SEVENTH FRAMEWORK PROGRAMME
Area ICT-2009.1.4 (Trustworthy ICT)

Visual Analytic Representation of Large Datasets for Enhancing Network Security

D6.1 Threat Landscape Identification Scenario

Contract No. FP7-ICT-257495-VIS-SENSE

Workpackage	WP 6 - Workpackage 6
Author	SYM
Version	1
Date of delivery	M36
Actual Date of Delivery	M38
Dissemination level	Public
Responsible	SYM
Data included from	UKON

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°257495.

The VIS-SENSE Consortium consists of:

Fraunhofer IGD	Project coordinator	Germany
Institut Eurecom		France
Institut Telecom		France
Centre for Research and Technology Hellas		Greece
Symantec Ltd.		Ireland
Universität Konstanz		Germany

Contact information:
Dr Jörn Kohlhammer
Fraunhofer IGD
Fraunhoferstraße 5
64283 Darmstadt
Germany

e-mail: joern.kohlhammer@igd.fraunhofer.de
Phone: +49 6151 155 646

Contents

Abstract

This deliverable presents the results of applying the integrated VIS-SENSE framework to various security data sets comprising several months of accumulated data. The goal is to evaluate the *usability* of the framework and its *capabilities* to identify, represent and explain various Internet threat phenomena, in which different security events have been clustered and attributed to a common *root cause*.

To this aim, different security analyses are presented in this deliverable in order to highlight as much as possible the extensive set of visual analytics capabilities of our framework, namely:

- Visual Analysis of Web Threats Dynamics
- Visual Analysis of Scammers Operations
- Visual Analysis of Spam Sent from Hijacked Networks

These very diverse applications are different instantiations of the first VIS-SENSE scenario as defined previously in D1.2 (threat landscape visualization), by which the users of the target groups of the project aim to get insights into various threat ecosystems using *visual analytics* tools.

1 Introduction

Security intelligence is the process of collecting information and applying the knowledge, expertise and skills of security analysts to derive business value. The objective of this application scenario consists in monitoring both known and unknown threats by performing security investigation and *root cause analysis* through mining very large security data sets.

Security companies have recently realized the business value that the analysis and visualization of massive amounts of data can bring thanks to an improved understanding of Internet attacks.

The central idea behind the application of the VIS-SENSE framework is thus to help security analysts understand the common root causes of attack phenomena observed in the Internet thanks to effective visual analytics tools. Analysts must be able to quickly attribute or classify any given set of attacks to the phenomenon that has likely generated them, and represent them graphically through aggregate summaries highlighting, in a visual and meaningful way, the modus operandi of the cybercriminals behind those attacks (*i.e.*, a task that is typically referred to as *strategic analysis).*

Ideally, analysts should be warned as soon as possible of any significant deviation observed in those cybercriminal behaviours. The visual analytics system should thus highlight the occurrence of a new phenomenon, or any significant changes in the behavior of malicious actors identified by the automated attack attribution methods (a task typically referred to as *tactical analysis*). Whenever possible, it should be possible to bridge the gap between the strategic and tactical analyses, *e.g.*, by leveraging behavioural models obtained from the strategic analysis component to correlate this high-level information about attackers behaviour with low-level intrusion or attack information collected on other networks, and vice versa.

As we will demonstrate in this document, the R&D efforts carried out by VIS-SENSE partners have led to the development of very effective visual analytics tools that can be used to perform intelligence analyses on security datasets, and quickly generate new insights. In fact, this deliverable presents the results of applying an integrated Web version of the VIS-SENSE framework to various security data sets comprising several months of accumulated data. Our goal was to evaluate the *usability* of the framework and its *capabilities* to identify, represent and explain various Internet threat phenomena, in which different security events are clustered and attributed to a common *root cause.*

More specifically, different applications of the framework are presented in this document:

- Visual Analysis of Web Threats Dynamics
- Visual Analysis of Scammers Operations
- Visual Analysis of Spam Sent from Hijacked Networks

These applications represent different instantiations of the first VIS-SENSE scenario that has been defined previously in D1.2. They are all related to *security intelligence*) use cases, by which users of the target groups identified in the project aim to get insights into various threat ecosystems.

The first application was formulated in the beginning of the project and described quite extensively as part of the VIS-SENSE use cases (see D1.2). Therefore, we use this previously-known application as *baseline* to evaluate the usability of the visualization tools developed in the project as well as their effectiveness. We achieve this by comparing the new results obtained with the VIS-SENSE framework, to the ones obtained previously without the framework (which were often published in previous papers).

To illustrate the *broad applicability* of the VIS-SENSE tools, another application was performed on a completely new and unknown dataset (made of *scam emails*) that was not initially integrated in our data collection infrastructure. This analysis was done in collaboration with a group of researchers external to the project. The insights generated by this analysis have been published in a joint paper and presented at the International Workshop on Cyber Crime (IWCC'13) – an IEEE Security & Privacy workshop [12].

Finally, the last application demonstrates the *cross-domain* correlation capabilities of the VIS-SENSE framework by performing visual analytics on spam sent from *hijacked networks*, as detected by our BGP analysis tools. This last analysis fulfils a key requirement of this project by establishing an explicit link between the two Internet planes: (i) the *control* plane (BGP infrastructure), and (ii) the data plane, *i.e.* in this case *spammers* activities. To the best of our knowledge, this is the first analysis of this kind that brings such clear evidence of correlation between these two aspects, but also highlights the modus operandi of so-called *fly-by spammers* that were until now only hypothetically described by previous research.

For every application, we try to emphasize the various capabilities and the novel aspects of the developed tools. We also evaluate to what extent they fulfil the user requirements laid out in the beginning of the project. Note that a more detailed evaluation is performed in D6.3 (VIS-SENSE Framework Evaluation).

1.1 The TRIAGE Web framework

Recall that *visual analytics* tools aim at integrating and establishing synergies between two main components: *data mining* algorithms and effective *visualizations.*

For this scenario, we have leveraged the TRIAGE data mining software framework, which was initially developed (in the WOMBAT project) to address a well-known security problem that is often referred to as *attack attribution*, or how to attribute (potentially) different attacks to a common *root cause*, based on the combination of all available evidence. Note that, by *root cause*, we do not refer to the identification of a given machine that has launched one specific, isolated attack (*i.e.*, IP traceback), but instead, we are more interested in having a better idea of the various individuals, groups or communities (of machines) that are responsible for large-scale attack phenomena. Also, as noted in previous work done by the WOMBAT consortium, the ultimate goal of doing intelligence analysis is not *per se* to offer names of individuals to law enforcement agencies. The goal is, instead, to provide models of the acting entities that we are facing. However, through generalization, these models can help in understanding the threats that every person or organization who connects to the Internet currently faces.

For the sake of clarity, and to make this deliverable as self-contained as possible, we briefly recall the main concepts of TRIAGE. This data mining framework combines clustering techniques with a data fusion process that is based on multi-criteria decision analysis (MCDA) algorithms [18]. Thanks to this, virtually any type of security events can be automatically grouped together based upon a number of common elements (also called *features*), hence generating clusters of events that are likely due to the same *root cause.* As a result, TRIAGE can identify and generate two types of Clusters:

- *1D Clusters*: these are simple data structures that are grouping events that are considered to be similar with respect to a given feature (or attribute), eventually across multiple analysis jobs (incremental analysis)

- *Multi-Dimensional Clusters* (or in MDC's in short), which are generated across *all dimensions* and are thus grouping events that share *more than one similar attribute* (according to the data fusion model defined by the analyst)

Furthermore, to address the inherent scalability issue of pairwise clustering approaches, a *prototype extraction* algorithm was designed and implemented as *preprocessing* step in order to automatically compress the raw data set before clustering the data objects, such that the TRIAGE algorithm could be applied to larger data sets (typically, 1 or 2 orders of magnitude larger). We have observed experimentally that we obtain usually a good compression ratio of the data set using this prototype extraction algorithm. In most cases, the size of the resulting set of prototypes is only 1-10% of the original dataset size (depending on the type of data and the algorithm parameters), which means that *medium-sized* datasets (about 100,000 objects) may be processed in a reasonable time and in a single batch processing job.

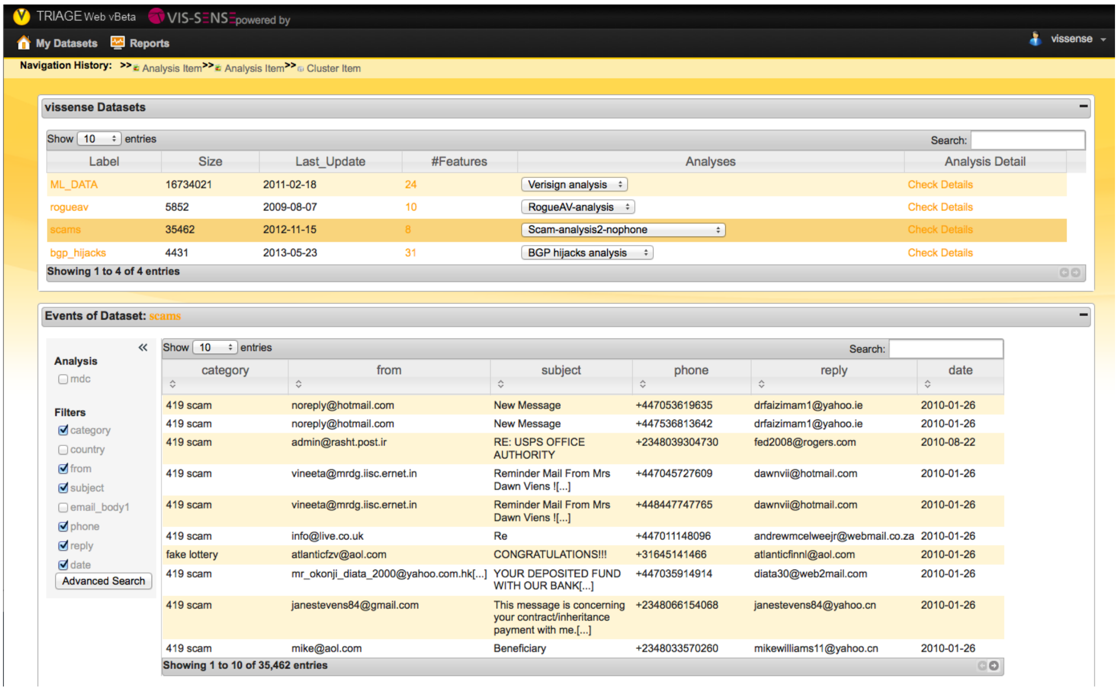

Navigation History: >> Analysis Item >> Analysis Item >> Cluster Item

vissense Datasets

Show 10 entries Search:

Label	Size	Last_Update	#Features	Analyses	Analysis Detail
ML_DATA	16734021	2011-02-18	24	Verisign analysis	Check Details
rogueav	5852	2009-08-07	10	RogueAV-analysis	Check Details
scams	35462	2012-11-15	8	Scam-analysis2-nophone	Check Details
bgp_hijacks	4431	2013-05-23	31	BGP hijacks analysis	Check Details

Showing 1 to 4 of 4 entries

Events of Dataset: scams

Analysis
□ mdc

Filters
☑ category
□ country
☑ from
☑ subject
□ email_body1
☑ phone
☑ reply
☑ date

Advanced Search

Show 10 entries Search:

category	from	subject	phone	reply	date
419 scam	noreply@hotmail.com	New Message	+447053619635	drfaizimam1@yahoo.ie	2010-01-26
419 scam	noreply@hotmail.com	New Message	+447536813642	drfaizimam1@yahoo.ie	2010-01-26
419 scam	admin@rasht.post.ir	RE: USPS OFFICE AUTHORITY	+2348039304730	fed2008@rogers.com	2010-08-22
419 scam	vineeta@mrdg.iisc.ernet.in	Reminder Mail From Mrs Dawn Viens ![...]	+447045727609	dawnvii@hotmail.com	2010-01-26
419 scam	vineeta@mrdg.iisc.ernet.in	Reminder Mail From Mrs Dawn Viens ![...]	+448447747765	dawnvii@hotmail.com	2010-01-26
419 scam	info@live.co.uk	Re	+447011148096	andrewmcelweejr@webmail.co.za	2010-01-26
fake lottery	atlanticfzv@aol.com	CONGRATULATIONS!!!	+31645141466	atlanticfinnl@aol.com	2010-01-26
419 scam	mr_okonji_diata_2000@yahoo.com.hk[...]	YOUR DEPOSITED FUND WITH OUR BANK[...]	+447035914914	diata30@web2mail.com	2010-01-26
419 scam	janestevens84@gmail.com	This message is concerning your contract/inheritance payment with me.[...]	+2348066154068	janestevens84@yahoo.cn	2010-01-26
419 scam	mike@aol.com	Beneficiary	+2348033570260	mikewilliams11@yahoo.cn	2010-01-26

Showing 1 to 10 of 35,462 entries

Figure 1.1: The *Dataset* page of the TRIAGEWEB interface, displaying details on the datasets available for analysis.

As demonstrated later in this document, this hierarchical data model (*Prototypes - Clusters - MDClusters*) can also be beneficial to interactive data visualization, as it provides a coarse clustering layer at the highest level (using data objects like Clusters and Prototypes, whose patterns are associated with aggregate data to be visualized but also less expensive queries). However, more details can still be provided on demand if the user wants to inspect results at the lowest level (data set entities).

We refer the interested reader to previous deliverables (in particular D3.3, D1.2) for a more extensive description of the TRIAGE workflow as well as the various algorithms upon which it is built.

The R&D efforts carried out by VIS-SENSE partners have thus led to the development

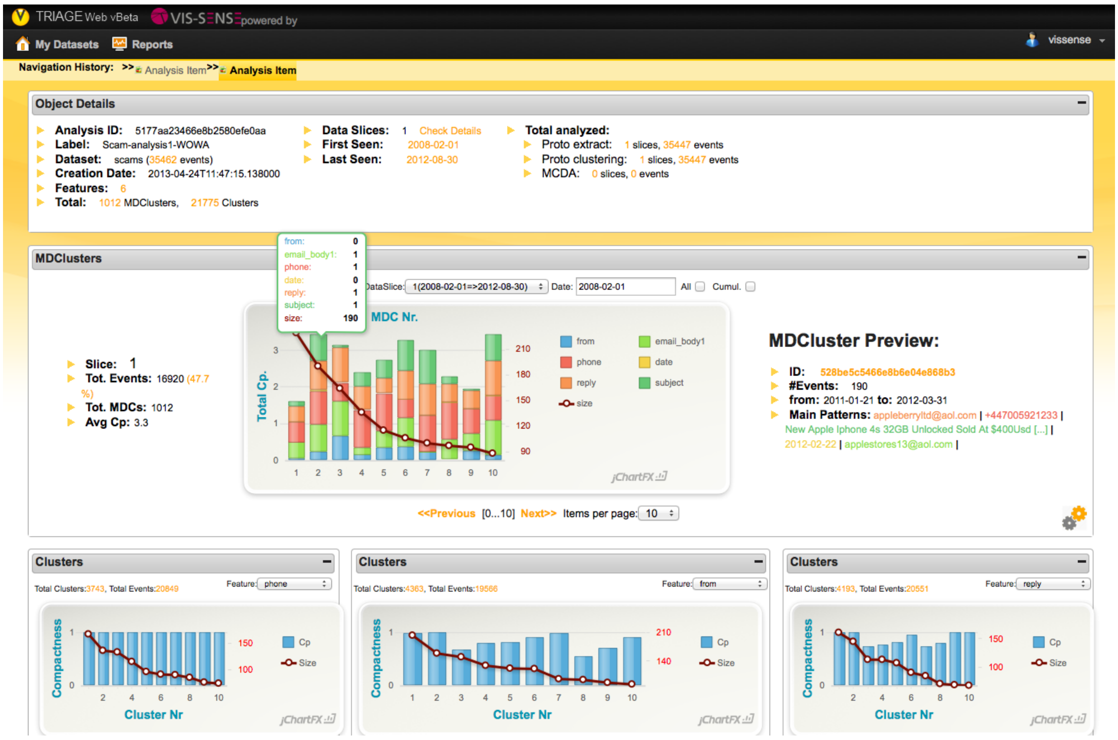

Object Details

- **Analysis ID:** 5177aa23466e8b2580efe0aa
- **Label:** Scam-analysis1-WOWA
- **Dataset:** scams (35462 events)
- **Creation Date:** 2013-04-24T11:47:15.138000
- **Features:** 6
- **Total:** 1012 MDClusters, 21775 Clusters
- **Data Slices:** 1 Check Details
- **First Seen:** 2008-02-01
- **Last Seen:** 2012-08-30
- **Total analyzed:**
  - Proto extract: 1 slices, 35447 events
  - Proto clustering: 1 slices, 35447 events
  - MCDA: 0 slices, 0 events

MDClusters

- **Slice:** 1
- **Tot. Events:** 16920 (47.7 %)
- **Tot. MDCs:** 1012
- **Avg Cp:** 3.3

DataSlice: 1(2008-02-01=>2012-08-30) Date: 2008-02-01 All □ Cumul. □

from: 0
email_body1: 1
phone: 1
date: 0
reply: 1
subject: 1
size: 190

<<Previous [0...10] Next>> Items per page: 10

MDCluster Preview:

- **ID:** 528be5c5466e8b6e04e868b3
- **#Events:** 190
- **from:** 2011-01-21 **to:** 2012-03-31
- **Main Patterns:** appleberryltd@aol.com | +447005921233 | New Apple Iphone 4s 32GB Unlocked Sold At $400Usd [...] | 2012-02-22 | applestores13@aol.com |

Clusters

Total Clusters:3743, Total Events:20849 Feature: phone

Clusters

Total Clusters:4363, Total Events:19566 Feature: from

Clusters

Total Clusters:4193, Total Events:20551 Feature: reply

Figure 1.2: The *Analysis* page of the TRIAGEWEB interface, displaying details on a specific analysis (in this case, a *scam emails* analysis).

of an integrated Web user interface that incorporates several visualization components developed by the partners and integrate them into a Web application framework called TRIAGEWEB. This was made possible thanks to the modular architecture of the VIS-SENSE framework. Figure 1.1 and 1.2 display some screenshots of the TRIAGEWEB user interface, showing details on the *Dataset* page (main starting page showing the available datasets and displaying events details via a dynamic table), and the TRIAGE *Analysis* page respectively (showing the overview of analysis results). More explanation on what is being displayed in these screens will be given later in the application Chapters, in which we will demonstrate the capabilities and utility of TRIAGEWEB through different analyses of security datasets.

Note that a short demonstration video is available on Youtube at: `http://youtu.be/gnAciFR9ANI`.

1.2 High-Level requirements

In D1.2 (Use Case Analysis and User Scenarios), some high-level requirements were defined for the security intelligence scenario according to three different categories, which can be summarized as follows:

- *Network analytics*:
  - support for the definition of MCDA and clustering algorithms
  - interactive visual analysis techniques to help the decision maker define or tune MCDA and clustering parameters
  - improved scalability of the framework
  - incremental analysis to highlight the occurrence of new phenomena
  - support for the detection of significant changes in the modus operandi of malicious actors
- *Visualization*:
  - advanced interactive visualizations for graph-based representations (*e.g.*, ability to select, zoom, or move around nodes and edges)
  - visual feature selection
  - analysis of feature interdependencies (for better understanding of events correlations within the same MDC)
  - suitability to represent very large and complex networks (scalability)
  - visual exploration at varying levels of detail
  - understanding of complex feature relationships and facilitating the interpretation of the results (interactivity & sense-making)
- *System integration*:
  - integration of remote API's (*e.g.*, REST, WAPI, etc)
  - provide interfaces for the tight coupling of modules available on the same hardware (to improve system responsiveness)

These high-level requirements will be considered in every application (whenever possible), together with *usability* aspects and the evaluation of the quality of the insights generated by the framework[1].

[1] Note that a more detailed evaluation of the VIS-SENSE framework in light of the design requirements set during the project is provided in D6.3 (Framework evaluation)

2 Visual Analysis of Web Threats Dynamics

In this Chapter, we demonstrate how TRIAGEWEB is used to visually analyse a type of threat that has started to emerge in the last years, namely *rogue security*software. This type of misleading application is distributed through large-scale web-hosted campaigns. A *rogue AV* software pretends to be legitimate security software, such as an antivirus scanner, but in reality, these programs provide little or no protection and, in fact, may actually install the very malicious code it purports to protect against.

In the following Sections, we describe how we leveraged our TRIAGEWEB framework (and the underlying TRIAGE multi-criteria clustering technology) to analyze the rogue campaigns through which this type of malware is distributed, *i.e.*, what are the techniques, server infrastructure and coordinated efforts employed by cyber-criminals to spread their rogue software. Actually we revisit the experimental results obtained previously by some of the partners and published in the Symantec Internet Security Threat Report in 2009 (special edition *Symantec Report on Rogue Security Software* [17]) and in a previous international academic publication [10].

In a nutshell, we demonstrate how the VIS-SENSE visual analytics tools integrated in TRIAGEWEB can deliver visual insights into different networks of rogue domains that are likely linked to the same campaign, helping analysts understand the modus operandi of the criminal organizations behind them in a reduced time and with much less efforts than previously.

2.1 Introduction

A rogue security software program is a type of misleading application that pretends to be a legitimate security software, such as an anti-virus scanner, but which actually provides the user with little or no protection. In some cases, rogue security software (in the following, more compactly written *rogue AV*) actually facilitates the installation of the very malicious code that it purports to protect against [17].

Rogue AV makes its way on victim machines in two prevalent ways. First, social engineering techniques, such as Web banner advertisements, pop-up windows and attractive messages on blogs or sent via spams, can be used to convince unexperienced users that a rogue tool is free and legitimate and that its use is necessary to remediate often non-existent threats found on the victim's computer (hence, the other name *scareware* given to those programs). A second, more stealthy technique consists in attracting victims to malicious web sites that exploit vulnerabilities in the client software (typically, the browser or one of its plugins) to download and install the rogue programs, sometimes without any user intervention (*i.e.*, via *drive-by* downloads).

Despite its reliance on relatively unsophisticated techniques, rogue AV has emerged as a major security threat, in terms of the size of the affected population (Symantec's sensors alone reported 43 million installation attempts over a one-year monitoring period, the number of different variants unleashed in-the-wild (over 250 distinct families of rogue tools have been detected by Symantec, and the volume of profits generated by cyber-crooks, as reported in a white paper published by Symantec in [17]

The prevalence and effectiveness of this threat has thus spurred considerable research in the security community. It can be reasonably assumed that malware code, the infrastructure used to distribute it, and the victims that encounter it do not exist in isolation, but are different aspects of the coordinated effort made by cyber-criminals to spread or distribute rogue AV. We will refer to such a coordinated activity as a rogue AV *campaign*, which is very likely managed by the same group of people, who are reusing, at various stages of the campaign, the same techniques, strategies, and tools (for obvious reasons of development cost).

The purpose of the analysis scenario presented in this Chapter is to identify any *emerging patterns* in the way rogue domains are created, grouped, and interconnected with each other, based upon common elements (*e.g.*, rogue AV-hosting sites, DNS servers, domain registration), and for these reasons, are very likely associated to the same rogue AV campaign.

2.2 Description of the Data Set

This analysis of the rogue AV threat has been performed by leveraging HARMUR, a **H**istorical **AR**chive of **M**alicious **UR**Ls, which was introduced in D2.2 (Data collection infrastructure). HARMUR [13] builds upon two types of information feeds: URL feeds that provide lists of fresh URLs likely to be of interest, and analysis feeds that build a wide range of contextual information around each URL introduced in the system by the URL feeds. To make our evaluation easier, we decided to reuse the very same dataset as the one used in previous publications such as in [10, 17], so that we could use our previous results as a baseline to evaluate the new results obtained with TRIAGEWEB.

HARMUR (and the associated WAPI-enabled interface) has been described extensively in other deliverables (mainly in D2.2 and D1.2). To make this deliverable as self-contained as possible, we give a very short summary of the dataset in this section. We invite the reader who is familiar with this dataset to skip this section and go directly to the analysis results.

HARMUR is a repository of information on the characteristics and the dynamics associated to web-related threats. HARMUR collects information on domains that are believed to be suspicious or malicious through a variety of security data sources. For each suspicious domain, HARMUR tries to look at the characteristics of the hosting infrastructure (web servers, DNS information, geographical location of the servers and hosting Autonomous System), at the domain registration information (WHOIS data) and at the security information (retrieved from Norton Safeweb and Google SafeBrowsing).

The HARMUR data set we have considered for this analysis was collected over a period of approximately two months, in July and August 2009. The rogue AV-hosting servers were identified through a variety of means, including automated and manual feeds.

To build our experimental data set, we have considered 5,852 DNS entries pointing to 3,581 distinct IP addresses of web servers that were possibly hosting rogue security software. After analysis, the 3,581 Web servers have been broken down into the following categories:

- 2,227 Web servers (as identified by their unique IP addresses) were hosting domains serving only rogue security software,
- 118 servers hosted rogue security software along with domains that served malicious code,
- the remaining IP addresses served malicious code along with innocuous domains.

From this dataset, a number of *domain features* that were considered to be relevant for analyzing rogue AV campaigns have then been extracted and uploaded to TRIAGEWEB by using the WAPI communication interface (see also Table 2.1 for an overview of these domain features):

- *domain*: the DNS name associated to the website URL on which the distribution of a rogue AV software has been detected (*e.g.*, `windowsantivirus2008.com`);
- *server_ip*: the IP address(es) of the web server(s) hosting a given rogue domain. Since multiple IP addresses may point to the very same domain, the corresponding feature vector is a *set* of server IP addresses;
- *class_c*: the Class C-subnet(s) of the web server(s) hosting a given rogue domain. The corresponding feature vector is thus a *set* of Class C-subnets;

- *class_b*: the Class B-subnet(s) of the web server(s) hosting a given rogue domain. The corresponding feature vector is thus a set of Class B-subnets;
- *ns*: the IP address(es) of the authoritative nameserver(s) for a given rogue domain. The corresponding feature vector is thus a set of nameserver IP addresses;
- *whois_reg*: refers to the name or email address of the *registrant* (as given in the *Whois* registration data);
- *whois_rar*: refers to the name of the Registrar (as given in the *Whois* regaistration data);
- *geo*: refers to the geolocation of the web server hosting a given domain (as obtained through IP-to-geolocation databases);
- *day*: refers to the creation date of the domain, as given in the *Whois* registration data.

Figure 2.1 displays a screen shot of the rogue AV dataset as it appears in the TRIAGEWEB interface, which provides a convenient data table that enables the user to:

- *browse* the data via page navigation;
- *sort* the data on any field being displayed in the table;
- perform quick *search* queries on specific patterns (*e.g.*, a given registrant email address, as illustrated in Figure 2.1);
- enable or hide specific data fields via the checkboxes in the left margin, which adds or removes columns and updates the data table on-the-fly (*dynamic field selection*)

As shown in Figure 2.2, simple statistics can be obtained for each dataset feature, such as the max and min values, the cardinality (number of unique values) and data type (*e.g.*, list, unicode, string, integer, IP address, etc).

The *Advanced Search* of the Events data table filter allows the user to perform a search in the data set based on a combination of different criteria (on multiple fields). *Wildcard* characters (*%*) enable the analyst to search for partial matches in string attributes.

Finally, the activation of the *mdc* checkbox (in the left margin) will add one more column within the data table to display the identifier of the *MDCluster* in which an entity (*i.e.*, a *rogue domain* in this case) has been attributed as a result of the TRIAGE clustering analysis. This enables the analyst to quickly search and find any *campaign* that was possibly associated to a given threat.

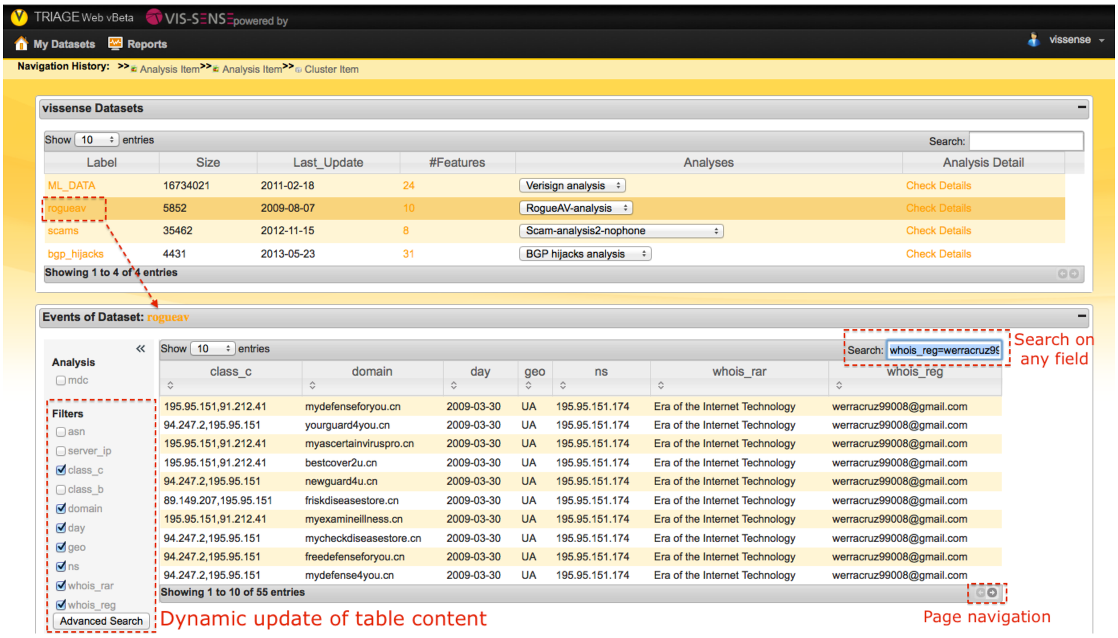

Navigation History: >> Analysis Item >> Analysis Item >> Cluster Item

vissense Datasets

Show 10 entries   Search:

Label	Size	Last_Update	#Features	Analyses	Analysis Detail
ML_DATA	16734021	2011-02-18	24	Verisign analysis	Check Details
rogueav	5852	2009-08-07	10	RogueAV-analysis	Check Details
scams	35462	2012-11-15	8	Scam-analysis2-nophone	Check Details
bgp_hijacks	4431	2013-05-23	31	BGP hijacks analysis	Check Details

Showing 1 to 4 of 4 entries

Events of Dataset: rogueav

Analysis
□ mdc

Filters
□ asn
□ server_ip
☑ class_c
□ class_b
☑ domain
☑ day
☑ geo
☑ ns
☑ whois_rar
☑ whois_reg

Advanced Search

Show 10 entries   Search: whois_reg=werracruz9

class_c	domain	day	geo	ns	whois_rar	whois_reg
195.95.151,91.212.41	mydefenseforyou.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
94.247.2,195.95.151	yourguard4you.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
195.95.151,91.212.41	myascertainviruspro.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
195.95.151,91.212.41	bestcover2u.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
94.247.2,195.95.151	newguard4u.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
89.149.207,195.95.151	friskdiseasestore.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
195.95.151,91.212.41	myexamineillness.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
94.247.2,195.95.151	mycheckdiseasestore.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
94.247.2,195.95.151	freedefenseforyou.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com
94.247.2,195.95.151	mydefense4you.cn	2009-03-30	UA	195.95.151.174	Era of the Internet Technology	werracruz99008@gmail.com

Showing 1 to 10 of 55 entries

Search on any field

Dynamic update of table content

Page navigation

Figure 2.1: *Dataset* page displaying details on the Rogue AV dataset through a flexible data table (which provides sorting, page navigation and search capabilities), as well as dynamic field selection and hiding (left margin of the table).

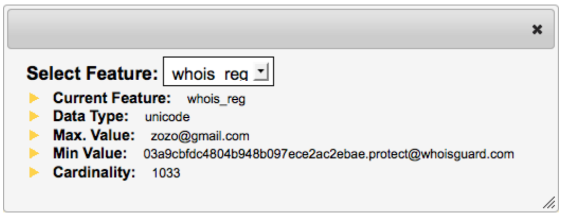

Select Feature: whois_reg

- Current Feature: whois_reg
- Data Type: unicode
- Max. Value: zozo@gmail.com
- Min Value: 03a9cbfdc4804b948b097ece2ac2ebae.protect@whoisguard.com
- Cardinality: 1033

Figure 2.2: Pop-up showing global statistics on the *Rogue AV* dataset.

In the table on top of the *Dataset* page, the user may then select a specific analysis performed on a dataset and click on `check details` to navigate to the page displaying the analysis results.

