

FTC Robocall Challenge Submission: Crowd-Sourced Call Identification and Suppression

Daniel V. Klein and Dean K. Jackson, Google¹

1. Overview

We recommend the creation of a system that allows users to report, to an online database system, the originating telephone number of unwanted solicitations, advertisements or robotically placed calls (henceforth called “spammers”). We also recommend that users’ telephones or external hardware may automatically query the database about the telephone number of an incoming call (before the call is answered, or *even before the telephone rings*) to determine if the caller has been flagged as a spammer by other users, and optionally block the call or otherwise handle it differently from a non-spam call.²

The recommended system thereby would provide a means whereby users can make reports of spam calls as well as ask if others have reported a caller as a spammer. While the first few people called would get spammed, after a sufficient number of reports are made, further calls would be blocked.

The recommended system would work on most types of telephonic platforms – smartphones, some feature phones, POTS³ lines, VoIP, PBX, and telephony providers – through the use of software and optional in-line hardware. In addition to crowd-sourced blacklisting, we also recommend a means to whitelist specific numbers so that, for example, emergency calls will always go through.

2. Background

The system would be comprised of globally accessible databases, possibly replicated on multiple servers worldwide, possibly locally cached, combined with call interceptors (combinations of hardware and/or software) for POTS lines, cell phones (including smartphones and some feature phones), PBXs and VoIP phones. The following overview describes the relevant existing technologies for establishing this system.

2.1 Introduction

Telephone calls that originate in the United States (or that are routed into this country from overseas) come with an identification code commonly known as Caller-ID. This information contains the telephone number (and optionally the name) of the caller, although Caller-ID may be

¹ This proposal does not reflect any internal Google anti-spam technology that is currently in place, other than known industry approaches such as DKIM.

² Because even if users cannot define spam, they can, as Supreme Court Justice Potter Stewart did, say “I know it when I see it”. http://en.wikipedia.org/wiki/I_know_it_when_I_see_it.

³ Plain Old Telephone Service – a wired non-PBX telephone instrument.

(in general) actively blocked by the caller using the vertical service code⁴ *67 (in which case the Caller-ID is usually shown as “Private”) and it may (rarely) be genuinely unknown to the telephone system. The actual Caller-ID information that is displayed depends on the type of equipment that originates the call.⁵

Smartphones (and some feature phones) provide both telephony and IP-based services. As of June 2012, there were more cell phones in service in the United States than the population of the country⁶, as of February 2012 more than half of these phones were smartphones, and with current trends will exceed 70% of all cell phones in 2013.⁷ Wired telephones (both PBX and POTS lines) also generally exist in concert with internet connectivity. As of June 2012, over 78% of the population of the United States were also Internet users⁸ and of those, over 90% had broadband.⁹ The combination of telephony and the Internet provide a means for cooperatively blocking robotically generated telephone phone calls and other forms of unwanted solicitation.

2.2 History and Related Systems

Cooperative crowd-sourced spam identification services for email filtering are well known (e.g., DCC¹⁰ or “Distributed Checksum Clearinghouse”). In this class of system, email “signatures” (sophisticated checksums of the filtered contents of the email) are collected and shared amongst participants. The more times an email “signature” is seen by the collected users of the system, the more times it has been sent, and thus the more likely it is to be spam. There are trusted ways in which legitimate, widely distributed emails such as for mailing lists can be whitelisted, but in general when an email is sent to thousands of users, it is probably spam. While spammers try to circumvent these features by subtly changing the contents of their messages for each recipient, the crowd-sourced signature collection systems attempt to generalize the signature analysis to avoid being fooled by the spammers.

If only a few users participate in crowd-sourced identification, it is easy for the spammers to win, since the sample-space for spam email checksums is small. However, when many thousands of users participate, then the spammers are only able to deliver their messages to the first few dozen recipients. After that happens, the crowd-sourced identification marks the email signatures(s) as “spammy”, and further attempts to send the message to other recipients will fail.

The DCC system is highly effective, and is designed so that each participant computes their own checksums, and reports them to the crowd-sourced servers. The servers rate-limit the signatures that can be reported to them, and also rate-limit the queries any one user may make. In this way, a user (or even a collection of users) cannot bombard the servers with false

⁴ http://en.wikipedia.org/wiki/Vertical_service_code contains a summary of currently defined codes.

