Understanding and Detecting International Revenue Share Fraud

Merve Sahin*
SAP Security Research
merve.sahin@sap.com

Aurélien Francillon
EURECOM
aurelien.francillon@eurecom.fr

***Abstract*—Premium rate phone numbers are often abused by malicious parties (e.g., via various phone scams, mobile malware) as a way to obtain monetary benefit. This benefit comes from the 'revenue share' mechanism that enables the owner of the premium rate number to receive some part of the call revenue for each minute of the call traffic generated towards this number.**

This work focuses on International Revenue Share Fraud (IRSF), which abuses regular international phone numbers as the so-called *International Premium Rate Numbers (IPRN)*. IRSF often involves multiple parties (e.g., a fraudulent telecom operator in collaboration with a premium rate service provider) who collect and share the call revenue, and is usually combined with other fraud schemes to generate call traffic without payment. Although this fraud scheme has been around for several years, it remains to be one of the most common phone fraud schemes, reportedly leading to billions of dollars of losses every year.

In this paper we explore the IRSF ecosystem from multiple angles, via: (i) A telephony honeypot that observes IRSF attempts towards an unused phone number range (i.e., *a phone number gray space*), (ii) A dataset of more than 3 Million test IPRNs and more than 206K test call logs we collected from several online IPRN service providers during 4 years, and finally, (iii) A real-world call data set from a small European operator, involving 689K call records, that we analyze to find IRSF cases. By leveraging our observations from (ii), we propose several machine learning features that can be used in IRSF detection. We validate our approach on the dataset in (iii), achieving 98% accuracy with a 0.28% false positive rate in detecting the fraudulent calls.

I. Introduction

Telephony is the largest and the oldest deployed network. Telephony networks carry a huge volume of call, messaging and data traffic every day. This is a complex and opaque ecosystem, which combines multiple technologies and involves various types of service providers and customers. Because calls can be expensive, and can be used to monetize third party services, telephony becomes a very profitable environment for many fraud schemes [57]. Among these, International Revenue Share Fraud (IRSF) is one of the most profitable for the fraudsters. According to the 2017 CFCA fraud loss survey, IRSF costs telecom operators $6.10B a year (roughly 20% of the estimated communication fraud) [12]. Moreover, it remains to be the most common fraud type reported by the telecom operators in 2019 [15]. IRSF can affect all users of the telephone network, both the individual users (fixed lines, prepaid and postpaid mobile subscribers) and the enterprise phone systems (that often include a Private Branch Exchange - PBX - to handle the internal and external communications of the enterprise). As an example, a recent report from Check Point Research [61] shows how the attackers target and exploit PBX systems worldwide, and use this access to generate calls to the so-called IPRNs for profit.

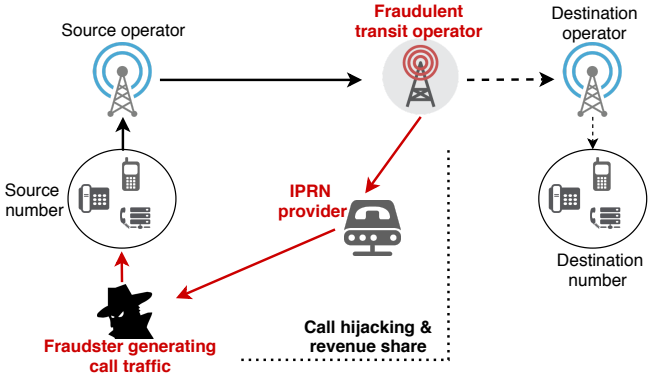

Fig. 1: IRSF Overview.

A. International Revenue Share Fraud Overview

IRSF often involves three types of fraudulent actors:

- A third party service provider who advertises regular phone numbers as the *so-called* International Premium Rate Numbers (IPRN). Throughout the paper, we will call these service providers as the *IPRN providers*.
- A fraudster who obtains IPRNs from an IPRN provider, and generate fraudulent call traffic to those phone numbers to receive monetary benefit.
- A fraudulent telecom operator who hijacks the phone calls and reroutes them to the IPRN provider, while collecting and sharing the call revenue. The hijacking of the calls is often performed during their transit, and it is often impossible to identify the fraudulent transit operator due to the opacity of international call routing [57].

An overview of the IRSF fraud scheme is presented in Figure 1. Next, we explain each of these actors in more detail.

*This author completed part of the work at EURECOM.

Network and Distributed Systems Security (NDSS) Symposium 2021
21-25 February 2021, Virtual
ISBN 1-891562-66-5
https://dx.doi.org/10.14722/ndss.2021.24051
www.ndss-symposium.org

Fig. 2: Example advertisements from IPRN providers.

1) IPRN providers: A simple online search for *international premium rate numbers* reveals many websites advertising them, and promising fast and easy money payout guarantee for the call traffic generated to these numbers. Some of the websites also provide easy-to-setup and ready-to-use Integrated Voice Response services (IVRs) that can be used for setting up the *premium rate service*, such as broadcasting audio books, live radio, live games or weather forecast. Figure 2 shows some example advertisements taken from IPRN provider websites.

In addition, such websites often provide specific web interfaces for testing purposes: they publish a set of *test numbers* that the fraudsters can call to check if the calls they generate will reach the IPRN provider and generate revenue. This ensures that the hijacking of the call is successful on this route, i.e., at some part of the call route, the involved fraudulent transit operator is able to hijack and re-route the call to the IPRN provider (which enables the revenue share). Throughout the paper, we will call these test numbers as *test IPRNs*.

The test interfaces also show the records of the calls that are initiated to the test IPRNs. Fraudsters can view the call records in real time, to check if the current call they are making has reached the IPRN provider. Note that, as we will discuss in detail later, reachability of these test IPRNs may not be consistent. That is why the IPRN providers often report some statistics about which countries/operators were recently reachable from which other countries/operators.

2) Fraudster who generate call traffic: Fraudsters aim to generate high volume of phone calls to the so-called IPRNs via compromised phone lines. To be able to avoid paying for the phone calls, they may use different fraud schemes [31]:

- **Fraudulently obtained SIM cards.** Fraudsters may use stolen SIM cards or, SIM cards that are obtained via subscription fraud (e.g., subscribed with fake identity and payment details). They often use the SIM card in a roaming network, abusing the fact that the call records may not be immediately available to the home operator [29], [65].
- **PBX hacking.** Fraudsters might gain unauthorized access to enterprise telephone systems (PBXs) (e.g., by abusing remote access interfaces, weak credentials or misconfigurations [39]) and use the phone lines to initiate outgoing phone calls (often, multiple calls in parallel) [6], [44], [48], [61].
- **Wangiri (One-ring/Callback) scam.** Fraudsters trick the phone users to call back certain international phone numbers, by ringing users' phones for only a few seconds, which results in missed call notifications [14], [50], [58], [66]. In recent years, a large number of countries have been target of Wangiri scam. Consumer protection bodies, regulatory bodies, and news agencies in various countries try to warn the users against this type of scam [51].
- **Mobile malware.** Malware that infects mobile phones might be able to initiate phone calls stealthily [7], [21], [35], [60].

In the case of PBX hacking or fraudulently obtained SIM cards, fraudsters aim to generate large volume of calls as quickly as possible (before they are detected). To stay under the radar of fraud detection mechanisms, they might generate calls to multiple IPRNs, while keeping the duration of each call low. Thus, they may use the test interfaces of IPRN providers to find the suitable destinations that will allow revenue share. They also often commit fraud during the weekend, when the fraud management team of the telecom operator is less likely to react quickly.

In the case of Wangiri scam, fraudsters would determine to use certain IPRNs belonging to the destination countries to which the calls are likely to be hijacked. It is possible that, some of the victims' calls will not generate revenue for the fraudster, because the victims may be using different originating operators. For mobile malware, it might be possible to remotely update the telephone numbers that will be used for fraud.

3) Fraudulent telecom operator: Since the deregulation of the telecommunication industry in most of the countries, lots of small or medium sized operators have emerged. Most of them do not own their own infrastructure, but resell the service they buy from other operators (e.g., Mobile Virtual Network Operators-MVNOs) [57]. Moreover, with the advances in VoIP technology, the number of transit operators that work in wholesale international market have also increased. Open source software and cheap equipment reduced the cost of creating a telecom operator [53]. While setting up a telecom operator became rather easy and low cost, the telecommunication network still remains opaque. Operators make bilateral agreements to buy and sell call traffic. However, due to the confidentiality of these agreements, it is often not possible to know the complete route that a call has taken. As a result, it is difficult to know if an operator manipulates the call route, and to pinpoint this operator.

IRSF requires the collaboration of a telecom operator who will route the calls to the IPRN provider. In some cases, this operator may own the destination number, in other cases it may be the same entity as the IPRN provider. In Section II we will explain in more detail the mechanisms enabling the hijack of the phone calls and variations of IRSF.

B. Our study and contributions

In this paper, we analyze the IRSF ecosystem from multiple perspectives and finally we propose a machine learning method for IRSF detection. We use three main data sources for our study (Figure 3 and Table I summarize the data sources and experiments):

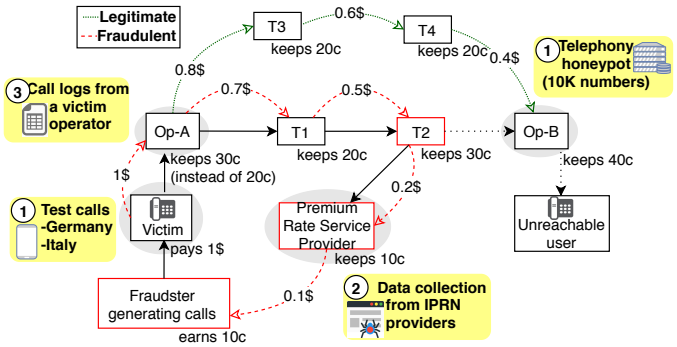

Fig. 3: Summary of experiments on International Revenue Share Fraud.

1) We present our findings from a telephony honeypot that was designed to detect the possible IRSF attempts on a specific phone number range of 10,000 numbers (Section III). This honeypot receives 259 international phone calls in around 2 years. In addition, we generate several calls from Germany and Italy to test the accessibility of the honeypot numbers during this experiment.
2) We present various insights about the IRSF ecosystem using the 3.14M *test numbers* and 206K *test call logs* that we collect from the public test interfaces of IPRN provider websites (Section IV). Combining these insights with our domain knowledge on how IRSF works, we then present a set of machine learning features that can be used in detection of IRSF (Section VI).
3) We use a real-world dataset of 689K Call Data Records (CDRs) from a telecom operator to evaluate the machine learning features that we propose.

