Overlay Network Measurements with Distribution Evolution and Geographical Visualization

Andri Lareida, Sebastian Schrepfer, Thomas Bocek, Burkhard Stiller
Communication Systems Group CSG@IfI, University of Zurich
Binzmühlestrasse 14, CH-8050 Zürich, Switzerland
Email: [lareida|bocek|stiller]@ifi.uzh.ch, sebschrepfer@hotmail.com

***Abstract*—Since Video-on-Demand (VoD) streaming and general overlay traffic, especially file sharing, are causing the main portion of today's Internet traffic, Internet Service Providers (ISP) face (a) challenges of network congestion during peak hours and (b) tussles between Video Content Providers (VCP) and ISPs, which can lead to congestion. However, dedicated traffic traces are not available today due to (a) the lacking interest of VCPs to share their internal data, *e.g.*, avoiding problems with privacy issues, and (b) the unavailability of measurement studies for overlay networks that are neither too specific nor cover more than a snap-shot of the network, which does reflect the essential traffic evolution over time. Thus, this paper bridges this gap between suitable traffic data and a video consumption analysis by (1) the development and prototyping of a new system to continuously monitor Video Consumption in Overlay Networks (VIOLA), resulting in an extensive and comprehensive measurement data set for detailed network evolutions over time in case of BitTorrent, and (2) a generic and reusable geographic visualization approach termed GeoChart, which aggregates selected data into a generic data structure to visualize them on a map.**

I. Introduction

Video-on-Demand (VoD) streaming is causing the main part of Internet traffic, but overlay traffic and especially file sharing and streaming applications still account for a large portion of the total Internet traffic [6], [24]. Thus, Internet Service Providers (ISP) face the challenges of network congestion during peak hours and tussles between Video Content Providers (VCP) and ISPs, which can lead to congestion as well [7]. Furthermore, overlay networks used by file sharing applications, such as BitTorrent (BT) [4] or the Akamai NetSession Interface, cause expensive inter-domain traffic due to the unawareness of the underlying physical network topology. To optimize caching solutions for VoD and locality mechanisms for file sharing detailed content demand traces are required, which are typically not publicly available today due to two main reasons. First, VCPs have no interest to share their internal data, as they could face privacy issues. Second, studies measuring overlay networks are either too specific or cover only a snap-shot of the network, which does not reflect evolution over time. Thus, to study video consumption in detail, an extensive and comprehensive measurement data set is required.

Since VCPs have no incentives to share their data sets the second best option is to collect data from the 2nd largest Internet traffic category [27], which is file sharing overlay networks. Since, these are mainly used for sharing video files and availability of content is equal or even better than in VoD services, leading to VoD providers analyzing P2P systems to find popular content [14]. Typically VoD services are free or available for a flat fee, *e.g.*, YouTube or Netflix, and, therefore, are comparable to file sharing systems. Overlay networks have been measured before, *e.g.*, [12][30], but these data sets only represent snap-shots that cannot be used to optimize time-dependent mechanisms, such as [10], [18], or caching and locality characteristics of overlay networks in general, which are strongly connected to time due to the global nature of those networks.

Furthermore, suitable visualizations, such as choropleth maps, assist in analyzing the potential of mechanisms, *i.e.* [9]. Typically, network measurement data contains IP addresses that can be mapped to a geographic location [22]. Overlay network measurement data is no different in this regard and can, therefore, be visualized on a map. Comparing measurement data categorized according to countries defines a very common and generic use case. However, appropriate tools that are simple to use and freely available are sparse.

To fill the gap of continuous P2P measurement data this paper presents and validates a new system to continuously monitor VIdeo cOnsumption in overLAy Networks (VIOLA). The data collected for the validation is available at [19]. VIOLA is targeting at the most commonly used overlay network today — BT [27]. As a second contribution this paper develops and investigates a generic geographic visualization approach, termed GeoChart. The data collected by VIOLA is aggregated into the generic and reusable data structure defined by GeoChart forming the basis for the interactive choropleth visualization, which is available here [20].

The remainder of this paper is structured as follows. Section II presents related works, Section III and IV explain the key design and implementation choices made for VIOLA and GeoChart. Finally, Section V presents key results of measurements taken between April 7 and 20, 2015 and adds examples on how this data is visualized and how it can help to draw conclusions.

