crude AV.

 ISBN: 978-1-61208-298-1

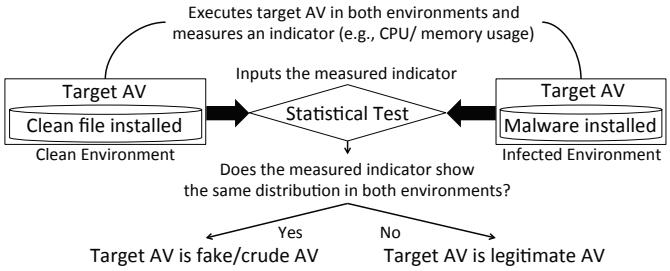

Fig. 1. Our basic approach

Second, it should be impractical for fake/crude AV to evade the indicator. The criminals that make use of fake/crude AV for their own profit do not want to spend a lot of money on development because it would reduce their revenues. A fake/crude AV indicator would thus be ideal since criminals would have to incorporate the same functionalities as legitimate AV in order to evade it and they would cost a fortune to develop software equivalent to legitimate AV.

Third, a fake/crude AV indicator should enable automatic distinction of fake/crude AV from legitimate AV. If the distinction process is automated, it can be incorporated into legitimate AV or deployed on software download sites [1][2]. The guidelines presented in Section II-C assume manual inspection of several aspects of suspicious behavior of AV-like software. In this paper, we seek a fake/crude AV indicator that does not require manual intervention. Rather than relying on visual inspection, we seek a fake/crude AV indicator from information that can be obtained in a systematic way. For example, we look for system call patterns or resource usage patterns that differentiate fake/crude AV from legitimate AV.

One approach to deriving a fake/crude AV indicator is to use binary analysis. We cannot rely on the source code because the source code of malware is not available in public. Binary analysis has the potential to identify a lot of operations (e.g., system calls and their arguments) that malware may do. However, it is not easy to apply binary analysis to fake/crude AV because malware can be obfuscated using a technique like binary obfuscation [17]. In addition, it is not straightforward to find a sequence of operations that can differentiate fake/crude AV from legitimate AV. Hence, we decide not to take this approach.

B. Basic Approach

Our basic approach is to *compare* potential indicators obtained in *clean* and *infected* environments. A clean environment is one in which no malware has been installed. We assume that an execution environment just after installing an operating system from read-only media such as DVD-ROM is clean. Therefore, an environment with harmless files is also clean. A clean environment can be prepared by using the recent technology of virtual machines. Once a clean environment has been prepared, we can reuse it by saving it as a virtual machine image. An infected environment is one in which malware has been installed. This environment is also saved as a virtual machine image for reuse.

Figure 1 illustrates the basic approach used in this paper. For each sample of AV-like software (it is unknown whether the sample is fake/crude AV or legitimate AV at this point of time), indicators are measured in clean and infected environments and *statistically* compared. The key insight behind this approach is that legitimate AV behaves *differently* in clean and infected environments while fake/crude AV behaves *similarly* in clean and infected environments, because legitimate AV thoroughly analyzes files suspected to be infected with malware. fake/crude AV *cannot* change its behavior depending on the presence of an infection because it does not detect malware infection.

This approach satisfies the three requirements described in Section III-A. First, the fake/crude AV indicator does not use features specific to each instance of fake/crude AV, and thus, should be applicable to a wide variety of fake/crude AV. As shown in Section IV, our indicator successfully identified all (39 out of 39) fake/crude AV samples and all (8 out of 8) legitimate AV samples. Second, it is impractical to try to evade the indicator. Since fake/crude AV must change its behavior depending on the presence of malware infection, it must be equipped with the detection facilities that are equivalent to legitimate AV; that is, criminals must develop a legitimate AV to evade the indicator. Finally, the distinction process can be automated. There is no need for manual intervention since the indicator can be used in a systematic way and indicator measurements are compared using a statistical test.

C. Examining Potential Indicators

To discover a good indicator, we examine three candidates that can be obtained systematically: 1) the file access pattern, 2) CPU usage, and 3) memory usage. To obtain them, we use system calls hook to get the file accesses and performance monitor, a default application of the Windows OS, to get the CPU usage and memory usage.