With these experiments and analyses, we make several contributions in this paper:

- Via the telephony honeypot, we provide empirical evidence on the hijack of an unused phone number range that belongs to a small European operator.
- By analyzing the phone number space used by IPRN providers, we demonstrate the extent and prevalence of IRSF.
- We point out the challenges in IRSF detection and analyze the existing detection techniques.
- We evaluate our machine learning approach via the real-world call dataset and achieve 98% accuracy (with 0.28% FPR) in experiments with 10-fold cross validation. Moreover, our machine learning model detects majority of the fraudulent calls (without any false positives) when applied to previously unobserved fraud cases.

II. Background: IRSF Mechanisms

International revenue share is a complex fraud scheme that manipulates certain weaknesses of the telephony networks, and it may come in many variations. In this section we will examine the types of abuse in the telephony network that would enable IRSF, and the different forms of this fraud.

A. Enablers of IRSF

1) Misuse of phone number allocations: Numbering plans enable the allocation of phone number ranges worldwide. First, each country is assigned a country code by the International Telecommunication Union (ITU) [2]. Then, the regulator in each country further allocates phone number ranges for landline, mobile or special purpose services, and assigns these numbers to telecom operators.

However, there are several issues with the numbering plans:

- There is no centralized numbering plan database that keep up-to-date information about all operators in every country.
- Numbering plans show the allocated phone number ranges, but do not give information on individual numbers (e.g., whether a given phone number is in use or not).

Due to these issues, the originating or transit operators often cannot know if the destination number is legitimately allocated and/or currently assigned to a user [4].

As we will show in upcoming sections, the phone numbers that are abused for IRSF are in fact regular phone numbers (landline or mobile), and they are not allocated for the purpose of international premium rate services. In fact, as stated in ITU guideline E.169.2 [3], the universal, legitimate number range allocated for international premium rate services consists of the 3-digit code *+979*, followed by a 9-digit phone number. Thus, any other number advertised as an International Premium Rate Number is not legitimate.

In particular, the ITU guideline E.156 [36] reports the misuse of regular phone numbers as premium rate numbers, stating that:

> [International country] codes are not designed to be used as charging band indicators for calls that are *terminated short* of the designated country. Furthermore *separate provision* is made within ITU-T Recommendations for designating International Premium Rate and Shared Cost Service. [...] E-series numbering resource will only be utilized by the assignee for the specific application for which they have been assigned.

2) Abuse of international call routing mechanisms: Routing of an international call from originating operator to the destination operator is usually enabled by several transit operators carrying the call in between. Each operator on the call route decides where to route the call next, depending on their peering agreements and routing algorithms. However, the operators have only partial visibility on the call route, and the agreements and algorithms they employ are often confidential [57]. As a result, it is often not possible to trace the route that a call has taken.

Fraudulent telecom operators can take advantage of this, and terminate the phone calls before they reach the actual destination operator (also called, *short-stopping* [36], [57]). In IRSF, the fraudulent transit operator can short-stop the calls by misrouting them to the IPRN provider. Due to the opaqueness of the call route, identifying the fraudulent party becomes almost impossible.

TABLE I: Summary of the datasets used in our study.

Dataset	Time interval	Nature of data	# of records	Data Fields	Way we obtain
Honeypot call logs (1)	Jan16 - Oct18	Call record	259	A number, B number, Date	Telephony Honeypot
Test IPRNs (2)	Jan16 - Nov19	Phone number	3.14 M	Phone Number, Collection date, Source Website	Crawling providers' sites
Test call logs (2)	Jan16 - Nov19	Call records	206 K	A number, B number, Date, Duration, Source Website	Crawling providers' sites
Real world call logs (3)	Oct14 - Dec14	Call record	689 K	A number, B number, Date, Duration, Localisation	Operator

3) Number range hijacking: Telecom operators employ complex algorithms to handle the interconnections with other operators. For each outgoing call, the operator decides which route it will send the call over: This can be another operator with a peering agreement, or through an interconnect broker that trades call traffic [36]. As depicted in Figure 3, Operator A may have multiple options (in this case, T1 and T3) to route the call to its final destination, Operator B. In our example, T1 advertises a lower price (70 cents) to route this call, compared to T3 (that advertises 80 cents). If Operator A, for instance, employs a *Least Cost Routing* [18] policy, it is likely to select the cheapest route (T1) to maximize its own profit (30 cents instead of 20). However, the low cost of this route is in fact due to T2, which is advertising a very cheap rate, but actually not terminating the call to its legitimate destination. Instead, T2 advertises a low cost route to attract traffic with the goal of hijacking (take over). The calls to this destination number range, will then be misrouted to an IPRN provider.

As routing decisions are dynamic, fraudsters need to constantly make sure that the call they generate will be routed through the involved fraudulent operator. This is the reason why IPRN providers employ the test interfaces (mentioned in Section I-A1). Note that, test interfaces often advertise unallocated or unused number ranges. Otherwise the test calls may ring the phones of genuine users when the hijack of the route fails.

A significant example of number range hijacking for IRSF occured in 2005, towards Pacific Islands [5], [25]. The originating operators who incurred losses due to this fraud started to block all the calls to the Pacific Islands to avoid fraud and protect their customers [59]. As a result, Pacific Islands were not reachable anymore from certain parts of the world, and their revenue from incoming international call traffic significantly dropped. Their reputation was damaged while they were not responsible for the fraud, nor they were able to do anything against it [5].

B. Variations of IRSF

Depending on the fraud agreement between the operator and the IPRN provider, the phone number space abused for fraud can either be allocated and in use (assigned to actual users), or can be unused (not assigned to any user) or unallocated. Moreover, the phone number space can be actually owned by the operator who is abusing it, or the operator might just be hijacking and short-stopping the calls that are intended to be terminated on this number range. In the latter case, the success of calls depends on the hijacking ability of the operator.

Although it is challenging to identify the exact relationship between the IPRN providers and the telecom operators who collaborate with them, it is possible that these two are in fact the same entity. It is also possible that there is a chain of premium rate number resellers, so the IPRN provider is not directly connected to the operator abusing/hijacking the number range.

Depending on the revenue share mechanism and fraud agreement, the owner of the misused number range may or may not be aware that its numbers are used for IRSF [4], [36]. In summary we can distinguish the following cases:

- *Terminating operator as part of the scheme:* The terminating operator (owner of the victim number range) resells its numbers to a premium rate service as IPRNs, and terminates the illegitimate traffic on its own network. The revenue of traffic is shared between the terminating operator, the premium rate service provider and the fraudster who is generating the calls.
- *Transit operator is rerouting calls illegitimately:* A transit operator can make an agreement with a premium rate service provider to misroute the hijacked phone calls. In this case, the fraudulent transit operator short-stops the calls to the victim number range, keeps the termination fee, and shares the benefit with the IPRN provider (as depicted in Figure 3). Thus, the owner of the victim number range may not be aware of the fraud scheme performed using its numbers [45], as the calls will never reach their intended destination.

III. Telephony Honeypots to Observe IRSF Attempts

Operators can setup telephony honeypots to observe hijacks on their own unused number ranges. Such a honeypot can collect calls that are aiming for revenue share, but were not properly hijacked by the fraudulent transit operators (i.e., failed IRSF attempts). In this section, we describe our findings from a similar honeypot that we employed for 2 years, between January 2016 and October 2018.

Nature of the honeypot numbers: Our honeypot consists of 10,000 phone numbers that are reserved as *technical numbers* in GSMA roaming database. This number range is essentially reserved for quality control and testing purposes and it is not supposed to be assigned to individual users. Thus, these numbers are not supposed to receive any calls. However, as the number range belongs to a small European country (a slightly high-cost destination), it is likely to be targeted by the fraudsters. In that sense, our honeypot acts like a network telescope [19] for telephony.

We set up the honeypot as an Asterisk PBX server located in the premises of the telecom operator. The operator directly routes the incoming calls to this PBX server, over a SIP trunk. In the honeypot, half of the honeypot numbers are configured to immediately decline (hangup) the received calls. For the other half, the server first emits a ring tone for 12 seconds and then emits a busy signal for 10 seconds before hangup. Note that, 10,000 numbers actually constitute a small number range when compared to the total allocated phone number space in this country. However, we still observe many IRSF attempts, as we will soon describe. As an additional experiment, starting from September 2016, we also generated periodic test calls to a few of the honeypot numbers from two mobile phones that we deployed in Germany and Italy.

Observations from the collected calls During the 2-year period, our honeypot received 259 international calls with caller IDs indicating that the calls were originated from 77 different countries. These calls come from 156 unique caller IDs, while 78 calls have their caller IDs anonymized. Figure 4 shows the monthly number of international calls received by our honeypot (excluding the test calls generated by us).

An empirical evidence of hijack on honeypot numbers As can be seen in Figure 4, our honeypot observed an unusual call traffic in January'17. More precisely, in 5 days (from the 11th to the 16th), 117 calls were received from 48 countries.

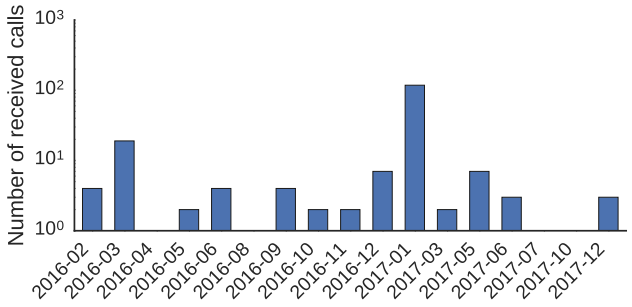

Fig. 4: Monthly number of received calls at honeypot.

Moreover, starting from the 6th of January'17 (12pm) to the 7th of January'17 (5am), the 30 test calls that we generated from Germany were all answered and billed for 1,5 minutes on average. In addition, these calls were not received by our honeypot; which indicates that they were terminated at a different location.

We looked up the origination numbers of these calls in a commercial numbering plan database [1]. For many of the origination numbers, we could only find a matching country code, which means they belong to unallocated number ranges, and indicates that they might actually be spoofed numbers. Also, it is interesting that some originating numbers belong to 'supplementary services', which are usually the premium rate number ranges in their corresponding countries.