Site Id	*domain*	*server_ip*	*class_c*	*class_b*	*ns*	*whois_reg*	*whois_rar*	*geo*	*day*
271665	windowsantivirus2008.com	74.54.82.219, 209.62.20.233	74.54.82 209.62.20	74.54 209.62	74.54.82.119	domadmin @privateregistrations.ws	DIRECTI	US	2008-06-04
271621	Xp-2008-Antivirus.com	208.73.210.27, 208.73.210.121	208.73.210	208.73	204.13.161.55, 204.13.160.55	-	-	US	-
272656	malwaredefender2009.com	67.43.237.75, 211.95.73.189	67.43.237 211.95.73	67.43 211.95	208.76.62.100, 75.102.60.66	jsfsl2341@googlemail.com	Regtime Ltd.	CN	2009-03-04
211552	anti-malware-2010.com	74.205.8.7	74.205.8	74.205	216.69.185.2, 208.109.255.2	ANTI-MALWARE-2010.COM @domainsbyproxy.com	GODADDY.COM	US	2009-05-31
122287	antivirus360remover.com	174.132.250.194	174.132.250	174.132	207.218.223.162, 207.218.247.135	ANTIVIRUS360REMOVER.COM @domainsbyproxy.com	GODADDY.COM	US	2009-02-22
272539	norton-antivirus-2010.com	74.208.42.60, 74.208.156.41, 82.165.245.27	74.208.42, 74.208.156, 82.165.245	74.208, 82.165	74.208.3.8, 74.208.2.9	proxy1994891 @1and1-private-registration.com	GODADDY.COM	US	2007-07-08
272540	nortonantivirus2010.com	69.64.145.229, 209.249.222.18, 208.116.34.163	69.64.145 209.249.222 208.116.34	69.64 209.249 208.116	209.249.221.130, 72.34.41.47, 74.81.64.51	support@NameCheap.com	ENOM	US	2006-08-13
334096	home-antivirus2010.com	72.52.210.132	72.52.210	72.52	76.73.35.154, 72.52.210.132	blair@8081.ru	ONLINENIC	US	2009-07-14
334091	homeanti-virus-2010.com	72.52.210.130	72.52.210	72.52	76.73.35.155, 72.52.210.130	blair@8081.ru	ONLINENIC	US	2009-07-14
389838	homeav-2010.com	72.52.210.133	72.52.210	72.52	76.73.35.158, 72.52.210.133	tours@infotorrent.ru	ONLINENIC	US	2009-07-14
465709	pc-anti-spyware-2010	174.139.5.50, 209.31.180.235	174.139.5, 209.31.180	174.139, 209.31	174.139.5.50, 209.31.180.235	argue@8081.ru	ONLINENIC	US	2009-07-29
465706	pc-anti-spy2010.com	174.139.243.45 209.31.180.233	174.139.243 209.31.180	174.139 209.31	174.139.243.45, 209.31.180.233	pixie@ml3.ru	ONLINENIC	US	2009-07-29
465710	p-c-anti-spyware-2010.com	174.139.243.46, 209.31.180.234	174.139.243 209.31.180	174.139 209.31	174.139.243.46, 209.31.180.234	kites@e2mail.ru	ONLINENIC	US	2009-07-29

Table 2.1: Network observables used as *domain features* for a set of rogue AV domains and associated web servers.

2.3 Overview of Analysis Results

Figure 2.3 shows a screenshot of the main *Analysis* page, which enables the user to perform a quick assessment of the clustering results for a given dataset. This page contains three main *panels*:

- **Panel 1**: gives some general details on the analysis, such as the total number of *Clusters* and *MDCs*, the number of *features* used for the analysis, some details on the *data slices* (used for larger data sets that need to be processed in several batches), and finally the `first_seen`, `last_seen` timestamps and total number of events that have already been processed.
- **Panel 2**: visual overview of MDClusters compactness, size, and patterns preview
- **Panel 3**: visual overview of 1D Clusters compactness, size, and patterns preview

The bar chart in Panel 2 is showing the overall MDC compactness, broken down by *feature*) for the top 10 MDCs (which can be adjusted to display more data). As we will describe in more details in Section 2.5, this enables the analyst to perform a quick evaluation of MCDA data fusion results, and identify for each MDC what set of features is mainly responsible the MDC formation. In this analysis, we can also see from this chart that 6 different features have been included in the clustering algorithm: *server_ip*, *class_c*, *class_b*, *domain*, *ns*, *whois_reg.*

However, before going into the details of these *MDClusters*, we will first turn our attention to the analysis of the *1D Clusters*, *i.e.*, Clusters of events that share only one particular feature.

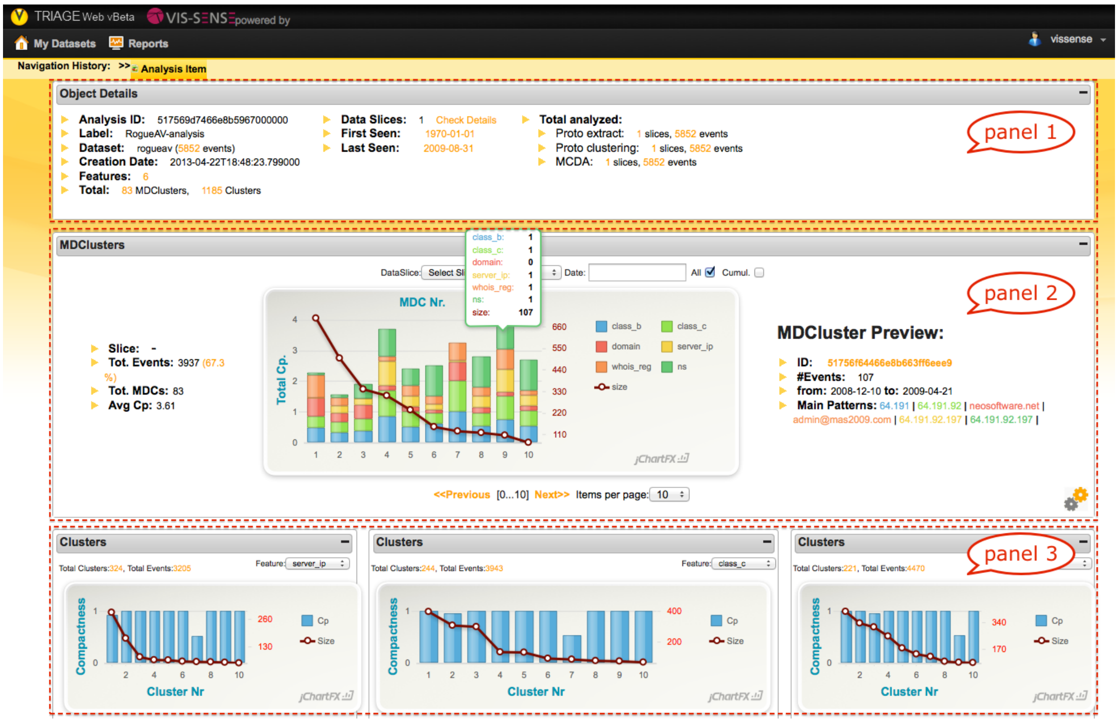

Navigation History: >> Analysis Item

Object Details

- Analysis ID: 517569d7466e8b5967000000
- Label: RogueAV-analysis
- Dataset: rogueav (5852 events)
- Creation Date: 2013-04-22T18:48:23.799000
- Features: 6
- Total: 83 MDClusters, 1185 Clusters

- Data Slices: 1 Check Details
- First Seen: 1970-01-01
- Last Seen: 2009-08-31

- Total analyzed:
- Proto extract: 1 slices, 5852 events
- Proto clustering: 1 slices, 5852 events
- MCDA: 1 slices, 5852 events

MDClusters

DataSlice: Select Sl Date: All ☑ Cumul. □

class_b:	1
class_c:	1
domain:	0
server_ip:	1
whois_reg:	1
ns:	1
size:	107

- Slice: -
- Tot. Events: 3937 (67.3 %)
- Tot. MDCs: 83
- Avg Cp: 3.61

MDCluster Preview:

- ID: 51756f64466e8b663ff6eee9
- #Events: 107
- from: 2008-12-10 to: 2009-04-21
- Main Patterns: 64.191 | 64.191.92 | neosoftware.net | admin@mas2009.com | 64.191.92.197 | 64.191.92.197 |

<<Previous [0...10] Next>> Items per page: 10

Clusters

Total Clusters:324, Total Events:3205 Feature: server_ip

Clusters

Total Clusters:244, Total Events:3943 Feature: class_c

Clusters

Total Clusters:221, Total Events:4470

Figure 2.3: *Analysis* page displaying details on the Rogue AV analysis results: analysis summary (Panel 1), MDClusters (panel 2) and 1D Clusters visual overview (Panel 3).

2.4 Analysis of 1D Clusters

In the lower part of the *Analysis* page (Panel 3 in Figure 2.3), the analyst can browse the one-dimensional *Clusters* that have been identified as a result of step 2 in the TRIAGE processing chain. Note that by default, a similarity metric along with an appropriate threshold will be automatically defined for every feature depending on the data type, unless this choice is overridden by the user. Some examples of default similarity measures implemented in TRIAGEWEB are given in Table 2.2. Quite obviously, the user may always override these default clustering parameters if needed. Figure 2.4 illustrates the typical dialog window in TRIAGEWEB for displaying to the user clustering parameters for a specific feature.

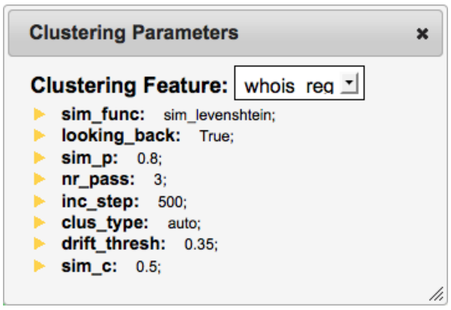

Clustering Feature: whois req

- **sim_func:** sim_levenshtein;
- **looking_back:** True;
- **sim_p:** 0.8;
- **nr_pass:** 3;
- **inc_step:** 500;
- **clus_type:** auto;
- **drift_thresh:** 0.35;
- **sim_c:** 0.5;

Figure 2.4: Clustering parameters for feature *whois_reg*.

As shown in Figure 2.5, the bar charts represent the *Compactness* (or *Cp* for short), which represents in fact the homogeneity of the data Clusters, as well as their total *Size*. This enables the user to do a quick visual assessment of clustering results. As illustration, Figure 2.5 shows the Cp and Size of *domain* (left) and *server_ip* Clusters (right). When the user points the mouse over one of the bar in the chart, a preview of the patterns of the associated Cluster is being displayed in a tooltip, along with a link to a page that contains the Cluster details.

These visual charts have enabled us to assess clustering results and to compare the results obtained with TRIAGEWEB to our previous results obtained in other analyses []. We observe that the number of Clusters, their Cp and average size are very close to the ones obtained before. However, the TRIAGEWEB interface now makes it much easier to analyze results and allows us to get insights much faster then before into this particular security dataset.

Next to the visual charts, the user interface provides also a dynamic *Clusters table*,

Table 2.2: Default similarity metrics used in TRIAGEWEB.

Data type	**Default similarity metric**	**Description**
list/set	*sim_tanimoto*	based on **Jaccard** index
string/unicode	*sim_levenshtein*	normalized **edit** distance
binary/categorical	*sim_generic*	binary comparison (1 or 0)
IP address	*sim_IP*	normalized diff. between two IP's (in decimal format)
email	*sim_Ngram*	**N-gram** similarity
date	*sim_date*	normalized diff. between 2 dates
datetime	*sim_datetime*	normalized diff. between 2 timestamps
distribution	*sim_jsdiv*	normalized similarity based on **Jensen-Shannon** divergence
dictionary	*sim_dict*	average Jaccard index
fuzzy hash	*sim_ssdeep*	based on **ssdeep** algorithm
2D coordinates	*sim_coords*	normalized Euclidean distance

in which every row displays details of a given 1D Cluster, as illustrated in Figure 2.6. Again, the table provides dynamic functionalities such as sorting, page navigation and searching for specific patterns. Two different drop-down lists enable the user to select a particular dataset feature (or *attribute*) to quickly switch to another set of Clusters (*e.g.*, switching from server_ip Clusters to domain Clusters). This enables the user to change perspective very easily and focus on particular groups of Events. The last column (named "pattern") will be updated automatically to give a preview of Clusters patterns regarding the selected feature. When a value is entered in the *search* filter, the entered pattern will then automatically be applied to the selected feature.

The user may want to dive into the details of a specific Cluster starting from either the Clusters visual charts or from the dynamic Clusters table, using the hyperlinks provided in each perspective. Figure 2.7 shows an example of page displaying the details a domain Cluster, which is automatically built when the user clicks on the associated Cluster identifier. In this Figure, the left frame displays some general details of the Cluster (*e.g.*, number of prototypes, total number of events, `first_seen` and `last_seen` timestamps, etc). In the right frame, a dynamic plotting tool enables the user to inspect the patterns of the Cluster. By default, the analysis feature associated to that Cluster will be selected and displayed in the chart. In the Cluster example of Figure 2.7, we observe that many rogue domains grouped in this Cluster have in common the same naming convention regarding the *Whois* registrant name: `[a-z]*@whoisprivacyprotect.com`.

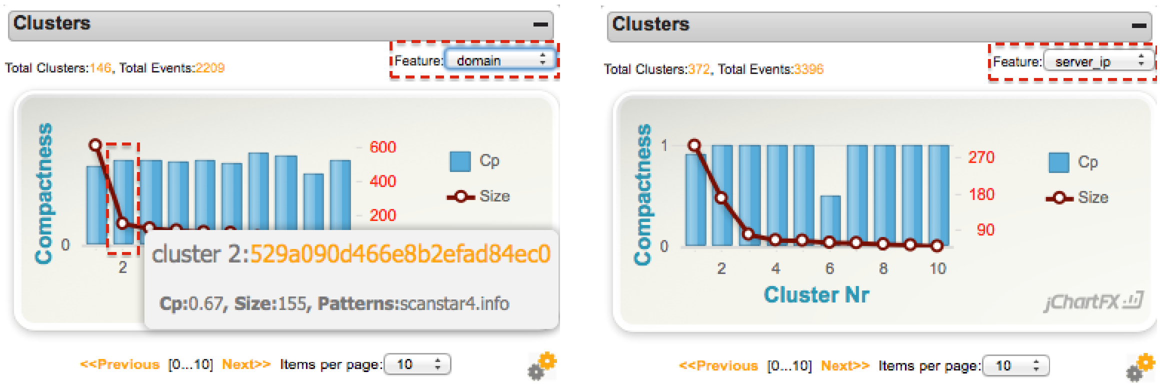

Total Clusters:146, Total Events:2209

Feature: domain

cluster 2:529a090d466e8b2efad84ec0

Cp:0.67, Size:155, Patterns:scanstar4.info

<<Previous [0...10] Next>> Items per page: 10

Total Clusters:372, Total Events:3396

Feature: server_ip

<<Previous [0...10] Next>> Items per page: 10

(a) *domain* Clusters.

(b) *server_ip* Clusters.

Figure 2.5: *Rogue AV* analysis: Bar charts showing Clusters *Compactness* and *Size*, enabling quick assessment of clustering results.

However, the analyst may change this view in order to analyse the distribution of Cluster events with respect to another feature. Actually, every feature can be seen as a viewpoint giving a certain *perspective* on a particular group of Events, which in turn can reveal interesting insights regarding the modus operandi of the miscreants or attackers. For example, we can look at the same *whois_reg* Cluster through the *server_ip* dimension in order to find out which server IP addresses were apparently associated to the rogue domains registered by these POC email addresses having apparently a similar naming scheme. This is illustrated in Figure 2.8 where we can observe the result of grouping the rogue domains of this particular *whois_reg* Cluster by server IP address. Three main server IP's (belonging to the same Class B-network) were apparently hosting the rogue domains associated to these registrants. By repeating this process for other Clusters, the analyst may quickly gain interesting insights into the dataset by simply inspecting various subsets of dataset entities across different dimensions.

Note also that the Cluster dynamic plotting tool provides some other controls to deal with possible *scalability* issues, such as:

- a *slider* for adjusting the *percentile* threshold on-the-fly, which enables the user to focus either on the top values of the distribution, or to look at the less frequent values as well (distribution tail)

- a selector for the *granularity* level, by which the user may choose between the *prototype* or *event* level of detail (where the *prototype* level may be more useful for

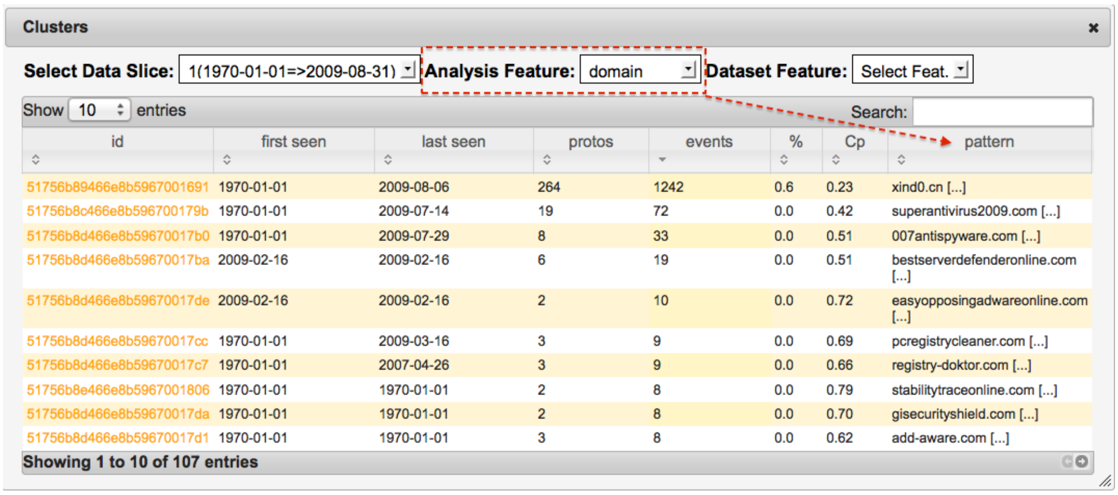

Clusters

Select Data Slice: 1(1970-01-01=>2009-08-31) Analysis Feature: domain Dataset Feature: Select Feat.

Show 10 entries Search:

id	first seen	last seen	protos	events	%	Cp	pattern
51756b89466e8b5967001691	1970-01-01	2009-08-06	264	1242	0.6	0.23	xind0.cn [...]
51756b8c466e8b596700179b	1970-01-01	2009-07-14	19	72	0.0	0.42	superantivirus2009.com [...]
51756b8d466e8b59670017b0	1970-01-01	2009-07-29	8	33	0.0	0.51	007antispyware.com [...]
51756b8d466e8b59670017ba	2009-02-16	2009-02-16	6	19	0.0	0.51	bestserverdefenderonline.com [...]
51756b8d466e8b59670017de	2009-02-16	2009-02-16	2	10	0.0	0.72	easyopposingadwareonline.com [...]
51756b8d466e8b59670017cc	1970-01-01	2009-03-16	3	9	0.0	0.69	pcregistrycleaner.com [...]
51756b8d466e8b59670017c7	1970-01-01	2007-04-26	3	9	0.0	0.66	registry-doktor.com [...]
51756b8e466e8b5967001806	1970-01-01	1970-01-01	2	8	0.0	0.79	stabilitytraceonline.com [...]
51756b8d466e8b59670017da	1970-01-01	1970-01-01	2	8	0.0	0.70	gisecurityshield.com [...]
51756b8d466e8b59670017d1	1970-01-01	1970-01-01	3	8	0.0	0.62	add-aware.com [...]

Showing 1 to 10 of 107 entries

Figure 2.6: Dynamic data table showing Clusters overview

larger Clusters)

- a selector for choosing the *data slice* (when applicable), as Clusters may span across multiple data slices (for larger datasets)
- a selector for choosing to look at *cumulative* patterns when the Cluster is spread over multiple slices

These controls, combined with the underlying TRIAGE data model, ensure a good responsiveness of the user interface as well as effective analysis capabilities even when dealing with large data sets, as the user may then focus on the important patterns first and ask for more details only when this is needed.

In conclusion, we observe that this *Cluster analysis* of the *rogue AV* dataset made through TRIAGEWEB has provided interesting results that are consistent with previous analyses of the same dataset [10, 17]. The key advantage of doing this cluster analysis using TRIAGEWEB is obviously the significant gain in time and efforts required to obtain the same insights. As demonstrated in the next Section, the rich and informative MDC visualizations integrated in TRIAGEWEB provide even more insights for the analyst, also in much less time than with previous analyses.

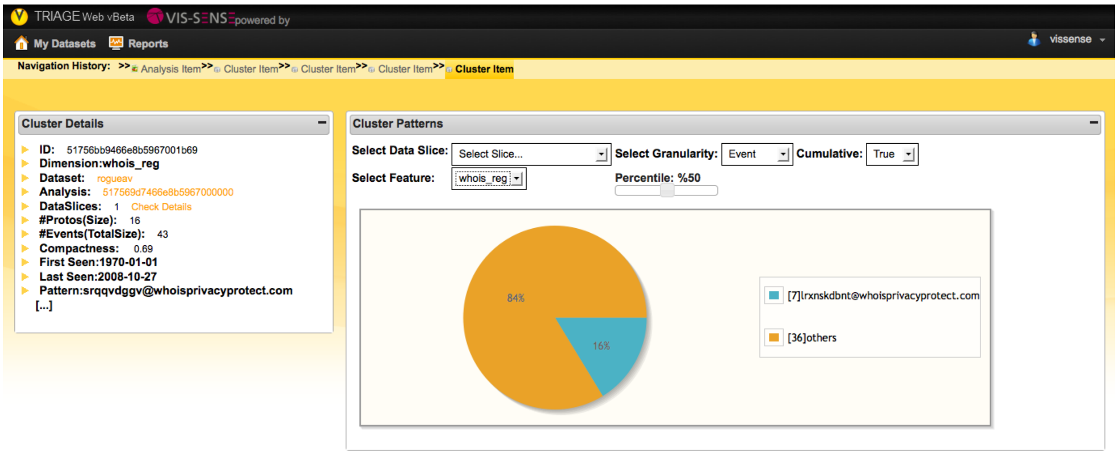

Navigation History: >> Analysis Item>> Cluster Item>> Cluster Item>> Cluster Item>> Cluster Item

Cluster Details

- ID: 51756bb9466e8b5967001b69
- Dimension:whois_reg
- Dataset: rogueav
- Analysis: 517569d7466e8b5967000000
- DataSlices: 1 Check Details
- #Protos(Size): 16
- #Events(TotalSize): 43
- Compactness: 0.69
- First Seen:1970-01-01
- Last Seen:2008-10-27
- Pattern:srqqvdggv@whoisprivacyprotect.com
  [...]

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True
Select Feature: whois_reg Percentile: %50

Figure 2.7: Cluster page showing the details of a specific *whois_reg* Cluster.

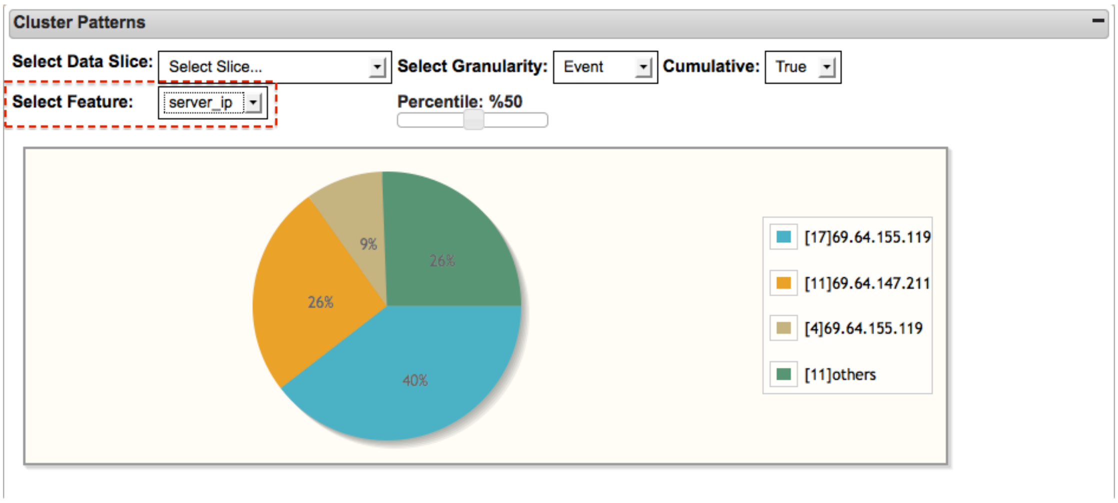

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True
Select Feature: server_ip Percentile: %50

Figure 2.8: Cross-feature pattern analysis of the *whois_reg* Cluster of Fig. 2.7, viewed here wrt *server_ip*.

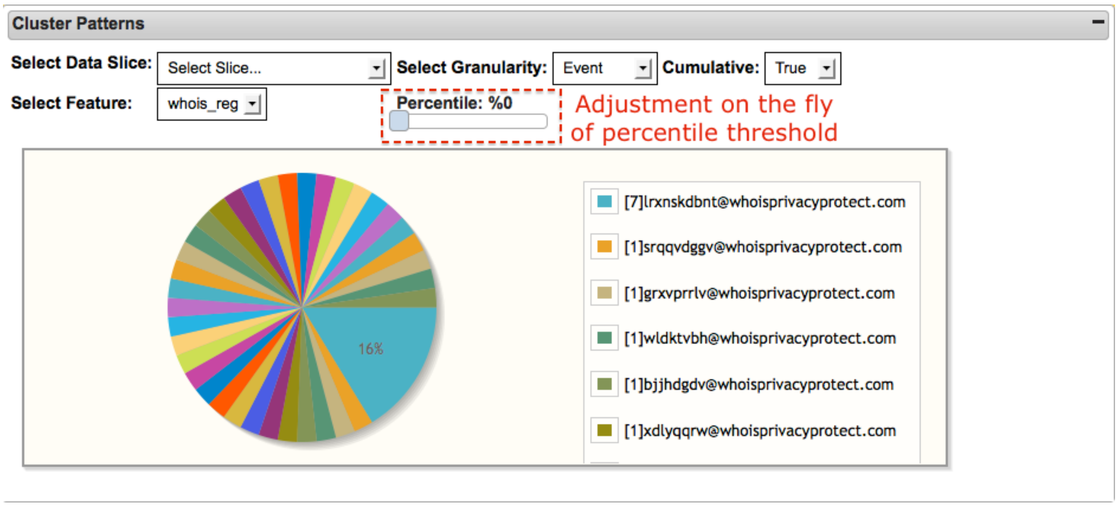

Cluster Patterns

Select Data Slice: Select Slice...

Select Granularity: Event

Cumulative: True

Select Feature: whois_reg

Percentile: %0

Adjustment on the fly of percentile threshold

16%
[7]lrxnskdbnt@whoisprivacyprotect.com
[1]srqqvdggv@whoisprivacyprotect.com
[1]grxvprrlv@whoisprivacyprotect.com
[1]wldktvbh@whoisprivacyprotect.com
[1]bjjhdgdv@whoisprivacyprotect.com
[1]xdlyqqrw@whoisprivacyprotect.com

Figure 2.9: Adjustment of the percentile threshold for more fine-grained pattern analysis.

2.5 Multi-Criteria Analysis Results

We have seen in previous Section that 1D Clusters of rogue domains may provide interesting patterns and reveal how these rogue websites are created and managed. In fact, each feature can be seen as a viewpoint giving a certain *perspective* on the underlying phenomenon, which in turn can highlight interesting aspects of the modus operandi of the miscreants.

However, the fact that groups of rogue domains are clustered w.r.t. a given aspect does not necessarily mean that these websites are linked to the same rogue AV campaign or the same group of individuals. Only one common feature can be merely due to a coincidence, to a commonly-seen pattern, or to a commonly-used technique.

To identify groups of rogue domains that are likely associated to the same campaign, in a more systematic and reliable manner, certain clusters are likely to be merged whereas some others may have to be split. To aid the analyst make such a decision, the multi-criteria aggregation component of TRIAGE may be used to effectively combine all these viewpoints, such that the final result reflects the expectations of the expert regarding the combination of features that must be satisfied in order to attribute two rogue domains to a specific campaign.

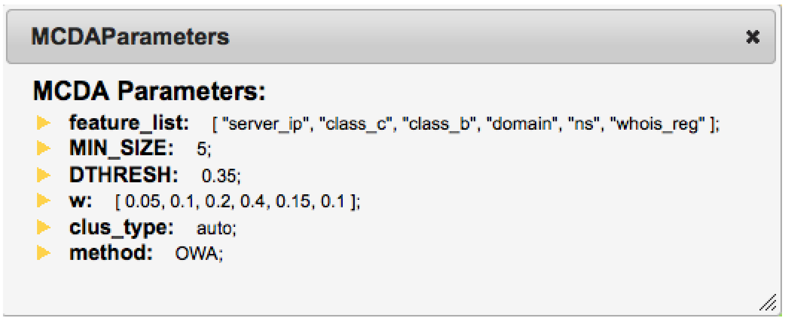

MCDAParameters

MCDA Parameters:

- **feature_list:** ["server_ip", "class_c", "class_b", "domain", "ns", "whois_reg"];
- **MIN_SIZE:** 5;
- **DTHRESH:** 0.35;
- **w:** [0.05, 0.1, 0.2, 0.4, 0.15, 0.1];
- **clus_type:** auto;
- **method:** OWA;

Figure 2.10: Window displaying OWA parameters for the *rogue AV* dataset

We have thus used TRIAGEWEB to perform the MCDA aggregation of rogue domains correlations using a data fusion model similar to the one defined previously in [10]. That is, we have used two different Ordered Weighted Averaging functions (OWA and Weighted OWA) as aggregation means, using the very similar weights and thresholds. Figure **??** represents the TRIAGEWEB window that displays the defined MCDA parameters for the OWA analysis. Note that the TRIAGEWEB interface is flexible in this regard, as the analyst can always decide to create a new analysis for the same dataset but using

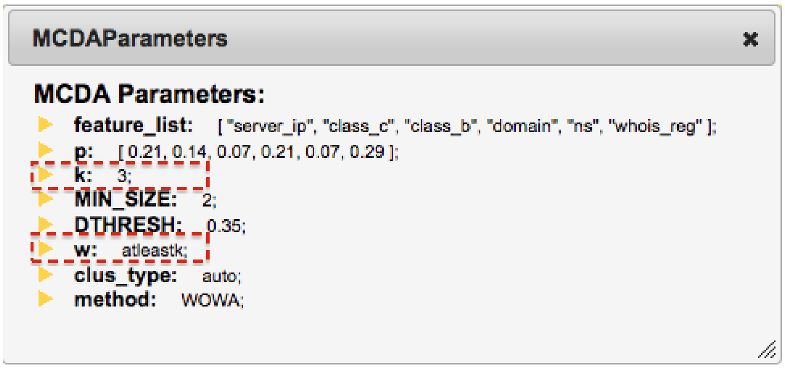

MCDA Parameters:

- feature_list: ["server_ip", "class_c", "class_b", "domain", "ns", "whois_reg"];
- p: [0.21, 0.14, 0.07, 0.21, 0.07, 0.29];
- k: 3;
- MIN_SIZE: 2;
- DTHRESH: 0.35;
- w: atleastk;
- clus_type: auto;
- method: WOWA;

Figure 2.11: Window displaying WOWA parameters for the *rogue AV* dataset. Some fuzzy quantifiers (*at least k*) facilitate the definition of weighting vectors by a human expert.

a different set of parameters, in order to facilitate the comparison of the results based on different data fusion models.

The weighting vector w is used in OWA to reflect some vague statement regarding the number of strong correlations that is required to generate a high aggregate value. By defining w like in Figure 2.10, we give more importance to strong correlations starting from the third highest position. In other words, it means that the two highest scores will have lower weights (0.05 and 0.10 respectively) and thus at least three strong correlations will be needed in order to have a global score above 0.35 (the sum of weights given to the first three higher similarity values).

As explained in previous deliverables (more specifically in D3.3 - Attack Attribution Module), we have improved the usability of the TRIAGE data fusion algorithms by providing various methods to help the analyst define the weighting vectors used in OWA and Weighted OWA aggregation operators. This is illustrated in Figure 2.11, where we can see that the data fusion model is defined more broadly as "*at least k*" correlations (with $k = 3$). Besides the data fusion model, the analyst must only be able to provide a broad quantification of the importances of individual features (*e.g.*, on a categorical scale going from "very high", "high", "neutral", "low", "very low"). The TRIAGEWEB interface will subsequently adjust the weighting vectors in order to reflect these importances together with the data fusion model that was defined (the result of which is reflected in vectors p and w as shown in Figure 2.11). We believe that it is indeed much easier for a human expert to express his domain knowledge using this type of *linguistic quantifiers* rather than providing precise numerical values. Note that if the analyst defines synergies or redundancies among pairs of features, the data fusion model will automatically be set to *Choquet* integral (which represents the most flexible but also the more complex aggregation function), and some helper methods will then define or adjust the weighting vectors and fuzzy measures accordingly.

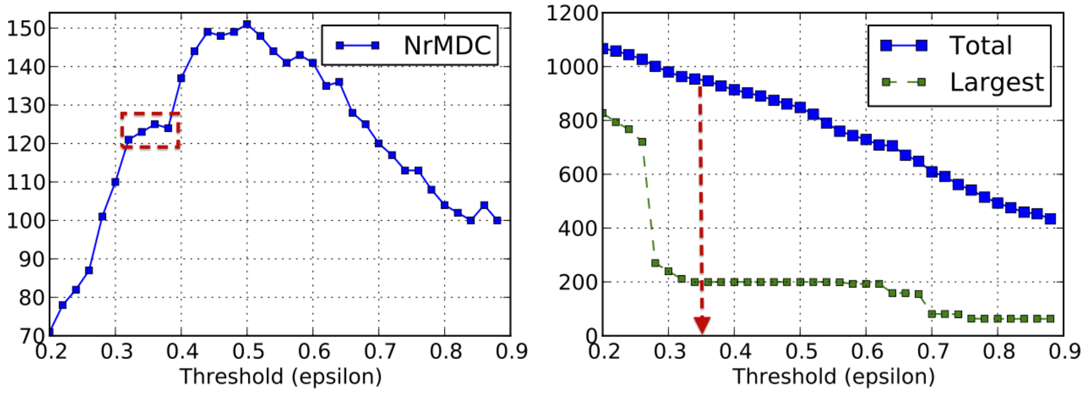

Figure 2.12: Sensitivity analysis to determine the most appropriate decision threshold in the MCDA fusion step.

As described in other deliverables and also in previous work like [18], another important parameter to define is the **decision threshold** ϵ (denoted by DTHRESH in the user interface). This threshold is used to eliminate irrelevant links in the combined graph that is built as a result of the MCDA fusion, so as to avoid unwanted or inconsistent linkage between dataset entities. It is thus recommended to let the system perform a *sensitivity analysis*, *i.e.*, to let ϵ increase from a very low to high value, and then observe the number of components (or MDCs) that are identified in the resulting graph, as well as the size of the largest MDC. This is illustrated in Fig. 2.12 where we can observe the impact of the threshold on the MCDA clustering results (similar diagrams can be obtained for the WOWA aggregation). We can observe different regions of interest in these plots.

In the first region of the plot, irrelevant edges having low weight in the aggregated graph will be removed, and large connected sub-graphs start to split into a number of more consistent sub-graphs.