⁵ http://en.wikipedia.org/wiki/Caller_ID

⁶ http://en.wikipedia.org/wiki/List_of_countries_by_number_of_mobile_phones_in_use.

⁷ <http://goo.gl/yUx1b> and <http://goo.gl/evSdo>.

⁸ <http://goo.gl/5T4C5> and <http://www.internetworldstats.com/america.htm>.

⁹ <http://gigaom.com/2012/09/04/90-of-us-households-with-computers-have-broadband/>

¹⁰ <http://www.rhyolite.com/dcc/>.

signatures, nor can they launch a denial of service attack.

3. Robocall and Spammer Identification

We recommend a similar crowd-sourced service based on Caller-ID (both phone number and Caller-ID text string) and/or ANI¹¹, where telephone subscribers may either ask the server (or local cache) “is this telephone number marked as a spammer?” Alternatively, subscribers may block and thereby report a telephone number as a spammer whenever they receiving a call from such a number. When a telephone number and, if available, the Caller-ID text string have been reported more than a given threshold number of times, that caller information would be marked as being associated with a spammer, and users’ telephone devices would have the option of automatically dropping the call (possibly before the phone is made to ring). Because both the telephone number and Caller-ID text would be logged as spam, spammers who merely pick new spoofed Caller-ID numbers will find it more difficult to evade the system over time, because Caller-ID text strings typically associated with spamming calls would also trigger the blocking mechanism.

Asking the question “is this telephone number marked as a spammer?” would be answered as quickly as possible, so that pending telephone calls can be routed quickly and efficiently.¹² However, providing a report to the system that “this number is a spammer” would require a measure of care (possibly with multiple steps in the process) so that it would not be easy to fraudulently report a number as a spammer, yet it must be sufficiently quick and easy that users would be encouraged to use the system. As with systems like DCC, the first few spam calls from any one source would go through to the user. However, as more calls are reported as spam, the system would be able to rapidly block the spam calls.¹³

In the ideal implementation, the only numbers which could be reported as “spam” are those associated with a call that has just been received by the user. In practice, the people who place robocalls would attempt to subvert the system by flooding the database with fraudulent reports. We will address security issues in a later section, below.

¹¹ http://en.wikipedia.org/wiki/Automatic_number_identification.

¹² The system is designed to be “fail safe”. In the event of lack of internet connectivity or a timeout when contacting the call-identification servers, the participating telephones behave as before. The users could also be provided with the option of downloading the current spammer database to their phones or call interceptor devices for quicker local access (although using a cached copy would trade greater algorithmic analysis and accuracy of centralized servers for access speed). A Bloom filter or blacklist-only database could also be used to reduce the size of the data needing to be downloaded.

¹³ In order to accurately determine which numbers are associated with spammers, the system could take into account much more than simply the raw number of spam reports on a specific number. An adaptive algorithm could be developed based on call histories, and refined as time goes on (since the effectiveness of any implemented system will be dependent on the number of participants and the number of calls processed). The basis for this algorithm is outlined in a later section.

3.1 Cell phone implementation

Presently, when a call is received by a cell phone, the phone rings and the user has an opportunity to answer or ignore the call. This decision is usually based in part on the identity of the caller.¹⁴ Once a call has been accepted, the user may terminate the call at any time.

For smart cell phones, we recommend an addition to the existing telephone user-interface (UI) which the user may opt in to. For example, the UI could be included as a part of the default-from-the-manufacturer telephone functionality, or it could be downloaded as a separate app that would work in concert with the cell phone interface. When a call is routed to a cell phone, but before the software in the phone causes it to ring, the recommended system would perform the following steps:

- If the caller is known to the user (because the caller's phone number is in the contact list of the user, or perhaps in the list of outgoing calls the user has made), then the phone would ring immediately.
- If the caller is unknown to the users, the phone would use its IP connectivity to query the network-based servers to determine if the incoming call is a crowd-source-identified spammer.
 - If the servers identify the caller as a spammer, the phone would (according to user-selected options for server-identified calls in the settings) either signal the user with a special ring, signal the user with a special "skin" on the call, mute the ringer (allowing transfer to voicemail), or drop the call entirely and never ring at all.
 - If the servers identify the caller as a non-spammer, the phone rings and the user could provide input after the call on its legitimacy. If the user then determines that the caller *is* a spammer, the phone would provide an option to tag the caller as a spammer through an easy-to-use button that surfaces after the caller hangs up. Pressing this button would then report the call to the servers as "spam". The interface could also provide the caller with a number of options for future calls from the reported phone number (block/mute/special ring).¹⁵

3.2 POTS Phones Implementation

For a POTS line, we recommend an Internet enabled device (an external "call interceptor") that would typically be interposed between the network interface device¹⁶ and the on-premises

¹⁴ There are also cell phone apps available which block private calls, similar to the *77 anonymous call block features provided by telephone companies, and others which use a database of known marketers (which is downloaded to the phone) and which feature personalized black/whitelists.