This incident is a strong evidence that the honeypot number range was advertised as an IRSF destination during this time period, and it attracted a lot of call traffic. Although the hijack of the number range was not always successful (as we observed from the 117 calls received by the honeypot), it was indeed successful on our test calls originated from Germany.

Note that, our honeypot is only able to observe the calls that were failed to be hijacked. If a fraudster makes an initial test call to an advertised IPRN, and the test call is successfully hijacked, only then he would start generating the actual fraud calls (in which case, we cannot observe in the honeypot). Thus, we did not expect to receive a very large number of calls in the honeypot.

We also note that, the calls received on the honeypot are not likely to be robocalls/telemarketing: Such unsolicited calls are often domestic, or spoofing domestic caller IDs, whereas our honeypot received international caller IDs. Unwanted calls will also generally target known telemarketing lists or commonly used number ranges, while our range has never been allocated to real users. Some of the calls could be misdials, however, most of the calls were received during the 5-day period, when our own test calls were also hijacked.

IV. A DEEPER LOOK INTO THE IRSF ECOSYSTEM

In this section, we present our observations on the IRSF ecosystem, using the data collected from online IPRN providers.

A. Data Collection

By making online searches for"international premium rate number" keywords, we identified 45 websites that provide IPRNs. We then searched for the names of these IPRN providers online and on social media. With this, we identified the test interfaces of 15 IPRN providers.

These test interfaces are often advertised on social media (e.g., Facebook, Twitter) or online forums, together with the credentials for a publicly available test account (e.g., username:*test*, password:*test123*). Those public test accounts are designed so anyone can access the test interface, list the advertised test IPRNs, and see the previous test calls recorded on the system. Typically, these accounts are used for testing IPRN success, and the fraudsters then obtain separate, private accounts to generate fraudulent calls and obtain the payout. We have never obtained such private accounts, but we aimed at collecting data from the publicly available test accounts that are explicitly advertised by the IPRN providers. Although automated data extraction was difficult on some of the interfaces, we were able to collect the test IPRNs from 12 of the websites, and from 6 of them, we could also collect the test call logs (which are evidence of successfully hijacked calls). Overall, we collected 3,149,793 distinct test IPRNs and 206,263 test call logs from January'16 to November'19. Moreover, using a commercial numbering plan database [1], we extracted further information on the test IPRNs and the source and destination numbers of the call logs. Table II summarizes the type of data we extracted using the numbering plan database. Note that, even though we use an up-to-date numbering plan database, our database may not be 100% accurate. For example, some recent number range allocations or modifications might be missing. However, such commercial databases provide the most comprehensive data available to us (and to many operators).

As another remark on accuracy: Some websites may obfuscate the source numbers in the test call logs, e.g., by removing the last few digits. However, as long as the country

TABLE II: Type of data extracted using the numbering plan database.

Data field	Explanation
Country Name	Name of the country that the phone number is allocated to.
Country Code (CC)	International country prefix as specified in ITU-T E.164 Recommendation [2].
National Destination Code (NDC)	A number prefix that identifies a geographic area or a service. (Some number allocations do not contain NDC, but only SN.)
Subscriber Number (SN)	Initial digits of the subscriber number that is used to distinguish between different areas prefixed by the same NDC. (Some number allocations do not contain SN, but only NDC.)
Subscriber Number Length (SNL)	Valid length (in terms of the number of digits) of the subscriber number. (Some number allocations do not specify the SNL.)
Number type	The type of the international CC-NDC-SN sequence (e.g., Mobile, Fixed, Special Service)
Network and operator name	Name of the operator holding the number range. (Not available for all number allocations.)
Number range validity	- Invalid CC: The number does not match any CC (No such instances in our dataset). - Unallocated number range: The number matches a CC, but does not match any allocated NDC-SN sequence inside the country code. - Valid number range: The number matches with an allocated CC-NDC-SN sequence.
Number length validity	- Valid length: The number has a valid length according to its matching number range. Note that for the unallocated number ranges (i.e., numbers matching only a CC), number length can still be valid. Also, if SNL is not specified for a number, we count it as valid length. - Invalid length: The number has an invalid length according to its matching number range.

code, national destination code and/or the subscriber number are available, we can extract data about those numbers using the numbering plan database. Moreover, some of the source numbers might be spoofed (e.g., by the fraudsters who generate the calls from a compromised PBX) but others cannot (e.g., calls generated from a stolen SIM card).

B. Analyzing IPRN providers

In Table III, we give an overview of the data that we collect from each of the 12 IPRN providers. In the *Entire data* column, we present the data collected during the total duration of January'16 to November'19. However, starting from February'18, we improved our data collection framework, and we were able to collect more fine-grained data from 9 of the providers. Thus, in the *Fine-grained data collection* column, we present more detailed statistics on the lifetime of IPRNs (i.e., the average number of days that an IPRN has been advertised on this website) and the number of new (unseen) IPRNs advertised per day, per provider.

As it can be observed in Table III, IPRN providers have different profiles: While some of them (P1, P2) update the advertised IPRNs every few days, and add thousands of new IPRNs per day; others may keep the same IPRNs for almost a year (P3, P4, P8) or not update any IPRN at all (P9).

Another observation is that, there are very few IPRNs that are advertised by multiple providers. Among 3.14M IPRNs, only 88,356 of them are repeated across multiple providers. In fact, except P1 & P3 that share around 80K IPRNs; and P9 & P12 that advertise very few IPRNs; the rest of the providers only have a small part of their numbers shared. We find that, a shared IPRN is advertised by 2 providers on average, and 4 providers maximum.

However, looking at the advertised number ranges (instead of distinct IPRNs) gives a different picture. According to Figure 5, when the last 4 digits of the IPRNs are ignored, ∼80% of the number ranges are shared across all providers. This shows that, the hijacked number ranges are shared between the IPRN providers in a fine-grained way. In other words, a fraudulent operator who hijacks a number range is likely to assign different portions of the range to different providers, or the hijack is done on very specific ranges.

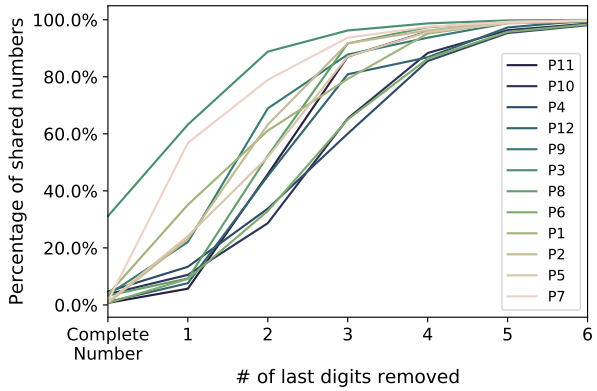

Fig. 5: Number sharing between IPRN sources.

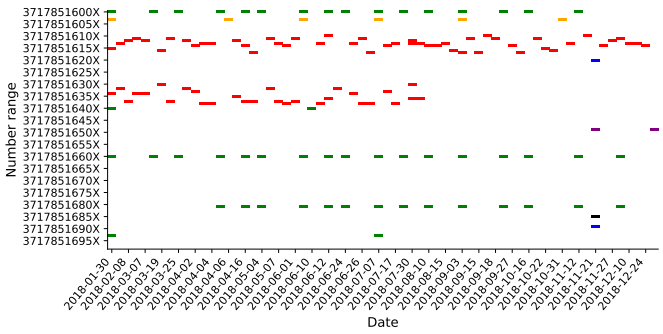

Fig. 6: IPRNs advertised by different providers, for a small number range from Latvia. (Red:P1, Green:P4, Orange:P3, Purple:P6, Black:P7, Blue:P8)

As an example, Figure 6 demonstrates the advertisement of a range of 1000 Latvian numbers by 6 different IPRN providers during 2018. We can easily observe that each IPRN provider advertises a separate sub-range. Furthermore, we can see that the number ranges are rotated (likely to avoid blacklisting) and that significant "holes" are present.

This figure shows the test IPRNs, but in fact after a successful test call, a fraudster is provided one or multiple dedicated IPRNs in a similar number range, to which he will

TABLE III: Summary of the data collected from 12 IPRN providers.

	Entire data Jan'2016 - Nov'2019		Fine-grained data collection Feb'2018 - Nov'2019			
Provider	# Distinct IPRNs	# Shared IPRNs	# Distinct IPRNs	# Data collection days	Lifetime of IPRNs (days) ($avg \pm stdev$)	# New IPRNs per day ($avg \pm stdev$)
P1	2,335,962	85,751	2,335,881	597	5 ± 7	$4,385 \pm 6,380$
P2	460,370	2,981	460,319	245	5 ± 5	$1,954 \pm 2,373$
P3	269,016	83,779	14,791	598	354 ± 191	22 ± 35
P4	39,775	1,865	39,775	602	398 ± 205	55 ± 144
P5	49,985	532	36,344	598	78 ± 130	106 ± 981
P6	21,432	144	21,432	86	48 ± 34	$188 \pm 1,055$
P7	7,603	101	7,603	355	102 ± 56	55 ± 84
P8	1,951	64	1,951	356	322 ± 87	2 ± 2
P9	254	254	216	608	564 ± 155	–
P10	47,611	1,794	–	–	–	–
P11	1,803	13	–	–	–	–
P12	774	774	–	–	–	–
Total	3,149,356		2,837,205			$4,539 \pm 6,306$

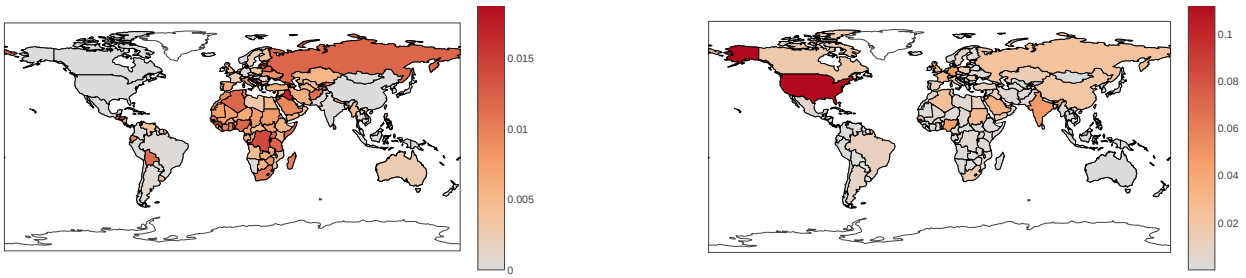

(a) Advertised test IPRNs (destinations).