In the second region, we can observe a first *plateau* between the values $\epsilon = 0.32$ and

0.38, which seems to indicate some reasonable starting point as values for the decision threshold (because the number of MDClusters becomes stable). At the same time, we should always consider the size of the largest MDC, which appears to be also pretty stable in this threshold interval.

Increasing further ϵ up to an excessive value can lead to a significant loss of nodes and edges in the aggregated graph (as isolated nodes will be disregarded by the graph clustering algorithm), which means that we may also loose too much information. We observe that an appropriate range of values for ϵ lies usually around k/n, with k the minimum amount of correlated features desired by the analyst to link two events. For example, if the analyst decides to define a model with at least 3 strong correlations, when the total number of features included in the analysis is 6, then a good starting point for the decision threshold will be 0.50. In our analysis, we have used a threshold of 0.35 (within the first plateau) to be consistent with the previous analyses made in [10, 18].

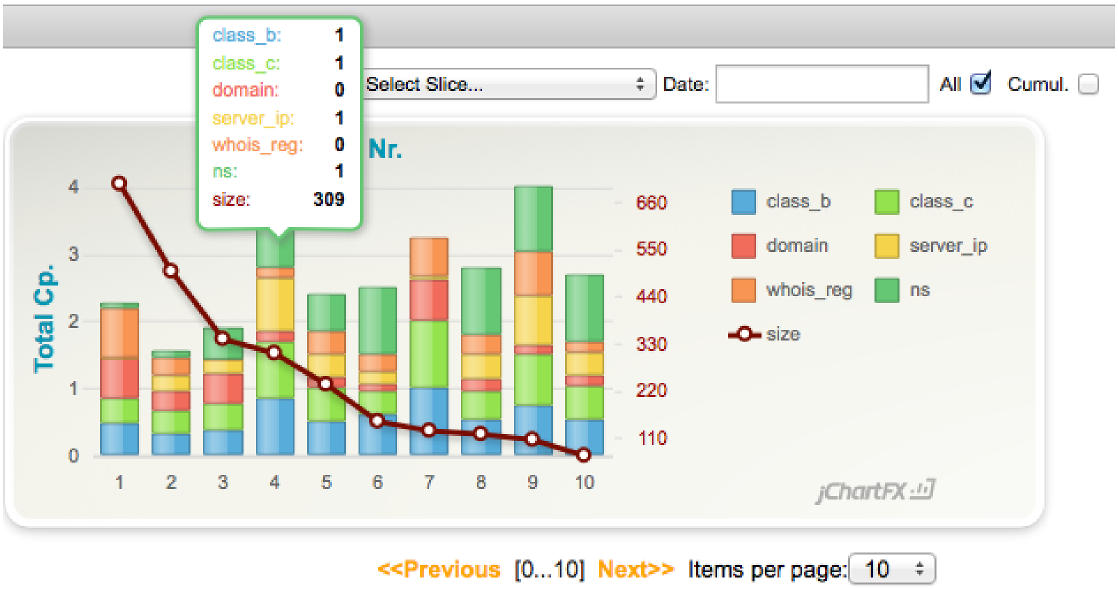

Select Slice... Date: All ☑ Cumul. □

<<Previous [0...10] Next>> Items per page: 10

Figure 2.13: *Rogue AV* analysis: stacked bar chart showing the total MDC compactness (broken down by *feature*) for the top 10 MDCs – enabling a quick evaluation of MCDA data fusion results.

2.6 Insights into Rogue AV Campaigns

We now turn our attention to the analysis of MDCs results. To quickly evaluate their consistency, the TRIAGEWEB interface provides the analyst with a visual stacked bar chart that displays the overall Compactness and Size of MDClusters (Figure 2.13). Controls are provided for quick navigation through MDC results (*e.g.*, 'next" and "previous", number of items per page, etc), and a *tooltip* showing the detailed values is displayed on top of each bar whenever the user points the mouse cursor on it. In Figure 2.13, we see that the average *Cp* value of an MDC is visually displayed with respect to each feature using a color map, which makes it easy to understand which features tend to correlate events (in general), but also what are the main reasons behind the identification of certain MDCs. Note that the average *Cp* value across all MDCs lies around 3.43, which is in adequacy with the MCDA model defined previously (*at least 3* strong correlations).

From Figure 2.13 we observe that most MDCs have globally a high compactness value (consistent with the data fusion model), except for a few ones (such as MDC_1 and MDC_3). A deeper inspection of those MDCs reveals that these are quite large connected components, which explains why they are less compact since they are made of several loosely connected subgraphs (probably reflecting larger groups of miscreants who are exchanging tools or using the same infrastructure).

We note also that IP-related features contribute the most to the correlation of rogue domains. In many cases, *whois_reg* seems to complete or reinforce those correlations. It is also interesting to see that some MDCs are correlated by *class_c* and *class_b*, but not by *server_ip* (such as MDC_7), which justifies the selection of those features.

MDCs have in general lower C_p values with respect to *domain* names, but we found a few MDCs in which domain name correlation plays a significant role (like for MDC_1, MDC_3 and MDC_7). Finally, we observe that every MDC is characterized by varying degrees of correlation regarding each feature, but overall there are always *at least three* different features having a high correlation (except maybe for the three largest MDCs).

Case study: Example of Rogue AV campaign

We now present a more in-depth analysis of an illustrative case study, in order to show the kind of visual insights we can get into the behavior of *Rogue AV campaigns*, which were identified, in a systematic and automated manner, by the TRIAGE multi-criteria clustering technique.

As illustrated in Figure 2.14, MDClusters can be visualized and explored via the TRIAGEWEB interface through a convenient dynamic table, which again provides *sorting* (on any column), page *navigation* and *searching* capabilities. Like for the Clusters table, two different drop-down lists (*Analysis* and *Dataset feature* DDLs) enable the user to select a particular feature (or *attribute*). However, the behavior of the MDC table is slightly different than for 1D Clusters: selecting one or another feature will only update the last table column (named "Feature val.") in which a preview of MDC patterns will be displayed regarding the selected feature. This enables the analyst to have a quick overview of MDC patterns. When a value is entered in the *search* filter, the entered pattern will also automatically be applied to the selected feature, which can be used to find a particular MDC.

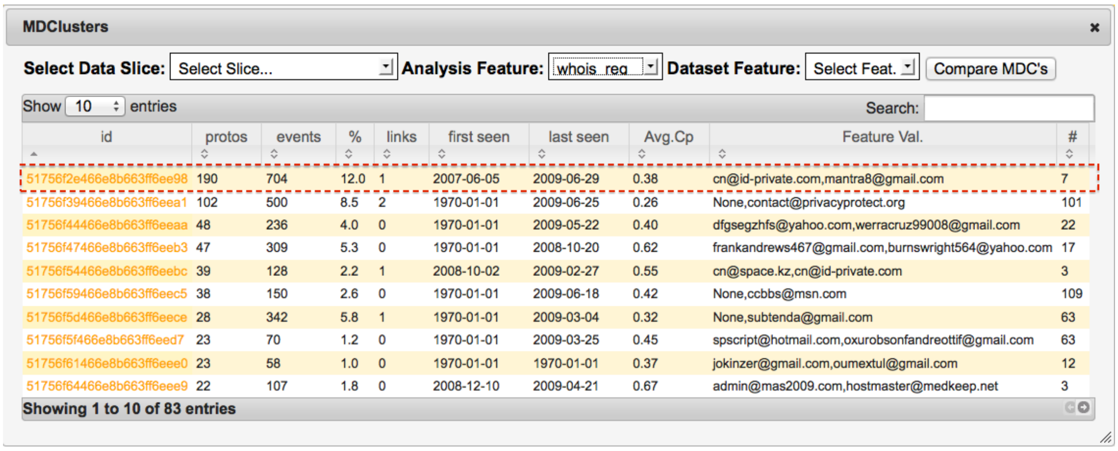

id	protos	events	%	links	first seen	last seen	Avg.Cp	Feature Val.	#
51756f2e466e8b663ff6ee98	190	704	12.0	1	2007-06-05	2009-06-29	0.38	cn@id-private.com,mantra8@gmail.com	7
51756f39466e8b663ff6eea1	102	500	8.5	2	1970-01-01	2009-06-25	0.26	None,contact@privacyprotect.org	101
51756f44466e8b663ff6eeaa	48	236	4.0	0	1970-01-01	2009-05-22	0.40	dfgsegzhfs@yahoo.com,werracruz99008@gmail.com	22
51756f47466e8b663ff6eeb3	47	309	5.3	0	1970-01-01	2008-10-20	0.62	frankandrews467@gmail.com,burnswright564@yahoo.com	17
51756f54466e8b663ff6eebc	39	128	2.2	1	2008-10-02	2009-02-27	0.55	cn@space.kz,cn@id-private.com	3
51756f59466e8b663ff6eec5	38	150	2.6	0	1970-01-01	2009-06-18	0.42	None,ccbbs@msn.com	109
51756f5d466e8b663ff6eece	28	342	5.8	1	1970-01-01	2009-03-04	0.32	None,subtenda@gmail.com	63
51756f5f466e8b663ff6eed7	23	70	1.2	0	1970-01-01	2009-03-25	0.45	spscript@hotmail.com,oxurobsonfandreottif@gmail.com	63
51756f61466e8b663ff6eee0	23	58	1.0	0	1970-01-01	1970-01-01	0.37	jokinzer@gmail.com,oumextul@gmail.com	12
51756f64466e8b663ff6eee9	22	107	1.8	0	2008-12-10	2009-04-21	0.67	admin@mas2009.com,hostmaster@medkeep.net	3

Showing 1 to 10 of 83 entries

Figure 2.14: Dynamic data table showing MDClusters overview

Let us focus now on a particular MDCluster, as highlighted in the MDC table shown in Figure 2.14. We will refer to it as "MDC 1", as this MDC seems to represent the largest MDC (and thus also the largest campaign). About 700 rogue domains have been grouped in this MDC, accounting for 12% of the dataset. To understand the reasons why these rogue AV domains have been grouped, the analyst may open the associated MDC page by clicking on the hyperlink built on its identifier. A screenshot of the page showing the details of MDC 1 is shown in Figure 2.15, where we can see three main parts:

- **Panel 1**: overview of MDC attributes (size, average Cp, first_seen, last_seen, etc)
- **Panel 2**: visual chart showing the distribution of *correlation combinations*
- **Panel 3**: dynamic plotting tool, to let the user inspect MDC patterns

For MDC 1, from the pie chart in Panel 2 we can deduce that the main reasons behind the formation of this cluster is a combination of *class_c*, *class_b*, *domain* name and *whois_reg* (for 56% of the rogue domains grouped in this MDC). Another visualization of *individual Cp values* (*i.e.*, the average compactness of the MDC by feature) can be obtained by clicking on the check details link next to the overall average Cp value in Panel 1. As illustrated in Figure 2.16, the analysis of these individual Cp values enables the user to understand which features are the most correlative ones, ans thus which characteristics of the rogue domains are linking them together within the MDC. For MDC 1, we clearly see that most of the rogue domains are linked by the same (or very similar) *whois_reg*, *domain*, *class_c* and *class_b* networks.

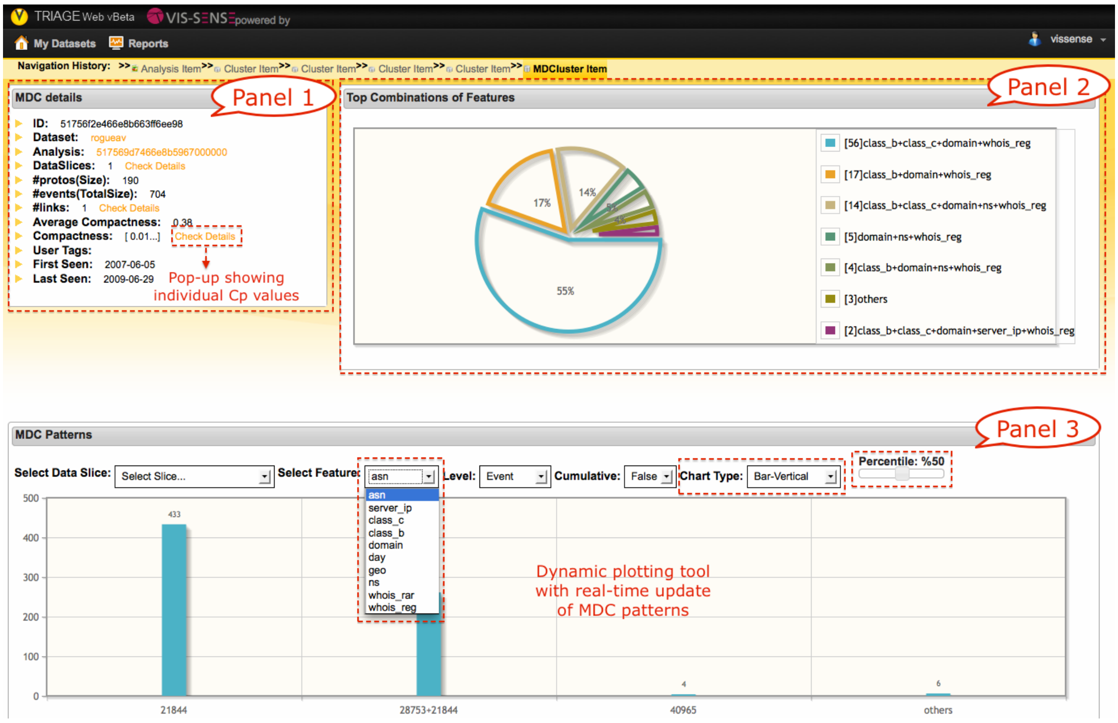

Navigation History: >> Analysis Item >> Cluster Item >> Cluster Item >> Cluster Item >> Cluster Item >> MDCluster Item

MDC details

- ID: 51756f2e466e8b663ff6ee98
- Dataset: rogueav
- Analysis: 517569d7466e8b5967000000
- DataSlices: 1 Check Details
- #protos(Size): 190
- #events(TotalSize): 704
- #links: 1 Check Details
- Average Compactness: 0.38
- Compactness: [0.01...] Check Details
- User Tags:
- First Seen: 2007-06-05
- Last Seen: 2009-06-29

Pop-up showing individual Cp values

Top Combinations of Features

- 🟦 [56]class_b+class_c+domain+whois_reg
- 🟧 [17]class_b+domain+whois_reg
- ▨ [14]class_b+class_c+domain+ns+whois_reg
- 🟩 [5]domain+ns+whois_reg
- 🟩 [4]class_b+domain+ns+whois_reg
- ▨ [3]others
- 🟪 [2]class_b+class_c+domain+server_ip+whois_reg

MDC Patterns

Select Data Slice: Select Slice... Select Feature: asn Level: Event Cumulative: False Chart Type: Bar-Vertical Percentile: %50

Dynamic plotting tool with real-time update of MDC patterns

Figure 2.15: *MDCluster* page showing the details of a specific MDCluster.

MDC patterns and correlations can be further inspected through the dynamic plotting tool of Panel 3 – as illustrated in Figure 2.17. Some flexible user controls are also provided with this dynamic plotting tool, such as:

- *data slice* selection (when applicable for larger data sets that are processed in different batches)
- drop-down list for dataset *feature selection*
- selection of *granularity* level (prototype or event level)
- selector for the *cumulative* flag (for MDCs that spread across multiple data slices)
- selector for the *orientation* of the chart (vertical or horizontal display), which can be useful to accommodate with long-tailed distributions
- slider control for controlling the number of elements being displayed via the *percentile* threshold

From Figure 2.17, we can now understand the reasons why MDC 1 has a high Cp value for the whois_reg feature, as it appears that a large majority of the domains have been registered using the very same POC email address (`cn@id-private.com`). However, it is interesting to note that some other email accounts have also been used to create domains apparently linked to the same campaign.

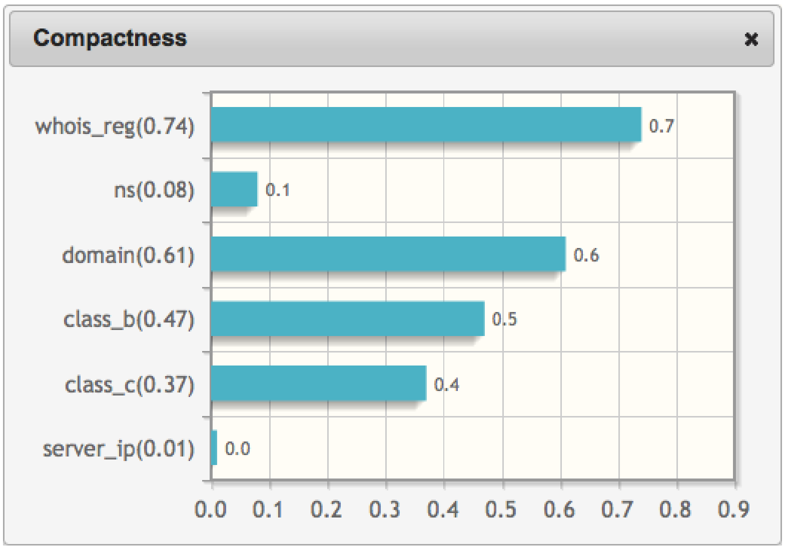

Figure 2.16: Analysis of individual Compactness values of an MDC (for better understanding of entity correlations within an MDC).

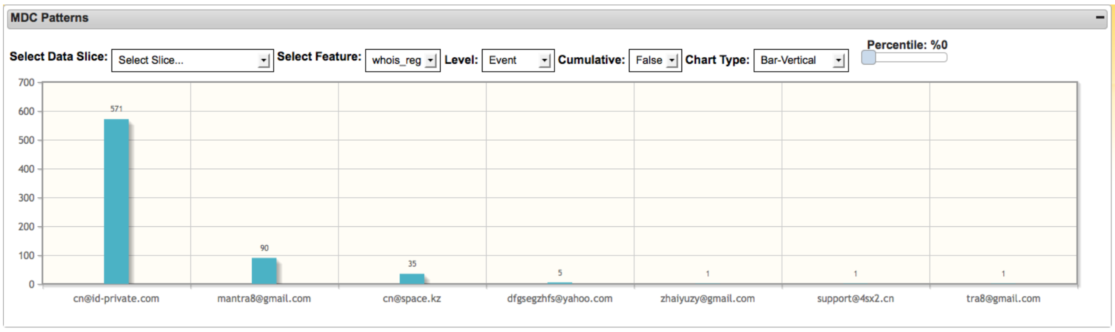

Select Data Slice: Select Slice...
Select Feature: whois_reg
Level: Event
Cumulative: False
Chart Type: Bar-Vertical
Percentile: %0

Figure 2.17: Analysis of MDC patterns via dynamic chart tool.

In order to understand now how all these domains are connected to each other, the MDC page provides different visualization tools that allow the analyst to explore interactively the MDC entities and their relationships. The first one is the interactive graph viewer, which is used to create a "big picture" of the campaign using a node-link diagram. As illustrated in Figure 2.18, we can now understand the TTP (tactics, techniques and procedures) of the miscreants behind this rogue AV campaign. About 700 rogue domains have been registered in the **.cn** top-level domain (resolving to 135 IP addresses in 14 subnets), on eight specific dates over a span of eight months. In fact, most domains were created using the same *Whois* email address (shown in red in the middle of the visualization) on only two consecutive days (nodes in purple on the timeline below the graph). While the same Chinese registrar (Era of the Internet Technology) was used by these cyber-crooks for the registration of all domain names, we can now observe that the campaign was initiated long before the two main dates on which the bulk of the campaign was done. Thanks to this graph diagram, the analyst can now click on different nodes to follow the different links. We can also observe common patterns in the domain names which are composed of exactly 5 alphanumeric characters, apparently chosen in a random fashion, which indicates the use of automated tools to create those domains. The interactive graph viewer provides various controls to make the exploration and navigation even easier:

- a control for zooming in and out
- a control for displaying or hiding edges, to avoid information overload in large MDCs
- a fish-eye lens control for increasing the node labels and thus improve their read-

ability

- a control for enabling the node rearrangement mode (*i.e.*, let the user move around nodes to improve the graph layout)
- a control for changing the edge thickness scale
- a control for changing the node size scale
- a control for changing the label text size scale
- a control to toggle between curved and straight edges
- a control for exporting the graph in various image formats

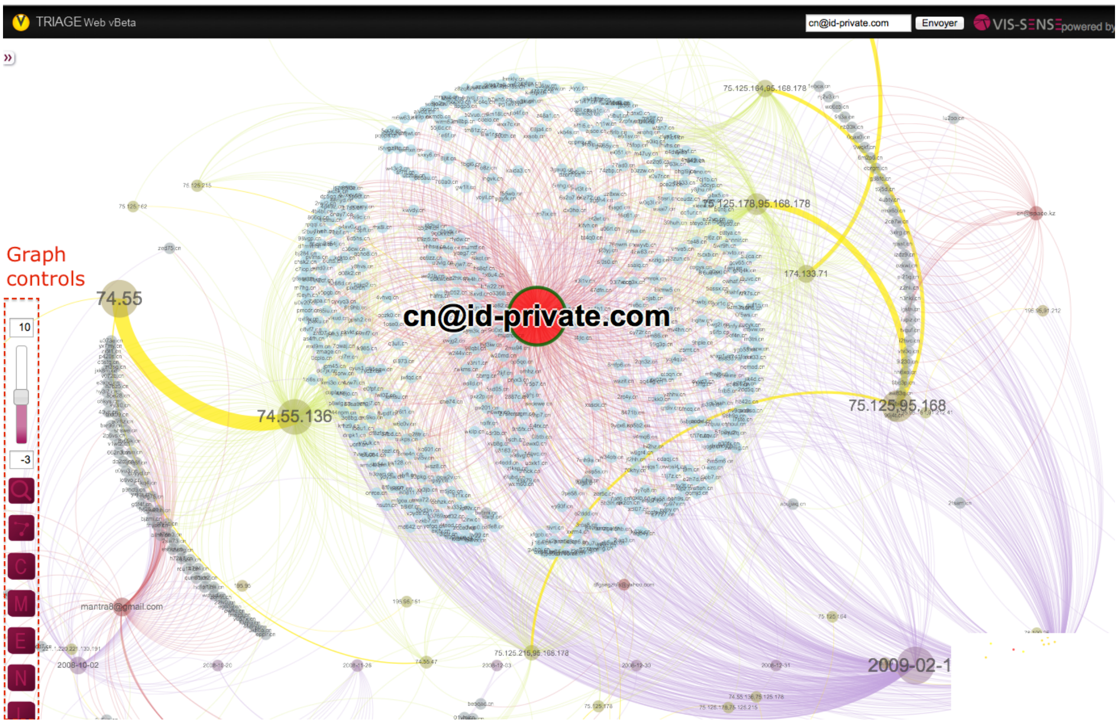

Figure 2.18: Visual analysis of an MDC through the interactive **Graph** viewer.

The MDC page of the TRIAGEWEB interface integrates several MDC visualizations developed by VIS-SENSE partners, which provides different perspectives of the clustered data. As described in other deliverables (mainly in D4.1, 4.2, 4.3 and 6.3), each visualization has different strengths or advantages in terms of scalability, ability to represent events relationships, its compactness and readability. Hence, the analyst may use the most appropriate visualization according to the MDC structure and size.

As an example, we have visualized MDC 1 using two other representations:

- Figure 2.19 visualizes all MDC events using a space-filling **Treemap**-based representation in which relationships between different feature patterns are represented with splines overlaid on the Treemap.

- in Figure 2.20, all MDC events are visualized using a so-called **Chord** diagram, in which the different feature patterns are laid on a circular basis and represented by different colors, and *ribbons* of different sizes are being drawn to represent the various relationships between these feature values (the size of the ribbon is proportional to the frequency of the association between two values).

In both visualizations, we use the 1D Clusters to represent the patterns of each dimension within the same MDC. The two visualizations provide also some *interaction* capabilities: by going over a particular feature value, both the *Treemap* and the *Chord* diagram will be updated on-the-fly to display only the relations associated to a specific pattern – as represented in Figure 2.20, where we can see the relationships between the most frequent *whois_reg* value (`cn@id-private.com`) and all other dimensions.

Furthermore, a dynamic *feature selection* tool enables the user to enable or disable certain features within the visualizations. The goal is to decrease the information overload in large MDCs, and let the analyst focus on particular aspects of an MDCluster when required. This is illustrated in Figure 2.21 in which we visualize MDC 1 using the Chord viewer and where the features *server_ip* and *ns* have been disabled to focus on the relationships between a particular *whois_reg* pattern and the rest of the diagram.

From our own experience, we note that the space-filling *Treemap* representation provides higher scalability and is better suited for generating an overall structural overview of an MDC, especially for large, highly variable MDCs. The *Graph* node-link representation is probably the less scalable visualization; however, it is also the most flexible one for exploring and understanding nodes relationships, as well as to discover "similarity paths" between two events. Finally, the *Chord* diagram provides a good trade-off between scalability and exploration of relationships within an MDC.

The last visualization component provides a classical geographical map in which the MDC events are represented according to their origin. This obviously requires the user to set in TRIAGEWEB one of the dataset features as the "country feature". Figure 2.22

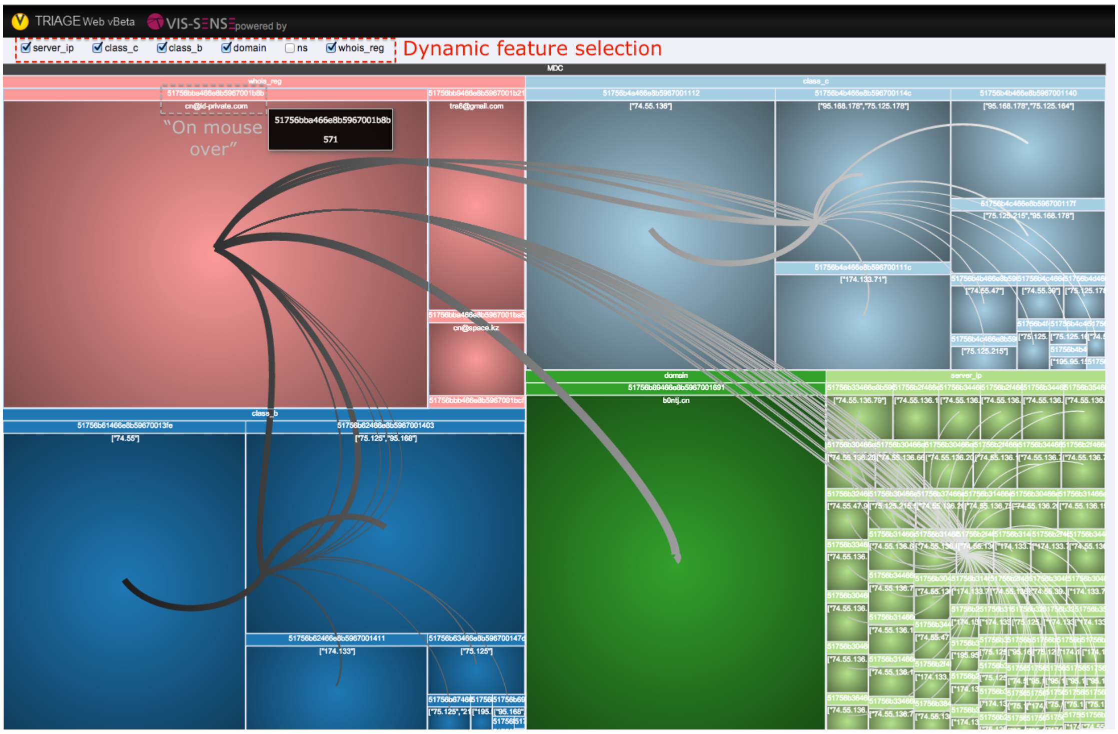

Figure 2.19: Visual analysis of an MDC through the interactive **Treemap**-based viewer.

represents the country origins of the rogue domains of MDC 1, as mapped from their associated *server_ip* address.

Fine-grained analysis of MDC events

Finally, the analyst may go one step further in the low-level analysis of an MDC by opening a dynamic table showing all events and their attributes – as illustrated in Figure **??**. A dynamic field selection tool lets the analyst select which features to be displayed in the data table. Like for all other data tables, the data can be sorted on any field or searched via the filter box. Page navigation controls are also provided to facilitate the data exploration.

The analyst can them click on any two events in the data table to trigger the similarity calculation and thus enable a detailed comparison of two events belonging to the same MDC – once again, to let the analyst get a better understanding of events correlations

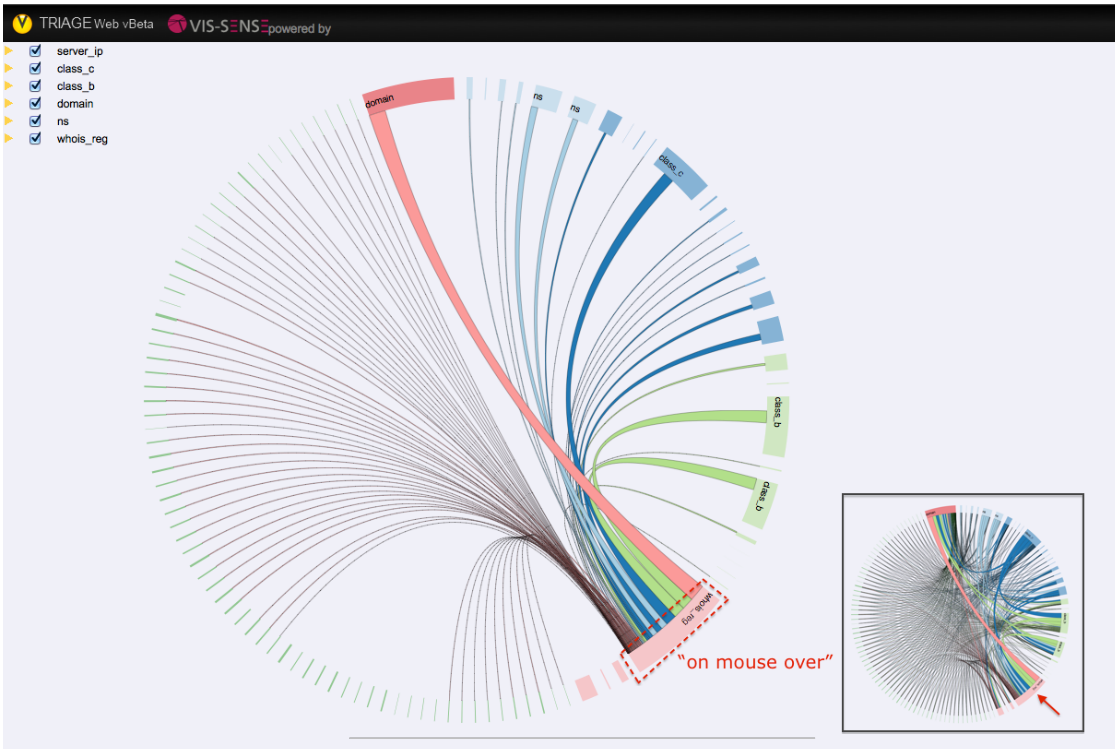

Figure 2.20: Visual analysis of an MDC through the interactive **Chord** viewer.

and the reasons why these events have been grouped within the same MDC. This Events comparison process is illustrated in Figure 2.24.

Because of the decision threshold used to identify MDCs, it may happen that two MDCs are loosely interconnected by certain correlations (which were, however, not sufficient for the algorithm to merge the two clusters). In that case, the two MDCs will appear as "linked" to each other, and an MDC comparison tool enables the analyst to gain insights into these weaker links and understand whether the two MDCs might be linked to the same phenomenon (campaign), or not.

As illustrated in Figures 2.25 and 2.26, in our case study it turned out that MDC 1 was indeed linked to another MDC having certain commonalities (*e.g.*, similar *whois_reg*, *server_ip* and *domain* names), which suggests that the two MDCs are likely associated to the same miscreants.

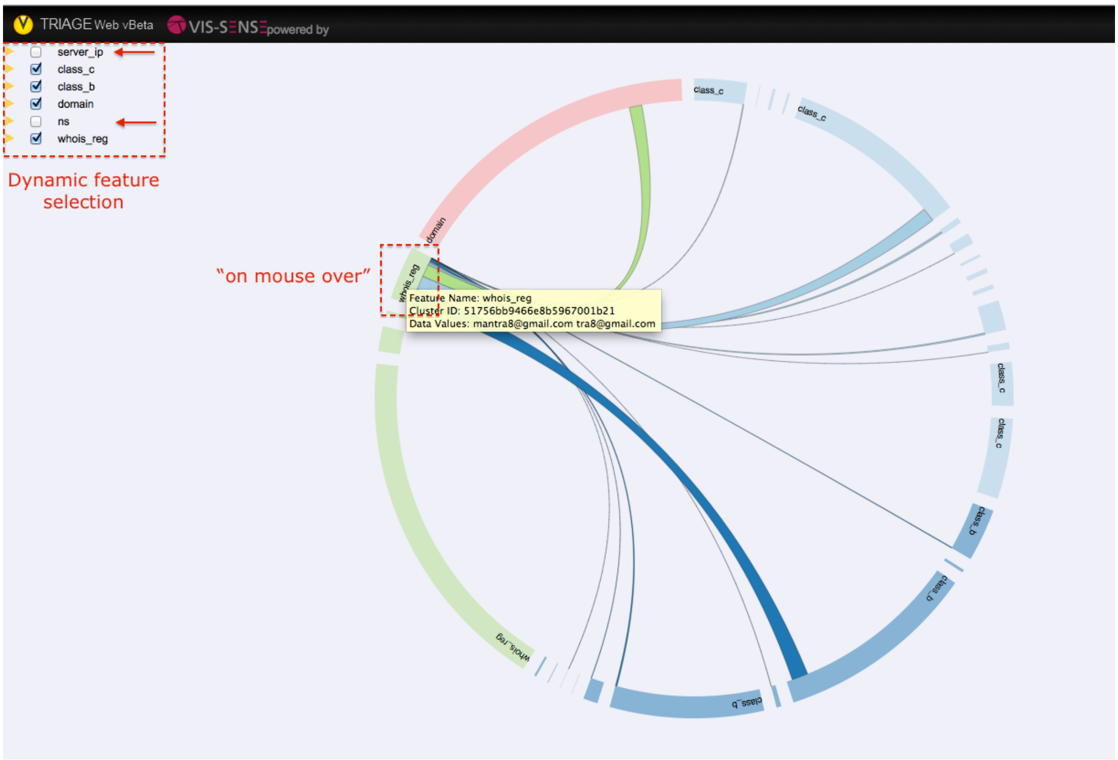

☐ server_ip
☑ class_c
☑ class_b
☑ domain
☐ ns
☑ whois_reg

Dynamic feature selection

"on mouse over"

Feature Name: whois_reg
Cluster ID: 51756bb9466e8b5967001b21
Data Values: mantra8@gmail.com tra8@gmail.com

Figure 2.21: Dynamic feature selection in the MDC **Chord** viewer.