¹⁵ The users' decision to mute, block, or create a special ring for a specific number could also provide an additional data point for the recommended algorithm. For example, numbers that are blocked are more likely to be spam, whereas numbers that are just associated with special rings are not.

¹⁶ The network interface jack is typically located in a box attached to an exterior wall or in an out-of-the-way location (e.g., in the walls, basement, closet). http://en.wikipedia.org/wiki/Network_interface_device

telephones - usually one per household.¹⁷ This call interceptor device would serve two purposes: automatic identification of incoming calls from crowd-sourced information and a means for providing identification to the crowd-sourced servers. The device would be designed to “fail safe”, so that in the event of a power failure or lack of IP connectivity, the household telephones would behave as before.

This call interceptor would be able to detect the Caller-ID of the incoming calls and query the network-based servers to determine if the incoming call is a crowd-source-identified spammer. If the servers identify the caller as a spammer, the call interceptor would (according to user-selected options) either cause the on-premises telephones to ring with a special ring or suppress¹⁸ the ringing of the telephones, essentially blocking the call.

However, if the servers do not identify the caller as a spammer, the call interceptor would cause the on-premises telephones to ring, and would pass the Caller-ID information onward to them. While it is expected that all incoming calls would be desired, it is possible that a new, as-yet unidentified spammer is calling. Upon receipt of the incoming call, if telephone subscribers determine that the call is robotically generated, an unwanted solicitation, or other form of spam, they would be able to dial a (newly proposed) vertical service code any time before the next call is received. This code would be received by the call interceptor (for example, *4SP for a spam call, or optionally more specific codes such as *4RO for a robocall, etc.), which would then cause the interceptor to pass the information about the caller via the internet connection to the crowd-sourced servers. Based on the user’s preferences, marking a call as a spammer could also simultaneously block that number for that user’s phone – this option would increase the incentive for users to mark calls as spam. The more users who mark a caller as a spammer, the more quickly that caller would be blocked from bothering others.

For calls that are historical in nature (e.g., that appear as messages on a user’s home answering machine), the call interceptor could optionally also maintain a limited history of incoming phone numbers, and allow the user to access this list (either from their telephone keypad or through a front-panel interface) to mark retroactively a call as spam.

3.3 IP PBX Systems and VoIP Gateways Implementation

A number of computer-based IP PBX systems and VoIP gateways provide a means for a user

¹⁷ A fully functioning ARM-based Linux system (e.g., Raspberry Pi, <http://www.raspberrypi.org/>) can be purchased for about \$35. We believe that a device with fewer features, dedicated to blocking robocalls and other solicitations, could easily be produced in bulk for sale at a substantially reduced cost. Eventually, call interceptors could be built into phones much as Caller-ID is built into phones today.

¹⁸ Caller-ID service on a POTS line is typically provided as a burst of data tones between the first and second rings of the phone. To suppress ringing completely, digital phone with ANI service are required. Alternatively, the call suppressor could wait until after the Caller-ID information is received before ringing the phone.

to inexpensively build a small office or in-home communications system (e.g., Asterisk¹⁹). Since these are often open-source, the call interceptor/blocking technology could easily be provided as a plug-in for these systems.

3.4 Web-based Reporting of Spammers

It is also desirable to provide a means for flagging spammers to cellular subscribers without a smartphone or to POTS line subscribers without a call interceptor. These users could still enter an offending telephone number to a web-based interface, similar to how the FTC currently accepts reports of robotic callers. This interface would likely be used by a minority of users, but would be available nevertheless.