1) File Access Pattern: While a legitimate AV has to investigate a file's content to determine whether it is infected with malware, fake/crude AV only traverses directories to obtain real pathnames; it does not access files as often as legitimate AV does, because fake/crude AV does not look hard for malware infection.

Unfortunately, file access patterns are not a good indicator of fake/crude AV because the patterns of some fake/crude AV samples are similar to those of legitimate AVs. Figure 2 shows the similarity of file access patterns between 8 legitimate AV products and 39 fake/crude AV samples. Each bar corresponds to one fake/crude AV sample, and shows the ratio of files accessed by each fake/crude AV sample to those commonly accessed by legitimate AV samples. (There are 10 bars in the figure because the remaining 29 samples do not access the files). Figure 2 shows that six fake/crude AV samples resemble legitimate AV products in terms of file access.

2) CPU Usage: Next, we investigate CPU usage. Since a malware scan such as signature matching is executed in user mode not kernel mode, we focus on the proportion of user mode time of CPU usage. User-time is expected to increase in legitimate AV if it is executed in infected environments because sophisticated scanning algorithms consume a lot of user-mode time. On the other hand, fake/crude AV is not expected to increase user-mode time in infected environments because it does not search for malware infection.

 ISBN: 978-1-61208-298-1

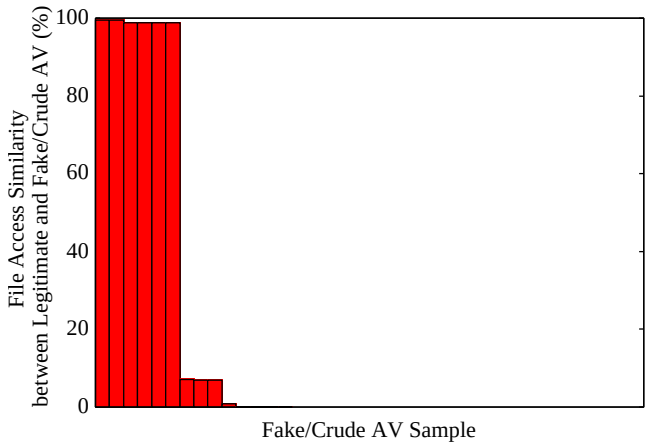

Fig. 2. Similarity of file access patterns between legitimate AV and fake/crude AV. The file access pattern is not be a good indicator because 6 out of 39 fake/crude AV samples access files accessed by legitimate AV products.

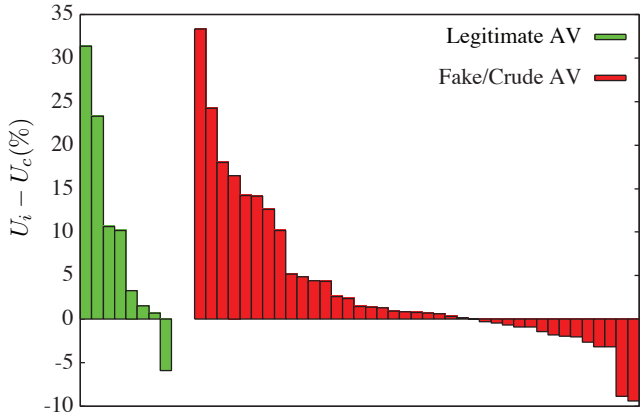

Fig. 3. Comparison of user-mode times in clean and infected environments. U_i and U_c represent user-mode time ratios in an infected and a clean environment, respectively. The green bars show $U_i - U_c$ of legitimate AV products, and the red bars show those of fake/crude AV samples. Regardless whether it is fake/crude or legitimate, $U_i - U_c$ ranges from -10% to 30%.

In spite of our expectations, the differences in user-mode time between clean and infected environments are not useful for distinguishing fake/crude AV from legitimate AV. Figure 3 shows the differences in user-mode time ratio between clean and infected environments. The y-axis shows $U_i - U_c$, where U_i stands for the user-mode time ratio measured in an infected environment and U_c stands for the user-mode time ratio measured in a clean environment. The green bars show $U_i - U_c$ of legitimate AV products, and the red bars show those of fake/crude AV samples. The overall trend in the user-mode time ratios is the same in the legitimate products and fake/crude samples. $U_i - U_c$ ranges from -10% to 30%.