II. Related Work

Overlay network measurements conducted in the past yielded valuable lessons learned, such as the aspect of random response of trackers or that not all nodes can be contacted because of NAT (Network Address Translation) problems, which have to be taken into account when designing a new measurement study targeting overlay networks. Specifically for BT measurements, respective techniques and pitfalls are well known [17]. P2P measurement studies have been conducted on different levels of abstraction, *e.g.*, as macroscopic, microscopic, and complementary ones [17]. Capturing the entire BT system in full detail is considered to be impossible due to the enormous number of peers and connections between them. BT measurements are always a tradeoff between resolution and coverage. Currently available data sets [30], [12], [26], [8], [5], [15] either have a high resolution or cover many swarms, but no data set exists, which covers the time, content, and user dimensions of BT.

On one hand, macroscopic studies of network traffic cover high level facets of participating nodes (or users) for large parts of, or even the whole overlay network. However, those measurements do not provide any details on overlay nodes' connections or performance. Technically, these studies mostly use crawlers for BT portals (such as Kickass Torrents [16]) and trackers, but do not connect to overlay nodes themselves to collect data. An example of distributed measurement of BT is [12] based on [11], which crawled trackers and index sites. However, snapshots or only high level data were collected. Only one of their experiments recorded detailed IP addresses over 88 h from 16 swarms. Another study [26] was measuring up to 100,000 swarms, in snapshots, which lack the time dimension. Although, these studies yield data sets containing millions of unique torrents and their nodes' IP addresses [30], those macroscopic studies only represent snapshots in time, which do not reflect any evolution of the network.

On the other hand, microscopic studies provide highly detailed insights into smaller parts of the overlay network, *e.g.*, inside a group or a swarm [8]. Unfortunately, those measurements do not document a complete picture of a network, mainly due to the effort that is required to collect those detailed insights for thousands of swarms. Microscopic studies typically follow a white box approach, in which the protocols under investigation need to be altered for explicit interactions with the measurement infrastructure. Thus, in case of overlay networks this approach involves a re-implementation of parts of the overlay management protocols, *e.g.*, the BT protocol. These altered clients are used to connect to overlay nodes and to receive internal information about them, such as client type, version, or bit-fields. Bit-fields are used to advertise what pieces of a file are available on a peer. With this technique connections between overlay nodes in traffic flows can be deduced [8]. Like their macroscopic counterparts, microscopic studies typically look at snapshots of the network only. When connecting to overlay nodes the problem of unreachable nodes arises due to NATs applied in many networks. Thus, even on a small scale it is very hard or even impossible to achieve a complete picture of connections and their performance between peers. This is not considered to be problematic for VIOLA, since it will collect data on a lower level of detail and, therefore, does not require to be able to connect to nodes at all.

Complementary approaches did augment those, micro- and macroscopic, measurements by adding DHT (Distributed Hash Tables) functionality to their crawlers [15] or by the use of plugins [5] for popular BT clients. A vantage point-based study [24] collected a large set of details about peers and how their interconnections, however, the collected data does not allow to identify the content shared, which is important information.

[8] performed a microscopic study in the BT network with a mechanism that infers connections between nodes. This approach delivers very detailed insights into the internals of a swarm, but is limited to a single swarm and covers only a time frame of 400 s, which is considered a snap-shot. A macroscopic study was conducted in [11] and later extended in [12], where torrents from the 3 most popular indexes at the time were collected and applied to query trackers with a distributed infrastructure to discover as many IP addresses as possible of a swarm. This process was repeated for every torrent found. The data is very detailed for a snap-shot of a swarm, but does not show its development over time or only for 16 swarms and not even 4 days.

In a similar study [30] indexes and related trackers were crawled to find IP addresses belonging to swarms. Additionally, the distributed trackers (Mainline DHT [21] and Azureus DHT [3]) were also queried. The study took 12 hours to complete one snap-shot, which is lacking the development over time dimension of a swarm's evolution again.

TABLE I. CLASSIFICATION OF EARLIER BT MEASUREMENT STUDIES AND THE NEW VIOLA APPROACH.

	Macro			Micro	Complementary		
Dimension	[30]	[12], [11]	[26]	[8]	[5], [24]	[15]	VIOLA
Content	✓	✓	✓	$\sim$	$\times$	✓	✓
Time	$\times$	$\sim$	$\times$	$\times$	✓	$\times$	✓
Peers	✓	✓	✓	✓	✓	✓	✓
Connections	$\times$	$\times$	$\times$	✓	✓	$\times$	$\times$

Table I provides an overview of measurement studies conducted in the past and VIOLA. The table shows that no study measured, so far, the dimensions content, time, and peers together, this gap is filled by VIOLA. Legend: ✓ fully covered, $\times$ not covered, $\sim$ limited coverage. Studies in the same column rely on the same data.