3.5 Telephony Providers Implementation

Telephony providers currently make a feature available to subscribers whereby private calls can be blocked. When a user selects this feature and a call is received that has blocked Caller-ID, the caller is informed with a recording that states that “the person you have dialed does not accept blocked Caller-ID calls.” In addition, the user’s phone does not ring, and the caller would need to redial in such a way that provides Caller-ID to the user.

A similar feature could be provided to the subscribers whereby the telephony provider queries the crowd-sourced spam database for each incoming call. If the caller is marked as a spammer, a recording could inform the caller that “the person you have dialed does not accept spam calls”, the user’s phone would not ring, and the caller would need to redial from a number that is not so marked in the database.

Additionally (and independently), for unblocked calls which are received by the user, telephony providers would have the ability to process the proposed vertical service codes *4RO, *4SP, etc. used to mark spammers, and transmit the cached ANI/Caller-ID information from the call just completed to the database servers on behalf of the user.

4. Effectiveness of Submission

The recommended system would work more efficiently as more telephone subscribers participate. All classes of telephone service would share the same information and the same reporting mechanisms. For POTS lines, the system would improve as each user places a new call interceptor into service, but the system would improve dramatically if cell phone manufacturers release new versions of their phones’ software that includes call interceptor software, and millions of users have the service now available to them. In the meantime, effectiveness would rapidly increase as smartphone users download and use an interceptor app. If telephony providers participate, even larger strides would be made, as many millions of users would have the interceptor service made available to them.

¹⁹ <http://www.asterisk.org/>.

To estimate the effectiveness of our submission, we used the 616,922 spam reports from the sample Robocall Complaint Data provided in five Excel spreadsheets²⁰ for the Robocall Challenge. If we assume that these reports are sampled from from the National Do Not Call Registry web-based complaint form, then we know that reporters must transcribe information from one medium (their phone) to another (a web form), a process which is both slow and error-prone (and one which our recommended system largely circumvents).

Even with that limitation, the data provides some stunning results. For example, one number (609-227-4849) was reported to the FTC 16,056 times from 257 area codes accounting for nearly 2% of the total reports! Many calls were reported with obviously bogus Caller-ID information (e.g., the number 111-111-1111), and we found 51,135 calls reported with no Caller-ID field, plus 45,487 Caller-ID fields reported only once. While these one-time-reported numbers may be due to user confusion or transcription errors on the report (such as users mistakenly entering their own number in the caller field), this problem would be mitigated in a system that is more tightly integrated with the telephones. Another explanation could be that spammers are spoofing Caller-ID information. This later problem is addressed below.

We eliminated the 51,135 “no originating number” reports from consideration, and of the remaining 565,787 reports calculate a spam probability computed solely based on the number of reported calls from a Caller-ID number²¹. As soon as a Caller-ID number is reported as a spammer by more than a threshold number of users, all subsequent calls from that number would be blocked in the recommended system. Based on this simple metric, our recommendation would be to block spammers from the sample data at the following rates:

| Call-Count Spam Threshold | Count of Blocked Calls | Percentage of Spam Blocked |
|---------------------------|------------------------|----------------------------|
| 10 | 473,789 | 83.7% |
| 25 | 430,543 | 76.1% |
| 50 | 391,617 | 69.2% |
| 100 | 347,434 | 61.4% |
| 200 | 299,841 | 53.0% |

²⁰ These files were chosen because they contained the complete Caller-ID number. The larger CSV file was not used in this analysis because it lacked complete traceback information.

²¹ In the proposed implementation below, the algorithm would be more sophisticated, and take many more factors into consideration, increasing both accuracy and blocking ability. Because we assume the data is user-transcribed (and because we note the variance in the Caller-ID Text field), we do not take into account the additional tracking information that would be otherwise available to the proposed system.

Note that due to the relatively high level of effort required to file a web-based complaint, the actual number of spammers and spam calls reported to the FTC is probably far lower than it would be if the reporting system were easy to use and tightly integrated with the telephone, as we recommend.

We also believe that for the web-based reports, a threshold of 10 for a caller to be marked as a spammer is reasonable for estimating the efficiency of a simple algorithm. In the proposed implementation, however, a higher threshold could probably be used with equal efficacy, as there would be a much larger complaint-base combined with other metrics available to track spammers.