Lessons learned

The knowledge we gained using the TRIAGEWEB visual analytics framework through this analysis has direct repercussions on nowadays security practices, and helps underlining weaknesses in currently employed techniques as well as potentials for new research avenues. Due to space limitations, we could not illustrate all the different types of rogue AV campaigns we have identified and visualized. However, the visual insights of this analysis revealed for example the characteristics of the infrastructure used to spread rogue AV, which can have important consequences on the effectiveness of countermeasures against this threat, more specifically, server IP and domain blacklisting, a technique commonly used to prevent end users from accessing malicious resources. In fact, we observed that the rogue AV infrastructure used in most MDCs (campaigns) comprises servers that not only host a large number of rogue AV sites but also servers where rogue AV sites coexist

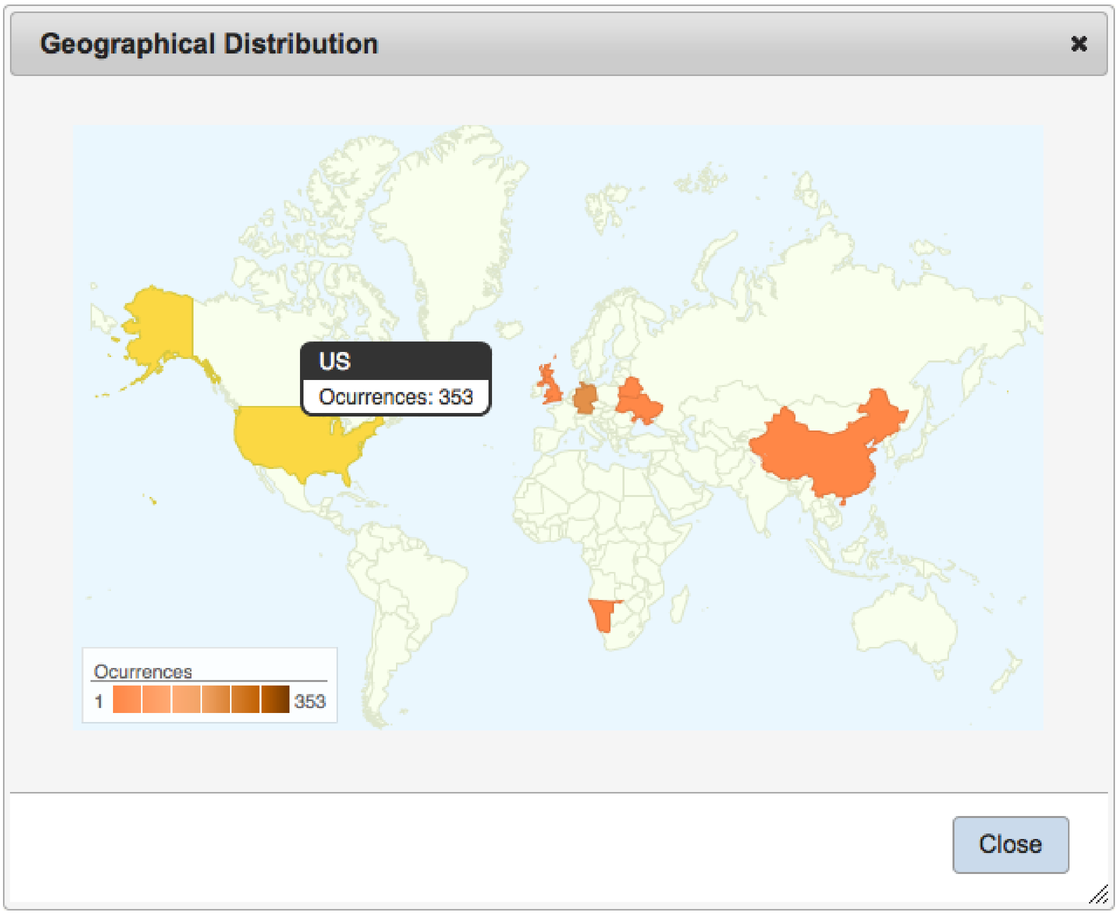

Figure 2.22: Visual analysis of an MDC via the Geographical map display.

with legitimate ones. This situation is a worst case for blacklisting: IP-based blacklisting (where access to a specific web server IP is blocked) is bound to generate many false positives, thus preventing users from visiting benign sites that happen to be hosted on server IPs that also serve malicious sites.

Conversely, domain name-based blacklisting (where access to a specific domain is blocked) is undermined by the easiness with which malicious actors can register large batches of domains. As we have seen with our case study on MDC 1, the registration of hundreds of automatically generated domain names is likely to be an active attempt to

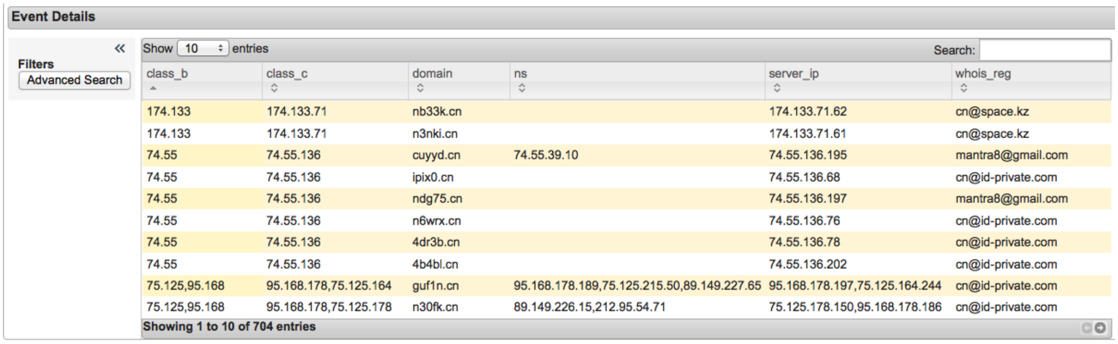

class_b	class_c	domain	ns	server_ip	whois_reg
174.133	174.133.71	nb33k.cn		174.133.71.62	cn@space.kz
174.133	174.133.71	n3nki.cn		174.133.71.61	cn@space.kz
74.55	74.55.136	cuyyd.cn	74.55.39.10	74.55.136.195	mantra8@gmail.com
74.55	74.55.136	ipix0.cn		74.55.136.68	cn@id-private.com
74.55	74.55.136	ndg75.cn		74.55.136.197	mantra8@gmail.com
74.55	74.55.136	n6wrx.cn		74.55.136.76	cn@id-private.com
74.55	74.55.136	4dr3b.cn		74.55.136.78	cn@id-private.com
74.55	74.55.136	4b4bl.cn		74.55.136.202	cn@id-private.com
75.125,95.168	95.168.178,75.125.164	guf1n.cn	95.168.178.189,75.125.215.50,89.149.227.65	95.168.178.197,75.125.164.244	cn@id-private.com
75.125,95.168	95.168.178,75.125.178	n30fk.cn	89.149.226.15,212.95.54.71	75.125.178.150,95.168.178.186	cn@id-private.com

Showing 1 to 10 of 704 entries

Figure 2.23: Fine-grained analysis of MDC events via dynamic data table.

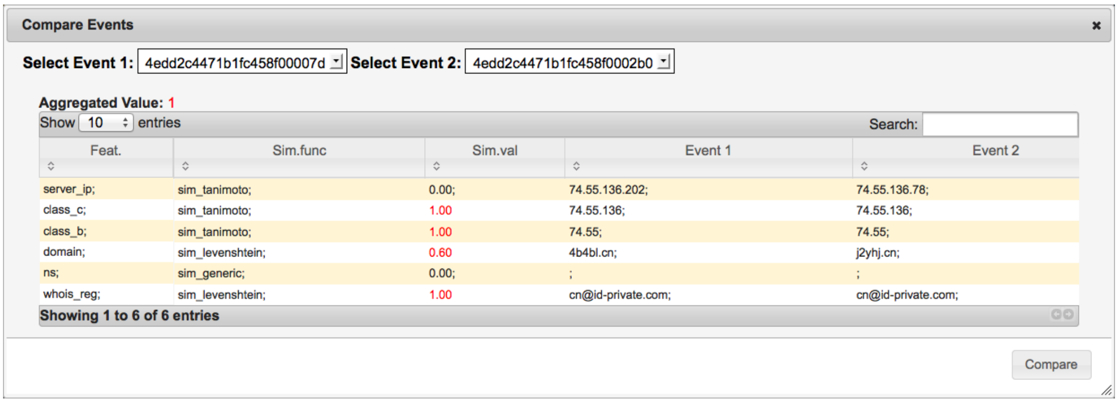

Select Event 1: 4edd2c4471b1fc458f00007d Select Event 2: 4edd2c4471b1fc458f0002b0

Aggregated Value: 1

Show 10 entries Search:

Feat.	Sim.func	Sim.val	Event 1	Event 2
server_ip;	sim_tanimoto;	0.00;	74.55.136.202;	74.55.136.78;
class_c;	sim_tanimoto;	1.00	74.55.136;	74.55.136;
class_b;	sim_tanimoto;	1.00	74.55;	74.55;
domain;	sim_levenshtein;	0.60	4b4bl.cn;	j2yhj.cn;
ns;	sim_generic;	0.00;	;	;
whois_reg;	sim_levenshtein;	1.00	cn@id-private.com;	cn@id-private.com;

Showing 1 to 6 of 6 entries

Figure 2.24: Comparison table showing correlations between two Events attributed to the same MDC.

evade such blacklists.

What would be a good strategy then to effectively fight rogue AV campaigns? Through an analysis of the victim access dataset performed in [10], it appears that taking down payment processing sites could help stop emerging rogue AV campaigns.

We also confirmed via this visual analysis that Rogue AV campaigns often rely on misleading DNS names to lure victims into trusting their products (*e.g.*, *pcsecurity-2009.com*). We have seen in our case study how such campaigns often lead to the automated deployment of large numbers of domains pointing to a few servers and following well-defined patterns in their naming schema. For all these reasons, as already

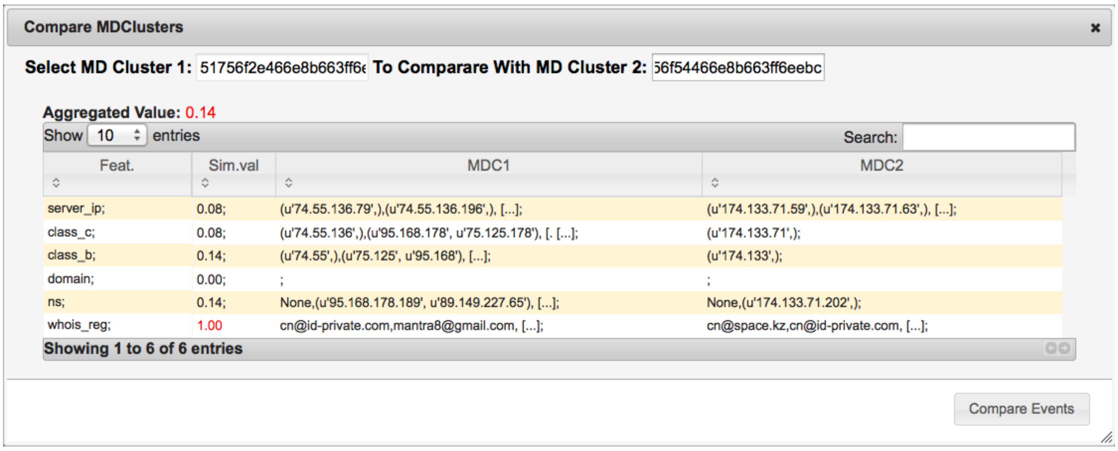

Select MD Cluster 1: 51756f2e466e8b663ff6ε **To Comparare With MD Cluster 2:** 56f54466e8b663ff6eebc

Aggregated Value: 0.14

Feat.	Sim.val	MDC1	MDC2
server_ip;	0.08;	(u'74.55.136.79',),(u'74.55.136.196',), [...];	(u'174.133.71.59',),(u'174.133.71.63',), [...];
class_c;	0.08;	(u'74.55.136',),(u'95.168.178', u'75.125.178'), [. [...];	(u'174.133.71',);
class_b;	0.14;	(u'74.55',),(u'75.125', u'95.168'), [...];	(u'174.133',);
domain;	0.00;	;	;
ns;	0.14;	None,(u'95.168.178.189', u'89.149.227.65'), [...];	None,(u'174.133.71.202',);
whois_reg;	1.00	cn@id-private.com,mantra8@gmail.com, [...];	cn@space.kz,cn@id-private.com, [...];

Showing 1 to 6 of 6 entries

Figure 2.25: Comparison table showing correlations between two linked MDCs.

noted in [15] for other type of threats, DNS seems to be a promising point of view for the detection of such anomalies.

2.7 Conclusions

In this application we have leveraged the visual analytics capabilties of the TRIAGEWEB framework to analyse real data and shed some light on the characteristics and dynamics of a specific threat landscape, that of rogue security software.

Our results can be leveraged in several ways. First, they give a more explanatory description of the rogue AV or client-side threats, in which, for example, individual, disconnected sites are substituted by sets of related sites in which time relationships (*e.g.*, creation dates) are more explicit thanks to very diverse cluster visualizations.

Second, campaign-level information – as obtained through the visual analysis of TRIAGE MDClusters – highlights the *modus operandi* of the criminals orchestrating the campaign, *i.e.*, how they register the domains, what are their hosting partners, the duration of their efforts, the sophistication of the tools available to them (*e.g.*, to automate the registration of domain names), and the countermeasures they employ against take-down efforts.

Finally, the patterns discovered by this analysis could yield means for identifying additional rogue AV sites pro-actively or reactively, for example through a closer monitoring of DNS registration patterns.

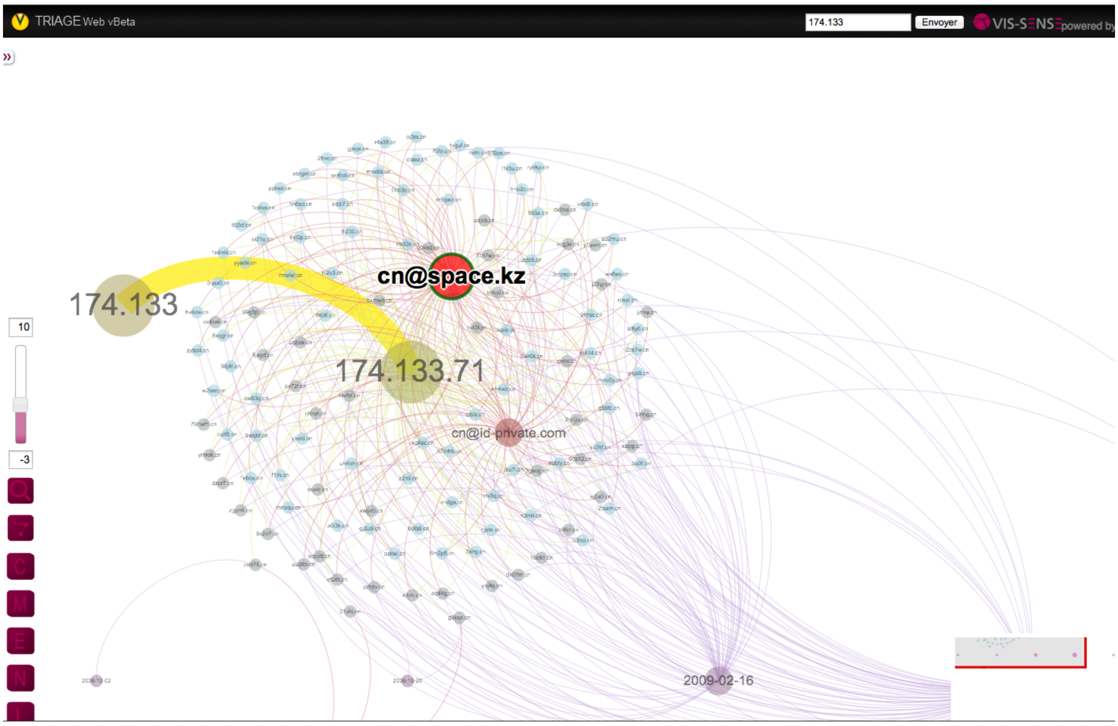

Figure 2.26: Graph visualization of MDC 2 (which appears to be linked to MDC 1).

In this application we have confirmed the specificities of Rogue AV campaigns and their foundations, as identified in previous work, such as in [10]. However, our visual analytics framework makes it much easier for an analyst to run such in-depth analyses and generate insights into complex threat phenomena, moreover in a significantly reduced amount of time.

3 Visual Analysis of Scammers Operations

419 scam (also referred to as Nigerian scam) is a popular form of fraud in which the fraudster tricks the victim into paying a certain amount of money under the promise of a future, larger payoff.

In this Chapter, we show how TRIAGEWEB is used to study how these forms of scam campaigns are organized and evolve over time. In particular, by applying the visual analytics tools developed in VIS-SENSE, we show that phone numbers and email addresses used by scammers are important identifiers to group messages together.

We detail several examples of 419 scam campaigns, some of which lasting for several years, and we illustrate the way scammers operate their campaigns by visualizing them with the VIS-SENSE visual analytics tools and by discussing their characteristics. The results of this analysis and the insights we gained into these scam operations have been published in a joint paper and presented at the International Workshop on Cyber Crime (IWCC'13) – an IEEE Security & Privacy workshop [12].

3.1 Introduction

Nigerian scam [3], also called "419 scam" as a reference to the 419 section in the Nigerian penal code [2], has been a known problem for several decades. The name encompasses many variations of this type of scam, like advance fee fraud, fake lottery, black money scam, etc.

Originally, the 419 scam phenomenon started by postal mail, and then evolved into a business run via fax first, and email later. 419 scam is a popular form of fraud in which the fraudster tricks the victim into paying a certain amount of money under the promise of a future, larger payoff. The prosecution of such criminal activity is complicated [8] and can often be evaded by criminals. As a result, reports of such crime still appear in the social media and online communities, e.g. *419scam.org* [1], exist to mitigate the risk and help users to identify scam messages.

Nowadays, 419 scam is often perceived as a particular type of *spam*. However, while most of the spam is sent today in bulk through botnets and by compromised machines, 419 scam activities are still largely performed in a manual way. Moreover, the underlying business and operation models differ. Spammers trap their victims through engineering effort, whereas scammers rely on human factors: pity, greed and social engineering techniques. Scammers use very primitive tools (if any) compared with other forms of spam where operations are often completely automated. A distinctive characteristic of email fraud is the communication channel set up to reach the victim: from this point of view, scammers tend to use emails and/or phone numbers as their main contacts [9], while other forms of spam are more likely to forward their victims to specific URLs. For instance, a previous study of spam campaigns [14] (in which scam was considered a subset of spam) indicates that 59% of spam messages contain a URL. However, even though 419 scam messages got eclipsed by the large amount of spam sent by botnets, they still pose a persistant problem that causes substantial personal financial losses for a number of victims all around the world.

The traditional spam and scam (not 419) scenarios have been already thoroughly studied (e.g. [14, 7]), where a big part of existing unsolicited bulk email identification techniques rely on high volumes of similar messages.

However, 419 messages are more likely to be sent in lower copies and from webmail accounts. Thus, criminals aim to stay unnoticed by the traditional spam filters and avoid drawing attention to abused webmail accounts. Although the exact distribution methods of 419 scam messages have not been studied as deeply as, for example, the distribution of botnet spam, based on Microsoft Security Intelligence Reports [6] 419 scam messages constitute on average to 8% of email spam traffic (based on the data over the last five years).

A recent study by Costin et al. [9] describes the use of phone numbers in a number of malicious activities. The authors show that the phone numbers used by scammers are often active for a long period of time and are reused over and over in different emails, making them an attractive feature to link together scam messages and identify possible campaigns. In this work, we test this hypothesis by using phone numbers and other email features to automatically detect and study scam campaigns by using a public dataset. To the best of our knowledge, this is the first in-depth study of 419 email campaigns.

Our goal is to study how scammers orchestrate their scam campaigns, by looking at the interconnections between email accounts, phone numbers and email topics used by scammers. To this aim, we use the TRIAGE multi-criteria decision algorithm to efficiently cluster scam emails that are sharing certain commonalities, even in the presence of more *volatile* features. Because of these commonalities, scam emails originating from the same scammer(s) can be grouped together, enabling us to gain insights about the scam campaigns. Additionally, we also evaluate the quality and consistency of the clustering results. For this, we perform threshold sensitivity analysis and we evaluate the homogeneity of clusters using *compactness* as a baseline metric.

In our analysis we have identified over 1,000 different campaigns and, for most of them, *phone numbers* represent the cornerstone that allows us to link the different pieces together. We also discovered some larger-scale campaigns (so-called “macro-cluster”), which are made of loosely inter-connected scam operations. We believe these are likely reflecting different scam runs orchestrated by the same criminal groups, as we observe the same phone numbers or email accounts being reused across different sub-campaigns.

As demonstrated by the experiments, our methods and findings could be leveraged to *pro-actively* identify new scam operations (or variants of previous ones) by quickly associating a new scam to ongoing campaigns. We believe that this would facilitate the work of law enforcement agencies in the prosecution of scammers. Our approach could also be leveraged to improve investigations of other cybercrime schemes by logging and investigating various groups of cybercriminals based on their online activities.

The rest of this Chapter is organized as follows. We start by describing the scam dataset (Section 3.2), to which we apply the TRIAGEWEB framework to identify scam campaigns. In Section 3.4 we present some insights gained from the analysis of one-dimensional *Clusters*. Then, in Section 3.5 we present an overview of MCDA analysis results and we focus on the identification of *MDClusters*. Finally, we describe a number of specific scam campaigns and present their characteristics in Section 3.6. We conclude and summarize our findings in Section 3.7.

3.2 Description of the Data Set

In this section we describe the dataset we used for analyzing 419 scam campaigns and provide some statistics of the scam messages. There are various sources of scam often reported by users and aggregated afterwards by dedicated communities, forums, and other online activity groups. The data chosen for our analysis come from `419scam.org` – a 419 scam aggregator – as it provides a large set of preprocessed data: email bodies, headers, and some already extracted emails attributes, like the scam category and the phone numbers. Note that IP addresses data are absent. We downloaded the emails for a period spanning from January 2009 until August 2012.

In our study we also exploited the fact that the phone numbers can indicate a geographical location, typically the country where the phone is registered. Although it does not prove the origin of the message or the scammer, still it references a country of a scam operation, and improves victim’s level of confidence in the received message. For example, receiving a new partnership offer from UK could seem suspicious if the phone contact has a Nigerian prefix, or a fake lottery notification with contact details originating from an African country while the victim being from Europe. Moreover, as shown in a previous study [9], 419 scam mobile phone numbers are precise in indicating the country of residence of the phone owner (scammer) as few roaming cases were found. Therefore, the phone attribute is precise enough to indicate geographical origins.

Table 3.1: General statistics table

Description	Numbers
Scam messages	36,761
Unique messages	26,250
Total email addresses	112,961
Unique email addresses	34,723
Total phone numbers	41,320
Total unique phone numbers	11,768
Number of countries	12

The resulting dataset consists of 36,761 messages with 11,768 unique phone numbers. The general statistics of the data are shown in Table 3.1. A first thing to notice is that the number of email addresses is three times bigger than the number of phone numbers, emphasizing the facility to acquire mailboxes for malicious purposes. However, this low ratio indicates a rather cheap and easy access to phone numbers.

Another specificity of these data is that each message can contain several email addresses and phone numbers, where a number of different email addresses can be as high as five per message: *From* email address, *reply* email, and other email addresses indicated in the body of the messages. Hence, although we collected only 36,761 messages, we extracted 112,961 email addresses from them.

In our dataset we did not notice any significant bursts of scam messages (verified on a monthly basis) during the three year span, suggesting that the email messages were constantly distributed over time. It is also important to note that the dataset is mostly limited to the European and African regions (with only a few Asian samples), which is due to the way the website owners are collecting and classifying the data. Nevertheless, the geographical distribution of the mentioned continents is reflected in our dataset, excluding only some minor actors.

To better understand the dataset, we look at the time during which emails and phones were advertised by scammers in scam messages. 71% of the email addresses in our dataset were used only during one day. The remaining were used for an average duration of 79 days each. Phone numbers have a longer longevity than email addresses: 51% of the phone numbers were used only for one day; the rest were used on average for 174 days (around 6 months). Hence, making it an important feature in our data clustering analysis.

Table 3.2 summarizes the phone number geographical distribution. UK numbers are twice as common as Nigerian, and three times more common than the ones from Benin, the third biggest group. Netherlands and Spain are the leading countries in Europe. Note that UK should be considered as a special case. As reported by *419scam.org* and Costin et al. [9], all UK phone numbers in this dataset belong to *personal numbering services* – services used for forwarding phone calls to other phone numbers and serving as a masking service of the real destination for the callee. In our dataset there are 44% of such phone numbers (all with UK prefix), another 44% are mobile phone numbers, 12% are fixed lines [9], and only less than 1% of the phones are non-existent.

Table 3.2: Phones by countries

Country	Total phones	Total in %
United Kingdom	4,499	43%
Nigeria	3,121	30%
Benin	1,448	14%
South Africa	562	5%
Spain	372	4%
Netherlands	263	3%
Ivory Coast	89	1%
China	68	1%
Senegal	47	0.5%
Togo	11	0.1%
Indonesia	1	0.01%

The initial *419 scam* messages are also labeled by the *419scam.org* [1] with a category. Around 64% of the emails are assigned to the category named “419 scam”, which is a subcategory of 419 scam and refers specifically to financial fraud types, e.g. transactions, lost funds, etc. As reported by [1], most of the remaining emails (24%) belong to “Fake lottery” category. However, this distribution has been changing over time as shown in Figure 3.1.

Data set in TRIAGEWEB

Figure 3.2 displays a screen shot of the *scam* dataset as it appears in the TRIAGEWEB interface. Like for previous analyses, the user interface provides an interactive data table with browsing, sorting, and advanced searching capabilities. A dynamic field selection filter enables the user to display or hide data fields via the checkboxes in the left margin, which adds or removes columns and updates the data table on-the-fly.

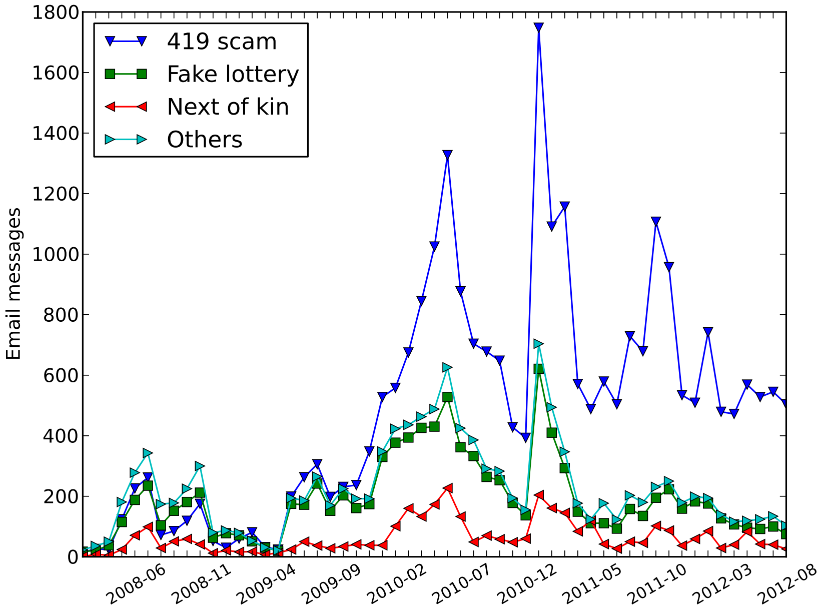

Figure 3.1: Scam email categories over time.

The activation of the *mdc* checkbox (in the left margin) will add one additional column within the data table to display the identifier of the *MDCluster* in which a scam email has been attributed as a result of the TRIAGE clustering analysis. This enables the analyst to quickly search and find any *campaign* that was possibly associated to a given scam pattern.

3.3 Overview of Analysis Results

Figure 3.3 shows a screenshot of the main *Analysis* page, which enables the user to perform a quick assessment of the clustering results for a given dataset. This page contains three subparts:

- **Top frame**: gives some general details on the analysis, such as the total number of *Clusters* and *MDCs*, the number of *features* used for the analysis, some details on the *data slices* (used for larger data sets that need to be processed in several

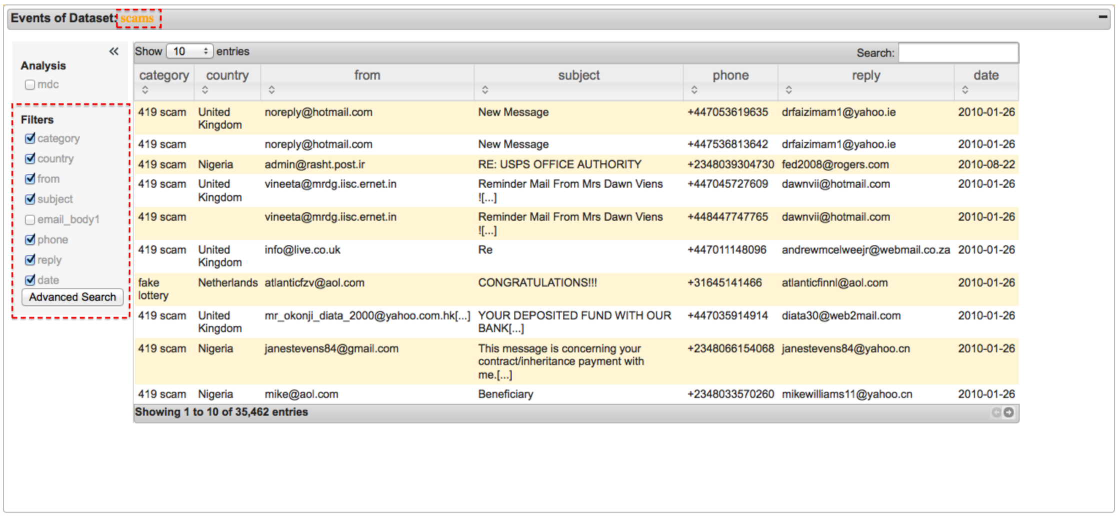

Events of Dataset: scams

Analysis
□ mdc

Filters
☑ category
☑ country
☑ from
☑ subject
□ email_body1
☑ phone
☑ reply
☑ date

Advanced Search

Show 10 entries

Search:

category	country	from	subject	phone	reply	date
419 scam	United Kingdom	noreply@hotmail.com	New Message	+447053619635	drfaizimam1@yahoo.ie	2010-01-26
419 scam		noreply@hotmail.com	New Message	+447536813642	drfaizimam1@yahoo.ie	2010-01-26
419 scam	Nigeria	admin@rasht.post.ir	RE: USPS OFFICE AUTHORITY	+2348039304730	fed2008@rogers.com	2010-08-22
419 scam	United Kingdom	vineeta@mrdg.iisc.ernet.in	Reminder Mail From Mrs Dawn Viens ![...]	+447045727609	dawnvii@hotmail.com	2010-01-26
419 scam		vineeta@mrdg.iisc.ernet.in	Reminder Mail From Mrs Dawn Viens ![...]	+448447747765	dawnvii@hotmail.com	2010-01-26
419 scam	United Kingdom	info@live.co.uk	Re	+447011148096	andrewmcelweejr@webmail.co.za	2010-01-26
fake lottery	Netherlands	atlanticfzv@aol.com	CONGRATULATIONS!!!	+31645141466	atlanticfinnl@aol.com	2010-01-26
419 scam	United Kingdom	mr_okonji_diata_2000@yahoo.com.hk[...]	YOUR DEPOSITED FUND WITH OUR BANK[...]	+447035914914	diata30@web2mail.com	2010-01-26
419 scam	Nigeria	janestevens84@gmail.com	This message is concerning your contract/inheritance payment with me.[...]	+2348066154068	janestevens84@yahoo.cn	2010-01-26
419 scam	Nigeria	mike@aol.com	Beneficiary	+2348033570260	mikewilliams11@yahoo.cn	2010-01-26

Showing 1 to 10 of 35,462 entries

Figure 3.2: *Dataset* page displaying details on the Rogue AV dataset through a flexible data table (which provides sorting, page navigation and search capabilities), as well as dynamic field selection and hiding (left margin of the table).

batches), and finally the `first_seen`, `last_seen` timestamps and total number of events that have already been processed.

- **Middle frame**: visual overview of MDClusters compactness, size, and patterns preview
- **Lower frame**: visual overview of 1D Clusters compactness, size, and patterns preview

As we can see, a list of six features was used for this analysis and include: the *from* email address, *phone* number, *reply* address, any email address found in the body (called *email_body1*), the sending *date*, as well as the email *subject*. Before going into the details of *MDClusters*, we will first turn our attention to the analysis of some *1D Clusters*, *i.e.*, Clusters of scam emails that share only one particular feature.