3) Memory Usage: Finally, we examine memory usage. Memory usage is expected to increase in infected environments in legitimate AV because it requires more memory to perform in-depth analyses of malware. On the other hand, fake/crude AV is not expected to increase memory usage because it should behave similarly in clean and infected environments.

Our preliminary results show memory usage is an effective indicator to distinguish legitimate AV from fake/crude AV.

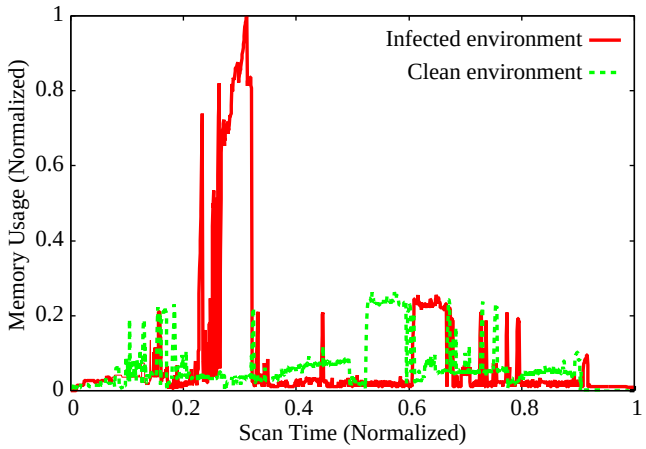

(a) McAfee (A legitimate AV product)

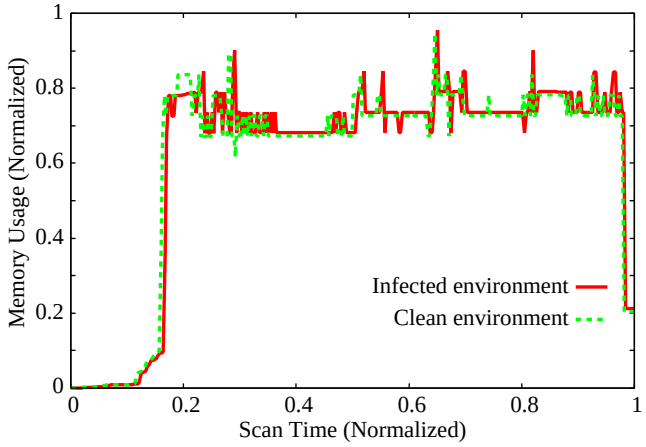

(b) Anti-Virus Elite (A fake/crude AV sample)

Fig. 4. Memory usages of legitimate AV and fake/crude AV. Memory usage and scan time are normalized. The legitimate AV uses a significant amount of memory when it detects malware. On the other hand, fake/crude AV hardly changes its usage in going from clean to infected environments.

Figure 4 shows the memory usages of one legitimate AV product and one fake/crude AV sample. It reveals memory usages differ in legitimate and fake/crude AV. In the legitimate AV product, the memory usage increases in infected environments. However, the fake/crude AV sample does not show such an increase in infected environments; it shows almost the same trend in both environments.

Figure 5 shows V_i/V_c for each legitimate AV product and fake/crude AV sample. V_i stands for the variance in an infected environment and V_c stands for the variance in a clean environment. As you can see from the figure, V_i/V_c shows different trends for legitimate AV and fake/crude AV. This suggests that memory usage is a good indicator of fake/crude AV.