(b) Test calls origins.

Fig. 7: Geographical representation of the source and destination of the IRSF fraud.

generate the actual fraud calls. We expect that the fraudster will have the successful calls logged on his own private account on the website, and we do not see these call logs in the public test account. The fraudster will get a share of the revenue generated from any call reaching the IPRNs that was assigned to him. This also suggests that the fraudulent operator does not assign the same IPRNs to multiple IPRN providers, because they wouldn't be able to know with whom they are supposed to share the revenue.[1]

C. Analyzing test IPRNs

Overall coverage. Our dataset includes test numbers targeting 238 out of 247 international destinations[2] and 1070 operators[3]. This shows that IRSF can target a large variety of countries, with varying call termination costs. Indeed, the whitepaper by TransNexus [9] analyzes payout rates from 193 countries and mentions that the fraudster's benefit can be as low as $0.00013 per minute. Another observation is that, none of the test numbers belong to the legitimate *Universal International Premium Rate Number* range (+979) specified by the ITU [3], [37]. Figure 7a shows the distribution of collected IPRNs across countries worldwide. African countries, Russia and part of Europe appear to be the most affected.

On the other hand, the popularity of the countries seem to change over time. Although some countries like Latvia, Lithuania and Iraq are always in the top 10 destinations, the other popular destinations keep changing over time (Table X in Appendix).

Validity. In Table IV, we present the validity of the test IPRNs, classified by the validity of the number range and number length. Number range validity checks if the number belongs to an allocated CC-NDC-SN (Country Code - National Destination Code - Subscriber Number) sequence defined in the numbering plan database. Overall, 70% of the numbers belong to a valid number range and have a valid length according to the numbering plan database we use. The remaining 30% either have an invalid length, or belong to an unallocated number range, or both. Note that the call route will initially be decided based on the CC or CC-NDC combination [4]. Thus, even if the rest of the number range is not allocated, or has invalid length, the call will be routed to the selected transit operator(s) which may hijack it at some point.

Number type. Next, we look at the number type information for the test IPRNs. Our numbering plan database specifies

TABLE IV: Validity of test IPRNs

	Valid length	Invalid length	Total
Valid range	70%	11.7%	81.7%
Unallocated range	8.7%	9.6%	18.3%
Total	78.7%	21.3%	100%

[1]It would be possible to use the caller ID for this, but this information is not reliable and is sometimes modified during an international call routing.

[2]Including countries but also, e.g., territories and satellite services.

[3]Note that operator information is only available for mobile numbers. Our numbering plan includes 1522 different mobile operators (including MVNOs).

a number type for each allocated number range (CC-NDC-SN sequence). However, for the 18.3% of the test IPRNs which do not match an allocated number range, number type information is not available. As we show in Table V, mobile number ranges are the most frequently abused. This may be because the mobile number ranges are usually more expensive, and therefore allow for a better revenue. Another possible explanation is that it is easier to check if a mobile number range is currently in use (assigned to someone) or not, by performing HLR lookups on the SS7 network.

TABLE V: Types of IPRN test numbers.

Number Type	%
Mobile	56.9
Fixed	15.9
Supplementary Services	7.8
Unallocated number range	18.3
Satellite	1.0

Dispersion of numbers. Figure 8 shows the cumulative distribution of the number of distinct test IPRNs by the percentage of countries. We find that, around 17% of countries have less than 1000 numbers advertised, whereas almost half of the countries have more than 10,000.

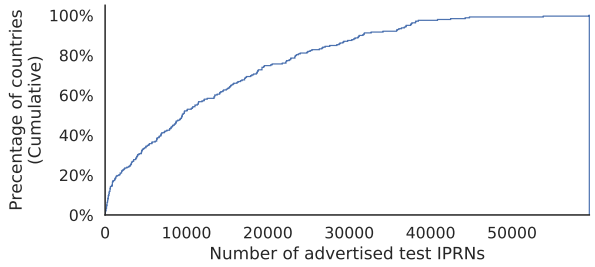

Fig. 8: Empirical cumulative distribution of the advertised test numbers by the % of countries.

It is likely that the numbers which are advertised as IPRNs are hijacked as part of a whole block of numbers. For example, a block of ten numbers in Latvia could be represented as +371xxxxxxxy, where the digits represented by an 'x' are fixed, and digits represented by a 'y' vary inside this range. Similarly a range of hundred numbers can be represented as +371xxxxxxyy. In our dataset, grouping the numbers in blocks of 10 (ignoring the last 1 digit) results in 2.4M number ranges, whereas grouping in blocks of 100 results in 1.6M ranges.

For the top 10 countries with the largest number of advertised test numbers, Figure 9 shows the number of collected test IPRNs, and the unique number ranges when the last digits of the number are ignored. As we can see from this figure, the quantity of test IPRNs is not always an indication of the *dispersion* of abused number ranges in that country. For instance, although Latvia has the largest quantity of test IPRNs, these numbers belong to a smaller set of number ranges, especially when compared to the countries like Iraq, Cuba or Guinea.

Next, we take a more detailed look into the dispersion of IPRNs in the phone number space of Latvia and Cuba. (We chose these countries as they both have 8-digit phone numbers and number spaces are easier to visualize.) Figures 10a and 10b shows the IPRNs (last digit ignored) together with the number types. In both countries, mobile number ranges are the most frequently abused. However, looking at the rest of the figures, the number ranges abused in Cuba are much more dispersed compared to Latvia.

We also observe that there is an accumulation of IPRNs in the beginning of the 4-digit number ranges. Figures 10c and 10d better demonstrate this accumulation in both countries. We believe that, when a 4-digit number range is hijacked, the initial numbers are advertised as test IPRNs, and the rest of the range is assigned to fraudsters for the actual IRSF call generation.

Rate of operators being involved in IPRNs. For the mobile test IPRNs, our dataset often contains the name of the mobile network and the operator. Moreover, from the numbering plan database, we can obtain the complete list of mobile network operators in a country. By combining these two sources, for each country, we can compute the ratio of mobile networks whose number ranges were involved in test IPRNs. Note that, the operators may or may not be a part of the fraud scheme, but for simplicity, we will call them the victim operators.

We find that, in 66.3% of countries, all mobile networks are victims of number range hijacking. In 86.5% of the countries, at least half of the mobile networks are hijacked. The average number of victim operators per country is 3.47. This observation is also in line with the data we present in Table VI, where we group the countries according to the total number of mobile networks per country.

TABLE VI: Ratio of affected operators, grouped by the number of mobile network operators per country.

	Number of such countries in dataset	Ratio of affected operators (avg.)
# Operators <3	98	96.9%
3 <= # Operators <6	62	85.4%
# Operators >=6	44	50.6%

In conclusion, our analysis shows that there is a large variety of number ranges that can potentially be used for IRSF: While almost all countries in the world are affected, IPRNs can belong to any number type and a large variety of operators.

D. Analyzing the test call logs

The 206,203 test call logs we collected contain calls from 248 origination countries to 199 destination countries. In particular, we observe 7082 distinct source-destination country pairs, which indicates that the test calls are widely disseminated.

Among these calls, 50,926 have invalid or anonymized caller IDs. The rest of the calls include 68,740 distinct origination numbers and 52,171 distinct destination numbers (i.e., test IPRNs). Table VII presents types of origination and destination numbers. A large number of test calls seem to be originated from mobile numbers, which possibly belong to stolen or abused SIM cards. Note that, we cannot completely trust the caller ID, as it can be spoofed by the fraudster (e.g., from a PBX) or during routing by an operator. On the other hand,

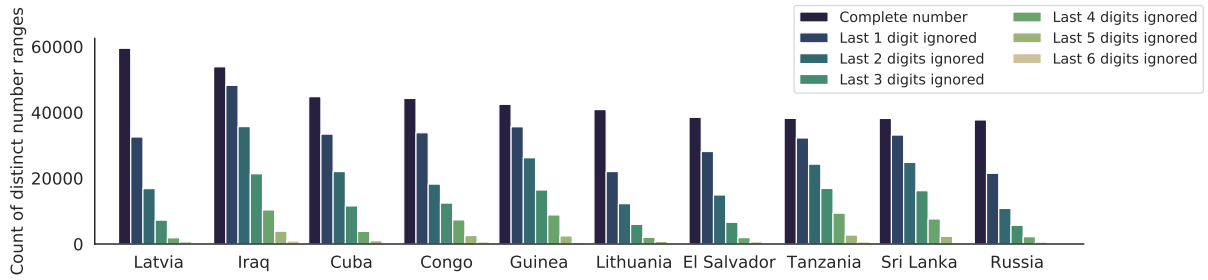

Fig. 9: Top 10 countries having test IPRNs advertised.

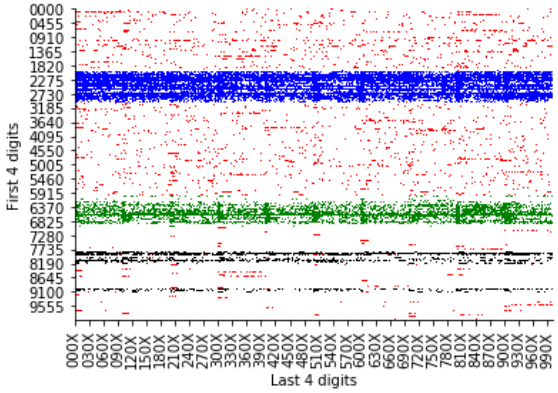

(a) Dispersion of IPRNs on the number range of Latvia.

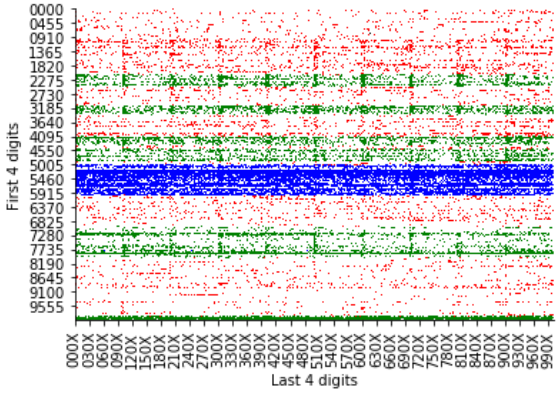

(b) Dispersion of IPRNs on the number range of Cuba.