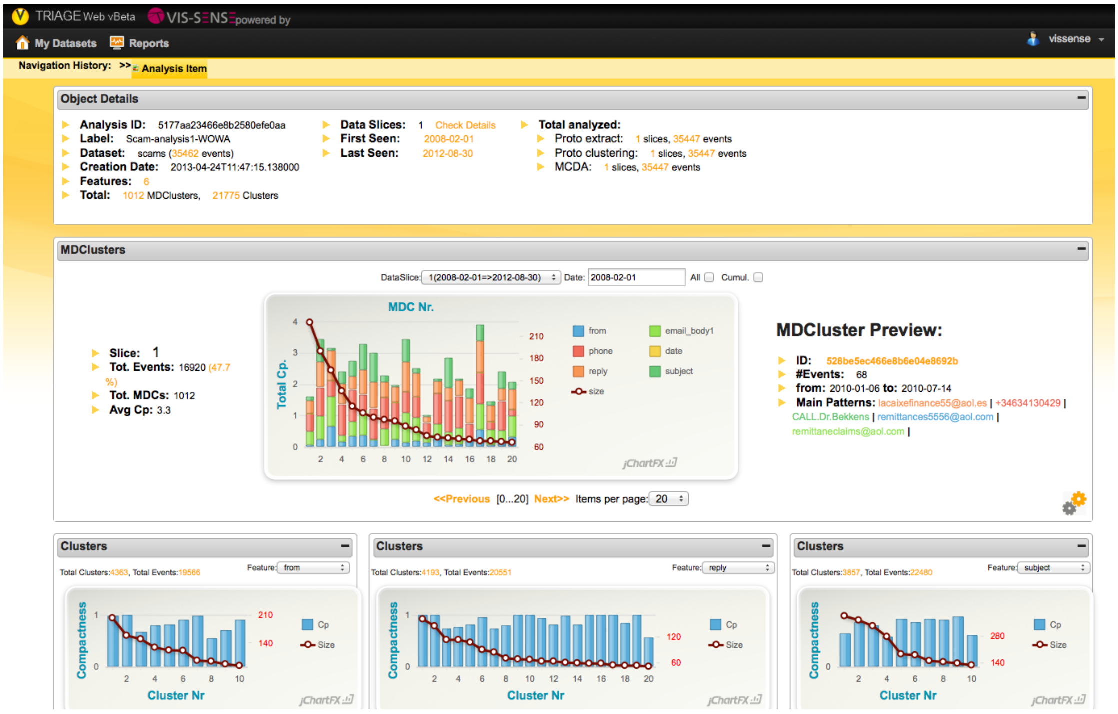

Navigation History: >> Analysis Item

Object Details

- **Analysis ID:** 5177aa23466e8b2580efe0aa
- **Label:** Scam-analysis1-WOWA
- **Dataset:** scams (35462 events)
- **Creation Date:** 2013-04-24T11:47:15.138000
- **Features:** 6
- **Total:** 1012 MDClusters, 21775 Clusters

- **Data Slices:** 1 Check Details
- **First Seen:** 2008-02-01
- **Last Seen:** 2012-08-30

- **Total analyzed:**
  - Proto extract: 1 slices, 35447 events
  - Proto clustering: 1 slices, 35447 events
  - MCDA: 1 slices, 35447 events

MDClusters

DataSlice: 1(2008-02-01=>2012-08-30) Date: 2008-02-01 All □ Cumul. □

- **Slice:** 1
- **Tot. Events:** 16920 (47.7 %)
- **Tot. MDCs:** 1012
- **Avg Cp:** 3.3

MDCluster Preview:

- **ID:** 528be5ec466e8b6e04e8692b
- **#Events:** 68
- **from:** 2010-01-06 **to:** 2010-07-14
- **Main Patterns:** lacaixefinance55@aol.es | +34634130429 | CALL.Dr.Bekkens | remittances5556@aol.com | remittaneclaims@aol.com |

<<Previous [0...20] Next>> Items per page: 20

Clusters

Total Clusters:4363, Total Events:19566 Feature: from

Clusters

Total Clusters:4193, Total Events:20551 Feature: reply

Clusters

Total Clusters:3857, Total Events:22480 Feature: subject

Figure 3.3: *Analysis* page displaying the analysis results for the *Scam* dataset: summary (top frame), MDClusters (middle frame) and 1D Clusters visual overview (lower frames).

3.4 Analysis of 1D Clusters

As a result of step 2 in the TRIAGE processing chain, a series of so-called 1D Clusters are created with respect to each feature selected for the analysis. To enable the user to perform a quick visual assessment of these clustering results, some visual bar charts are provided showing the *Compactness* (*Cp* for short) – or homogeneity – of the Clusters, as well as their total *Size*. Figure **??** illustrates this by showing the Cp and Size of *from* (left) and *subject* Clusters (right). When the user points the mouse over one of the bar in the chart, a preview of the patterns of the associated Cluster is being displayed in a tooltip, along with a link to a page that contains the Cluster details. Note that the similarity metric used to identify these Clusters has been set to *sim_Ngram*, *i.e.*, a similarity function that is based on the calculation and comparison of *N-grams* (which

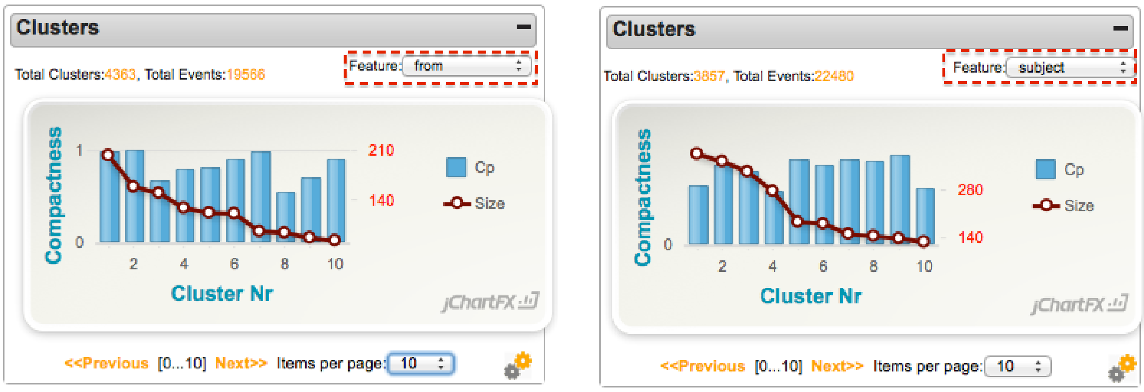
Clusters

Total Clusters:4363, Total Events:19566

Feature: from

(a) *from* Clusters.

Clusters

Total Clusters:3857, Total Events:22480

Feature: subject

(b) *subject* Clusters.

Figure 3.4: *Scam* analysis: Bar charts showing Clusters *Compactness* and *Size*, enabling quick assessment of clustering results.

turns out to be more reliable than the *edit* or *Levenshtein* distance, in particular for short strings such as email addresses).

Regarding the *from* Clusters, we observe that about 55% of the emails have been clustered in 5,129 Clusters, with a maximum size of 203 and an average size of only 4.5. This confirms the intuition that scammers are changing very often the mailboxes used to send their scam emails. The average Cp across all *from* Clusters is still 0.99 – indicating that the N-gram similarity and default thresholds used in TRIAGEWEB works well in practise.

Regarding subject Clusters, about 63% of the emails have been clustered in 5,242 Clusters, with a maximum size of 386 and an average size 5.8. Here too, this indicates the use of very diverse subjects in the scam campaigns (even though a deeper inspection of subject Clusters reveals that only a limited number of topics seem to be reused over and over again by scammers). The average Cp for all *subject* Clusters is also very high (0.98), showing again the effectiveness of the N-gram similarity method.

Let us now focus on a particular Cluster to illustrate the kind of insights that one can get by inspecting these data structures. As illustrated in Figure 3.5, starting from the dynamic *Clusters table*, we select the subject feature and obtain a quick overview of the Clusters. In this table, every row is displaying some details of a given 1D Cluster and the table provides dynamic functionalities such as sorting, page navigation and searching for specific patterns to make the cluster navigation easier. A particular *subject* Cluster seems to present interesting patterns related to some hypothetical *iPhone* offers. This

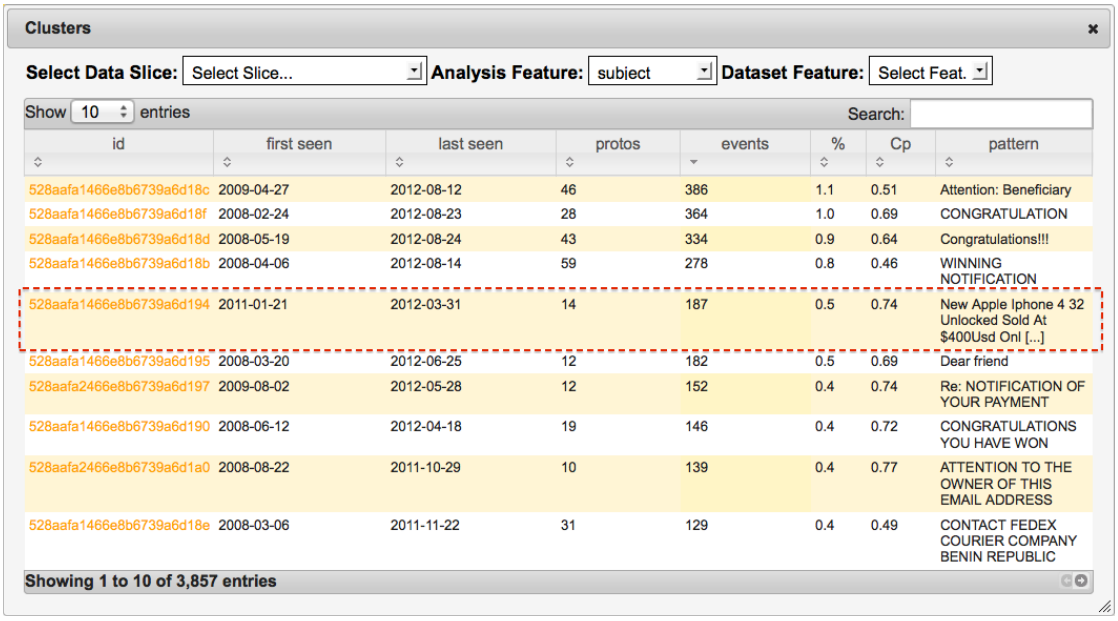

Clusters

Select Data Slice: Select Slice... Analysis Feature: subject Dataset Feature: Select Feat.

Show 10 entries Search:

id	first seen	last seen	protos	events	%	Cp	pattern
528aafa1466e8b6739a6d18c	2009-04-27	2012-08-12	46	386	1.1	0.51	Attention: Beneficiary
528aafa1466e8b6739a6d18f	2008-02-24	2012-08-23	28	364	1.0	0.69	CONGRATULATION
528aafa1466e8b6739a6d18d	2008-05-19	2012-08-24	43	334	0.9	0.64	Congratulations!!!
528aafa1466e8b6739a6d18b	2008-04-06	2012-08-14	59	278	0.8	0.46	WINNING NOTIFICATION
528aafa1466e8b6739a6d194	2011-01-21	2012-03-31	14	187	0.5	0.74	New Apple Iphone 4 32 Unlocked Sold At $400Usd Onl [...]
528aafa1466e8b6739a6d195	2008-03-20	2012-06-25	12	182	0.5	0.69	Dear friend
528aafa2466e8b6739a6d197	2009-08-02	2012-05-28	12	152	0.4	0.74	Re: NOTIFICATION OF YOUR PAYMENT
528aafa1466e8b6739a6d190	2008-06-12	2012-04-18	19	146	0.4	0.72	CONGRATULATIONS YOU HAVE WON
528aafa2466e8b6739a6d1a0	2008-08-22	2011-10-29	10	139	0.4	0.77	ATTENTION TO THE OWNER OF THIS EMAIL ADDRESS
528aafa1466e8b6739a6d18e	2008-03-06	2011-11-22	31	129	0.4	0.49	CONTACT FEDEX COURIER COMPANY BENIN REPUBLIC

Showing 1 to 10 of 3,857 entries

Figure 3.5: Dynamic data table showing *subject* Clusters overview for the Scam analysis.

specific Cluster is grouping 187 scam emails having a subject line similar to the ones displayed in the table.

To get more details about this Cluster, we can follow the hyperlink provided on top of its identifier, which leads us to the dedicated Cluster page. The Cluster has a compactness of 0.74, which indicates some good similarity among the email subjects. Indeed, as illustrated in Figure **??**, the email *subject* patterns of this Cluster are apparently consistent with each other and have many keywords in common or seem to reflect the use of common email templates.

However, the analyst may change this viewpoint in order to analyse the patterns of the very Cluster with respect to another feature, which can be seen as a different viewpoint giving another *perspective* on a particular group of Events. This in turn can reveal interesting insights regarding the modus operandi of the miscreants or attackers. In this example, the analyst may want to know the distribution of phone numbers that are associated with the use of these particular email subjects. In Figure **??**, we observe that only three different phone numbers were apparently used in association with these iPhone scams.

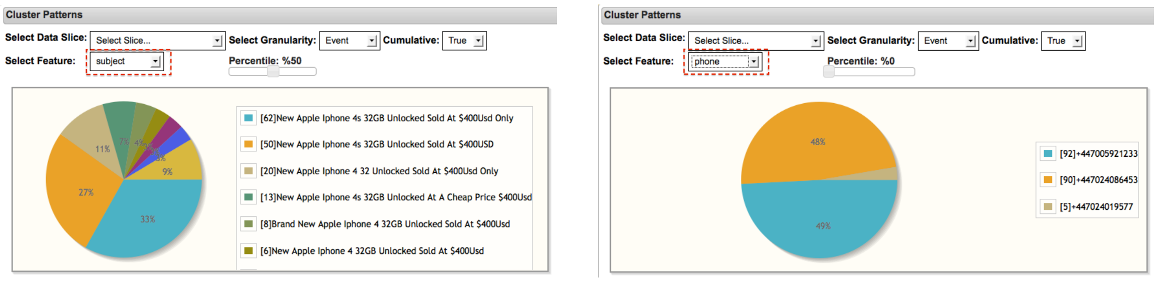

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True

Select Feature: subject Percentile: %50

- 🟦 [62]New Apple Iphone 4s 32GB Unlocked Sold At $400Usd Only
- 🟧 [50]New Apple Iphone 4s 32GB Unlocked Sold At $400USD
- ▨ [20]New Apple Iphone 4 32 Unlocked Sold At $400Usd Only
- 🟩 [13]New Apple Iphone 4s 32GB Unlocked At A Cheap Price $400Usd
- 🟩 [8]Brand New Apple Iphone 4 32GB Unlocked Sold At $400Usd
- 🟩 [6]New Apple Iphone 4 32GB Unlocked Sold At $400Usd

(a) *subject* distribution.

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True

Select Feature: phone Percentile: %0

- 🟦 [92]+447005921233
- 🟧 [90]+447024086453
- ▨ [5]+447024019577

(b) *phone* number distribution.

Figure 3.6: *Scam* analysis: visualizing the patterns of a specific *subject* Cluster.

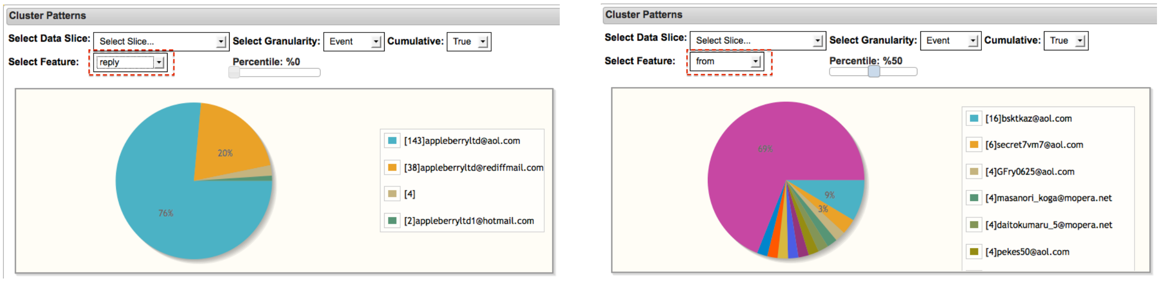

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True

Select Feature: reply Percentile: %0

- 🟦 [143]appleberryltd@aol.com
- 🟧 [38]appleberryltd@rediffmail.com
- ▨ [4]
- 🟩 [2]appleberryltd1@hotmail.com

(a) *subject* distribution.

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True

Select Feature: from Percentile: %50

- 🟦 [16]bsktkaz@aol.com
- 🟧 [6]secret7vm7@aol.com
- ▨ [4]GFry0625@aol.com
- 🟩 [4]masanori_koga@mopera.net
- 🟩 [4]daitokumaru_5@mopera.net
- 🟩 [4]pekes50@aol.com

(b) *phone* number distribution.

Figure 3.7: *Scam* analysis: visualizing the patterns of a specific *subject* Cluster, as viewed w.r.t. different features.

A similar reasoning holds for other features as well, and the analyst may now look at the same group of scam emails across any other dimension, such as the *reply* or *from* email addresses – as illustrated in Figure 3.7.

In conclusion, we observe that this *Cluster analysis* of the scam dataset already provides direct insights that are quite valuable to the analyst, as it can help him understand some emerging patterns in the way scam campaigns are operated. However, to identify these campaigns in a more reliable way, one still needs to combine the different features in an intelligent way, *i.e.*, using a data fusion model that reflects somehow the modus operandi of scammers.

3.5 Multi-Criteria Analysis Results

To identify groups of scam emails likely associated to the same campaign in a more systematic and reliable manner, we have used the multi-criteria aggregation component of TRIAGE to effectively combine all viewpoints provided by 1D Clusters. To do this, an MCDA fusion model was defined as illustrated in Figure 3.8. To this aim, the only input required from the user is the fuzzy linguistic quantifier (*e.g.*, "*at least k*"), the number of correlations k desired (in this case, $k = 3$) and the relative *importances* of individual features (*e.g.*, in this analysis, we have set a high importance on *phone* numbers, a medium importance on *reply* and *email_body1* addresses and a lower importance on *subject* and *date*). The TRIAGEWEB interface will then adjust the weighting vectors according to this model definition.

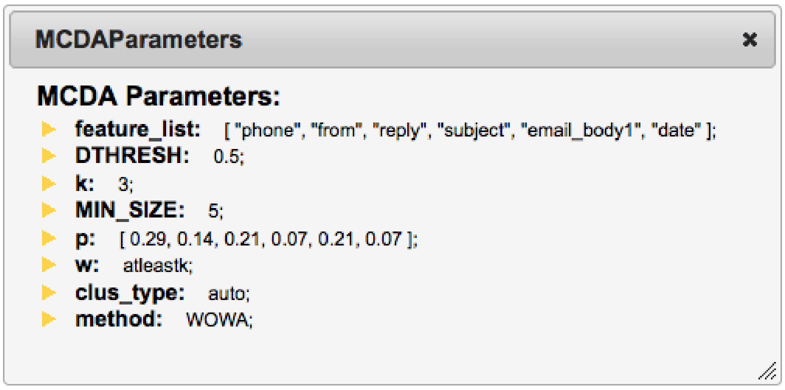

MCDA Parameters:

- feature_list: ["phone", "from", "reply", "subject", "email_body1", "date"];
- DTHRESH: 0.5;
- k: 3;
- MIN_SIZE: 5;
- p: [0.29, 0.14, 0.21, 0.07, 0.21, 0.07];
- w: atleastk;
- clus_type: auto;
- method: WOWA;

Figure 3.8: WOWA parameters used for the *scam* analysis

As described earlier, it is important to let the system run a sensitivity analysis on the **decision threshold** (DTHRESH), so as to set an appropriate thresholding value for the identification of MDCs. This threshold is used to eliminate irrelevant links in the combined graph that is built as a result of the MCDA fusion, and thus avoid inconsistent links between dataset entities. the rersult of this sensitivity analysis is illustrated in Fig. 2.12 where we can observe the impact of increasing the threshold on the clustering results (total number of MDCs and size of largest MDC). As explained before, a good region for choosing the decision threshold lies usually in the range of values for which the total number of MDCs is maximum. For this reason, we choose for our analysis a value of 0.35 as DTHRESH.

To quickly evaluate the consistency of MDC results, the TRIAGEWEB interface provides

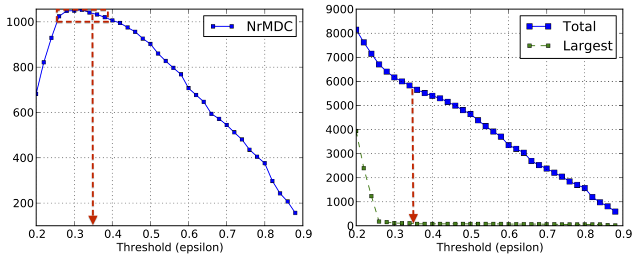

Figure 3.9: Sensitivity analysis to determine the most appropriate decision threshold in the MCDA fusion step.

the analyst with a visual stacked bar chart that displays the overall Compactness and Size of MDClusters (Figure 3.10). The algorithm found in total about 1,000 MDCs accounting for 48% of the dataset, and the average compactness of the clusters lies around 3.3, which is consistent with the MCDA data fusion model defined previously.

Thanks to this chart visualizing the average *Cp* value of every MDC (broken down by feature), it makes it easier to understand which features tend to correlate events (in general), but also what are the main reasons behind the identification of certain MDCs. For example, we observe that the *from* address (in blue) is usually not a strong correlation feature, whereas the *phone* number and *reply* address seem to be strong elements to correlate scam emails and identify campaigns. In some cases, the email subject seems to contribute also to the total compactness of MDCs, and the email_body address tend reinforce the correlations among emails with respect to scammer contact information. We can also observe that MDCs are on average pretty small, which indicates that a large number of small criminal groups seem to operate in an isolated and uncoordinated fashion, and are apparently using very diverse sets of email addresses, subjects and even phone numbers.

3.6 Insights into Scam Campaigns

This section provides now deeper insights into 419 scam campaign orchestration. We present some typical scam campaigns and show the connections between clusters, which are possibly run by the same group of scammers due to multiple strong interconnections among scam emails belonging to the same cluster.

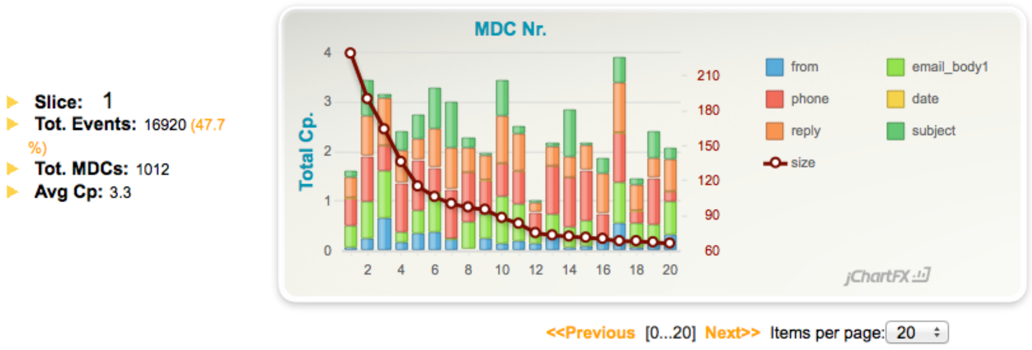

- Slice: 1
- Tot. Events: 16920 (47.7 %)
- Tot. MDCs: 1012
- Avg Cp: 3.3

<<Previous [0...20] Next>> Items per page: 20

Figure 3.10: Scam analysis: stacked bar chart showing the total MDC compactness (broken down by *feature*) for the top 10 MDCs – enabling a quick evaluation of MCDA data fusion results.

3.6.1 Scam campaign examples

To characterize 419 scam campaigns by looking at how they are operated, we have used the VIS-SENSE data visualization tools (and in particular, the TRIAGEWEB interface) to plot the clustered data in an organized manner and look at the "big picture". MDC visualizations are likely reflecting the organization of scam campaigns and their maintenance over time. Interestingly, various campaigns have different operational structures and manage resources differently, as depicted by the examples in Figure 3.19.

Figures 3.14, 3.19, 3.20 and 3.21 show examples of different scam campaigns corresponding to various MDCs, as identified using the TRIAGEWEB interface. The visualizations are drawn using a circular graph layout that represents the various dates on which scam messages were sent. The dates are laid out starting from 9 o'clock (far left in the graph) and are growing clockwise. The other cluster nodes, which highlight other email features and their relationships, are drawn using a force-directed node placement algorithm. The big nodes in the graphs refer to *phone* numbers and *from* addresses. The smaller nodes represent mostly *subjects* and email addresses found in the *reply* and *from* fields, or in the message content.

ESKOM campaign

Figure 3.14 is an example of a 419 scam campaign quite likely orchestrated by the same cyber criminals. From Figure 3.11 we can infer that 38% of the emails of this campaign are correlated by the *from* address and *phone* number, and on top of these two attributes, 26% of emails also have the *reply* and *email_body* address in common.

As we can observe from the graph visualization in Figure 3.14, this ESKOM campaign actually consists of two sub-campaigns: first, a one-year *fake lottery* campaign located in the upper-left part of the graph (Figure 3.14); secondly, a 1,5 year campaign impersonating *ESKOM Holdings*, an electricity company in South Africa.

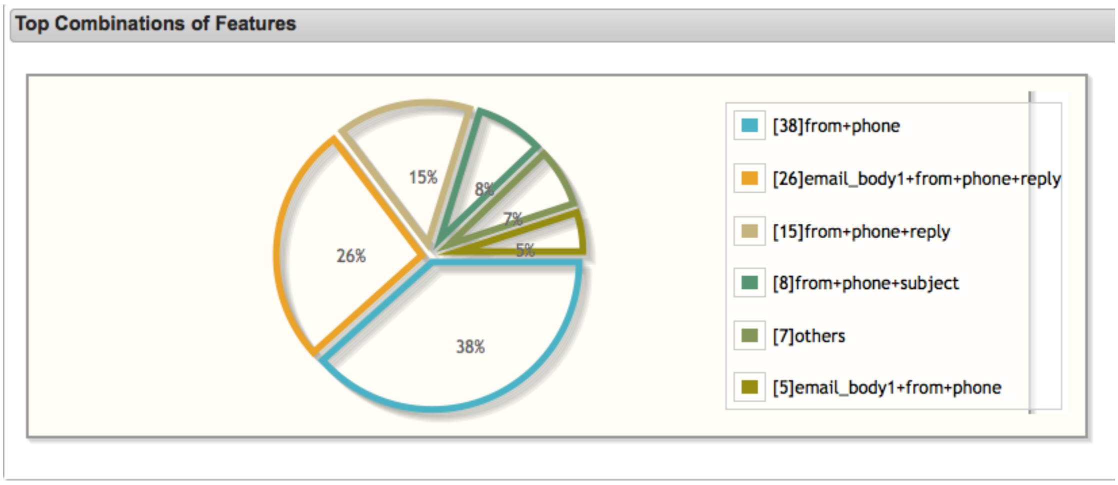

Figure 3.11: ESKOM campaign: distribution of correlation combinations among scam emails.

Figure 3.12 represents the same graph visualization as obtained in TRIAGEWEB using the interactive graph viewer. Even though scammers changed the topic of their scam, they kept re-using the very same phone number (represented in the center of the diagram). A noteworthy aspect of this campaign, shared with other campaigns we found, is that it relies on a few *from* email addresses (*i.e.*, the bigger nodes in the figure). A set of email addresses for *reply* and *body* was used in this campaign, however, after switching to other scam topics a new set of mailboxes and subjects was apparently used. Also, we observe that the load of the scam campaign is well distributed over time, and does not exhibit very high peaks on specific dates, hence keeping very low volumes of emails sent. Finally, the *from* email accounts used by scammers in this case are mostly Gmail accounts. As we have no sender IP information, we could not verify if these were spoofed or not. However, in case these are genuine email accounts, this suggests that scammers use such webmail accounts for long periods of time while staying unnoticed by the email providers.

Figure 3.13 provides another informative perspective on this ESKOM campaign using the *Chord* diagram. While the complete diagram (lower left frame) is pretty crowded and does not really help understand the overall structure of the MDC, the selection of a particular pattern in the circular layout highlights immediately all relationships towards other feature values, which can help us understand certain connections. In this case, Figure 3.13 displays the connections between the main from address used by scammers (`eskomholdings@gmail.com`) and all other patterns identified in this MDC (email *subject*, sending *dates*, *reply* address , etc). Similarly, the analyst may want to focus on less frequent patterns, such as another *from* address used within the same campaign (shown in lower right frame), in order to understand the connections to other more frequent values.

Sify-Rolex campaign

A similar campaign, presented in Figure 3.19a, illustrates the roles of email addresses and phone numbers in 419 scam. This campaign, which lasted for 1,5 year, changed topic 5 times at a frequency of 1 to 2 months, which is visible in the Figure by looking at the larger subgroups placed around the circle. These shorter sub-campaigns were most probably run by the same group of scammers as the same phone number was reused over all campaigns. Inversely, we observe that the email addresses and subjects were completely changed as scammers were moving from one campaign to another. Moreover, these email addresses were often selected to match the campaign topic and subjects, probably to make the scam messages appear authentic.

iPhone campaign

While we observed a large number of easily distinguishable campaigns (having clear patterns and not too many interconnections), we also identified a very different, more “chaotic” type of patterns reflecting thus a different modus operandi, as demonstrated with the graphical illustration in Figure 3.19b. This diagram shows a cluster representing a recent campaign of iPhone-related scams, which lasted for over 1,5 years. The communication infrastructure of the scammers operating this campaign is much more diverse – around 85 unique email accounts were used in this campaign. Moreover, it relies on a large number of “disposable” email addresses, which are typically used only once and seldom reused for long period of time. As opposed to previous examples, however, the same or quite similar subjects and *from* email address were often reused, as well as the very same two phone numbers.

Even though this iPhone campaign may look at first sight somehow “chaotic”, it might be due the fact that many scam emails belonging to this MDC are tightly interconnected by at least 3 or 4 different features, as highlighted by the individual compactness values shown in Figure 3.15. This fact is confirmed with the visualization of the distribution of features coalitions in Figure 3.16, from whihch we can conclude that 64% of the scam emails in this campaign are correlated by the features: *email_body1*, *phone*, *reply*, and *subject*.

Finally, Figure 3.17 and 3.18 illustrate two interactive visualizations as shown to the analyst in TRIAGEWEB, which are representing this iPhone campaign using a graph node-link visualization and a Chord diagram respectively. In both visualizations, we can clearly observe that a specific *reply* address (`appleberryltd@aol.com`) is the cornerstone of this campaign, as it appears to be a central element which is linking a majority of emails.

Internationally-operated campaigns campaign

Some scam campaigns still lack organization and do not always exhibit very clearly separated patterns, as illustrated by two campaigns depicted in Figures 3.19c and 3.19d. Both seem to use fewer email addresses but many different phone numbers that are changing over time. The first one is an international campaign that started with Chinese phone numbers (top-left part), then moved to UK-based anonymous proxy numbers (top-right part), and ended with Dutch-based phone numbers. Almost all analyzed features changed over time, except a single *from* address. Interestingly, most of the phone numbers were regularly switched over time.

We also note that both campaigns exploit *fake lottery* topics. For example, the second one represents a Spanish fake lottery campaign that uses topic-related email addresses (again, in order to look more legitimate), and scammers leverage up to 11 different Spanish phone numbers, which are also regularly changed over time in the campaign. In the middle of the diagram, we can see a larger node representing an email *subject* that has been reused in a large number of scam emails during this campaign, showing that scammers were probably reusing the same fake lottery email template for all these emails. However, this type of scam clusters illustrates well the challenge of identifying such dynamic campaigns, in which the links between scam emails originating from the same criminal group are constantly changing over time. This supports our choice of using a multi-dimensional clustering tool, which can not only take into account multiple features but can also identify groups of emails that are linked by *varying* sets of commonalities (*i.e.*, more *volatile* features). These complex patterns and this volatility in email attributes can also suggest that cyber criminals operate in separate groups, where each

Table 3.3: Macro-clusters, mean values of attributes

Macro-cluster	Nr. of campaigns	Phones	Mailboxes	Subjects	Duration	Countries	Topics
1	14	44	677	223	4 years	4	Lottery, lost funds, investments
2	43	163	1,127	463	4 years	7	Lottery, banks, diplomats, FBI
3	6	18	128	80	4 years	4	Lottery
4	5	8	111	51	3,5 years	2	Packaging, Guiness lottery, loans
5	6	7	201	96	1 year	1	Microsoft lottery, UPS & WU delivery, lost funds
6	4	7	82	33	2 years	1	Lottery, lost payments

group manages its own set of mailboxes and phone number(s), however these groups are somehow federated and are collaborating with each other, for example by sharing the same email templates, same distribution lists or exchanging new scam topics.

3.6.2 Macro clusters: Connecting sub-campaigns

As a follow-up of this analysis, we looked at scam campaigns from a broader perspective: by searching for loosely interconnected MDClusters. The goal was to pinpoint possibly larger-scale campaigns, which are made of weakly interconnected scam operations (*i.e.* different scam *runs*). For this purpose, we only used email addresses and phone numbers, since the other attributes are not considered as personally identifiable information. In fact, we looked for clusters that share at least one email address and/or phone number, and use this information to build so-called *macro-clusters*.

As a result, we identified a set 845 isolated clusters, and another set of 195 connected clusters, where the latter consists of 62 macro-clusters. The characteristics of the top 6 macro-campaigns are shown in Table 3.3. These macro-clusters are particularly interesting as they consist of a set of scam campaigns that appear to be loosely interconnected and therefore could be also orchestrated by the same cybercriminals. In fact, the links between different scam clusters were considered too weak by the clustering algorithm, because of the decision scheme and thresholds set as parameters, and thus these various scam runs were eventually grouped into separate clusters. However, these weak links can be easily recovered, and it is then up to the analyst to investigate how meaningful these interconnections really are. Indeed, we believe that it is much easier for a cyber investigator to start from a set of really meaningful scam clusters, to gradually increase the decision thresholds up to the point where she can decide herself to stop merging data clusters, as it might not be meaningful any more to attribute further different campaigns to the same group due to a lack of evidence.