D. Memory Usage as an Indicator

This section describes a concrete method to distinguish fake/crude AV from legitimate AV. On the basis of examinations in the previous sections, we choose memory usage as an indicator of fake/crude AV. A sample of AV-like software is installed in clean and infected environments and the

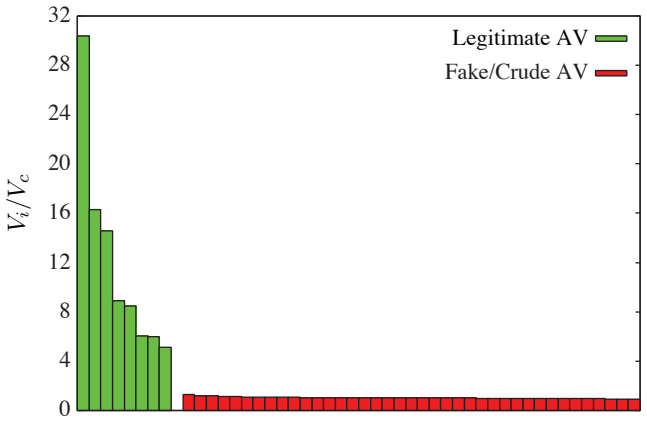

Fig. 5. Comparison of variances in clean and infected environments. V_i and V_c represent the variances of memory usage distributions in an infected and a clean environment, respectively. All legitimate AV products significantly increase the variances in the infected environment. However, fake/crude AV samples hardly change variances in these environments.

memory usage is measured in each environment. An infected environment is prepared by installing 905 unique instances of malware (about 500 MB in total), and a clean environment is by installing 500 MB of clean files. The installed malware instances and clean files are the same size and have the same directory structure. Note that we do not have to prepare these environments every time a test is performed, because a virtual machine image can be copied for reuse.

Memory usage is measured every second in each environment. For ease of mathematical formalization, the clean and infected environments are numbered 1 and 2. The measured memory usage values are grouped into a sequence for each environment and represented by Y_i, where i denotes the number of group ($1 \leq n \leq 2$).

If the distributions in Y_1 (clean env.) and Y_2 (infected env.) are not statistically different, we conclude that the tested sample of AV-like software is fake/crude AV. Otherwise, we conclude that the tested sample is legitimate AV.

To compare the memory usage distributions in each environment, we use the *Levene Test* [10], a well-known inferential statistic used to assess the equality of variances in different samples. It tests the null hypothesis that the population variances are equal. The Levene Test compares the distributions of two sequences Y_i and Y_j. If the results of the test are less than the significance level (0.05 in this paper), the difference is statistically significant. In our method, memory usage is measured M times to mitigate fluctuations and the Levene Test is repeated. If the M results of the Levene Test are all less than 0.05, we consider that the distributions are different. Since it is time-consuming to measure indicators M times, we measure indicators $\lceil\sqrt{M}\rceil$ times and perform the Levene Test on any pair of the measured indicators.

IV. Experiments

This section shows that memory usage can be used to distinguish fake/crude AV from legitimate AV. We collected 39 fake/crude AV samples listed in Table I from a malware collection site [12] and Malware Domain List [11]. In addition

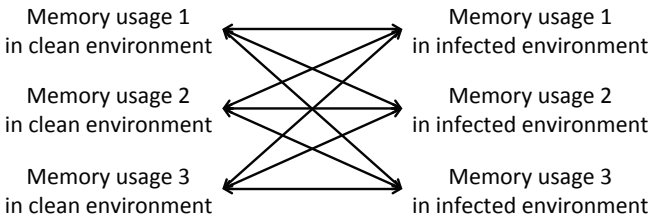

Fig. 6. Memory usages in different environments were measured for 3 times and the Levene test was performed on any pair of memory usages. If all pairs showed statistical significance, the tested AV sample was identified as legitimate. Otherwise, it was deemed fake/crude.

TABLE I. Fake/Crude AV samples

XP Internet Security 2011	XP Internet Security 2012 6.0.2900.2180
XP Home Security 2011	XP Home Security 2012 6.0.2900.2180
XP Anti Spyware 2011	XP Anti Spyware 2012 6.0.2900.2180
XP Antivirus 2011	XP Antivirus 2012 6.0.2900.2180
XP Security 2011	XP Security 2012 6.0.2900.2180
XP Total Security 2011	PC Privacy Cleaner 1.0.22.4
Patchup Plus	Virus Remover 2008 1.0.15.2
Security Tool	Virus Remover 2009 1.0.9.0
System Security	Anti Spy Safeguard 1.0.0.0
XL Guarder	Security Antivirus 2.0.2.18
Security Shield	Major Defense Kit 1.0.0.0
Protect Code	Anti Spyware Bot 9.6.9
Adware Bot 12.0.6	Security Defender 1.6.812.0
Reg Clean 1.0.0.1	Malware Removal Bot 12.0.6
Onescan 1.0.0.1	Anti Spyware Expert 1.0.22.2
Anti-Spyware 12.0.6	Anti-Virus Elite v5.0
Error Sweeper 2.8.0	Pest Detector 1.0.0.0
Registry Smart 2.10.0	Netcom3 PC Cleaner 9.1.10
Red Cross 1.0.0.0	Peak Protection 1.0.0.0
Privacy Control 2.6.0.0	