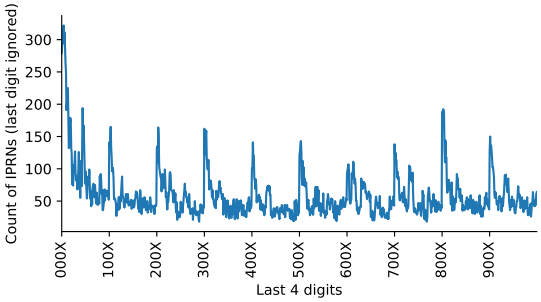

(c) Last digits of Latvia.

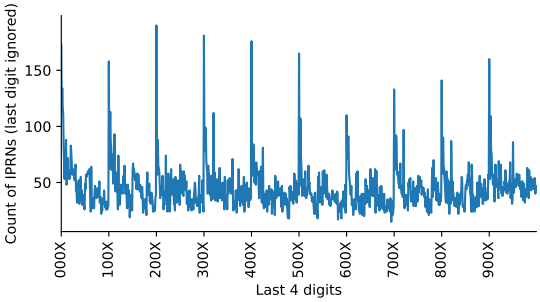

(d) Last digits of Cuba.

Fig. 10: Dispersion of IPRNs in Cuba and Latvia (Blue: Mobile, Green: Fixed, Black: Supplementary Services, Red: Unallocated).

TABLE VII: Types of originating and destination phone numbers observed in test calls.

Number Type	Originating Numbers	Destination Numbers
Mobile	62%	71%
Fixed	19%	13%
Supplementary Services	5%	5%
Unallocated number range	14%	8%
Satellite	-	6%

calls that originate from mobile networks are less likely to be spoofed, as the caller ID cannot be easily modified by the caller (unlike the calls from a PBX with a SIP trunk), due to the SIM card being authenticated to the mobile network.

We also observe that many of the test calls were repeated more than once. To analyze the actual number of fraud cases per country, we remove the recurring calls and only consider the unique source and destination phone number pairs. This leaves us with 119,684 call logs. In Figure 11, we present the cumulative distribution of the number of unique test calls by the percentage of originating countries. We can see that, half of the countries have less than 100 unique test calls originated, whereas approximately 10% of countries have more than 1000 unique test calls. In particular, the US, India and Germany were the top 3 most frequent test call originators. Only from the US, there were test calls to 155 different countries. An interesting point is that, our dataset contains few test IPRNs from these countries: The US has only 260, India has 1174, Germany has 2026 test IPRNs advertised. This might show that the fraudsters are likely to manipulate separate

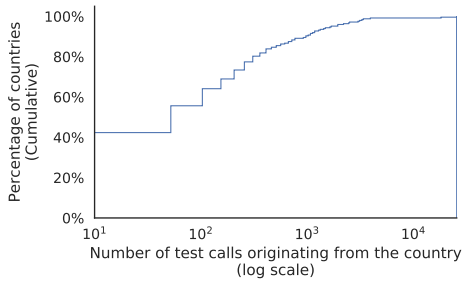

Fig. 11: Empirical cumulative distribution of the number of test calls by the % of originating countries.

sets of countries for originating and terminating the fraudulent calls. This phenomenon can also be observed in geographical distribution of call originations (Figure 7b), compared to the geographical distribution of test IPRNs (Figure 7a). In fact, these two figures seem to be almost the opposite of each other in terms of the distribution ratios. (Some obvious examples are the US, Russia, Bolivia, Syria, Congo, India and China.)

V. IRSF Detection: Existing Techniques

A. Blocking number ranges

A basic approach to prevent IRSF is to block all the calls to frequent IRSF destinations, or number ranges. However, this type of extensive interference may lead to unreachability of genuine users in the destination country, and may result in significant disruption of service and user complaints [59]. As we show in the previous section, the abused phone numbers may belong to a large variety of countries and operators, which makes it unrealistic to block all the suspicious number ranges.

B. Crowdsourced Blocklists and Hotlists

Another approach is to use list of previously identified destination numbers to block calls. Organizations like GSMA and CFCA share such lists with their members [11], [68]. There are also other industry initiatives (such as [54]) to collect and share the phone numbers identified to be used in Wangiri scams. However, our analysis shows that the phone number space that can be abused for IRSF is quite large: Fraudsters can easily circulate or renew the IPRNs, which will make such lists outdated. Moreover, it is difficult to know the efficiency of this approach, as a number range hijack that works from one country, may not work from another (similar to what we observed in Section III with test calls from Germany and Italy). At this point, we also note that the hijack of a number range towards a destination can be detected by the commercial Test Call Generation (TCG) platforms[4] that provide call initiation points worldwide. However, this process can get costly if the hijack occurs and the phone calls are answered and billed in transit (as we also experienced in our experiment in Section III). Currently, to the best of our knowledge, there is no commercial TCG platform that provide such a service to detect number range hijacks.

[4]e.g., sigos.com, araxxe.com, revector.com

C. Monitoring for test IPRNs for early detection

Leveraging the test IPRNs as an early detection mechanism has been proposed by a few commercial services [10], [13] that provide databases of test IPRNs (similar to the data we have collected). As the fraudsters are likely to make few test calls before starting the actual fraud, monitoring for the test IPRNs or close number ranges may indeed be an efficient approach. However, the efficiency of this approach relies on the test IPRN database being up-to-date: When the fraudsters start to abuse a different number range, or a new IPRN provider enters into the market, the test IPRN database should be immediately updated. Another disadvantage of this approach is that, monitoring for large number ranges may result in a high number of false positives. We will refer to this as the *naive approach* in the next section.

D. Call record analytics and anomaly detection

Applying analytics and anomaly detection techniques based on historical call data records is another method to detect IRSF. Commercial fraud management systems such as [31], [62] claim to integrate call analytics and machine learning techniques together with hotlists from different sources, and achieve near real time detection in different types of IRSF calls. However, the inner details and the accuracy rates of such solutions are not available to us.

There are also a few academic studies that use clustering and other anomaly detection techniques to identify IRSF calls in different contexts. However, these studies report low accuracy and high false positive rates. In particular, [38] focuses on mobile networks and proposes to generate an international outgoing call graph (from domestic users to foreign destinations), and use a Markov Clustering based algorithm to the isolate fraudulent activities. Authors obtain a 6-months call dataset from a large mobile operator, and use the customer complaints (i.e., fraud cases that was reported to the operator) and other online complaints (from public forums or social media) as a ground truth on fraudulent destination numbers. Overall, their clustering method was able to detect the destination numbers that correspond to 78% of the IRSF calls in the dataset. On the other hand, this method seems to yield a high ratio of false positives: Out of the 24K *candidate* phone numbers that they found suspicious, only 9.3% was actually associated with IRSF or other fraud activities in the ground truth dataset.

In a more recent work, Meijaard et al. [42] analyze the use of Isolation Forest algorithm to identify IRSF calls as anomalies. To be able to detect IRSF before the fraudulent calls are complete, they use only pre-call features such as source country, destination country, anonymized representations of the source and destination numbers, timestamp and whether the call was answered, canceled or busy. They avoid post-call features like duration and cost of the call. They use an industrial anti-fraud solution as the ground truth to compute the accuracy of the results. On a dataset of 10K phone calls from 9 users, the industrial solution detects 150 IRSF calls. Unfortunately, isolation forest does not seem to perform well for this task: The paper reports to detect 45% of fraudulent calls with 2% false positive rate, while up to 87% of detection is achieved at the cost of 5% false positives. This rate of false positives is likely to clog the fraud management team with

false alerts, when the real-world volume of phone calls is considered. One problem with this approach (as also discussed in the paper) is that fraudulent calls can be distributed over many countries that also involve legitimate calls. As a result, relying on the country names and phone numbers might misguide the algorithm. Moreover, anomaly detection techniques like Isolation Forest may not perform well, because there can be multiple types of anomalies in the dataset, and the algorithm may isolate the wrong type of anomaly. As also mentioned in [46], the *diverse types of anomalies* and the *class imbalance* between the normal and anomalous instances pose challenges to anomaly detection, especially for unsupervised techniques that *do not have prior knowledge of the true anomalies*. In conclusion, those previous approaches have too high false positive rates, making them impossible to use in practice.

VI. Our approach: Leveraging test IPRNs for machine learning features

In this section, we propose several features that can be used with supervised Machine Learning algorithms to detect potential IRSF calls. Our features aim to leverage the intelligence we gathered from test IPRNs, and also the domain knowledge on how IRSF works. Next we explain each feature in detail. Note that, we often refer to the originating number of a phone call as the *A-Number*, and the destination phone number as the *B-Number* in the rest of the paper.

A. Features related to call destination

Distance to test IPRNs. This feature computes the proximity of the destination number to the closest known test IPRN. For instance, the number can match an exact test IPRN (in which case the distance will be 0), or it can be in the same number range with a few digits of difference.

Dispersion digit of test IPRNs in the destination country. This feature reflects our analysis (Section IV) that the test IPRNs can be dispersed over the phone number space in a country-specific way. For example, if the test IPRNs are widely dispersed over the number range allocation of the country, we would have less confidence in labeling a phone number as suspicious, even if it is very close to a known test IPRN. In other words, this feature aims to provide a *spreading factor* based on how spread the test IPRNs are in the destination country's numbering plan. For this, we simply propose to use the analysis in Figure 9, looking at the number of digits that give the highest variance to the test IPRNs. For instance, the *dispersion digit* for Latvia is 11 (which corresponds to the last digit in international format), as most of the variety of test IPRNs comes from this digit (i.e., ignoring the last digit results in losing almost half of the distinct number ranges). On the other hand, the dispersion digit for Guinea is the 9th digit (out of 12 digits), which shows that the test IPRNs are more widely dispersed in country's number space.

IRSF likelihood of the destination number. This feature combines the two previous features, in a way that will better reflect the IRSF likelihood of the destination number. The idea is that, if the dispersion digit is low, we would have less confidence to label a B-Number as suspicious. On the other hand, the length of the number and the distance to test IPRNs is inversely correlated to the suspiciousness of the number. Combining these factors, we compute a IRSF likelihood score for the B-Number as:
(Dispersion digit / (Number length x Distance to test IPRNs)).