III. SYSTEMS DESIGN

Based on the design decisions taken for VIOLA and GeoChart, this section presents the use cases of both parts of the system followed by the key design descriptions.

A. Use Cases

VIOLA measures the BitTorrent Network in a distributed manner. In contrast to existing BitTorrent measurement systems (cf. Section II), which typically take snap-shots of the overlay network, VIOLA is able to monitor a large number of swarms over an extended period of time. VIOLA is deployed on one master node, which gives instructions to slave nodes deployed on smaller machines. The data gathered by slaves is returned to the master and stored in a database. VIOLA discovers torrents from a torrent portal and starts to measure those torrents discovered.

GeoChart visualizes country-referenced numerical data by relying on a well-defined relational data model depicted in Figure 1. GeoChart takes arbitrary data and produces a world map, where each country will be colored relatively to the value assigned to it. The main strength of GeoChart is the interactive interface that allows a user to select the data to be displayed and also lets the user change the color function used to determine the color of countries. The first option to use GeoChart is to feed the data directly into the view which will render the map with it. The second option is to use the

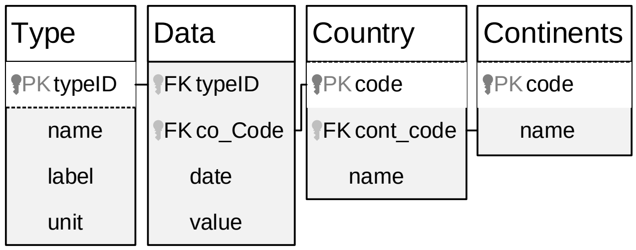

Fig. 1. GeoChart's data model. Country and Continents are static data.

provided back-end to read the data from a database following the relational data model.

B. VIOLA Systems Design

The BitTorrent network is orchestrated by trackers, which serve a maximum of 50 randomly selected IP address and port pairs of peers that share the torrent specified in the request. Additionally, BitTorrent has distributed trackers in the form of DHTs (Mainline and Vuze), which also return random IP address and port pairs. Therefore, measuring BitTorrent swarms is possible by querying trackers and DHTs. However, the random peer selection results in a lot of redundancy of peers discovered, especially for large swarms. Discovering all peers of a single swarm is an instantiation of the coupon collector problem [29]. To create a scalable system, a distributed approach was chosen, which consists of one master and an arbitrary number of slaves. The master is responsible for collecting newly published torrents, communicating relevant torrents to all slaves, and storing the results received in a data base. A slave is responsible for registering with the master, querying trackers and DHT for torrents received, and sending back results. While VIOLA has a specialized data model for its purpose GeoChart uses a generic data model that allows different data types to be entered.

The data model used in persisting the collected data follows the actual objects measured. The relational database consists of three tables *TORRENTS*, *ANNOUNCE RESULT*, and *PEERS*. The *TORRENTS* table contains information about the torrent itself, such as title and size. Furthermore, it contains meta data used in measurement, such as the *ACTIVE* flag defining if a torrent should still be measured. The *ANNOUNCE RESULT* stores announce meta data from announces executed by slaves, *e.g.*, the IP address of the slave executing the announce, the tracker the reply was received from, or the number of seeders and leechers reported by that tracker. Since a torrent is identified by the info hash, it is used as a foreign key to link the announce data to the torrent meta data. Finally, actual IP addresses of peers returned from the tracker are stored in the *PEERS* table, which contains among others IP address, port number, AS number, and the country code, which are resolved through geo IP databases from Maxmind [22].

C. GeoChart

GeoChart makes it possible to create individual, interactive choropleth world maps (*i.e.* Figure 3). The map is depicted with an equirectangular projection. Choropleth maps visualize one value per country. The color is determined by calculating the relative position of the value in the range of *[0, maximum value]*. A specific range of two predefined colors will represent the boundaries of this range. The data point will be converted into the matching color on the correct position in that range by a function.

The choropleth map uses color functions instead of classifications for the representation of values. Color functions do not suffer from some of the problems of fixed color classes, like loss of precision [2]. It is not required to define fixed ranges, but rather one function, which defines the complete course of the data set. To avoid a biased representation of the map, multiple color functions are provided. The color functions influence the relative position in the above-mentioned range of values. They were selected from the basic set of mathematical functions and can be extended easily. The logarithmic scale was chosen as the standard color function for the map since it does not react as strong to outliers than linear functions, leading to a general improvement of the color distribution. The user is able to select the desired color function in the interactive user interface.