to the fake/crude AV samples, 8 legitimate AV products listed in Table II were collected. Clean and infected environments were prepared by using KVM with Qemu 1.1.1, in which Windows XP SP3 was installed and 1GB of memory is allocated. Although we used Windows XP in the experiments, our approach was not limited to a specific OS. These environments were prepared as described in Section III-D. The Levene Test was repeated 9 times. In other words, M equaled 9. Since $\lceil\sqrt{M}\rceil$ was 3 in this case, memory usage in each environment was measured 3 times.

The Levene Test was performed on pairs of measured memory usages in clean and infected environments as shown in Figure 6, and counts of less than significance level (0.05) were gathered. If the count reached 9, the tested AV sample was identified as legitimate. Otherwise, it was identified as fake/crude.

Our method identified all 39 fake/crude AV samples. The results are shown in Table III. The second row in Table III shows that none of fake/crude AV samples had positive counts of 9. In particular, the fake/crude AV could not evade detection by changing memory usage at random. For example, the two fake AV samples in Figures 7 and 8 change their memory usage at random. However, our indicator detects them correctly, because the Levene Test is performed on any pair of the measured indicators. Moreover, Figure 9 shows the memory usage of Anti-Virus Elite, a crude AV. Surprisingly, Anti-Virus Elite hardly changes its usage between clean and infected environments despite that it has a function to detect malware.

The rate of false positives is low; our method correctly

 ISBN: 978-1-61208-298-1

TABLE II. LEGITIMATE AV PRODUCTS

Avast Pro Antivirus 7.0.1426	G Data Antivirus 2011 21.1.0.1
AVG Antivirus 2012.0.1913	Kaspersky Anti-Virus 2011 11.0.2.556
McAfee VirusScan 15.0.294	ESET NOD32 Antivirus 4.2.71.2
Norton AntiVirus 18.7.0.13	Panda Antivirus Pro 2011 10.00.00

TABLE III. RESULTS OF LEVENE TEST

Name	# of Positive Results	Result
Adware Bot	0	Fake/Crude
Anti Spy Safeguard	0	Fake/Crude
Anti-Spyware	5	Fake/Crude
Anti Spyware Bot	0	Fake/Crude
Anti-Virus Elite	1	Fake/Crude
Anti Spyware Expert	3	Fake/Crude
Error Sweeper	3	Fake/Crude
Major Defense Kit	0	Fake/Crude
Malware Removal Bot	0	Fake/Crude
Netcom3	5	Fake/Crude
Onescan	0	Fake/Crude
Patchup Plus	0	Fake/Crude
PC Privacy Cleaner	5	Fake/Crude
Peak Protection	0	Fake/Crude
Pest Detector	3	Fake/Crude
Privacy Control	3	Fake/Crude
Red Cross	0	Fake/Crude
Reg Clean	0	Fake/Crude
Registry Smart	0	Fake/Crude
Security Antivirus	0	Fake/Crude
Protect Code	2	Fake/Crude
Security Defender	6	Fake/Crude
Security Shield	3	Fake/Crude
Security Tool	0	Fake/Crude
System Security	5	Fake/Crude
Virus Remover 2008	0	Fake/Crude
Virus Remover 2009	0	Fake/Crude
XL Guarder	0	Fake/Crude
XP AntiSpyware 2011	5	Fake/Crude
XP AntiSpyware 2012	4	Fake/Crude
XP AntiVirus 2011	6	Fake/Crude
XP AntiVirus 2012	2	Fake/Crude
XP HomeSecurity 2011	2	Fake/Crude
XP HomeSecurity 2012	4	Fake/Crude
XP InternetSecurity 2011	2	Fake/Crude
XP InternetSecurity 2012	2	Fake/Crude
XP Security 2011	4	Fake/Crude
XP Security 2012	3	Fake/Crude
XP TotalSecurity 2011	2	Fake/Crude
Avast	9	Legitimate
AVG	9	Legitimate
McAfee	9	Legitimate
NOD32	9	Legitimate
G Data	9	Legitimate
Norton	9	Legitimate
Kaspersky	9	Legitimate
Panda	9	Legitimate