In addition, other information could also be used to increase the success-rate of correctly identifying incoming phone calls as spam, even when the calling numbers themselves have never been tagged as spam. For example, large non-carrier entities (such as Google Voice, large banks, etc.) have large quantities of unallocated numbers. These numbers can be used as low-interaction honeypots.²² If a sufficient quantity of these unallocated numbers are called from any given Caller-ID, that entity is calling telephones which are known to not belong to anyone. Anyone “hunting for targets” like this is very likely to be a spammer.

5. Additional Implementation Details

The following section details some of the specifics of the implementation of the call interceptor technology that is recommended here.

5.1 Whitelisting and Removal From Blacklist

Some telephone numbers would have to be whitelisted in the recommended system. For example, emergency alerting systems must be whitelisted. In another example, while the bill collection numbers for the local power or gas company might be marked as a spammer, the emergency reporting systems of those entities must be whitelisted so that their calls cannot be blocked.

Similarly, because the crowd-sourced spam identification system is largely automated, some legitimate entities may accidentally or inappropriately get blacklisted. For example, the local museum may call past donors (an action protected by 47 U.S.C. § 227(a)(2) as a “established business relationship”) and still be flagged as a spammer. Additionally, although telephone bill collectors are hugely unpopular, they are performing a service on behalf of an entity that has an established business relationship with the person being called. These types of callers should be protected from being marked as spam in the database.

Therefore, an easy-to-use appeals process would also be needed wherein callers can request a review of their status. This process could be streamlined for first-time offenders, but would require human reviewers to evaluate repeated claims on a case-by-case basis, supported by

²² [http://en.wikipedia.org/wiki/Honeypot_\(computing\)](http://en.wikipedia.org/wiki/Honeypot_(computing)).

tools to access detailed database information about the caller. In the case of “legal but unwanted calls”, a user could block a specific phone number they have personally tagged as spam. The appeals process would also allow certain callers to be reclassified as “not spam” for anyone who had not personally chosen to block that caller. While this compromise position may make legal callers unhappy, we view the spirit of this challenge as being aimed at protecting the *recipients* of unwanted calls.

Additionally, a protected API would need to be provided to allow telephony providers to notify the system when a number has been canceled or reissued, so that the number can be removed from (or have their spam score markedly lowered in) the database.

5.2 Detecting Forged Caller-ID (“spoofing”)

Although forging/spoofing Caller-ID is illegal in the United States,²³ we still expect it to occur, especially since the robotic calls the recommended system would be seeking to block are also illegal according to 47 U.S.C. § 227(b)(1). It is not possible to prevent spoofing now²⁴ - but it is possible to detect it!

The more people who participate in the system, the more phone numbers and their corresponding text strings would be in the database. As would be clearly and conspicuously described for users before they adopt the technology, the reporting user’s phone number and text string would be needed in the database for reasons described later in this submission (for example, to limit submissions from a single user that could indicate a malicious spam report). In addition, other online databases where users enter their own name and phone number(s) could provide for a clearly disclosed opt-in for that data to be included in the database for the prevention of spoofing (e.g., Google+ profiles, Facebook profiles, etc.) In order for a spoofer to impersonate someone, they must send a fraudulent number and a string.²⁵ If the number and string do not match what is already in the database, we know there is a spoofer (either the original caller or the subsequent one). By its nature, the recommended system would be able to check and see if the number and text are as expected; if not, the system could optionally escalate the caller information to the FTC and FCC, especially if the number is reported by the user as being spam.

Another feature that could further strengthen the system would be an additional mechanism that would compare received phone call identifications to other participating users’ contact lists (who

²³ <http://www.govtrack.us/congress/bills/111/s30/text> and <http://www.fcc.gov/guides/caller-id-and-spoofing>.

²⁴ Cryptographic technology (similar to DKIM) could be used to enable telephony providers to authenticate Caller-ID information and guarantee the accurate functioning of the call blockers. However, this requires a change to the telephony infrastructure, and would take much longer to implement and roll out nationwide.

²⁵ As it is, many people will not answer their phone if they see a name and number they do not recognize. So “Jane Doe” is not a convincing text string, and “Internal Revenue Service” would be avoided. However, “Publishers Clearinghouse” would probably be a strong inducement to answer the phone, even though the Caller-ID proves to be fraudulent!

also opt-in to this feature). Consider for example, that a number of people in this opt-in community have 412-555-1111 listed in their electronic address books as “Joe Smith” or “J. Smith”, etc. Then if anyone else in the opt-in community receives a call putatively from 412-555-1111 that does not also have a text field similar to “Joe Smith”²⁶, the system can immediately determine that the Caller-ID information is fraudulent because the name and number do not match. Using this method, the call can be more quickly classified as spam (without unfairly marking the forged Caller-ID as a spammer).