Test IPRN ratio of the destination country. This feature computes the number of distinct test IPRNs that belongs to the destination country, to all of the distinct test IPRNs collected. If this ratio is low, this destination country would be less suspicious.

Test call ratio of the destination country. To compute this feature, we take the test calls with distinct caller ID - IPRN pairs. We then compute the ratio of the number of test calls directed to the destination country (from distinct source numbers), over all of the distinct test calls.

Test call spreadness This feature computes the number of distinct origination countries that called this destination in the test call logs we collected. If a destination country has received test calls from many different originations, it is likely to be a popular IRSF destination (i.e., the hijacking of its number space affects multiple originations).

Number length validity. This is a boolean feature that checks if the destination number has a valid length. We gather this information from the commercial numbering plan database [1].

Number type. We map this categorical feature to numerical values, based on our observations in Section IV-D. The values are correlated to the likelihood of the number type to be used for IRSF: $Fixed \mapsto 1$, $Mobile \mapsto 2$, $Supplementary\ services \mapsto 3$, $Satellite\ numbers \mapsto 4$, $Unallocated\ numbers \mapsto 5$.

B. Features related to historical call records

These features relate to the call history of the originating number (A-Number). Note that, for the transit operators it may not be possible to compute these features, as they do not always carry the calls to/from a specific set of users like the retail operators.

Call frequency metric. This feature first computes the duration between the initial call related to this A-Number and B-Number pair, and the current call. We then divide this value by total number of calls in between. For instance, if the A-Number has called the B-Number 3 times in the last 60 minutes, this value would be 20. If we observe the first call between A and B Numbers, we assign the value of 0.

Time to previous call. This feature is based on the observation that the fraudsters will try to create as many calls as possible in a small amount of time. Thus, this feature computes the number of seconds that has passed since the last outgoing international call originating from this user (i.e., the A-Number).

Hour. We use the hour of the current call as another feature. As the fraudsters might prefer the night time or non-business hours to avoid being detected, this can be an important feature that distinguishes the fraudulent calls.

We highlight the fact that all our features are pre-call features, which means they do not require the call to be answered or completed. Thus, they can be used to detect IRSF

or to identify suspicious calls, before the calls go through. Moreover, most of our features do not directly rely on the test IPRNs or A/B Numbers. By integrating the intelligence gathered from the cumulative data, we aim to detect fraudulent calls even if the destination number range does not exactly match the test IPRNs.

C. Other possible features

Depending on the data available to the telecom operator, additional features can help in detecting IRSF. For instance, the payout rate (pay-back, cash back) advertised by the IPRN providers for this particular destination can relate to the likelihood of fraud. In addition, the difference between the average price of a high quality route and a standard (often low quality) route from the originating country/operator to the destination country/operator can be used as a feature: If the difference is too large (i.e., the price of the low-cost route is 'too cheap'), the route is likely to include fraud. For mobile destination numbers, real time HLR lookups can be used to check if the number is assigned to a real user, and currently in use. Note that, we do not use these additional features in our evaluation section, due to the lack of (timely) data. Depending on the usage of the fraud detection system, post-call features (such as call duration and cost of the call) and behavioral patterns (such as multiple simultaneous calls) can be combined with the pre-call features.

VII. Evaluation on a Real-World Dataset

In this section, we use the features proposed in Section VI-A and VI-B to detect IRSF on a real world Call Data Record (CDR) dataset that we obtained from a small retail telecom operator in Europe.

A. Challenges related to real world data

First, we want to emphasize the difficulty of obtaining real world CDR datasets for research purposes. Indeed, CDR is very sensitive data, which makes it difficult for operators to share CDRs with third parties. Anonymizing the call records would not be a solution, as we need to parse the phone numbers to be able to compute most of our features. To overcome this challenge, we ran all our experiments involving CDRs under NDA, and on a server hosted on the telecom operator's premises.

A second challenge is obtaining the ground truth on fraudulent calls. While the best way to validate our approach would be to experiment on a pre-labeled dataset, telecom operators often does not have such data, unless they employ a very accurate fraud detection solution with manual verification. Thus, we first needed to create a ground truth, as we will explain in the following sections.

The third challenge is to obtain timely data. Unless there is a near real time data sharing framework with the telecom operator, the obtained dataset is likely to be old. In our case, the CDRs were collected between 10/01/2014 and 12/31/2014. Note that, this time period is even before we started to collect the test IPRNs. However, we see this as a positive aspect: In a real world deployment it is also likely that the time period of analyzed phone calls will not exactly overlap with the collection period of test IPRNs. In addition, note that we have at most 3 months of historical data of the users.

B. Details of the dataset

Our CDR dataset contains 3 months of outgoing international call records from the users of this small operator. This corresponds to 689,015 phone calls from 6,289 users (both mobile and fixed lines). Each call record contains the A-Number and B-Numbers in international format, the date and time of the call, duration, and the localization data. The localization corresponds to the country from which the call is initiated. While most of time this is the country of the home operator, it can also be a different country if the user is roaming.

C. Addressing the Machine Learning Challenges

1) Lack of ground truth: Telecom operators carry a huge volume of call traffic to various destinations, for a large number of users. Although they may employ automated fraud detection systems or human analysts, some of the fraud might go unnoticed, or the cost of detecting some fraud might be higher than the cost of fraud. Moreover, fraud detection systems may have large number of false alerts and it may not be feasible to manually investigate all of them. As a result, it is often not possible to obtain a 100% accurate ground truth on what is legitimate and what is fraudulent.

Solution. As our CDR dataset did not have any labels on fraudulent and benign calls, we tried to use the collected test IPRNs as an indication to detect the possible fraud cases. We found calls that hit 12 of the test IPRNs. Upon further manual investigation, we realized that these calls belong to 4 different possible IRSF cases, involving many other calls. Table VIII gives the details about these cases, such as the total duration of the fraud, number of calls, number of distinct destinations, and how many test IPRNs were matched (either exactly, or with 2 or 4 digit distance). We find that in three of the cases, victim numbers were fixed lines, possibly belonging to compromised PBX systems of certain enterprises. In fact, for the *PBX Hack#1* case, the originating number belongs to an outgoing call center that makes large number of international calls. Only in one case (*SIM theft*) we found that the victim was using a mobile number roaming in Spain (i.e., localization was Spain). This case is likely to involve a stolen mobile phone, and the 2 day duration of fraud is possibly because the victim did not cancel his SIM card after the phone was stolen. In fact, the organized crime in Spain is well known to use IRSF with stolen phones [22], [29], [64].

Our fraud findings for the ground truth were later approved by the telecom operator. Thus, we can be sure that the fraud-labeled data is indeed fraud. However, we cannot be sure that the rest of the data does not include any undetected IRSF calls. Yet, we treat all the rest of the calls from these 4 users as benign calls. While this is not a perfect solution, it is realistic: In real world, telecom operators can never have 100% certainty about the legitimacy of each call, and some of the fraud goes unnoticed. Thus, this was the best approach available to us.

Combining the calls from the users involved in those four fraud cases, our labeled dataset contains 3084 fraudulent and 158,703 legitimate calls. Looking at the fraudulent calls, we observe that multiple IPRNs were used (at least 40 distinct numbers) in every fraud case. We also observe that each fraud case involves a different set of IPRNs, meaning that different

fraud cases did not have any destination numbers in common. (Note that, the count of distinct IRSF phone numbers does not correspond to the number of bad actors, as one fraudster can easily obtain multiple numbers, or a single IRSF number can be recycled and assigned to multiple fraudsters in time.) Moreover, multiple destination countries were involved in each fraud case. The average duration of fraudulent calls is 10 minutes (602 ± 445 seconds), probably because the fraudsters try to stay under the radar of fraud detection systems by initiating short duration calls to several different numbers.

2) Imbalanced dataset: Imbalanced datasets is a common problem when applying machine learning for fraud detection [20], [46]. For instance, considering the four fraud cases we detected in our dataset with 689K calls, the ratio of fraudulent calls to all calls is 0.0046. Because there are too few fraud calls compared to the benign data, unsupervised learning methods such as clustering and isolation forest often result in a large number of false positives [24], [42]. On the other hand, supervised machine learning algorithms, in particular the Random Forest classifier is found to achieve the best results for fraud detection in previous work [24], [41]. In addition, the classification problem in such extremely imbalanced datasets is addressed in two ways in the literature: Assigning a higher cost to the misclassification of the minority class, or using a sampling technique [23], [26], [27].

Solution. Following the similar previous work that aims at fraud detection [24], [41], we chose to use the Random Forest algorithm and to down-sample the majority class (i.e., the legitimate calls) by randomly selecting the quantity equal to the fraudulent calls. Thus, our final dataset contains 3084 fraudulent calls and 3084 legitimate calls. To avoid the bias of the sampling procedure, we have run our experiment 10 times (with different re-sampling of the benign calls) and averaged the results.

D. Evaluation

For the experiment, we used the *sklearn* [47] implementation of the Random Forest algorithm [17][5]. To evaluate the accuracy, we first trained and tested the model on the balanced dataset with 10-fold cross validation, re-sampling the benign calls 10 times and taking the average.

In addition, we compute the accuracy of the *naive approach*, i.e., using the number ranges of known test IPRNs for detection. In particular, we identified the calls that would be labeled as fraudulent, if we were monitoring for the number ranges of test IPRNs (e.g., ignoring the last 2 digits or the last 4 digits). Table IX shows the True Positive Rate (TPR), False Positive Rate (FPR) and accuracy of all methods in comparison.

Compared to the naive approach of monitoring -2 digit number ranges, our method achieves 33% higher TPR, with only a slight increase in FPR. We can also see that, the naive approach with -4 digits may significantly increase the FPR, even though it improves detection rate compared to -2 digits. Overall, our method achieves 98% accuracy, which is much higher than both of the naive approaches. Note that, although the naive approach can be useful to detect IRSF cases in historical CDRs, our experiment shows that it may not be reliable for automated and real time detection of fraud, with high accuracy.

To gain a better insight into the Random Forest model, we demonstrate the feature importances (based on the impurity-based importance provided in sklearn [17]) when trained on a balanced dataset, in Figure 12. Although this measurement has some disadvantages such as favoring the features with high cardinality [8], [16], overall we can see that the features related to the destination country have the highest importance, followed by the features related to the destination number, and the features related to the call history at last.