identified the 8 legitimate AV products listed in Table II, i.e., no false positives in this experiment. All of these samples show statistical differences in clean and infected environments. Table III shows that the Levene Test reveals statistical significances with all of the legitimate AV products. As a result, our method judges these to be legitimate.

V. DISCUSSION

A. Evasion

1) Random Memory Usage: It is useless to change memory usage at random to evade our indicator. If memory usage is measured only once in clean and infected environments, the randomly changing memory usage could evade our indicator. However, as explained in Section III-D, memory usage is measured M times in our approach and the Levene Test is performed multiple times. To deceive our approach, fake/crude AV samples must change their memory usage based on the presence of malware. This means that the fake/crude AV would have to act as legitimate ones; that is, they would correctly detect malware infections.

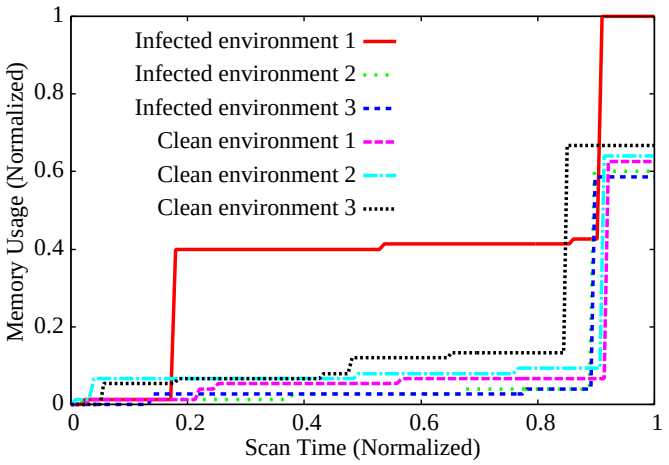

Fig. 7. Memory usages of Anti Spyware Expert, a fake AV sample. The memory usage in infected environment 1 is obviously different. However, since the other memory usages are similar, our indicator identifies this AV sample as fake. Memory usage and scan time are normalized.

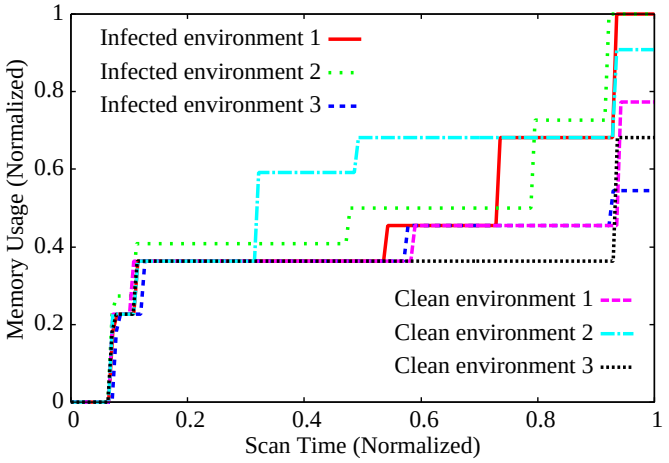

Fig. 8. Memory usage of System Security, a fake AV sample. Although the memory usages are different, the Levene Test does not show statistical significance for some pairs (e.g., infected environment 3 and clean environment 1). Memory usage and scan time are normalized.