Users would need to affirmatively opt-in to allow the system to use data from their online profiles or from their phone contact lists. This opt-in data would not be retained in any way by the system except anonymized and in aggregate (so that it would be possible to check credentials by asking “is 412-555-1111 Joe Smith a valid combination?”, but not to look up Joe Smith’s phone number, nor to see who has Joe in their contact list, nor to read Joe’s contact list if he participates in this opt-in component). The information would also be encrypted and otherwise securely protected.

For a spoofer who fraudulently uses one number and calls hundreds of people, this could be handled later in the chain; basically, if someone asks to be whitelisted after the system has put them on the spammer list as a result of their number being spoofed, that data might be very useful to the FTC and FCC to see who/what/when a number was spoofed. When the “this is spam” interface is engaged, it should yield to the system the rights necessary to record all available data about the call, and should also allow the system to share that data with law enforcement entities who can help track down offenders.

Ultimately, however, a determined spoofer would be hard to defeat – they simply need to provide forged Caller-ID credentials that match valid credentials for a valid user. In order for this (or any) system to work in a foolproof manner, there needs to be foolproof Caller-IDentification, and we do not have that today. ANI provides a theoretically non-spoofable identity, but ANI is not currently available to the end-user, so a foolproof system must be implemented by providers. For any system to be able to block all spammers, 100% provider participation in both database query and support of the proposed vertical service codes is required – and we will not have that any time soon, either. Consequently, we are providing our best recommendation based on what exists in the world today, by providing an easy-to-implement and easy-to-rollout solution. We acknowledge that spoofing is a problem, and propose a system that seeks to block as many spam calls as possible in spite of this fact.

5.3 Preventing Denial of Service Attacks

While it would be easier to allow arbitrary reports of spamming and robotic callers from any putative call interceptor device, the reality is that malicious individuals could attempt to disrupt the service (and, for example, mark an unfavored public figure or a competitor as a spammer,

²⁶ The number would need to match exactly, but because the name may vary somewhat in address books, fuzzy-matching algorithms would be used for validation of the name.

preventing their calls from going through). To prevent this, we recommend that the spammer database record tuples of (IP address, EMEA or MAC address, reported number) to detect unique reports²⁷ to prevent malicious users from artificially elevating a telephone number's spam rating, and to rate limit reports from single sources.²⁸

In all cases, the recommended software or hardware call interceptor should require that to report a phone number as a spammer, that number must have called in to the device making the report. Additionally, reporting a spammer would need to be easy, but not trivial. On cell phones, a multi-step confirmation process could be sufficient. For a web-based report, a CAPTCHA or similar device could be used to verify that a human is submitting the report.

The recommended system could also make it difficult to mark a number as a spammer that a user has called first (or that they have in their contact list), as this indicates a prior relationship with the caller.

5.4 Spam Detection Algorithm

In order to accurately determine which numbers are associated with spammers, the recommended system would take into account much more than simply the raw number of spam reports on a specific number. We recommend an algorithm that would be developed based on aggregated call histories, and refined as time goes on (since the effectiveness of any implemented system would depend on the number of participants and the volume of calls that have been processed). When used below, “periodically” may mean daily or monthly; “daily at midnight” is a likely configuration.

At the high level, the recommended algorithm would be used to calculate a “spam score” for each originating name/number pair. The spam score is a probability value that expresses whether calls from a source are junk or spam – the higher the score, the more likely a call is spam. To compute a score, we recommend an algorithm that would:

- Collect data on all calls monitored by system, including information about:
 - Originating number and name (caller)
 - Destination number (reporter)
 - Local time call originated and ended
 - Timezone (i.e., offset from UTC), including Daylight Savings data
 - Whether the call was reported as spam
 - (optionally) the type of spam reported.
- Wherever possible, anonymize all data stored in the database (see “Privacy” below).

²⁷ Since it is unlikely that a spammer will repeatedly call a single number, it is likewise unreasonable for a single user to report a spammer's number repeatedly.

²⁸ Since a telephone can only receive so many calls per day, making more than a reasonable number of spammer reports in a single day (or hour) from a single IP address should be greeted with some skepticism.