Macro-clusters usually span across long time periods and exhibit various bursts of emails reflecting different campaigns, which use various topics and can even be operated in different countries. An example of a macro-campaign is illustrated in Figure 3.20, where it consists of 6 different scam campaigns of various sizes that include UK and Nigerian phone numbers. We can easily distinguish them in the diagram as they appear as separate subgroups, each one having one or two bigger nodes (representing phone numbers reused multiple times) and a tail of connected nodes representing a series of *from* email addresses. Notice that campaigns in this case are well separated with respect to phone numbers and emails, which are dedicated to each campaign (or operation), and the overlaps between campaigns are quite limited. However, there is a small node just in the center that indicates how these are interconnected (through a common *from* email address). Some contact details were also reused and we used that for grouping them together. All together, these campaigns lasted for almost 3,5 years. Over this rather long time period, scammers have sent emails using 51 distinct subjects and 8 different phone numbers. This diversity of the topics suggests that there might be some competition among them, as they try to cover different online trick schemes instead of specializing in a single one.

Another example of a macro-campaign is illustrated in Figure 3.21, which consists of 14 sub-campaigns that can be more or less identified in the diagram as separate groups revolving around different phone numbers. Each one has a few dedicated phone numbers (44 in total) and its own set of *from*, *reply* and embedded email addresses. However, in this case it appears that scammers were operating these different scam runs sequentially, sometimes reusing certain resources of previous campaigns. Hence, in forensic investigations it might be necessary to look sometimes at weaker links that may possibly connect together some individuals or criminal groups that could be crime associates.

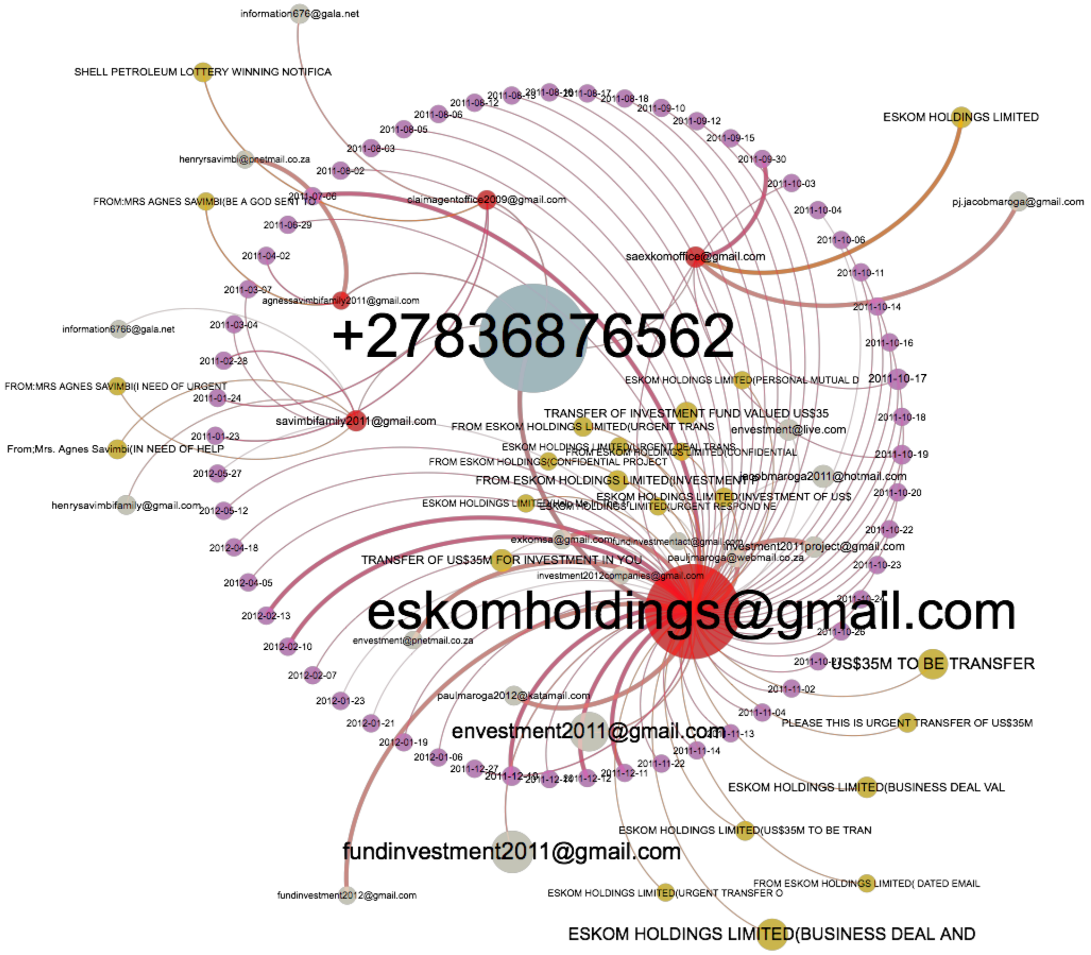
information676@gala.net
SHELL PETROLEUM LOTTERY WINNING NOTIFICA
henryrsavimbi@pnetmail.co.za
FROM:MRS AGNES SAVIMBI(BE A GOD SENT TO
claimagentoffice2009@gmail.com
agnessavimbifamily2011@gmail.com
information6766@gala.net
FROM:MRS AGNES SAVIMBI(I NEED OF URGENT
savimbifamily2011@gmail.com
From;Mrs. Agnes Savimbi(IN NEED OF HELP
henrysavimbifamily@gmail.com
+27836876562
2011-08-02
2011-08-03
2011-08-05
2011-08-06
2011-08-12
2011-08-13
2011-08-16
2011-08-17
2011-08-18
2011-09-10
2011-09-12
2011-09-15
2011-09-30
2011-10-03
2011-10-04
2011-10-06
2011-10-11
2011-10-14
2011-10-16
2011-10-17
2011-10-18
2011-10-19
2011-10-20
2011-10-22
2011-10-23
2011-10-24
2011-10-26
2011-07-06
2011-06-29
2011-04-02
2011-03-07
2011-03-04
2011-02-28
2011-01-24
2011-01-23
2012-05-27
2012-05-12
2012-04-18
2012-04-05
2012-02-13
2012-02-10
2012-02-07
2012-01-23
2012-01-21
2012-01-19
2012-01-06
2011-12-27
2011-12-19
2011-12-20
2011-12-12
2011-12-11
2011-11-22
2011-11-14
2011-11-13
2011-11-04
2011-11-02
ESKOM HOLDINGS LIMITED
pj.jacobmaroga@gmail.com
saexkomoffice@gmail.com
ESKOM HOLDINGS LIMITED(PERSONAL MUTUAL D
TRANSFER OF INVESTMENT FUND VALUED US$35
FROM ESKOM HOLDINGS LIMITED(URGENT TRANS
envestment@live.com
ESKOM HOLDINGS LIMITED(URGENT DEAL TRANS
FROM ESKOM HOLDINGS LIMITED(CONFIDENTIAL
FROM ESKOM HOLDINGS(CONFIDENTIAL PROJECT
FROM ESKOM HOLDINGS LIMITED(INVESTMENT P
jacobmaroga2011@hotmail.com
ESKOM HOLDINGS LIMITED(INVESTMENT OF US$
ESKOM HOLDINGS LIMITED(Help Me In The
ESKOM HOLDINGS LIMITED(URGENT RESPOND NE
exkomsa@gmail.com
fundinvestmentact@gmail.com
investment2011project@gmail.com
TRANSFER OF US$35M FOR INVESTMENT IN YOU
pauljmaroga@webmail.co.za
investment2012companies@gmail.com
eskomholdings@gmail.com
envestment@pnetmail.co.za
US$35M TO BE TRANSFER
paulmaroga2012@katamail.com
PLEASE THIS IS URGENT TRANSFER OF US$35M
envestment2011@gmail.com
ESKOM HOLDINGS LIMITED(BUSINESS DEAL VAL
ESKOM HOLDINGS LIMITED(US$35M TO BE TRAN
fundinvestment2011@gmail.com
fundinvestment2012@gmail.com
FROM ESKOM HOLDINGS LIMITED(DATED EMAIL
ESKOM HOLDINGS LIMITED(URGENT TRANSFER O
ESKOM HOLDINGS LIMITED(BUSINESS DEAL AND

Figure 3.12: ESKOM campaign: interactive **Graph** visualization in TRIAGEWEB interface.

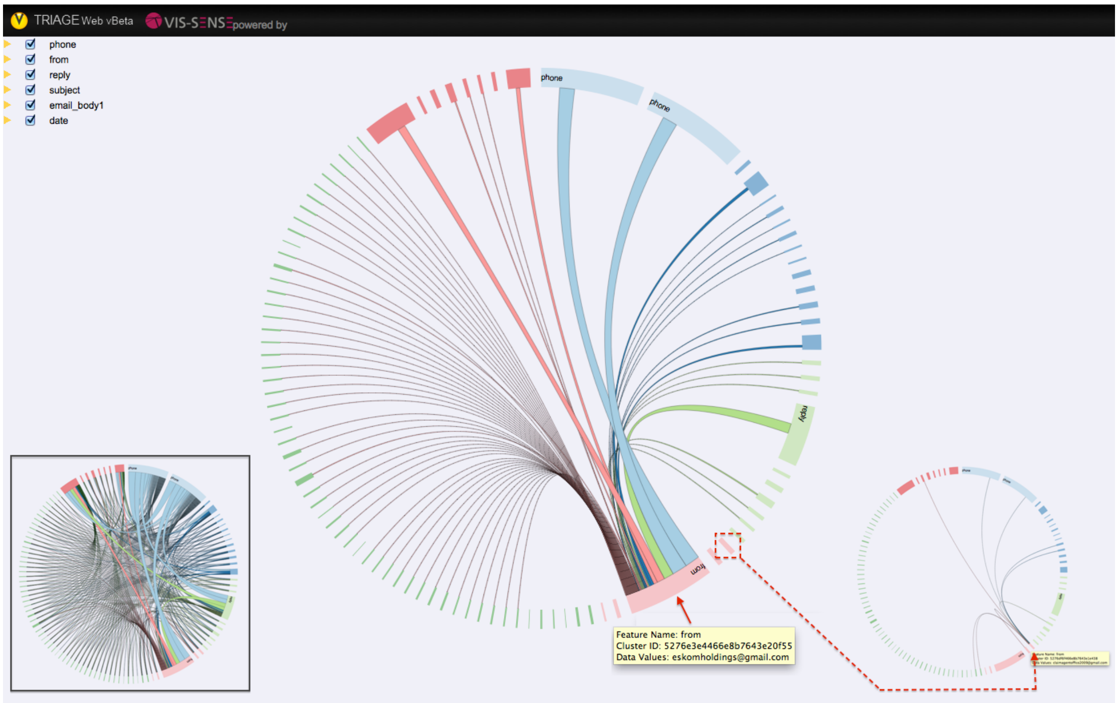

Feature Name: from
Cluster ID: 5276e3e4466e8b7643e20f55
Data Values: eskomholdings@gmail.com

Figure 3.13: ESKOM campaign: interactive **Chord** viewer in TRIAGEWEB interface.

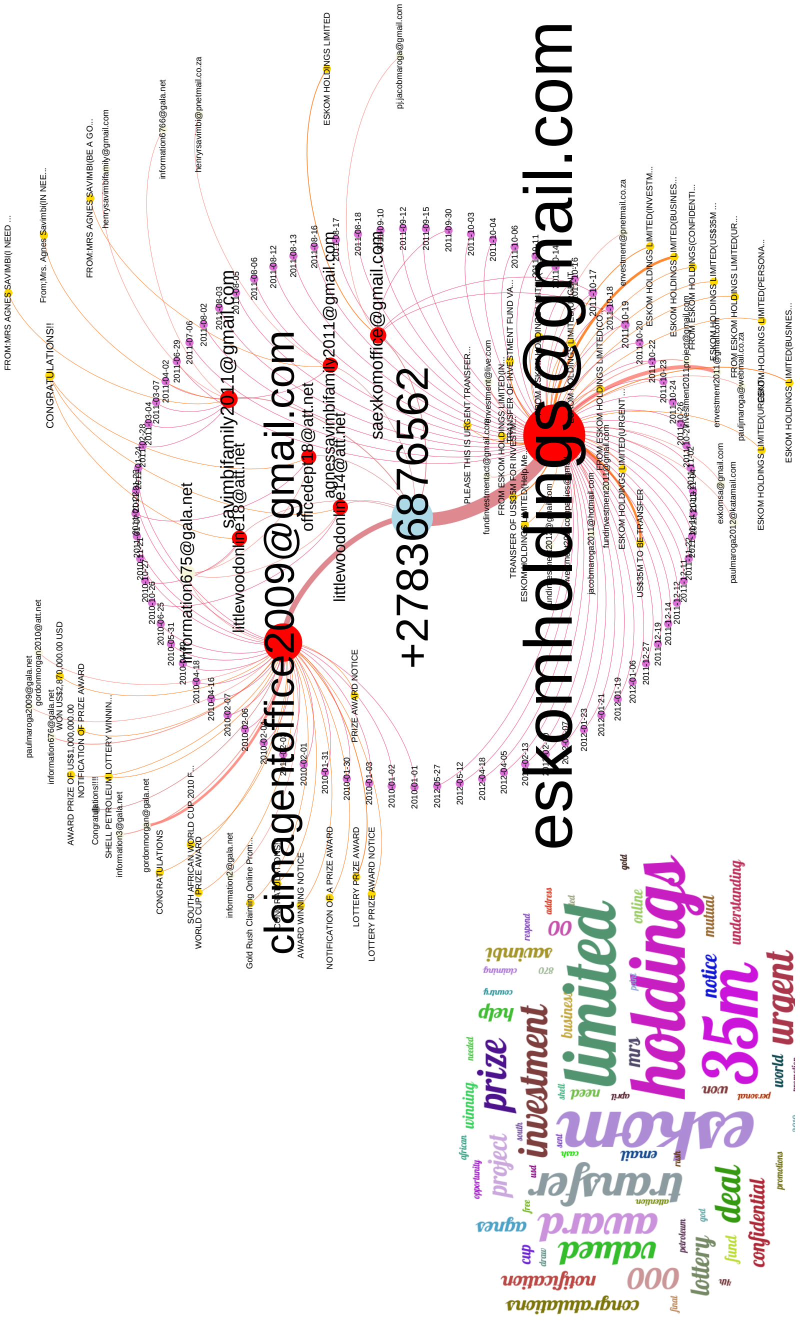
claimagentoffice2009@gmail.com
eskomholdings@gmail.com
+27836876562
information675@gala.net
littlewoodonline18@att.net
savimbifamily2011@gmail.com
officedept18@att.net
agnessavimbifamily2011@gmail.com
littlewoodonline14@att.net
saexkomoffice@gmail.com
pj.jacobmaroga@gmail.com
ESKOM HOLDINGS LIMITED
henrysavimbi@pnetmail.co.za
information6766@gala.net
henrysavimbifamily@gmail.com
FROM:MRS AGNES SAVIMBI(BE A GO...
From;Mrs. Agnes Savimbi(IN NEE...
FROM:MRS AGNES SAVIMBI(NEED ...
CONGRATULATIONS!!
PRIZE AWARD NOTICE
LOTTERY PRIZE AWARD
LOTTERY PRIZE AWARD NOTICE
NOTIFICATION OF A PRIZE AWARD
AWARD WINNING NOTICE
Gold Rush Claiming Online Prom...
information2@gala.net
SOUTH AFRICAN WORLD CUP 2010 F...
WORLD CUP PRIZE AWARD
CONGRATULATIONS
gordonmorgan@gala.net
information3@gala.net
SHELL PETROLEUM LOTTERY WINNIN...
AWARD PRIZE OF US$1,000,000.00
NOTIFICATION OF PRIZE AWARD
WON US$2,870,000.00 USD
paulmaroga2009@gala.net
gordonmorgan2010@att.net
information676@gala.net
PLEASE THIS IS URGENT TRANSFER...
TRANSFER OF US$35M FOR INVESTMENT
ESKOM HOLDINGS LIMITED(URGENT ...
US$35M TO BE TRANSFER
jacobmaroga2011@hotmail.com
fundinvestment2011@gmail.com
fundinvestmentact@gmail.com
investment@live.com
einvestment@pnetmail.co.za
ESKOM HOLDINGS LIMITED(INVESTM...
ESKOM HOLDINGS LIMITED(BUSINES...
ESKOM HOLDINGS LIMITED(CONFIDENTI...
ESKOM HOLDINGS LIMITED(US$35M ...
FROM ESKOM HOLDINGS LIMITED(UR...
ESKOM HOLDINGS LIMITED(PERSONA...
eskomsa@gmail.com
paulmaroga2012@katamail.com
paulmaroga@webmail.co.za
einvestment2011@gmail.com
eskom
holdings
limited
transfer
investment
award
prize
35m
urgent
deal
valued
notification
congratulations
confidential
lottery
000
project
agnes
help
savimbi
notice
mutual
understanding
world
email
winning
business
won
need

Figure 3.14: Lotteries (between 9 and 12 o'clock) and *ESKOM Holdings* impersonation.

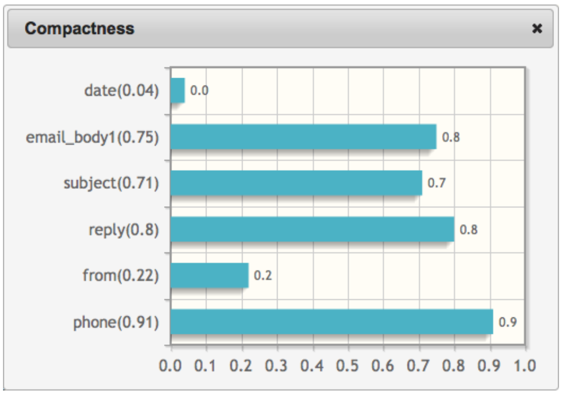

Figure 3.15: iPhone campaign: Compactness values for individual features.

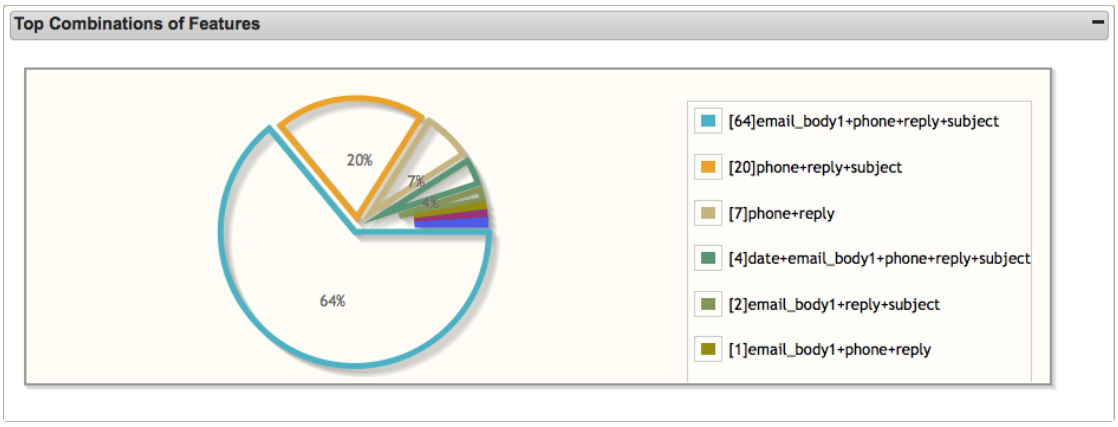

Figure 3.16: iPhone campaign: distribution of features coalitions among scam emails.

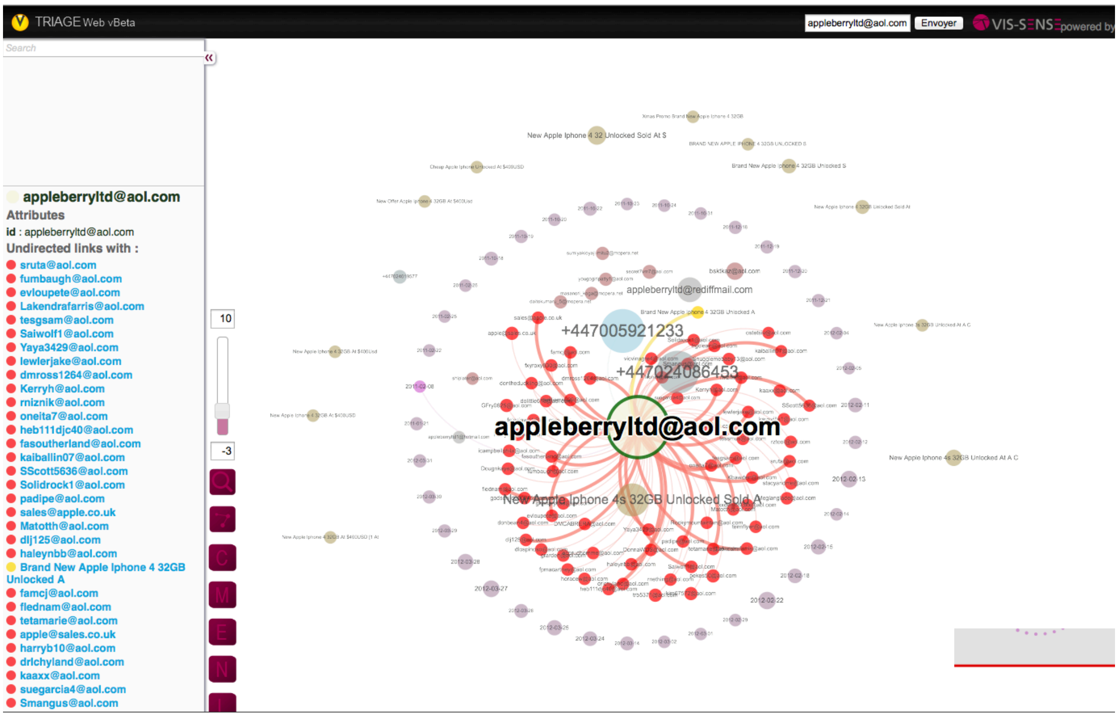
TRIAGE Web vBeta
appleberryltd@aol.com
Envoyer
VIS-SENSE powered by
Search
«
10
-3
C
M
E
N
+447005921233
+447024086453
appleberryltd@aol.com
appleberryltd@rediffmail.com
New Apple Iphone 4s 32GB Unlocked Sold A
New Apple Iphone 4 32 Unlocked Sold At $
Brand New Apple Iphone 4 32GB Unlocked A
New Apple Iphone 4s 32GB Unlocked At A C
2012-03-27
2012-02-22
2012-02-13

appleberryltd@aol.com

Attributes

id : appleberryltd@aol.com

Undirected links with :

- sruta@aol.com
- fumbaugh@aol.com
- evloupete@aol.com
- Lakendrafarris@aol.com
- tesgsam@aol.com
- Saiwolf1@aol.com
- Yaya3429@aol.com
- lewlerjake@aol.com
- dmross1264@aol.com
- Kerryh@aol.com
- rniznik@aol.com
- oneita7@aol.com
- heb111djc40@aol.com
- fasoutherland@aol.com
- kaiballin07@aol.com
- SScott5636@aol.com
- Solidrock1@aol.com
- padipe@aol.com
- sales@apple.co.uk
- Matotth@aol.com
- dlj125@aol.com
- haleynbb@aol.com
- Brand New Apple Iphone 4 32GB Unlocked A
- famcj@aol.com
- flednam@aol.com
- tetamarie@aol.com
- apple@sales.co.uk
- harryb10@aol.com
- drlchyland@aol.com
- kaaxx@aol.com
- suegarcia4@aol.com
- Smangus@aol.com

Figure 3.17: iPhone campaign: interactive **Graph** visualization in TRIAGEWEB interface.

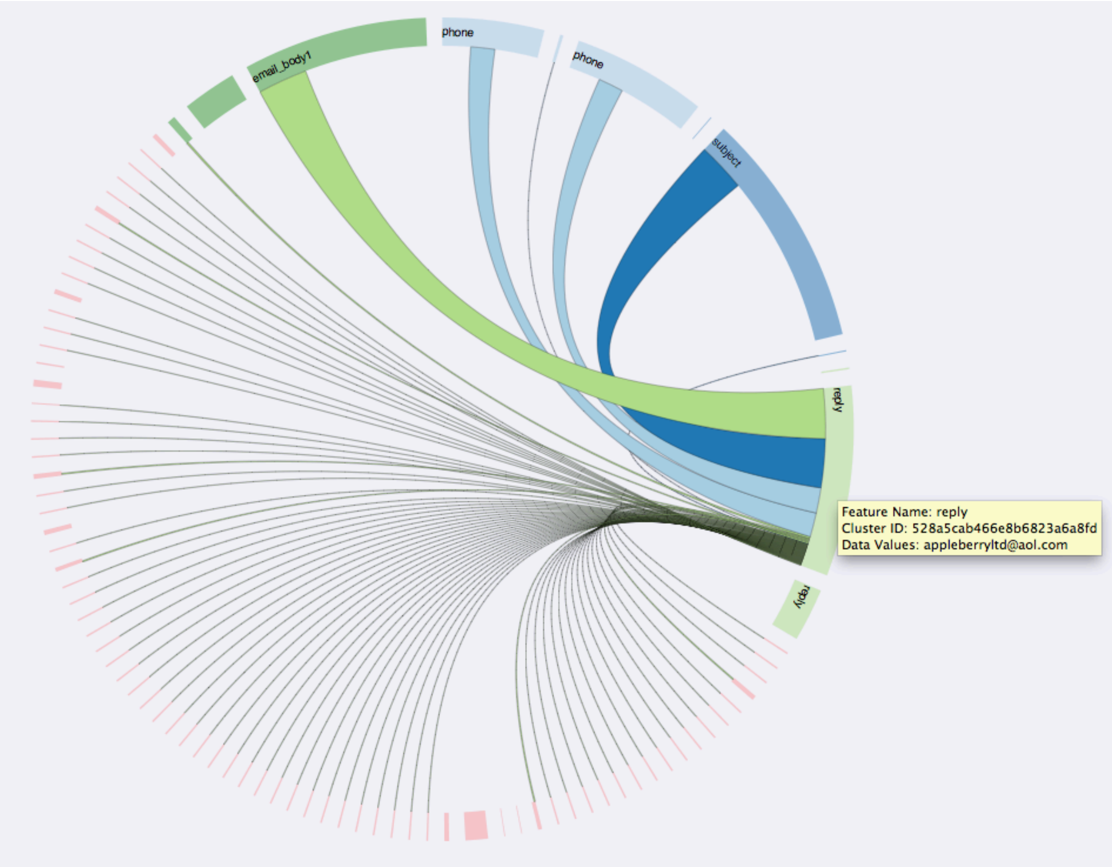

Feature Name: reply
Cluster ID: 528a5cab466e8b6823a6a8fd
Data Values: appleberryltd@aol.com

Figure 3.18: iPhone campaign: interactive **Chord** viewer in TRIAGEWEB interface.

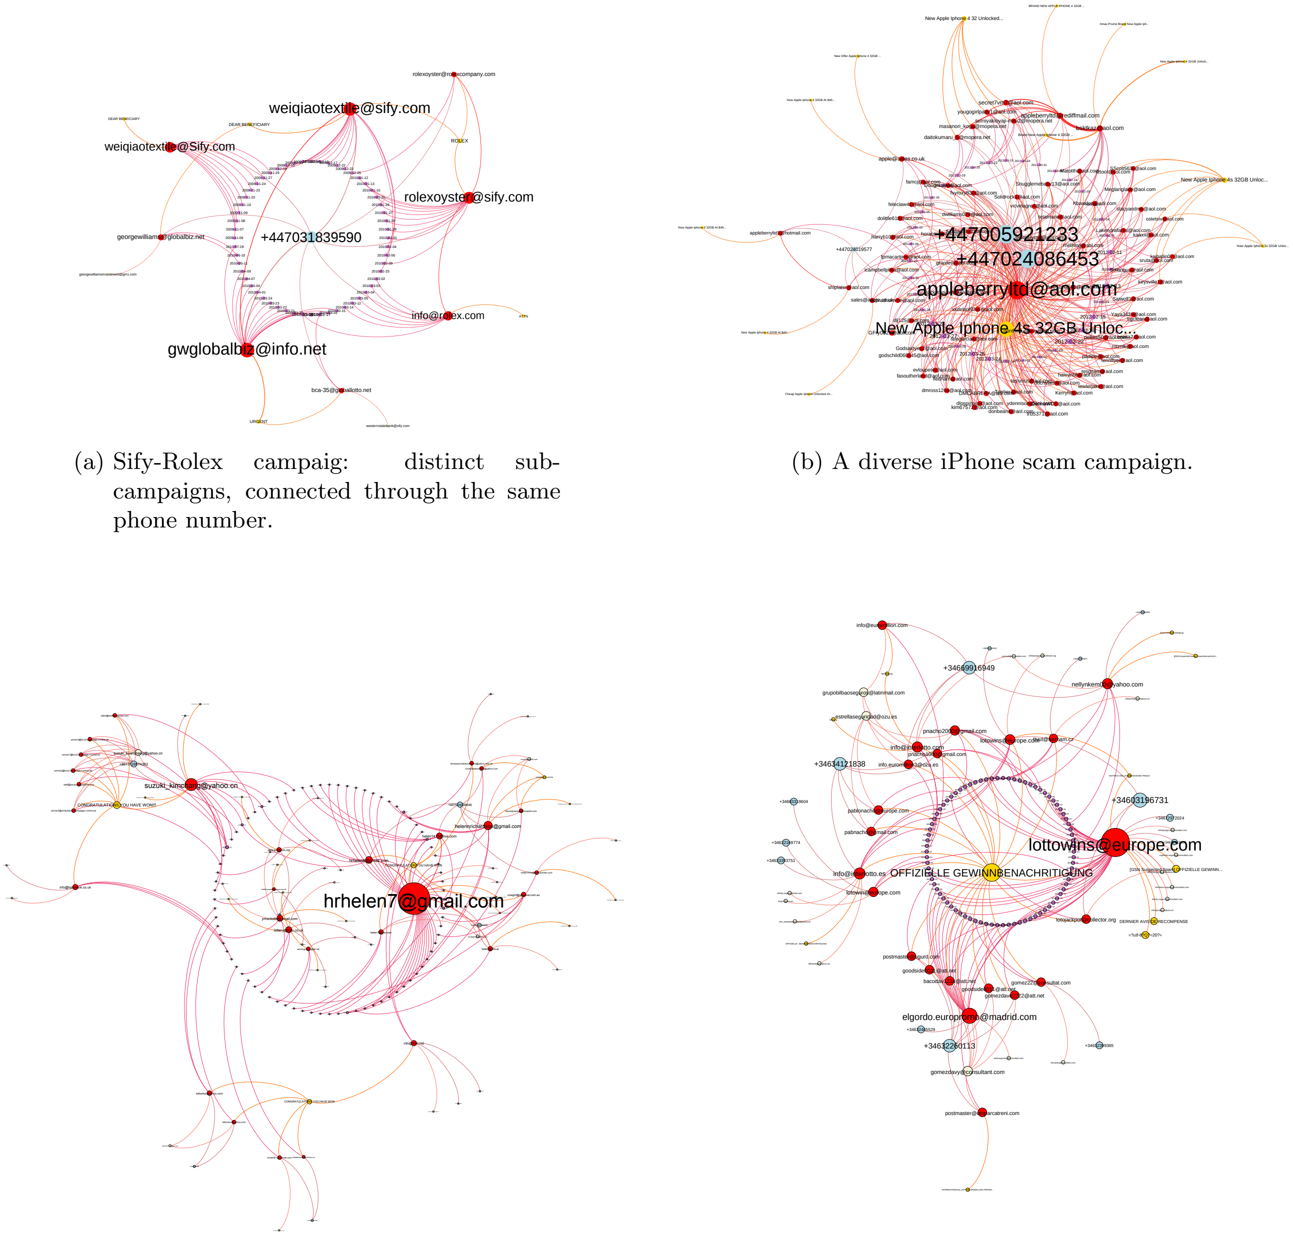

(a) Sify-Rolex campaig: distinct sub-campaigns, connected through the same phone number.

(b) A diverse iPhone scam campaign.

(c) International campaign operated in China, UK and Netherlands.

(d) Spanish lottery campaign changing often phone numbers and email addresses.

Figure 3.19: Examples of other scam campaign structures.

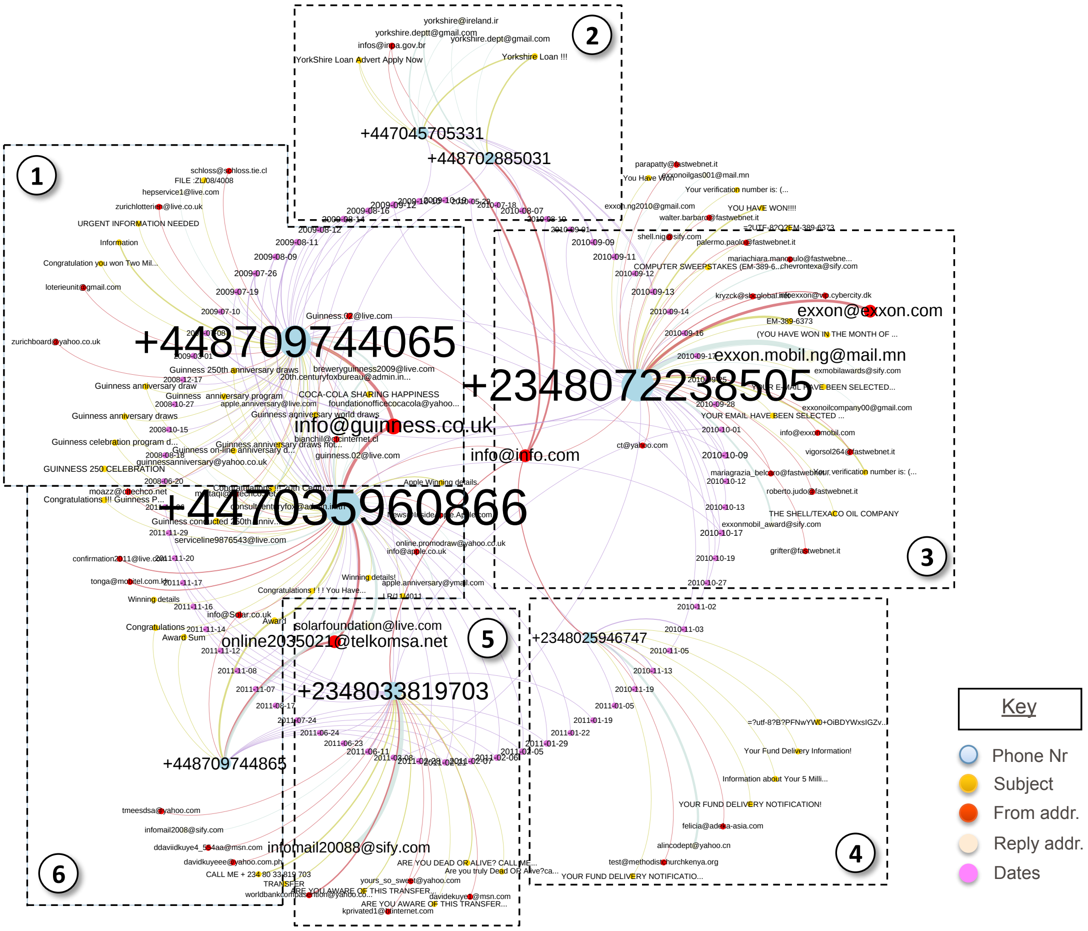

Figure 3.20: Example of macro-cluster #1. Date nodes (in purple) are laid in clock-wise fashion reflect the timeline of the campaigns.