2) Exploiting Open Source AV and Source Code of Legitimate AV: One possible approach to evade our indicator is to use open source AV or leaked source code of legitimate AV. Fake/crude AV samples based on legitimate AV could evade our indicator because their behavior is similar to legitimate AV. However, we believe our indicator raises the bar to developing fake/crude AV because fake/crude AV developers require the source code of product-quality legitimate AV. We also hope vendors can quickly develop effective signatures to detect fake/crude AV based on their products, since the legitimate vendors have the source code of their products and deeply understand the internal behavior of their products.

 ISBN: 978-1-61208-298-1

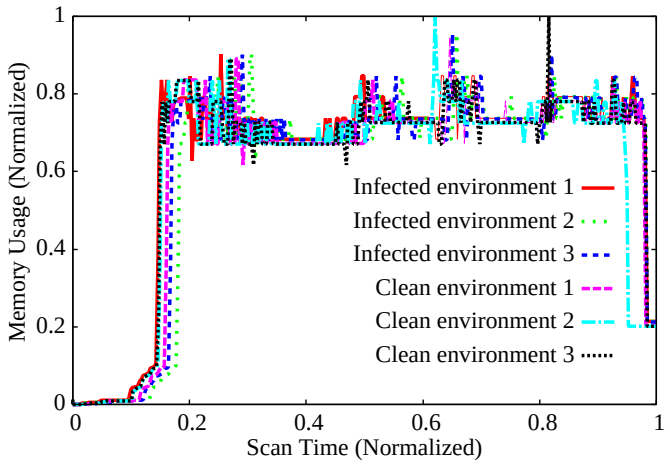

Fig. 9. Memory usages of Anti-Virus Elite, a crude AV sample. Although it can detect malware, all memory usages are almost same. Memory usage and scan time are normalized.

B. Deployment scenario

Our approach can serve to prevent software download sites form distributing fake/crude AV. Software download sites such as CNET [1] and PCMAG [2] should not distribute fake/crude AV. In spite of the careful management, though, CNET distributed a sample of fake AV in the middle of September 2012 [5]. A discrimination system based on our indicator can prevent the users of those sites from downloading fake/crude AV. Since AV software is usually indexed by tags such as "antivirus" in software download sites, all the pieces of software indexed by "antivirus" can be tested to decide if they are legitimate or fake/crude AV.

Our approach is ineffective at finding distribution sites for fake/crude AV. Since it uses the difference between legitimate and fake/crude AV, tested samples should be legitimate AV or fake/crude AV. However, it is difficult to automatically collect only AV-like software on the web. In this case, our approach cannot correctly classify non-AV samples. As a result, it reports a lot of false negatives and false positives.

C. Other Possible Indicators

We investigated three indicators, file access patterns, CPU usage and memory usage, and found that memory usage is a good indicator to distinguish legitimate AV and fake/crude AV. Although we have not sought other possible indicators in this paper, there may be other indicator to distinguish them. In the future, we plan to explore other possible indicators and determine which indicator is the best.

VI. Related Work

Recently, two studies have reported long-term analyses of fake/crude AV threat ecosystems. They show the traditional signature-based and blacklist-based approaches are useless against fake/crude AV. Rajab et al. show it is practically impossible to keep signatures with a high detection rate against fake/crude AV [15]. The detection rate rises and falls frequently. Cova et al. show that neither IP nor domain-based blacklists are effective on fake/crude AV [4]. Legitimate web sites are often blocked in IP-based blacklists, and domain-based blacklists are evaded by rotating short-lived domains.

Stone-Gross et al. suggest that credit-card companies should endeavor to identify fake/crude AV companies [16]. Fake/crude AV companies monitor the refunds that customers demand from their credit card providers, and they control these refunds so as to keep the chargeback rates low. However, this behavior leads to unusual patterns in chargebacks, which may be leveraged by credit-card companies to identify and ban fraudulent companies.

A white paper has been published to identify fake/crude AV by visual inspection [7]. It provides diverse characteristics about fake/crude AV. By using it, computer users can identify fake/crude AV by visual inspection. As described in Section II-C, some fake/crude AV samples do not have such characteristics.