- Periodically calculate a confidence weight for each destination number/spam reporters. This would be done to avoid abuse of the system.
 - Reporters with a single report are weighted as 1.0.
 - Find the percentage of calls reported as spam by the reporter. Graphing all reporters, sorted by percentage of calls reported as spam, if the graph exhibits “hockey stick” behavior,²⁹ then all reports past the “bend in the stick” for that reporter are weighted exponentially lower based on the distance between the bend and the end of the graph (i.e., we would attempt to detect abuse of the system for users who send in too many spam reports).
 - Find the raw number of calls reported as spam by the reporter. This also exhibits exponential decline past the “bend in the hockey stick”, but may require further tuning if legitimate high-use cases are identified.
- Aggregate repeated reports with matching origination/destination numbers.
 - If the decision is made to block originating numbers to a destination (that were previously marked as spam by that destination), this would indicate either an error in the implementation and/or abuse of the system.³⁰
 - This should be monitored, and should trigger an alert for manual review.
 - When the alert is triggered, all reports from the destination number would be weighted to 0.
- Periodically calculate a weight for each spam report, based on:
 - The confidence weight of the destination number/reporter, as discussed above.
 - Repeated reports, if they weren’t permanently blocked as discussed above.
 - The current age of the most recent report.
 - Exceedingly old reports are weighted as zero, but are not removed, so that someone reporting a spammer may personally block them.
 - The length of the call before the spam report was filed. If the call was over a threshold of time (e.g., 5 minutes), weight the spam report lower, as it might have been an accidental filing.
- Drop originating numbers with fewer than a threshold of calls. For example, if the number 301-555-1212 only made 10 calls this month (i.e., if only 10 queries about that number were made by spam interceptors), that number is not a significant number for purposes of spam computation. Note: this does not apply to originating *names*.
- Build a weighted mean³¹ of the caller information, such that high-volume originators are

²⁹ This would be where the slope of the graph rapidly changes from a gentle slope to a steep slope (i.e., the first derivative undergoes a rapid transformation, or the second derivative is not close to a constant).

³⁰ To clarify, if 123-456-7890 calls 123-555-1212 and 123-555-1212 marks that call as spam, we would always count this towards 123-456-7890’s spam score. We could also (optionally) immediately block 123-456-7890 from ever calling 123-555-1212 again. If this latter measure is taken, yet we continue to see calls from 123-456-7890 to 123-555-1212 being marked as spam, then something is broken (either there is abuse of the system or we are failing to properly block spam calls for a given to/from number pair).

³¹ http://en.wikipedia.org/wiki/Weighted_mean.

weighted more heavily than low-volume originators.

- This gives us a mean spam confidence score.
 - If an originating number has more than one originating name (possible Caller-ID spoofing) perhaps trigger an alert for manual review, and for alerts above a set frequency threshold, possibly automatically send the data to the FCC/FTC to report potentially fraudulent spoofing.
 - One (manual) alert response for this would be to add the originator to a greylist. For this greylist, the originating number would not trigger future alerts unless it's behavior significantly shifted.

5.5 Algorithm Refinements

Once the proposed preliminary system has been built, one could envision the following possible enhancements and refinements to the submission:

1. "Type of spam" refinements -- allow the user to report different types of spam. This could include "robotic caller", "recorded voice", "unwanted charity solicitation", "advertiser", "telemarketer" compared to block list, etc. and would allow users to filter differently depending on type of caller (for example, some users might not mind charities calling them, but would reject all advertisers).
2. Location correlation -- allow a 'heatmap' data source of spam reports. When a new destination number makes a request to the database, weight nearby (same area code?) spam reports significantly more heavily.
3. Name merging -- if an organization is using Caller-ID Spoofing legitimately (e.g., to unify the name field for each of their outgoing numbers), but doesn't merge those numbers, each number would be treated separately. For matching names in the same area code, especially ones that don't match to a database of names (the "phone book"), consider merging the spam reports for all originating numbers with matched names.
4. The interceptor apps in cell phones and the interceptor devices for POTS lines could have an option for the user to express how much that user trusts the interceptor system. This would establish an inverse relationship with spam score: low trust means a spammer must have a very high spam score to be blocked, whereas a high trust means all spammers are blocked. There could also be a "report-only" mode, where no calls are blocked, but the system would report the spam score of each incoming call.
5. Wherever possible, the system could limit cell phone-based reporting to WiFi (to conserve the user's data plan). We might hope that cell phone providers would allow database queries be free of roaming charges, although if not, users might reasonably be willing to offset their peace of mind against the small cost of querying the database when calls are received.