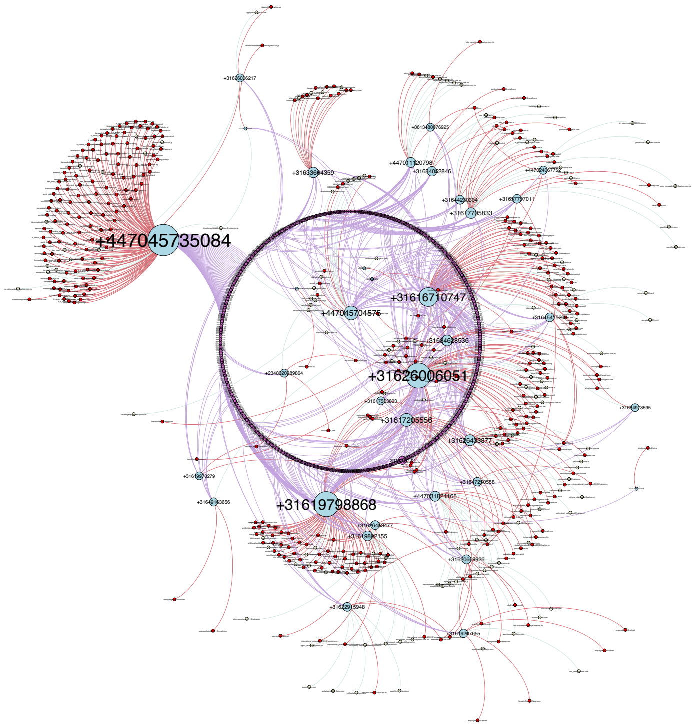

Figure 3.21: Example of macro-cluster #2. Date nodes (in purple) are laid in clock-wise fashion reflect the timeline of the campaigns.

3.7 Summary

In this Chapter, we have demonstrated how TRIAGEWEB was used to study how scam campaigns are organized and evolve over time. Our goal was to study how scammers orchestrate their scam campaigns, by looking at the interconnections between email accounts, phone numbers and email topics used by scammers. By applying the visual analytics tools developed in VIS-SENSE, we could generate many insights into these scam campaigns, for example that that phone numbers and email addresses used by scammers are important identifiers to group messages together.

We have detailed some typical examples of 419 scam campaigns and their characteristic, and we have illustrated the way scammers operate their campaigns by visualizing them using the arsenal of visual analytics tools developed by the VIS-SENSE partners. The results of this analysis have been published in a joint paper and presented at the International Workshop on Cyber Crime (IWCC'13) – an IEEE Security & Privacy workshop [12], which demonstrates once again the usefulness and effectiveness of our visual analytics tools.

4 Visual Analysis of Spam Sent from Hijacked Networks

4.1 Introduction

In 2006 Ramachandran et al. claimed in [16] to have observed spammers exploiting the routing infrastructure (control plane) of the Internet to steal blocks of IP addresses from their legitimate owner for a very short time and send spam using the stolen IP addresses. We refer to those spammers as *fly-by spammers.* In theory, if such sophisticated spammers would carry out short-lived hijacks of IP address blocks this would allow them to hop between those IP blocks fast enough to *circumvent* blacklists of spam senders still heavily used as a first layer of defence in spam filters. Moreover hijacking blocks of IP addresses is no different from stealing the IP identity of the addresses legitimate owner and thus hinder traceability of attackers which can lead to *misattributing* attacks when responding with possibly legal actions.

Despite anecdotal evidence mainly reported in [16, 11] and on some mailing lists [4, 5] the fly-by spammer phenomenon remained a open conjecture. As one of the tasks of the VIS-SENSE project aimed at studying at large scale this phenomenon to confirm or not this conjecture, the framework SpamTracer [19] has been developed to monitor the routing behavior of spam networks and evidence possible IP block hijacks performed by spammers. Recently, several cases of IP address blocks hijacked for a very short period of time during which they were used to launch spam campaigns, in other words fly-by spammers, were uncovered by SpamTracer. In this chapter we report on the analysis of the spam emails sent from those hijacked networks using triage. In particular we are interested in gaining more insights into the spam operations of fly-by spammers, e.g., what characterises spam operations run by fly-by spammers? Can we find some links between *different* apparently unrelated hijacked networks based on the spam emails received from them?

Answering those questions is actually important as it may improve our understanding of fly-by spammers and their "ecosystem". It may also help uncover some of the motivations for those spammers to carry out such sophisticated hijacking attacks prior to sending spam. Eventually such analysis could also help us refine the techniques we use to detect fly-by spammers and identify the most appropriate measures to defeat their operations.

In the following Sections, we start by describing the spam emails data set used in this analysis. In Sections 4.3 and 4.4 we leverage the TRIAGE and TRIAGEWEB applications to provide details on (i) per-feature patterns uncovered in spam emails and (ii) spam campaigns identified by TRIAGE. To illustrate the results, we analyze one such campaign and explain some of the insights into fly-by spammers operations we gained from it. We conclude and summarise our findings in Section 4.5.

4.2 Description of the Data Set

The analysis of spam campaigns launched from hijacked networks has been performed by extracting from the SPAMCLOUD data set all spam emails received from networks identified as hijacked by SPAMTRACER. The SPAMCLOUD data set has already been described in the VIS-SENSE deliverable 2.2 ”Data collection and infrastructure”. This spam feed is collected at Symantec.cloud spamtraps and consists of about 4M emails per day.

The spam data set that we have considered for this analysis consists of 4,431 spam emails received from 29 different hijacked networks over a period of 7 months, from January to July 2013. Normally the SPAMCLOUD data set includes mostly features related to the sender and the content of spam emails. In an effort to perform a comprehensive analysis of fly-by spammers operations we decided to enrich the emails sent from hijacked networks with information related to the spam advertised content hosting infrastructures, i.e., the servers hosting content advertised in spam emails.

In order to make this deliverable self-contained we provide in Table 4.1 a brief description of the spam email features available in the data set. We group them in three classes: (i) sender-related features, (ii) message-related features and (iii) scam hosting infrastructure-related features. A *timestamp* for each spam email is also available. Figure 4.1 shows the BGP hijacks data set page provided by TRIAGEWEB allowing to easily browse and search in the data.

4.3 Analysis of 1D Clusters

In this section we present a preliminary analysis of the data performed to have a better idea of the structure of spam emails with respect to each feature individually. In particular, this analysis allows to identify possible patterns (or clusters) of values within each feature and look at a pattern (or cluster) for a given feature with respect to all other features.

Spam sender feature	**Description**
sourceIp	the IP address of the spam sender host
ipCountry	the country part of the spam sender host geolocation
ipCity	the city part of the spam sender host geolocation
ipRegion	the region part of the spam sender host geolocation
ipLatitude	the latitude part of the spam sender host geolocation
ipLongitude	the longitude part of the spam sender host geolocation
prefix	the hijacked IP address block which the spam sender host IP address belongs to (in CIDR notation)
x_bot	the spam bot this spam sender host was found to belong to, based on signatures of known spam bots built from the SMTP communications and spam email messages sent by those bots
x_p0f_genre	the operating system family of the spam sender host, inferred using P0f (e.g., Windows)
x_p0f_detail	the operating system version details of the spam sender host, inferred using P0f (e.g., XP SP1+, 2000 SP3)
x_p0f_distance	the topological distance from the recipient server to the spam sender host, inferred using P0f
x_p0f_link	the type of Internet link the spam sender host is connected to, inferred using P0f (e.g., ethernet/modem)
x_p0f_signature	the TCP/IP stack signature of the spam sender host, inferred using P0f (e.g., S10:54:1:60:M1460,S,T,N,W8:.)
x_p0f_uptime	the uptime of the spam sender host
Message feature	**Description**
subject	the subject of the spam email
charset	the character-set of the spam email (e.g., us-ascii)
cte	the content-transfer-encoding of the spam email (e.g., 7bit)
messageSize	the size (in bytes) of the spam email message
uris	the URIs found in the spam email message
uri_domains	the domain name of the URIs found in the spam email message
Scam hosting infrastructure feature	**Description**
uri_server_ip	the IP addresses of the hosts to which the URI domain names of a given spam email resolved to
uri_whois_creation	the creation dates of the URI domain names of a given spam email, retrieved using WHOIS
uri_whois_expiration	the expiration dates of the URI domain names of a given spam email, retrieved using WHOIS
uri_whois_ns	the NS servers of the URI domain names of a given spam email, retrieved using WHOIS
uri_whois_rar	the WHOIS registrars of the URI domain names of a given spam email
uri_whois_reg	the WHOIS registrant email addresses of the URI domain names of a given spam email

Table 4.1: Features of the BGP hijacks data set, i.e., spam emails sent from hijacked networks.

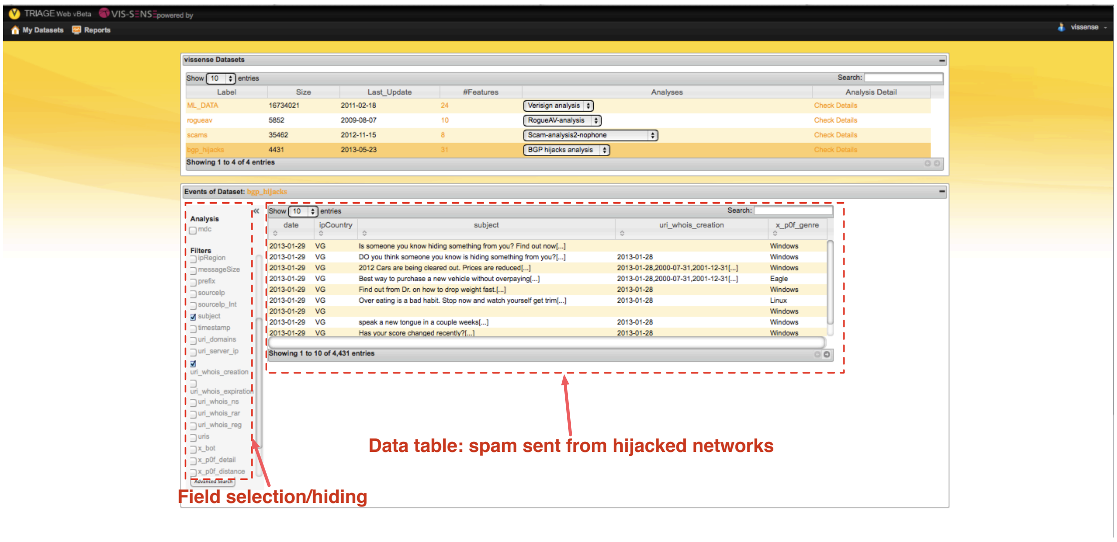

vissense Datasets

Show 10 entries Search:

Label	Size	Last_Update	#Features	Analyses	Analysis Detail
ML_DATA	16734021	2011-02-18	24	Verisign analysis	Check Details
rogueav	5852	2009-08-07	10	RogueAV-analysis	Check Details
scams	35462	2012-11-15	8	Scam-analysis2-nophone	Check Details
bgp_hijacks	4431	2013-05-23	31	BGP hijacks analysis	Check Details

Showing 1 to 4 of 4 entries

Events of Dataset: bgp_hijacks

Show 10 entries Search:

date	ipCountry	subject	uri_whois_creation	x_p0f_genre
2013-01-29	VG	Is someone you know hiding something from you? Find out now[...]		Windows
2013-01-29	VG	DO you think someone you know is hiding something from you?[...]	2013-01-28	Windows
2013-01-29	VG	2012 Cars are being cleared out. Prices are reduced[...]	2013-01-28,2000-07-31,2001-12-31[...]	Windows
2013-01-29	VG	Best way to purchase a new vehicle without overpaying[...]	2013-01-28,2000-07-31,2001-12-31[...]	Eagle
2013-01-29	VG	Find out from Dr. on how to drop weight fast.[...]	2013-01-28	Windows
2013-01-29	VG	Over eating is a bad habit. Stop now and watch yourself get trim[...]	2013-01-28	Linux
2013-01-29	VG			Windows
2013-01-29	VG	speak a new tongue in a couple weeks[...]	2013-01-28	Windows
2013-01-29	VG	Has your score changed recently?[...]	2013-01-28	Windows

Showing 1 to 10 of 4,431 entries

Figure 4.1: Data set page displaying details on the BGP hijacks data set through a convenient data table (providing sorting, navigation and search capabilities), as well as dynamic field selection or hiding (left pane).

TRIAGEWEB provides specific visualisations for the analysis of 1D clusters. Figure 4.2 depicts bar charts showing the compactness and size of clusters related to the scam hosting infrastructure features *uri_whois_reg* (on the left) and *uri_server_ip* (on the right). The clustering results overview provided by those charts allows to quickly see if there are interesting clusters within each feature that are worth taking a closer look at. For example, feature *uri_whois_reg* (left bar chart) contains a large cluster (cluster #1) of spam emails with a compactness of 1.0. This reveals that a *strong email address pattern* was found in the WHOIS registrant email addresses of URIs advertised in spam emails.

In order to have more information on the clusters created for each feature, a clusters table (Figure 4.3) is provided by TRIAGEWEB. To obtain more information on the clusters for the feature *uri_whois_reg* we can select the feature *uri_whois_reg* in the clusters table (Figure 4.3 (1)) and sort the clusters according to their size (in number of spam emails) (Figure 4.3 (2)). We can now see that, for example, cluster #2 contains 120 spam emails, has a compactness of 0.54 and has an interesting pattern of WHOIS registrant email addresses ”HLVCFWRRR@WHOISPRIVACYPROTECT.COM”. The pattern and the compactness both suggest that the WHOIS registrant emails addresses in this cluster probably share a invariable part. For example, as it has already been seen in the

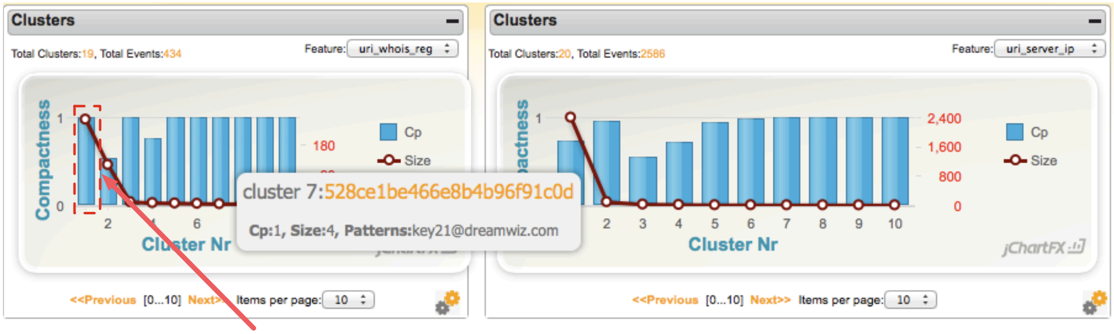

Clusters

Total Clusters:19, Total Events:434

Feature: uri_whois_reg

cluster 7:528ce1be466e8b4b96f91c0d

Cp:1, Size:4, Patterns:key21@dreamwiz.com

Clusters

Total Clusters:20, Total Events:2586

Feature: uri_server_ip

Compactness
1
0
2
4
6
Cluster Nr
180
Cp
Size
<<Previous [0...10] Next>> Items per page: 10
1
2
3
4
5
6
7
8
9
10
Cluster Nr
2,400
1,600
800
0
Cp
Size
jChartFX
<<Previous [0...10] Next>> Items per page: 10

Cluster #1 for feature "uri_whois_reg"

Figure 4.2: Bar chart showing Clusters Compactness and Size (quick evaluation of clustering results)

analysis of Rogue AV campaigns (see chapter 2) the *user* part of the email address could be different for the registration of several different domain names while the *domain* part of the email address would remain unchanged.

To obtain further details on a given cluster TRIAGEWEB provides a special analysis page, as depicted in Figure 4.4 for the example cluster #2. The pie chart shows the distribution of values of the cluster with respect to the feature *uri_whois_reg*. We can immediately see the *wide variety* of email addresses which still ended up in a single cluster.

It is interesting to look at the content of a cluster with respect to the feature that was used to create it but it is also interesting to look at it with respect other features. For our example cluster #2, Figure 4.5 shows the distribution of values for the feature *uri_domains*. Similar to the distribution of WHOIS registrant email addresses, the cluster contains a lot of different spam advertised domain names. All this suggests that all the domain names were actually registered using different automatically generated WHOIS registrant email addresses sharing a common pattern.

Looking at the distribution of values of cluster #2 with respect to other features further reveals (i) a lot of different spam email subjects (Figure 4.6), (ii) 8 different IP prefixes for the spam sender hosts (Figure 4.7) and (iii) no spam bot identified (Figure 4.8). The pattern in the WHOIS registrant emails addresses exhibited by the 120 spam emails in this cluster may already suggest that the apparently unrelated 8 hijacked prefixes from which those spam emails were sent might actually be linked to each other.

In this section we showed how TRIAGEWEB can be leveraged to visually identify and analyze per-feature patterns in TRIAGE 1D clusters. In the next section we present the

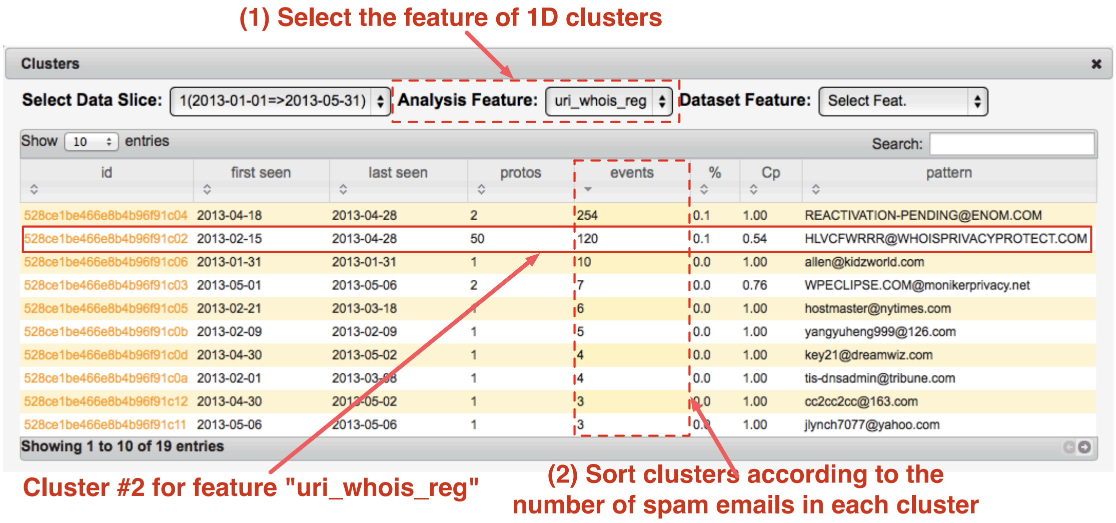

(1) Select the feature of 1D clusters

Clusters

Select Data Slice: 1(2013-01-01=>2013-05-31) Analysis Feature: uri_whois_reg Dataset Feature: Select Feat.

Show 10 entries Search:

id	first seen	last seen	protos	events	%	Cp	pattern
528ce1be466e8b4b96f91c04	2013-04-18	2013-04-28	2	254	0.1	1.00	REACTIVATION-PENDING@ENOM.COM
528ce1be466e8b4b96f91c02	2013-02-15	2013-04-28	50	120	0.1	0.54	HLVCFWRRR@WHOISPRIVACYPROTECT.COM
528ce1be466e8b4b96f91c06	2013-01-31	2013-01-31	1	10	0.0	1.00	allen@kidzworld.com
528ce1be466e8b4b96f91c03	2013-05-01	2013-05-06	2	7	0.0	0.76	WPECLIPSE.COM@monikerprivacy.net
528ce1be466e8b4b96f91c05	2013-02-21	2013-03-18	1	6	0.0	1.00	hostmaster@nytimes.com
528ce1be466e8b4b96f91c0b	2013-02-09	2013-02-09	1	5	0.0	1.00	yangyuheng999@126.com
528ce1be466e8b4b96f91c0d	2013-04-30	2013-05-02	1	4	0.0	1.00	key21@dreamwiz.com
528ce1be466e8b4b96f91c0a	2013-02-01	2013-03-08	1	4	0.0	1.00	tis-dnsadmin@tribune.com
528ce1be466e8b4b96f91c12	2013-04-30	2013-05-02	1	3	0.0	1.00	cc2cc2cc@163.com
528ce1be466e8b4b96f91c11	2013-05-06	2013-05-06	1	3	0.0	1.00	jlynch7077@yahoo.com

Showing 1 to 10 of 19 entries

Cluster #2 for feature "uri_whois_reg"

(2) Sort clusters according to the number of spam emails in each cluster

Figure 4.3: Dynamic data table showing Clusters overview

insights into fly-by spammers operations uncovered from the analysis of TRIAGE multi-dimensional clusters taking advantage of other TRIAGEWEB visualisations.

4.4 Multi-Criteria Analysis - Insights into Fly-by Spammers Operations

In this section we leverage TRIAGEWEB to perform a visual analysis of TRIAGE multi-dimensional clusters of spam emails sent from hijacked networks and shed some light on the modus operandi of fly-by spammers.

Prior to describing the multi-dimension clustering results, we briefly look at the MCDA parameters used via TRIAGEWEB as depicted in Figure 4.9. Without entering into too many details, we can first notice that 7 features were used in the MCDA (feature_list) and include: (i) the date, (ii) the sourceIp, (iii) the subject, (iv) the uri_domains, (v) the x_p0f_detail, (vi) the uri_server_ip and (vii) the uri_whois_reg. Furthermore, by looking at the importance factors (importance_factors) set for the analysis, the highest importance was set to the features sourceIp and uri_whois_reg while the lowest importance was set to the features date and x_p0f_detail. These importance factors really reflect the relevance of each feature when it comes to grouping together spam emails in a cluster (or spam campaign). In our case, grouping together spam emails having a similar sourceIp

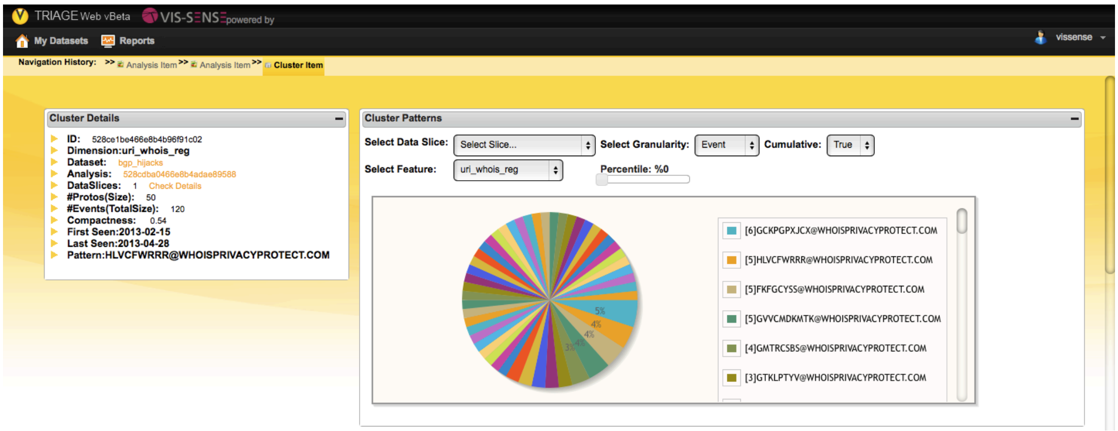

Cluster Details

- ID: 528ce1be466e8b4b96f91c02
- Dimension:uri_whois_reg
- Dataset: bgp_hijacks
- Analysis: 528cdba0466e8b4adae89588
- DataSlices: 1 Check Details
- #Protos(Size): 50
- #Events(TotalSize): 120
- Compactness: 0.54
- First Seen:2013-02-15
- Last Seen:2013-04-28
- Pattern:HLVCFWRRR@WHOISPRIVACYPROTECT.COM

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True

Select Feature: uri_whois_reg Percentile: %0

- [6]GCKPGPXJCX@WHOISPRIVACYPROTECT.COM
- [5]HLVCFWRRR@WHOISPRIVACYPROTECT.COM
- [5]FKFGCYSS@WHOISPRIVACYPROTECT.COM
- [5]GVVCMDKMTK@WHOISPRIVACYPROTECT.COM
- [4]GMTRCSBS@WHOISPRIVACYPROTECT.COM
- [3]GTKLPTYV@WHOISPRIVACYPROTECT.COM

Figure 4.4: Cluster page showing the details of a specific Cluster (in this case, a uri_whois_reg Cluster)

and uri_whois_reg is more relevant for our analysis then grouping together spam emails sharing a common date and x_p0f_detail (operating system of spam sender host).

Figure 4.10 shows the MDCluster preview panel provided by TRIAGEWEB to have a quick overview of the multi-dimension clustering results, e.g., the number of MDClusters, the size of MDClusters, for a given MDCluster the compactness of each feature. The MDClusters table depicted in Figure 4.11 allows to look at the list of MDClusters, sort them according to, for instance, the number of spam emails they contain (Figure 4.11 (1)) or their average compactness (Figure 4.11 (2)), see a sample of per-feature values for each MDCluster (Figure 4.11 (3)). There is a total of 29 MDClusters. Looking at the clusters and displaying sample values for the feature uri_whois_reg, we observe the pattern in the WHOIS registrant email addresses we already found while looking at the 1D clusters in the previous section. It thus appears that the MDCluster #2 contains spam emails sharing this pattern in the WHOIS registrant email addresses.

We can take a closer look at the MDCluster #2 using the TRIAGEWEB MDCluster page as depicted in Figure 4.12. This page provides a pie chart with the distribution of combinations of features in the MDCluster (Figure 4.12 (1)) and a dynamic per-feature chart (Figure 4.12 (2)). The pie chart tells us that no less than 59% of spam emails were included in this group because they share at least similar values for the combination of features (i) sourceIp, (ii) uri_server_ip, (iii) uri_whois_reg and (iv) x_p0f_detail. This combination of features is thus quite strong in this MDCluster and indicates that spam emails were sent from hosts with similar IP addresses and operating systems and advertised URIs with domain names hosted on shared servers and registered using similar

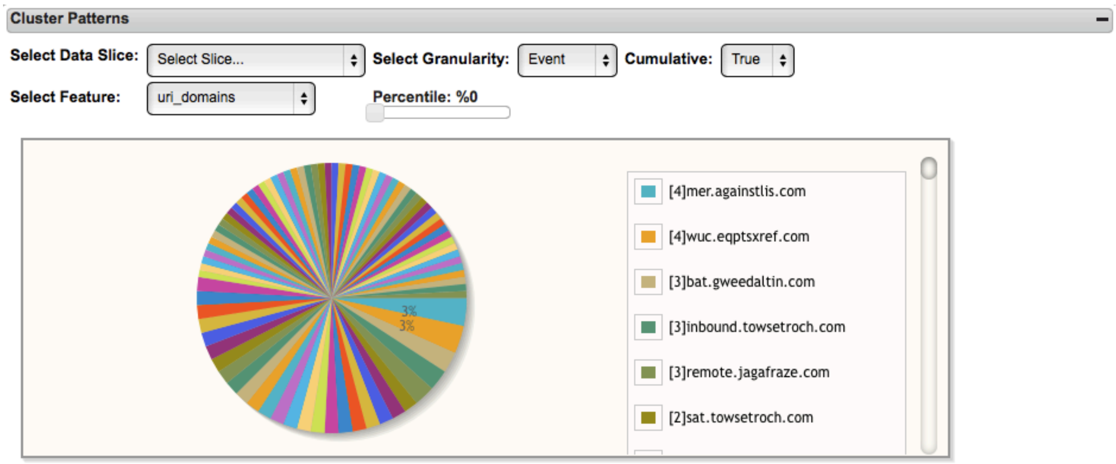
Cluster Patterns

Select Data Slice: Select Slice...

Select Granularity: Event

Cumulative: True

Select Feature: uri_domains

Percentile: %0

Figure 4.5: Cross-feature pattern analysis of a uri_whois_reg Cluster (viewed here wrt uri_domains)

email addresses (the WHOIS registrant email address pattern already observed before). TRIAGEWEB provides another way of looking at spam emails in our MDClusters, this time by looking at the distribution of spam emails per feature using dynamic charts (Figure 4.12 (2)). Figure 4.12 (2) shows the distribution of spam emails per spam sender IP prefix. We can see that a total of 7 prefixes are included in this MDCluster and that some sent much more spam that others. Figure 4.13 shows the distribution of spam emails per date. Interestingly, we can see that the spam campaign described by our MDCluster #2 was active everyday for 14 days from 2013-02-15 to 2013-02-28 and that more spam was received in the beginning of the campaign.

We now leverage the graph-based visualisation of MDClusters provided by TRIAGEWEB to look at the content of our spam campaign (MDCluster #2) and what insights into fly-by spammers operations we can uncover from it. From the graph in Figure 4.14 we can see that

- the campaign shared **very few** IP addresses of servers hosting spam advertised content (grey dots (1) in Figure 4.14);
- spam emails were sent from **7 different** hijacked prefixes (yellow dots (2) in Figure 4.14);
- **a lot of different** spam advertised URI domain names were used in spam sent from the different IP prefixes (blue dots (3) in Figure 4.14);

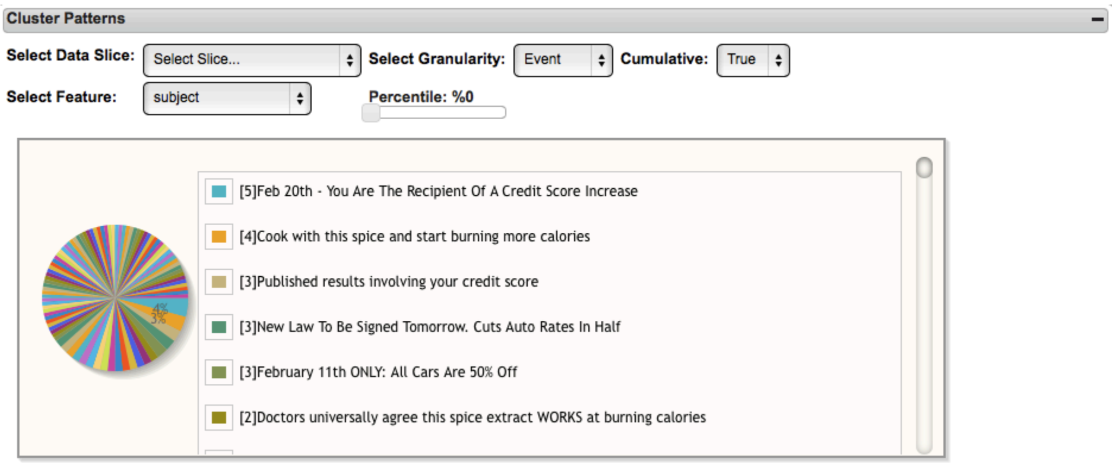

Cluster Patterns

Select Data Slice: Select Slice... Select Granularity: Event Cumulative: True

Select Feature: subject Percentile: %0

🟦 [5]Feb 20th - You Are The Recipient Of A Credit Score Increase

🟧 [4]Cook with this spice and start burning more calories

🟫 [3]Published results involving your credit score

🟩 [3]New Law To Be Signed Tomorrow. Cuts Auto Rates In Half

🟩 [3]February 11th ONLY: All Cars Are 50% Off

🟨 [2]Doctors universally agree this spice extract WORKS at burning calories

Figure 4.6: Cross-feature pattern analysis of a uri_whois_reg Cluster (viewed here wrt subject)

- URI domain names advertised in spam emails were registered using **many different** email addresses sharing the pattern [7-10 capital letters] @ WHOISPRIVACYPROTECT.COM (red dots (4) in Figure 4.14);

- **one main** WHOIS registrar was used to register all domain names of spam advertised URIs in the campaign (pink dot (5) in Figure 4.14);

- the campaign was active **everyday** for 14 days from 2013-02-15 to 2013-02-28 (purple dots (6) in Figure 4.14).

A closer look at some uri_whois_reg values in the graph clearly shows the pattern in the WHOIS registrant email addresses as shown in Figure 4.15. TRIAGEWEB provides another visualisation for the content of MDClusters, this time in a treemap fashion. Figure 4.16 shows the Treemap view of our MDCluster #2 for the spam campaign we are studying. This visualisation is complementary to the others and allows to very quickly have an idea of the diversity of values for each feature in an MDCluster. Thus, for our MDCluster #2, we can easily observe (i) the many diverse spam sender IP addresses, (ii) the limited number of shared IP addresses of servers hosting spam advertised content, (iii) the large number of spam email subjects and URI domain names and (iv) the prominent pattern in the WHOIS registrant email addresses used to register the different URI domain names.