Evaluation on previously unobserved data. As our dataset contains only 4 IRSF cases with a limited number of fraudulent calls, there is a possibility of over-fitting in the classifier, even if we apply 10-fold cross validation. Thus, we made a further evaluation by training the Random Forest classifier with three of the fraud cases, and testing it on the fourth, previously unseen case.

When we train the model on *PBX Hack #1*, *PBX Hack #2* and *SIM theft* cases, and test it on the *PBX Hack #3* case; we achieve 100% TPR and 0% FPR. This means, all of the 110 fraudulent calls in *PBX Hack #3* were correctly identified, without any false positives. Of course the performance on a larger dataset with a higher variety of calls might not be as good, but unfortunately we do not have access to such a large dataset.

When we train the model on *PBX Hack #1*, *PBX Hack #2* and *PBX Hack #3 cases*, and test it on the *SIM theft* case; we achieve 57% TPR and 0% FPR. In this case, only 308 of the 539 fraudulent calls are detected, but there are still no false positives. Note that, the lower TPR is possibly due to the fact that the fraud patterns of PBX hacking is different from SIM theft. This result shows the importance of the variety of fraud cases in the training data.

Finally, in Figure 13 we give an overview of our approach, demonstrating the requirements (such as the test IPRN database, numbering plan database, and historical call records) and the feedback loop to keep the classifier up-to-date.

E. Limitations

This first experiment demonstrates how the test IPRNs combined with machine learning approaches can help in IRSF detection. Note that, our dataset only contains examples of IRSF conducted via compromised PBX and stolen SIM cards. Thus, some of the CDR related parameters such as *call frequency* and *time to previous call* is designed to detect these fraud types. On the other hand, we did not have any example of Wangiri fraud. In fact, Wangiri fraud is different in nature: It involves a fraudster initiating calls to a large population and expecting the calls to be returned (rather than compromising a phone system and directly initiating the fraudulent calls). In terms of machine learning approach, additional features can be computed to detect Wangiri fraud such as the number of distinct users contacted by a certain A-Number. However, to compute this feature, the operator would need to access all its call records and maybe multiple operators would need to exchange data. We believe that our *IRSF likelihood* feature and

[5]Modified parameters: n_estimators=200, class_weight='balanced'.

TABLE VIII: Summary of the IRSF cases identified in the CDR dataset.

Fraud Case	Fraud Duration	# Calls	Total Call Dur. (min)	# Distinct Dest. Numbers	# Distinct Dest. Countries	Match with test IPRNs			Naive Detection Rate	
						Exact	Upto 2	Upto 4	2 digit	4 digit
PBX Hack #1	9h 59m 43s	2,263	22,152	59	11	1	39	52	81%	82%
PBX Hack #2	22h 31m 42s	172	1,670	44	8	2	11	16	19%	35%
PBX Hack #3	5h 7m 59s	110	563	91	2	1	32	90	4%	99%
SIM theft	2d 11h 24m 15s	539	6,577	40	16	8	24	35	8%	95%

TABLE IX: Evaluation of our method in comparison to the naive approaches.

	TPR	FPR	Accuracy
Naive method: Last 2 digits ignored	63.2%	0.12%	81.5%
Naive method: Last 4 digits ignored	82.4%	5.82%	88.3%
Our method (w. 10-fold cross valid.)	96.4%	0.26%	98.1%

other features related to the destination country may still be useful to assign a risk score for Wangiri fraud.

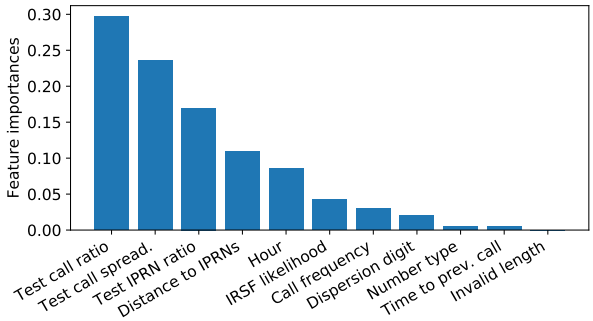

Fig. 12: Feature importances of the random forest algorithm.

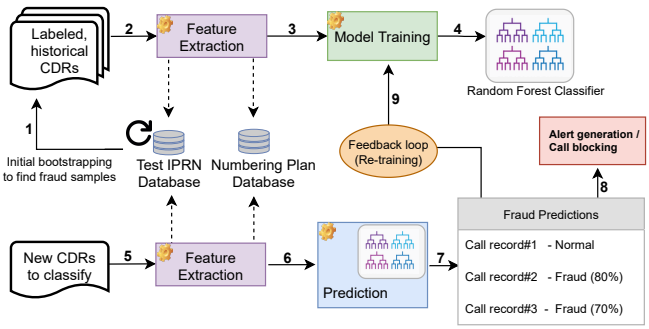

Fig. 13: Overview of the proposed machine learning approach.

Finally, fraudsters might try to avoid detection by changing their call patterns: Instead of generating large volume of calls in a few hours, they can spread out the calls over a longer time period, especially in the case of compromised PBX. This could make our CDR-related features less useful and make the calls more difficult to detect. On the other hand, it would mean less profit for the fraudsters, as they would need to slow down their operations instead of making the most profit in the shortest amount of time.

VIII. Related Work

In recent years, telecommunication fraud attracted a lot of attention in the research community. As the topic is very broad, covering many different types of fraud, some studies try to systematize the information. In particular, Sahin et al. [57] presents a taxonomy of telephony fraud, including different types of revenue share fraud, abuses of call routing and signaling, and fraud schemes that rely on bypassing call routes in interconnection. Moreover, Tu et al. [63] explores the voice spam ecosystem, describing the key challenges in fighting it, and evaluating the existing detection techniques. In addition, several studies focus on certain types of interconnect bypass fraud. For instance, Reaves et al. [52] presents a technique to detect illegitimate VOIP-GSM gateways in the international call routes, by analyzing the quality degradations in call audio. Another study [56] analyzes the impact of OTT bypass fraud, where the fraudulent transit operators hijack international phone calls and route them through Over-The-Top service providers, bypassing the destination operator.

To the best of our knowledge, there is no study that explores the IRSF ecosystem, in particular the test numbers and test call logs available on IPRN providers' websites. In a whitepaper from TransNexus, authors contact 121 premium rate number providers and gather information on the advertised number ranges (countries), payout rates and location of the company [9]. However, they do not look into the test IPRNs, but only analyze the most frequently advertised countries and compare the payout rates. Thus, our study is the first to shed light on the ecosystem of test IPRNs.

Data mining and machine learning techniques have been frequently used for fraud detection in different domains, including telecommunications [40], [49]. In fact, telecommunications was one of the first industries that adopted machine learning technologies due to the huge amount of high-quality data they store [67]. Most of the academic work in this field focus on applying machine learning on certain behavior patterns extracted from CDRs to detect: subscription fraud [30], [32], account takeover [33], [34], [55], simbox fraud [28], [43] and voice spam [41]. As each fraud type exhibit different call patterns, features used in these previous work varies. As we mentioned in Section V-D, [42] attempts to detect IRSF with anomaly detection, however this method suffers from a high rate of false positives.

IX. Conclusion

This paper studies the long-standing, yet unsolved problem of International Revenue Share Fraud (IRSF). IRSF can take various forms depending on the fraud agreement, collaborating parties, and the traffic generation methods employed by fraudsters. The complexity of IRSF increases even further with the third party providers operating online. In this paper, we first

analyzed the data we collected from such online providers to understand how they operate and how they abuse international phone numbers. By taking advantage of this information, we then proposed a set of machine learning features that can be used to detect IRSF calls. We finally validated this approach on a real-world call dataset. We believe that continuously monitoring the online IPRN providers would be very useful in fighting IRSF proactively.

REFERENCES

[1] "National Numbering Plans Collection," https://numplans.com/.

[2] "The international public telecommunications numbering plan," ITU-T Recommendation E.164, 1997.

[3] "Application of Recommendation E.164 numbering plan for universal international premium rate numbers for the international premium rate service ," ITU-T Recommendation E.169.2, 2000.

[4] "Consumer abuses and fraud issues relating to high tariff services," Electronic Communications Committee (ECC) report 86, 2006.

[5] "International revenue share fraud: Are we winning the battle against telecom pirates?" Black Swan Telecom Journal, November 2012.

[6] "Phone-hacking scam costs uk businesses estimated 1bn," http://www.channel4.com/news/phone-hacking-scam-costs-uk-businesses-estimated-1bn, May 2012.

[7] "You will be billed $90,000 for this call 3: F-secure discloses mobile app virus attacks," Privacy-PC.com news, March 2012.

[8] "Selecting good features Part III: random forests," https://blog.datadive.net/selecting-good-features-part-iii-random-forests/, 2014.

[9] "International premium rate number market," https://transnexus.com/whitepapers/international-premium-rate-number-market/, 2015.

[10] "Fraudstrike data sheet," http://www.fraudstrike.com/New_FraudStrike_Data_Sheet_v3_1_.pdf, 2016.

[11] "CFCA International Revenue Fraud Number Database," www.cfca.org, 2017.

[12] "Communications Fraud Control Association (cfca), 2017 global fraud loss surveys," 2017. [Online]. Available: http://v2.itweb.co.za/whitepaper/Amdocs_LINKED_2017_CFCA_Global_Fraud_Loss_Survey.pdf

[13] "International and fraud databases," http://www.ccmi.com/international-tariff-databases, 2017.

[14] "Are You Getting Unexpected Overseas Calls? Here's Why," https://miami.cbslocal.com/2019/04/08/are-you-getting-unexpected-overseas-calls-heres-why/, April 2019.

[15] "Communications Fraud Control Association announces results of 2019 global telecom fraud survey," 2019. [Online]. Available: https://cfca.org/sites/default/files/Fraud%20Loss%20Survey_2019_Press%20Release.pdf

[16] "Permutation Importance vs Random Forest Feature Importance (MDI)," https://scikit-learn.org/, 2019.

[17] "Random Forest Classifier," https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html, 2019.

[18] "Least Cost Routing (LCR)," https://www.bandwidth.com/glossary/least-cost-routing-lcr/, July 2020.

[19] "The ucsd network telescope," https://www.caida.org/projects/network_telescope/, 2020.

[20] I. Amit, J. Matherly, W. Hewlett, Z. Xu, Y. Meshi, and Y. Weinberger, "Machine learning in cyber-security - problems, challenges and data sets," *CoRR*, vol. abs/1812.07858, 2018.