5.6 Privacy and Security

Privacy and security are really important to Google, and our proposed method for combating unwanted robocalls takes that into account. In the recommended system, whenever possible

data could be anonymized with cryptographic hashes. In general, the system should hash non-“public” data (i.e., the information other than that which is provided by Caller-ID when a call is received). Thus, privacy and security is addressed by de-identifying information where possible, aggregating data, and other techniques. In addition, the materials describing the recommended technology (such as a pop-up description of the cell phone software and the packaging material for the POTS interface) would make clear that a key feature of the technology includes the reporting of anonymized information as described in this submission for the purposes of crowd-sources spam detection.

5.6.1 Spam reports indexed by reporter

For each reporter (destination number), the recommended system would only need to know which numbers were marked as spam; it would not need to retain detailed data on all calls that were received. For non-spam call data used in the algorithm, information would only be retained in aggregate. This spam reporting would be part of the clearly disclosed terms of service of the recommended system to keep the list reported spammers publically queryable by the system; this information would be necessary for the system to function at all.

For each reporter, the recommended system would store a list of originating numbers and numbers/text-strings that have been marked as spam, a timestamp of when each was marked as spam, and what type of spam each was marked as (if applicable).

5.6.2 Spam reports indexed by both originator and destination

To know if a spam report is a duplicate, the recommended system would need both the originating number (from Caller-ID) and the destination number (from the user making the report). Each U.S. telephone number can be one of approximately 10^{10} possible values. To build a unique anonymized key, the system would concatenate the originating and destination numbers (e.g., if 412-555-1212 calls 724-555-1234, the key becomes 41255512127245551234), which gives us 10^{20} possible unique identifiers, which can then be hashed for privacy. A rough estimate might be on the order of a million dollars of electricity to build a rainbow table to crack this data, without utilizing other methods to further obfuscate the data and increase the complexity of cracking the database.

Spam reports would be indexed by this unique identifier, and have the complete call data stored for them (caller and called number, caller text string, and time of day). The system may also keep copies of this data keyed by originating area code + originating name + destination, to enable deduplication of multiple-number originators (i.e., a call center or PBX which places outgoing calls from many different numbers).

5.6.3 Spam reports indexed by originator

For each originator (person making the calls that were checked by any recipient), the recommended system would not need to store unique call data; instead it would simply need aggregated, anonymized counts. The system would keep:

- The number of outgoing calls, recorded as:
 - The number of outgoing calls per day
 - The number of outgoing calls in one-hour blocks (e.g., calls made between 1PM-2PM, date is ignored).
 - Number of outgoing calls to each given area code.
 - Average of the length of time of the longest outgoing call to any unique recipient
- The number of outgoing calls marked as spam by the recipient, indexed as above, also indexed by “type of spam reported”.

The recommended system may also key a copy of this aggregated data by originating area code + originating name, to enable de-duplication of multiple-number originators (i.e., a call center or PBX which places outgoing calls from many different numbers with a single name).

5.6.4 Additional privacy and security consideration

All other information that is collected for the system, such as data from users’ contact lists or online profiles, would also be anonymized and securely stored using the principles outlined above. When choosing to opt-in to providing this personal contact information, the nature of the data collected and the purpose for the collection would be clearly and conspicuously disclosed.

6. Summary and Conclusion

The system we recommend will be highly effective, easy to implement, and easy to use. It would address nearly the complete range of telephonic devices, and would be available to nearly the entire population of the United States for a negligible cost. All of the infrastructure required to implement the submission already exists (in the telephone system and Internet), and a number of companies have the requisite experience needed to create the required software and optional hardware in short order (so this need not be a single-source solution, although for maximum effectiveness the databases should all be exact replicas of each other). People with disabilities can use it as easily as they use their existing telephones. The system can easily be rolled out incrementally. It is not a perfect system, but the imperfections are well known and accounted for.

The solution requires no legislative changes, and is 100% opt-in – users would freely choose to use the spam blocking service, or report spammers, or neither or both. The only parties who would suffer are the spammers themselves.