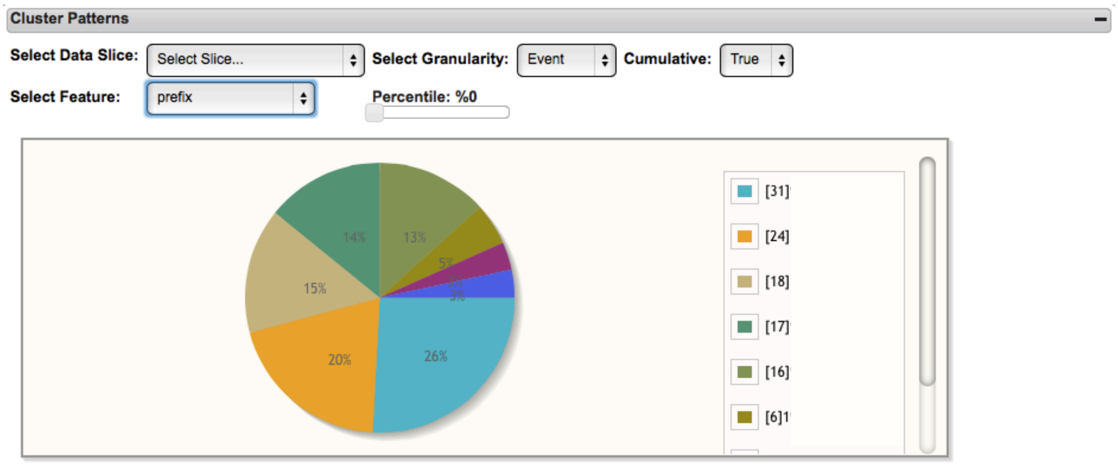

Cluster Patterns

Select Data Slice: Select Slice...

Select Granularity: Event

Cumulative: True

Select Feature: prefix

Percentile: %0

Figure 4.7: Cross-feature pattern analysis of a uri_whois_reg Cluster (viewed here wrt prefix)

In conclusion, from the evidence gathered during the analysis of this spam campaign, we strongly believe that those spam emails likely belong to a *single* spam campaign which is linked to the hijack of *seven* prefixes out of the 29 identified hijacked prefixes, possibly performed by a single cybercriminal organisation. Even though hijacked prefixes are apparently unrelated to each other, similar WHOIS registrant email addresses and shared servers hosting content advertised in spam emails appear to link some of them together. Fly-by spammers thus seem to hijack several prefixes over time and use each of them for a short period of time in order to circumvent spam sender blacklists and be able to continuously send spam and run their spam campaigns. It is very likely that other hijacked prefixes are linked in some way and TRIAGE combined with TRIAGEWEB will certainly be of great hep to uncover those other cases. In fact, without such an analysis we would have never been able to link some hijacks together based solely on the routing data that is used to detect the hijacks.

4.5 Summary

In this chapter we have presented an application of TRIAGE on spam emails sent from hijacked networks. By taking advantage of the myriad of visualisations provided by the TRIAGEWEB application we have been able to identify a spam campaign launched from several, apparently unrelated hijacked prefixes suggesting the hijacks were actually linked

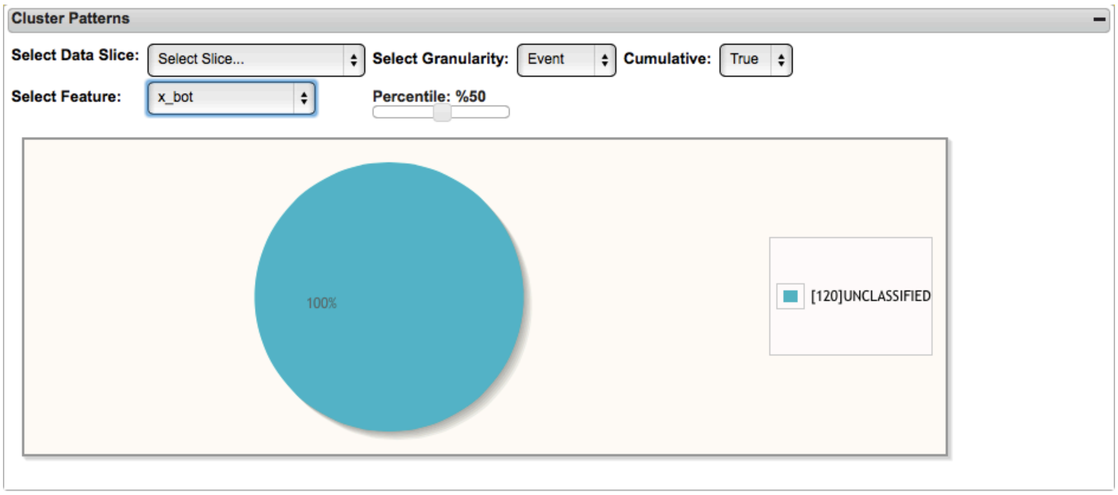

Cluster Patterns

Select Data Slice: Select Slice...

Select Granularity: Event

Cumulative: True

Select Feature: x_bot

Percentile: %50

Figure 4.8: Cross-feature pattern analysis of a uri_whois_reg Cluster (viewed here wrt x_bot)

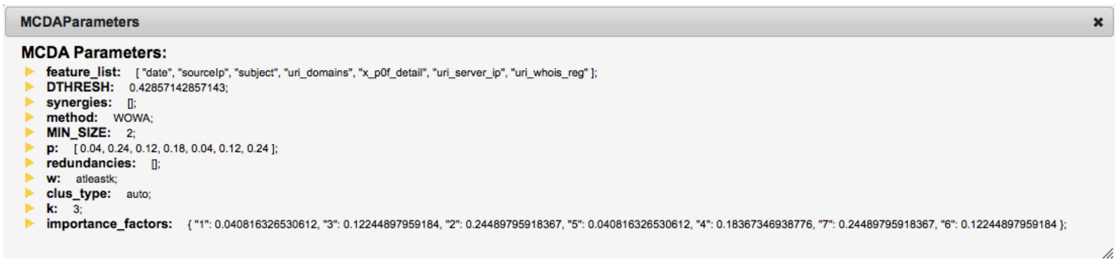

MCDAParameters

MCDA Parameters:

- **feature_list:** ["date", "sourceIp", "subject", "uri_domains", "x_p0f_detail", "uri_server_ip", "uri_whois_reg"];
- **DTHRESH:** 0.42857142857143;
- **synergies:** [];
- **method:** WOWA;
- **MIN_SIZE:** 2;
- **p:** [0.04, 0.24, 0.12, 0.18, 0.04, 0.12, 0.24];
- **redundancies:** [];
- **w:** atleastk;
- **clus_type:** auto;
- **k:** 3;
- **importance_factors:** { "1": 0.040816326530612, "3": 0.12244897959184, "2": 0.24489795918367, "5": 0.040816326530612, "4": 0.18367346938776, "7": 0.24489795918367, "6": 0.12244897959184 };

Figure 4.9: Window showing analysis parameters for the BGP hijacks data set

together and likely performed by the same cybercriminal organisation. Even though other findings and more insights into fly-by spammers operations could probably be uncovered by analyzing other MDClusters we presented another case where visualisations developed in the context of VIS-SENSE definitely help in analyzing security events.

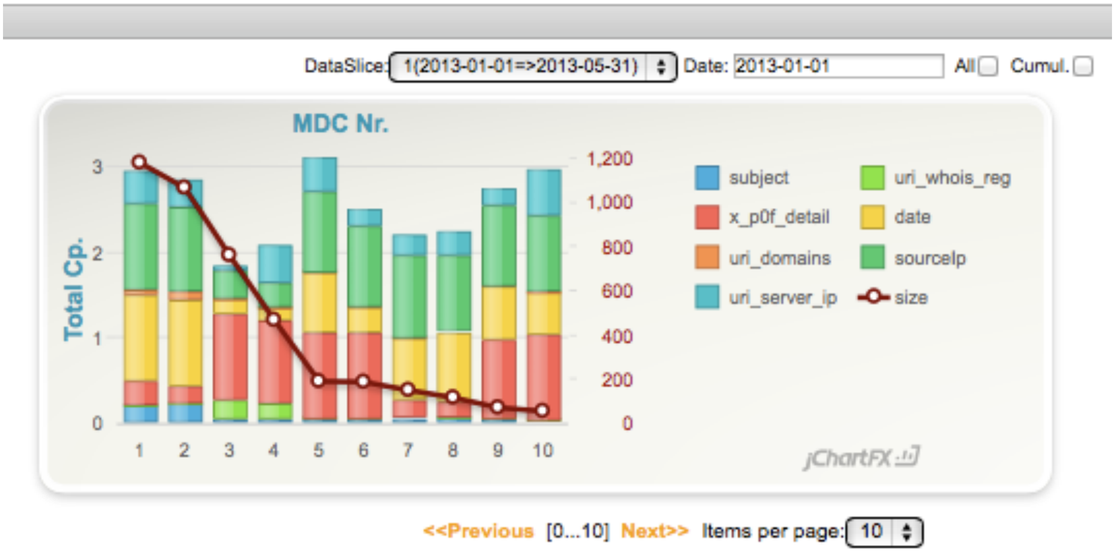

DataSlice: 1(2013-01-01=>2013-05-31) Date: 2013-01-01 All □ Cumul. □

<<Previous [0...10] Next>> Items per page: 10

Figure 4.10: The stacked bar chart showing the total MDC compactness (broken down by feature) for the top 10 MDCs (enabling a quick evaluation of MCDA data fusion results)

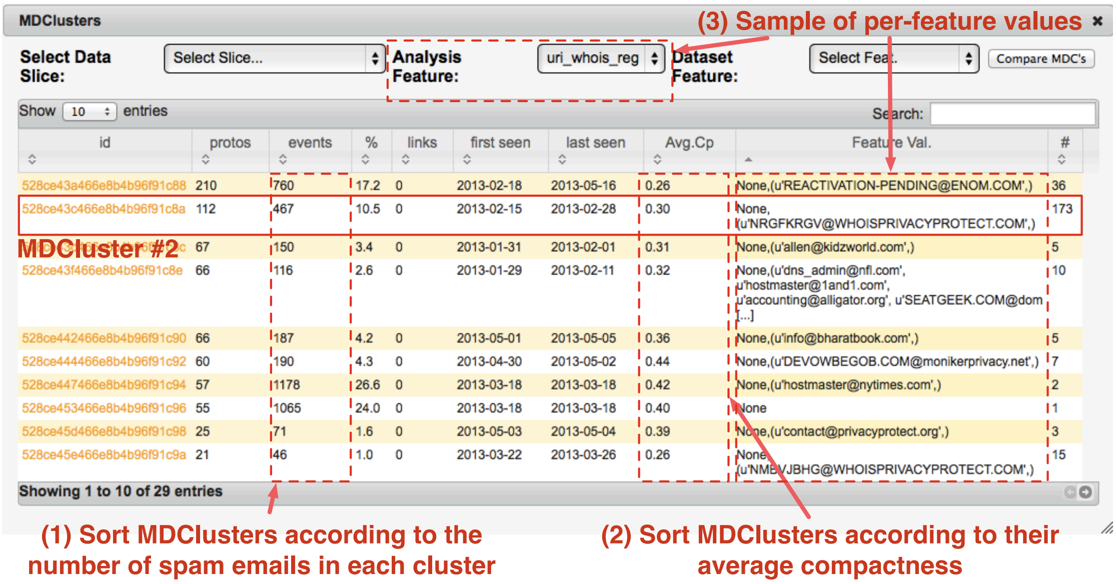

MDClusters

Select Data Slice: Select Slice...

Analysis Feature: uri_whois_reg

Dataset Feature: Select Feat..

Compare MDC's

Show 10 entries

Search:

id	protos	events	%	links	first seen	last seen	Avg.Cp	Feature Val.	#
528ce43a466e8b4b96f91c88	210	760	17.2	0	2013-02-18	2013-05-16	0.26	None,(u'REACTIVATION-PENDING@ENOM.COM',)	36
528ce43c466e8b4b96f91c8a	112	467	10.5	0	2013-02-15	2013-02-28	0.30	None, (u'NRGFKRGV@WHOISPRIVACYPROTECT.COM',)	173
528ce43[illegible]c	67	150	3.4	0	2013-01-31	2013-02-01	0.31	None,(u'allen@kidzworld.com',)	5
528ce43f466e8b4b96f91c8e	66	116	2.6	0	2013-01-29	2013-02-11	0.32	None,(u'dns_admin@nfl.com', u'hostmaster@1and1.com', u'accounting@alligator.org', u'SEATGEEK.COM@dom [...]	10
528ce442466e8b4b96f91c90	66	187	4.2	0	2013-05-01	2013-05-05	0.36	None,(u'info@bharatbook.com',)	5
528ce444466e8b4b96f91c92	60	190	4.3	0	2013-04-30	2013-05-02	0.44	None,(u'DEVOWBEGOB.COM@monikerprivacy.net',)	7
528ce447466e8b4b96f91c94	57	1178	26.6	0	2013-03-18	2013-03-18	0.42	None,(u'hostmaster@nytimes.com',)	2
528ce453466e8b4b96f91c96	55	1065	24.0	0	2013-03-18	2013-03-18	0.40	None	1
528ce45d466e8b4b96f91c98	25	71	1.6	0	2013-05-03	2013-05-04	0.39	None,(u'contact@privacyprotect.org',)	3
528ce45e466e8b4b96f91c9a	21	46	1.0	0	2013-03-22	2013-03-26	0.26	None, (u'NMBVJBHG@WHOISPRIVACYPROTECT.COM',)	15

Showing 1 to 10 of 29 entries

(3) Sample of per-feature values

MDCluster #2

(1) Sort MDClusters according to the number of spam emails in each cluster

(2) Sort MDClusters according to their average compactness

Figure 4.11: Dynamic data table showing MDClusters overview

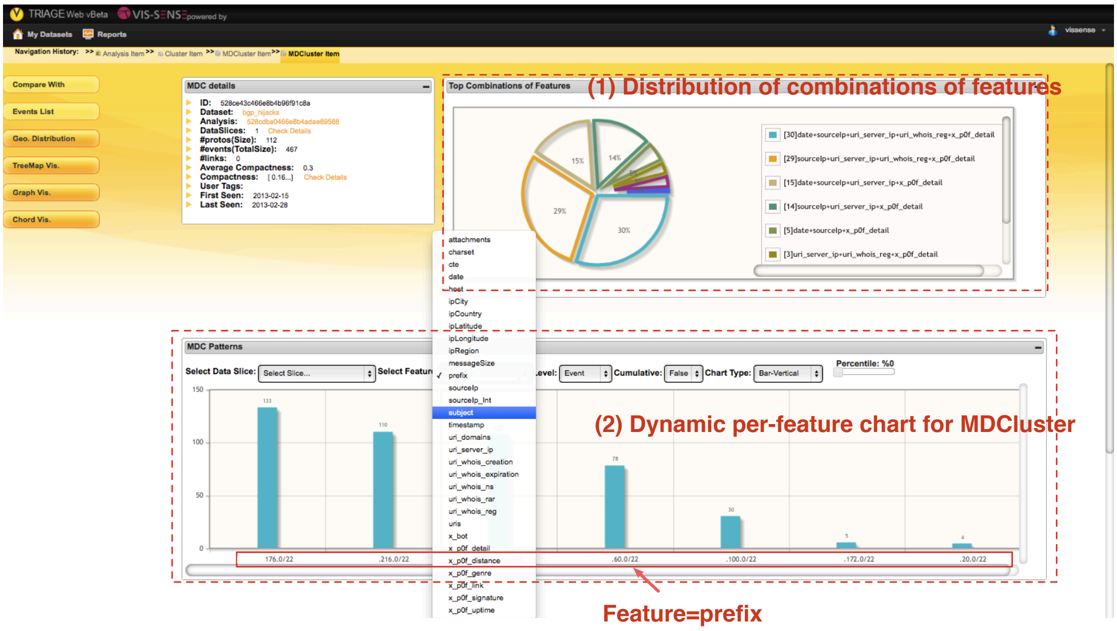

TRIAGE Web vBeta VIS-SENSE powered by

My Datasets Reports

Navigation History: >> Analysis Item >> Cluster Item >> MDCluster Item >> MDCluster Item

MDC details

- ID: 528ce43c466e8b4b96f91c8a
- Dataset: bgp_hijacks
- Analysis: 528cdba0466e8b4adae89588
- DataSlices: 1 Check Details
- #protos(Size): 112
- #events(TotalSize): 467
- #links: 0
- Average Compactness: 0.3
- Compactness: [0.16...] Check Details
- User Tags:
- First Seen: 2013-02-15
- Last Seen: 2013-02-28

MDC Patterns

Select Data Slice: Select Slice... Select Feature: prefix Level: Event Cumulative: False Chart Type: Bar-Vertical Percentile: %0

(1) Distribution of combinations of features

(2) Dynamic per-feature chart for MDCluster

Feature=prefix

Figure 4.12: MDCluster page showing the details of a specific MDCluster.

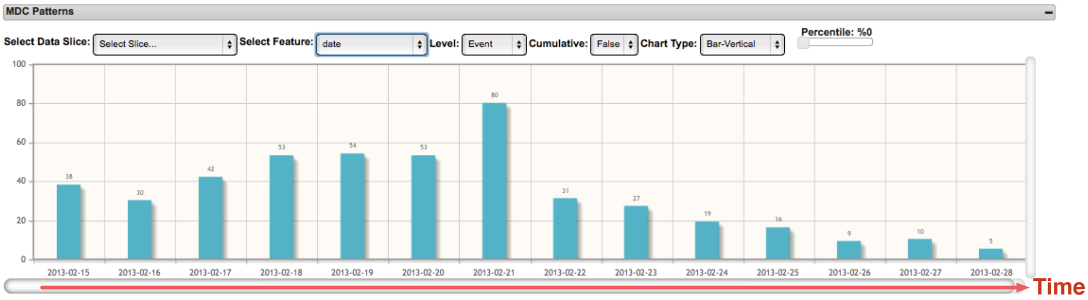

MDC Patterns

Select Data Slice: Select Slice... Select Feature: date Level: Event Cumulative: False Chart Type: Bar-Vertical Percentile: %0

Figure 4.13: Analysis of MDC patterns via dynamic chart tool.

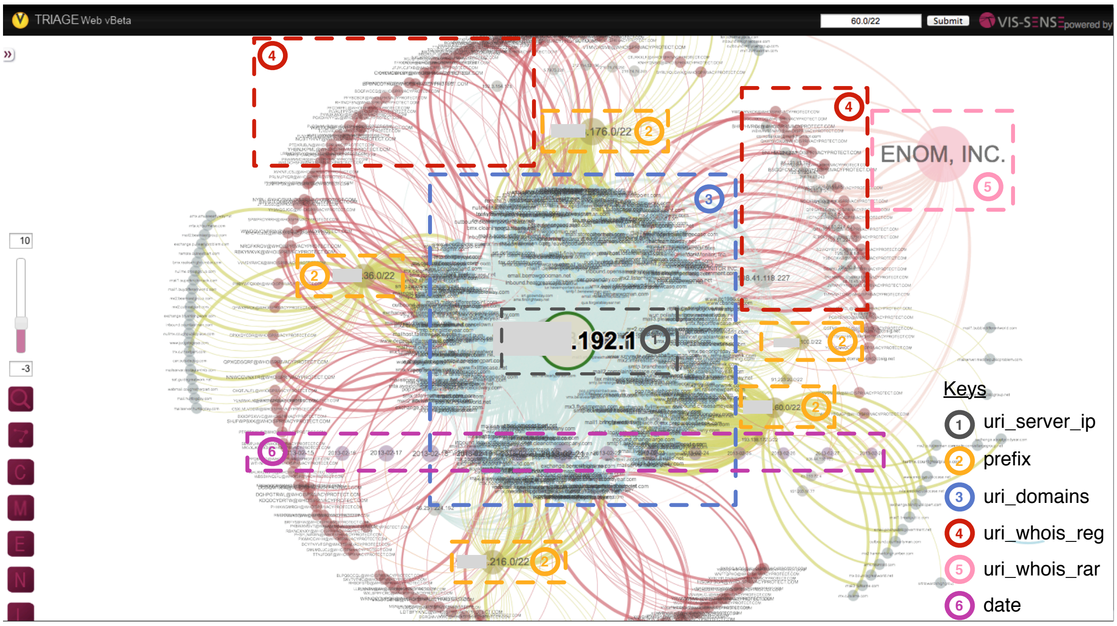

Figure 4.14: Visual analysis of an MDC through the interactive **Graph**-based viewer.

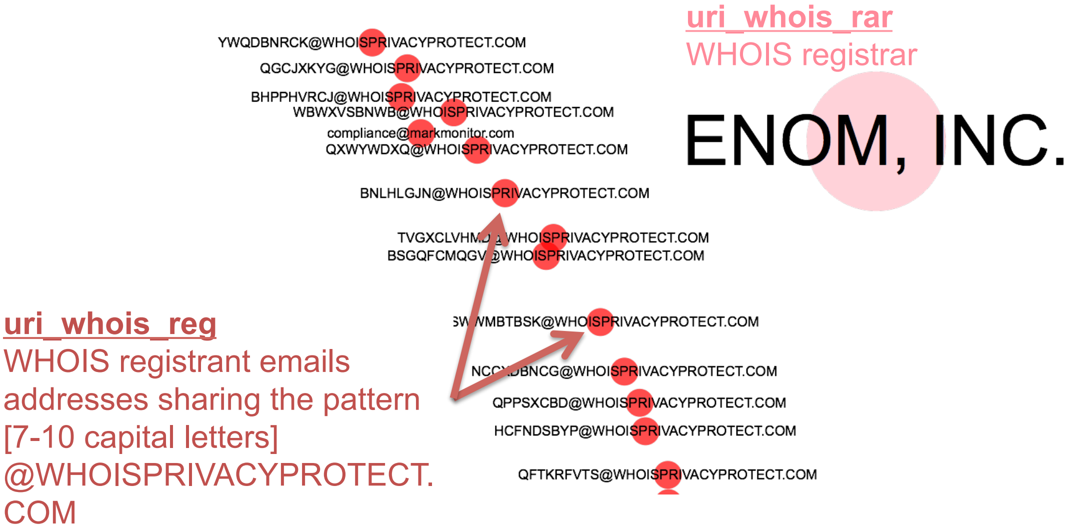
YWQDBNRCK@WHOISPRIVACYPROTECT.COM
QGCJXKYG@WHOISPRIVACYPROTECT.COM
BHPPHVRCJ@WHOISPRIVACYPROTECT.COM
WBWXVSBNWB@WHOISPRIVACYPROTECT.COM
compliance@markmonitor.com
QXWYWDXQ@WHOISPRIVACYPROTECT.COM
BNLHLGJN@WHOISPRIVACYPROTECT.COM
TVGXCLVHMD@WHOISPRIVACYPROTECT.COM
BSGQFCMQGV@WHOISPRIVACYPROTECT.COM
SWWMBTBSK@WHOISPRIVACYPROTECT.COM
NCCXDBNCG@WHOISPRIVACYPROTECT.COM
QPPSXCBD@WHOISPRIVACYPROTECT.COM
HCFNDSBYP@WHOISPRIVACYPROTECT.COM
QFTKRFVTS@WHOISPRIVACYPROTECT.COM
ENOM, INC.

uri_whois_rar
WHOIS registrar

uri_whois_reg
WHOIS registrant emails addresses sharing the pattern [7-10 capital letters] @WHOISPRIVACYPROTECT.COM

Figure 4.15: Visual analysis of an MDC through the interactive **Graph**-based viewer (snippet of a graph).

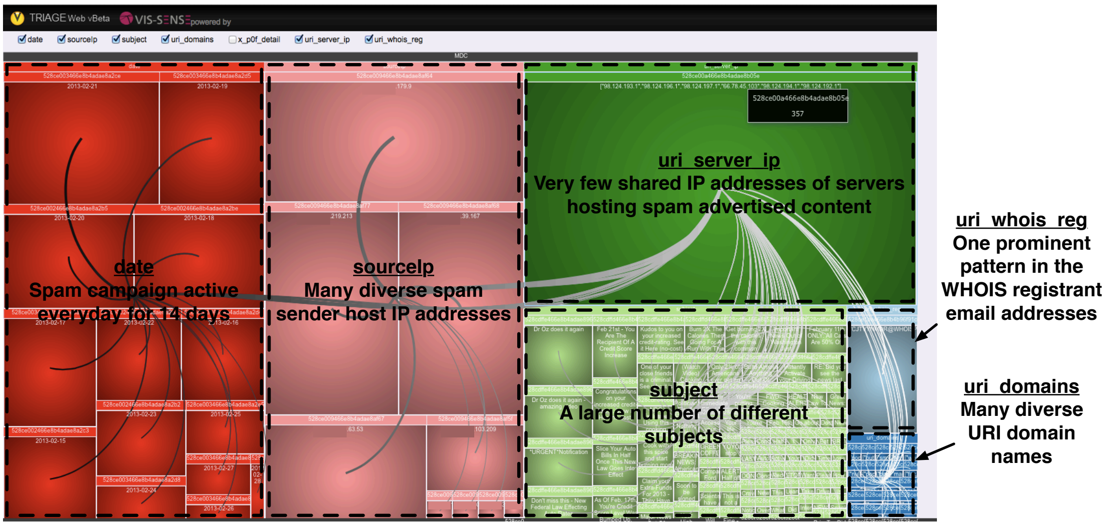

date
Spam campaign active everyday for 14 days

sourceIp
Many diverse spam sender host IP addresses

uri_server_ip
Very few shared IP addresses of servers hosting spam advertised content

subject
A large number of different subjects

uri_whois_reg
One prominent pattern in the WHOIS registrant email addresses

uri_domains
Many diverse URI domain names

Figure 4.16: Visual analysis of an MDC through the interactive **Treemap**-based viewer.

5 Conclusions

In this deliverable we have demonstrated how the R&D efforts carried out by VIS-SENSE partners have led to the development of visual analytics tools that can be used to perform intelligence analyses on security datasets, to generate new insights in a faster and more effective way. In particular, a web framework called TRIAGEWEB has been developed, as a fork of the full-fledged VIS-SENSE framework. We have shown how this web-based visual analytics framework helps security analysts in their quest to understand the common root causes of attack phenomena observed in the Internet, but also to generate insights into *modus operandi* of cyber criminals.

Our goal was also to evaluate the *usability* of the TRIAGEWEB framework and its *capabilities* to identify, represent and explain various Internet threat phenomena. to this aim, different applications of the framework were presented in this document and represent different instantiations of the first VIS-SENSE scenario that was defined in the start of the project (D1.2 - Use Case Analysis and User Scenarios):

- Visual Analysis of Web Threats Dynamics
- Visual Analysis of Scammers Operations
- Visual Analysis of Spam Sent from Hijacked Networks

The second application emphasizes the *broad applicability* of the VIS-SENSE tools, as it was performed on a completely new and unknown dataset made of *scam emails*, which was not initially integrated in our data collection infrastructure. Furthermore, this analysis was done in collaboration with a group of researchers external to the project and the insights generated by this analysis have been published in a joint paper and presented at the International Workshop on Cyber Crime (IWCC'13) – an IEEE Security & Privacy workshop [12].

The last application demonstrates also the *cross-domain* correlation capabilities of the VIS-SENSE tool, as our visual analytics tool were used successfully to analyze spam sent from *hijacked networks*, as detected by our BGP analysis tools – hence, fulfilling one of the key requirements of this project by establishing an explicit link between the *control* plane (BGP), and the *data* plane of the Internet. To the best of our knowledge, this is the first analysis of this kind that brings such clear evidence of correlation between these two aspects, but also highlights the modus operandi of so-called *fly-by spammers*.

Through these various applications, we could also evaluate the degree of fulfilment of the high-level functional requirements laid out in the beginning of the project. An overview of this evaluation is given in Table 5.1. Note that a more detailed evaluation of the requirements in the light of the design specifications defined in the project is available in deliverable D6.3 (VIS-SENSE Framework Evaluation).

From the requirements evaluation provided in Table 5.1, but also in light of the results shown in this deliverable, we can conclude that:

(i) most of the requirements have been successfully fulfilled to a high degree;

(ii) in some cases, the VIS-SENSE partners have gone beyond requirements definition and user expectations, by developing and integrating additional functionalities that were not envisaged or specified at the design stage.

Evaluation of High-Level Requirements Fulfilment

In the table below we evaluate the degree of fulfilment of every functional requirement set in the project start during the design and specifications phase. To this aim, we use an evaluation scale as defined here after:

- **None**: requirements have not been met at all
- **Low**: partially meets user requirements and expectations, but still lacking important functionalities
- **Satisfactory**: meets the minimal set of user expectations
- **High**: fulfils completely all critical user requirements
- **Very high**: fulfils all requirements and exceeds user expectations

Table 5.1: Evaluation of functional requirements fulfilment

Category	**Requirement**	**Degree of fulfilment**	**Comments**
Network analytics	support for the definition of MCDA and clustering algorithms	*High*	TRIAGE was enriched with various methods to help defining MCDA parameters
	improved scalability of the framework	*High*	A *prototype*-based clustering technique added to TRIAGE
	incremental analysis	*High*	Ability to run incremental analysis across data slices, on both on Clusters and MD-Clusters
	support for the detection of changes in the modus operandi	*Satisfactory*	Added patterns *search* functionalities in all types of TRIAGE objects - Change detection can be performed via identification of new MDCs (*e.g.*, in more recent data slices)

Category	**Requirement**	**Degree of fulfilment**	**Comments**
Visualization	advanced interactive visualizations for graph-based representations	*Very high*	The web-based Graph Viewer provides an extended set of functionalities to interact with graph visualizations
	visual feature selection	*High*	Visual charts and statistics are provided for 1D Clusters, helping the analyst understand the underlying dataset structure
	analysis of feature interdependencies	*Satisfactory*	Enabled through the cross-feature analysis of 1D Clusters and the analysis of features coalitions for MDClusters
	suitability to represent very large and complex networks	*High*	More scalable visualizations are provided besides the Graph-based viewer (Treemap-based viewer, Chord diagram, Matrix visualizations)
	visual exploration at varying levels of detail	*High*	All visualizations can be generated at two different levels: *prototype* or *event* level
	understanding of complex feature relationships (interactivity & sense-making)	*Very high*	The interactive visualizations tools, combined with the visual charts and dynamic plotting tools, enable the analyst to comprehend the *root causes* of MDC creation in much less time than before
System integration	integration of remote API's (*e.g.*, REST, WAPI)	*High*	WAPI v2 provides a RESTful interface and has been integrated in the framework
	provide interfaces for the tight coupling of software modules	*High*	Standard interfaces and protocols have been successfully used to provide a robust and efficient integration of all components

Bibliography

[1] 419 Scam Fake Lottery Fraud Phone Directory. http://www.419scam.org/419-by-phone.htm.

[2] Definition of 419 scam. http://www.419scam.org/419scam.htm.

[3] Definition of Nigerian scam. http://en.wikipedia.org/wiki/Nigerian_scam.

[4] NANOG (North American Network Operators' Group) Mailing List. http://www.merit.edu/nanog/.

[5] RIPE Mailing Lists. http://www.ripe.net/ripe/mail/ripe-mailing-lists/.

[6] Microsoft Security Intelligence Report. http://www.microsoft.com/security/sir/archive/default.aspx, 2008-2012.

[7] D. S. Anderson, C. Fleizach, S. Savage, and G. M. Voelker. *Spamscatter: Characterizing internet scam hosting infrastructure.* PhD thesis, University of California, San Diego, 2007.

[8] J. Buchanan and A. J. Grant. Investigating and Prosecuting Nigerian Fraud. *High Tech and Investment Fraud*, 2001.

[9] A. Costin, J. Isacenkova, M. Balduzzi, A. Francillon, and D. Balzarotti. The role of phone numbers in understanding cyber-crime. In *PST 2013, 11th International Conference on Privacy, Security and Trust*, Tarragona, Spain, 2013.

[10] M. Cova, C. Leita, O. Thonnard, A. D. Keromytis, and M. Dacier. An Analysis of Rogue AV Campaigns. In *Proceedings of the 13th international conference on Recent advances in intrusion detection*, RAID'10, pages 442–463, Berlin, Heidelberg, 2010. Springer-Verlag.

[11] X. Hu and Z. M. Mao. Accurate Real-time Identification of IP Prefix Hijacking. In *SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy*, pages 3–17, Washington, DC, USA, 2007. IEEE Computer Society.

[12] J. Isacenkova, O. Thonnard, A. Costin, D. Balzarotti, A. Francillon, and F. Eurecom. Inside the scam jungle: A closer look at 419 scam email operations. In *International Workshop on Cyber Crime (IWCC 2013)*, 2013.

[13] C. Leita and M. Cova. Harmur: Storing and analyzing historic data on malicious domains. In *Proceedings of the First Workshop on Building Analysis Datasets and Gathering Experience Returns for Security*, pages 46–53. ACM, 2011.

[14] A. Pathak, F. Qian, Y. C. Hu, Z. M. Mao, and S. Ranjan. Botnet spam campaigns can be long lasting: Evidence, implications, and analysis. *SIGMETRICS*, pages 13–24, 2009.

[15] M. A. Rajab, J. Zarfoss, F. Monrose, and A. Terzis. A multifaceted approach to understanding the botnet phenomenon. In J. M. Almeida, V. A. F. Almeida, and P. Barford, editors, *Internet Measurement Conference*, pages 41–52. ACM, 2006.

[16] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 291–302, New York, NY, USA, 2006. ACM.

[17] Symantec Corporation. Symantec Report on Rogue Security Software. Available online at `http://www.symantec.com/threatreport/archive.jsp`, October 2009.

[18] O. Thonnard. *A multi-criteria clustering approach to support attack attribution in cyberspace.* PhD thesis, École Doctorale d'Informatique, Télécommunications et Électronique de Paris, March 2010.

[19] P.-A. Vervier and O. Thonnard. SpamTracer: How Stealthy Are Spammers? In *The 5th IEEE International Traffic Monitoring and Analysis Workshop (TMA)*, Turin, Italy, Apr. 2013.