VII. Conclusion

In this paper, we have searched for an indicator that captures behavioral differences in legitimate AV and fake/crude AV. We have conducted experiments showing that memory usage would be a good indicator, and developed a systematic method, based on a statistical test, to distinguish fake/crude AV from legitimate AV. To demonstrate the effectiveness of our method, we collected and tested 39 real fake/crude AV samples and 8 legitimate AV products. According to our experiments, our method correctly identified all fake/crude AV samples and all legitimate AV products.

References

[1] "CNET | Download.com," http://download.cnet.com [retrieved: Jul, 2013].

[2] "PCMAG.COM," http://www.pcmag.com/downloads [retrieved: Jul, 2013].

[3] "Virus Total," https://www.virustotal.com [retrieved: Jul, 2013].

[4] M. Cova, C. Leita, O. Thonnard, A. D. Keromytis, and M. Dacier, "An Analysis of Rogue AV Campaigns," in *Proceedings of the 13th International Symposium on Recent Advances in Intrusion Detection (RAID '10)*, Sep. 2010, pp. 442–463.

[5] S. Doyle, "How To Remove RegGenie Rogue Antivirus Software – Uninstall RegGenie Malware (Identity Theft Protection)," http://botcrawl.com/how-to-remove-reggenie-rogue-antivirus-software/ [retrieved: Jul, 2013], Sep. 2012.

[6] M. Fossi, D. Turner, E. Johnson, T. Mack, T. Adams, J. Blackbird, M. K. Low, D. McKinney, M. Dacier, A. D. Keromytis, C. Leita, M. Cova, J. Orbeton, and O. Thonnard, "Symantec Report on Rogue Security Software," http://www4.symantec.com/Vrt/wl?tu_id=TeCm125590003756772344 [retrieved: Jul, 2013], Oct. 2009.

[7] A. Karnik, J. Avelino C. Rico, A. Prakash, and S. Honjo, "Identifying Fake Security Products," http://www.mcafee.com/us/resources/white-papers/wp-identifying-fake-security-products.pdf [retrieved: Jul, 2013], 2009.

[8] U. Kiguolis, "Remove Security Antivirus," http://www.2-spyware.com/remove-security-antivirus.html [retrieved: Jul, 2013], Mar. 2010.

[9] U. Kiguolis, "Remove Anti-Virus Elite," http://www.2-spyware.com/remove-antivirus-elite.html [retrieved: Jul, 2013], Feb. 2012.

[10] H. Levene, *Robust Tests for Equality of Variances*, Ingram Olkin and Sudhish G. Ghurye and Wassily Hoeffding and William G. Madow and Henry B. Mann, Ed. Stanford University Press, 1960.

[11] MDL, "Malware Domain List," http://www.malwaredomainlist.com/ [retrieved: Jul, 2013].

[12] Open Malware, "Open Malware — Community Malicious code researh and analysis," http://offensivecomputing.net [retrieved: Jul, 2013].

 ISBN: 978-1-61208-298-1

[13] F. Paget, "Running Scared: Fake Security Software Rakes in Money Around the World," http://www.mcafee.com/us/resources/white-papers/wp-running-scared-fake-security-software.pdf [retrieved: Jul, 2013], 2010.

[14] N. Provos, P. Mavrommatis, M. A. Rajab, and F. Monrose, "All Your iFRAMEs Point to Us," in *Proceedings of the 17th USENIX Security Symposium*, Jul. 2008, pp. 1–16.

[15] M. A. Rajab, L. Ballard, P. Mavrommatis, N. Provos, and X. Zhao, "The Nocebo Effect on the Web: An Analysis of Fake Anti-Virus Distribution," in *Proceedings of the 3rd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '10)*, Aug. 2010.

[16] B. Stone-Gross, R. Abman, R. A. Kemmerer, C. Kruegel, D. G. Steigerwald, and G. Vigna, "The Underground Economy of Fake Antivirus Software," in *Proceedings of the 10th Workshop on Economics of Information Security (online) (WEIS '11)*, Jun. 2011.

[17] Z. Wu, S. Gianvecchio, M. Xie, and H. Wang, "Mimimorphism: A New Approach to Binary Code Obfuscation," in *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS '10)*, Oct. 2010, pp. 536–546.

 ISBN: 978-1-61208-298-1