[21] A. Apvrille, "WinCE/Terdial or Impunity for Dialers," Fortinet blog, May 2010.

[22] M. Brignall, "Victory against Vodafone for schoolteacher billed 15,000," https://english.elpais.com/elpais/2015/10/08/inenglish/1444315248_950843.html, 2014.

[23] C. Chen and L. Breiman, "Using random forest to learn imbalanced data," *University of California, Berkeley*, 01 2004.

[24] A. Cidon, L. Gavish, I. Bleier, N. Korshun, M. Schweighauser, and A. Tsitkin, "High precision detection of business email compromise," in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1291–1307. [Online]. Available: https://www.usenix.org/conference/usenixsecurity19/presentation/cidon

[25] S. Davies, "Numbering misuse and fraud," ITU Seminar, May 2012.

[26] P. Domingos, "Metacost: A general method for making classifiers cost-sensitive," in *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD 99. New York, NY, USA: Association for Computing Machinery, 1999, p. 155164.

[27] C. Drummond and R. Holte, "C4.5, class imbalance, and cost sensitivity: Why under-sampling beats oversampling," *Proceedings of the ICML'03 Workshop on Learning from Imbalanced Datasets*, 01 2003.

[28] A. Elmi, S. Ibrahim, and R. Sallehuddin, "Detecting sim box fraud using neural network," in *IT Convergence and Security 2012*, K. J. Kim and K.-Y. Chung, Eds.

[29] Europol press release, "Europol supports Spanish police to dismantle serious cybercriminal group," https://www.europol.europa.eu/newsroom/news/europol-supports-spanish-police-to-dismantle-serious-cybercriminal-group, July 2015.

[30] H. Farvaresh and M. M. Sepehri, "A data mining framework for detecting subscription fraud in telecommunication," *Eng. Appl. Artif. Intell.*, vol. 24, no. 1, pp. 182–194, Feb. 2011.

[31] S. GmbH, "IRSF Detection," https://www.sigos.com/fraud-management/sigos-analytics/irsf-detection/, 2020.

[32] W. Henecka and M. Roughan, "Privacy-preserving fraud detection across multiple phone record databases," *IEEE Transactions on Dependable and Secure Computing*, 2015.

[33] C. S. Hilas, "Designing an expert system for fraud detection in private telecommunications networks," *Expert Syst. Appl.*, vol. 36, no. 9, pp. 11 559–11 569, Nov. 2009.

[34] C. S. Hilas and J. N. Sahalos, "User profiling for fraud detection in telecommunication networks," in *5th Int. Conf. technology and automation*, 2005, pp. 382–387.

[35] J. Huang, X. Zhang, L. Tan, P. Wang, and B. Liang, "Asdroid: Detecting stealthy behaviors in android applications by user interface and program behavior contradiction," in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: Association for Computing Machinery, 2014, p. 10361046.

[36] International Telecommunication Union, "Supplement 1: Best practice guide on countering misuse of e.164 number resources," 2007.

[37] ——, "Uiprn," https://www.itu.int/en/ITU-T/inr/unum/Pages/uiprn.aspx, 2017.

[38] N. Jiang, Y. Jin, A. Skudlark, W.-L. Hsu, G. Jacobson, S. Prakasam, and Z.-L. Zhang, "Isolating and analyzing fraud activities in a large cellular network via voice call graph analysis," in *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys 12. New York, NY, USA: Association for Computing Machinery, 2012, p. 253266. [Online]. Available: https://doi.org/10.1145/2307636.2307660

[39] P. Jourdan, "Common Mistakes Which Lead to Call Fraud ," https://www.3cx.com/blog/voip-howto/call-fraud/, February 2019.

[40] Y. Kou, C.-T. Lu, S. Sirwongwattana, and Y.-P. Huang, "Survey of fraud detection techniques," in *IEEE International Conference on Networking, Sensing and Control, 2004*, vol. 2, 2004, pp. 749–754 Vol.2.

[41] H. Li, X. Xu, C. Liu, T. Ren, K. Wu, X. Cao, W. Zhang, Y. Yu, and D. Song, "A machine learning approach to prevent malicious calls over telephony networks," in *2018 IEEE Symposium on Security and Privacy (SP)*, 2018, pp. 53–69.

[42] Y. J. Meijaard, B. C. M. Cappers, J. G. M. Mengerink, and N. Zannone, "Predictive analytics to prevent voice over ip international revenue sharing fraud," in *Data and Applications Security and Privacy XXXIV*, A. Singhal and J. Vaidya, Eds. Cham: Springer International Publishing, 2020, pp. 241–260.

[43] I. Murynets, M. Zabarankin, R. P. Jover, and A. Panagia, "Analysis and detection of simbox fraud in mobility networks," in *INFOCOM*, 2014.

[44] S. Networks, "PBX Hacking: How it Works ," https://tollshield.com/news/pbx-hacking-how-it-works, September 2016.

[45] E. N. Nielsen and France Telecom-Orange, "Draft report on the article 28(2) universal service directive: A harmonised berec cooperation process," BEREC, November 2012.

[46] G. Pang, C. Shen, L. Cao, and A. van den Hengel, "Deep learning for anomaly detection: A review," 2020.

[47] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in python," *J. Mach. Learn. Res.*, vol. 12, no. null, p. 28252830, Nov. 2011.

[48] N. Perlroth, "Phone Hackers Dial and Redial to Steal Billions," https://www.nytimes.com/2014/10/20/technology/dial-and-redial-phone-hackers-stealing-billions-.html, October 2014.

[49] C. Phua, V. C.-S. Lee, K. Smith-Miles, and R. W. Gayler, "A comprehensive survey of data mining-based fraud detection research," *CoRR*, vol. abs/1009.6119, 2010.

[50] E. PRIEZKALNS, "Massive Wangiri Attack Hits Angola," Commsrisk Blog, July 2020.

[51] ——, "Wangiri Warnings This Year: Belgium, Costa Rica, Finland, Ireland, Kenya, Lithuania and Poland," Commsrisk Blog, May 2020.

[52] B. Reaves, E. Shernan, A. Bates, H. Carter, and P. Traynor, "Boxed out: Blocking cellular interconnect bypass fraud at the network edge," in *USENIX Security*, 2015.

[53] E. REED-SANCHEZ, "How to Start Your Own Cellular Network: Micro Telcos and the Future of Communication," 2020. [Online]. Available: https://edwin.atavist.com/microtelco

[54] Risk & Assurance Group, "RAG Wangiri Blockchain Consortium," https://riskandassurancegroup.org/rag-wangiri-blockchain-ledger/, 2020.

[55] S. Rosset, U. Murad, E. Neumann, Y. Idan, and G. Pinkas, "Discovery of fraud rules for telecommunications- challenges and solutions," in *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '99. New York, NY, USA: ACM, 1999, pp. 409–413.

[56] M. Sahin and A. Francillon, "Over-the-top bypass: Study of a recent telephony fraud," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16. Association for Computing Machinery, 2016, p. 11061117.

[57] M. Sahin, A. Francillon, P. Gupta, and M. Ahamad, "Sok: Fraud in telephony networks," in *Proceedings of the 2nd IEEE European Symposium on Security and Privacy (EuroS&P'17)*, ser. EuroS&P'17. IEEE, April 2017.

[58] A. Sakounthong, "Wangiri (One Ring) Scam Grows 98% in Q1," https://hiya.com/blog/2019/05/03/wangiri-one-ring-scam-grows-98-in-q1/, May 2019.

[59] M. Sanford, "A coordinated approach on the international telecommunications scene between governments from the pacific and a non-governmental organisation," *Information and Communication Technologies in the South Pacific: Selected Issues*, vol. 9, 2009, available at: https://www.wgtn.ac.nz/law/research/publications/about-nzacl/publications/special-issues/hors-serie-volume-ix,-2009/sanford.pdf.

[60] B. Shrestha, M. Mohamed, A. Borg, N. Saxena, and S. Tamrakar, "Curbing mobile malware based on user-transparent hand movements," in *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2015, pp. 221–229.

[61] I. Solomon, O. Hamama, and O. Ventura, "INJ3CTOR3 Operation Leveraging Asterisk Servers for Monetization," https://research.checkpoint.com/2020/inj3ctor3-operation-leveraging-asterisk-servers-for-monetization/, November 2020.

[62] SUBEX, "IRSF Fraud," https://www.subex.com/fraud-management/irsf-fraud/, 2020.

[63] H. Tu, A. Doup, Z. Zhao, and G. Ahn, "Sok: Everyone hates robocalls: A survey of techniques against telephone spam," in *2016 IEEE Symposium on Security and Privacy (SP)*, 2016, pp. 320–338.

[64] D. VERDU, "Why Barcelona is Spain's stolen cellphone capital," https://english.elpais.com/elpais/2015/10/08/inenglish/1444315248_950843.html, 2015.

[65] L. Watson, "Five men jailed for 4.5m worldwide premium phone number scam," http://www.dailymail.co.uk/news/article-2048273/iPhone-fraud-gang-tricked-O2-paying-4-5m-premium-rates-scam.html, October 2011.

[66] A. Weckler, "'Wangiri' phone scam sweeping across Ireland is 'unprecedented' say operators," https://www.independent.ie/business/technology/wangiri-phone-scam-sweeping-across-ireland-is-unprecedented-say-operators-36240323.html, October 2017.

[67] G. M. Weiss, *Data Mining in Telecommunications*. Boston, MA: Springer US, 2005, pp. 1189–1201.

[68] T. Wilson, "Network Fraud - Bypass/Premium Rate Number - IRS," https://www.slideshare.net/teralight/network-fraud-bypass-premium-rate-number-irs.

X. Appendix

TABLE X: Top 10 countries for the number of newly advertised IPRNs, per semester, in 2018 and 2019.

2018 1st Half	2018 2nd Half	2019 1st Half	2019 2nd Half
Latvia	Iraq	Iraq	Congo
Russia	Latvia	Latvia	Iraq
Cuba	Guinea	Guinea	Latvia
Lithuania	Sri Lanka	Congo	Bolivia
Somalia	Somalia	El Salvador	Honduras
Guinea	Saudi Arabia	Tanzania	Nigeria
Belarus	Algeria	Nigeria	Algeria
Algeria	Lithuania	Bolivia	El Salvador
Madagascar	Tanzania	Lithuania	Italy
Tunisia	Uganda	Sri Lanka	Cuba