GeoChart supports multiple dimensions of data. It allows to represent a list of different values on the world map. Individual data sets can be easily exchanged without reloading the map. It can for example visualize the distribution of response times of peers as well as the sum of peers in the respective country all in the same graph. The current dimension's details are listed in a sidebar, which overlays the map. Each country is listed with the relative value, the absolute value, and its color. The data can be downloaded into a CSV (Comma Separated Value) format to further process it. Dimensions can be switched within the sidebar overlay.

IV. Key Implementation Details

Driven by the design of VIOLA and GeoChart, selected and important samples of the implementation of the GeoChart visualization and the VIOLA measurement approach are discussed, to provide insights into this approach's scalability and performance.

A. Geochart

GeoChart was split in to parts, the front end (view) termed GeoChart.js and the back end termed GeoChartServer. The key details of GeoChart.js are presented before those of GeoChartServer. GeoChart.js is a library, which is based on Web technologies, *e.g.*, HTML, CSS, and JavaScript. It relies mainly on the JavaScript framework D3.js and the geo-data storing format TopoJSON. D3.js is a JavaScript framework which is used to create interactive graphs. D3.js is based on the SVG (Scalable Vector Graphics) standard. In order to create a map with D3, it is necessary to provide the map data in a JSON (Java Script Object Notation), either GeoJSON or TopoJSON, format. The JSON input has to be generated by conversion from a shape file [28], which contains the country shapes. These world-wide shape files are provided by Natural Earth Data [23]. The conversion of the data is done with the TopoJSON command-line tool. It allows to remove unnecessary meta data and creates a valid TopoJSON file. The TopoJSON file, which is required by GeoChart.js, needs to provide the map data and the country identifier code (ISO-3166-1) with the keyword *iso_a2*.

The GeoChartServer acts as the interface between a MySQL database and the JavaScript based view and is written in Java. The GeoChartServer-lib offers the basic functionality required to connect to a MySQL database and serve the content to GeoChart.js. It offers the *IDataAdaptor* and *IDataAdaptorFactory* interfaces, which have to be implemented to support other data sources. The GeoChart-lib implements all Web services required by GeoChart.js. Therefore, no additional programming is required to use GeoChart and it is easily extendable to integrate custom data sources by implementing the provided interfaces.

B. VIOLA

VIOLA relies on java sockets for communications. The master is implemented according to the "Architecture of a Highly Scalable NIO-Based Server" [1] to handle many slaves simultaneously. A custom solution was chosen to minimize overhead and gain full control of the socket handling.

The slave uses the tTorrent Java library [25] which offers full BitTorrent client functionality. VIOLA does only reuse torrent parsing and tracker querying functionality from tTorrent, since slaves do not actively participate in file sharing activities. Due to trackers being offline often or blocking slaves after too many queries, only UDP trackers are being queried to avoid timeout issues with TCP trackers.

BitTorrents Mainline DHT offers a very convenient interface for discovering peers of a swarm since it does not block IPs like central trackers. To connect to Mainline DHT the JKad library [13] is used. With JKad it is possible to receive over 500 IP addresses in response of one query, which is a tenfold increase compared to the central tracker's response being limited to 50 addresses.

C. Data Aggregation

The data collected by VIOLA needed to be aggregated to be feed into the GeoChart database. This meant transferring the BT data into the model shown in Figure 1. Since both are MySQL databases they were hosted on the same server and therefore the aggregation could be done by using SQL scripts, which were scheduled as daily tasks. To generate the view over all unique IP addresses on one day a query using the "distinct" aggregation function could be used and the results was grouped by country. The result set could then be inserted into the *Data* table of GeoChart using a predefined type. For the daily top 5 torrents the process was more complex. First, a similar query as for the overall addresses was made, which grouped the result after info hash and country. Based on this result the top 5 torrents had to be found by summing up all countries per torrent. Using that ranking the top 5 torrents' details were inserted into the *Data* table. If the torrents were not already in the *Type* table they had to be added.

V. Results

Since the VIOLA approach had been instantiated, the first set of results are presented, including key parameters and initial outcomes of a VIOLA measurement executed in April 2015. Based on these data collected an instance of the GeoChart approach was applied to visualize and perform an evaluation. In turn, the BT data newly collected is analyzed initially and compared to previous work.

TABLE II. Measurement parameters.

Start	07.04.2015 19:00
End	20.04.2015 11:00
Slaves	10
Interval	20 min
Portal	kickasstorrents.to
Category	Movies

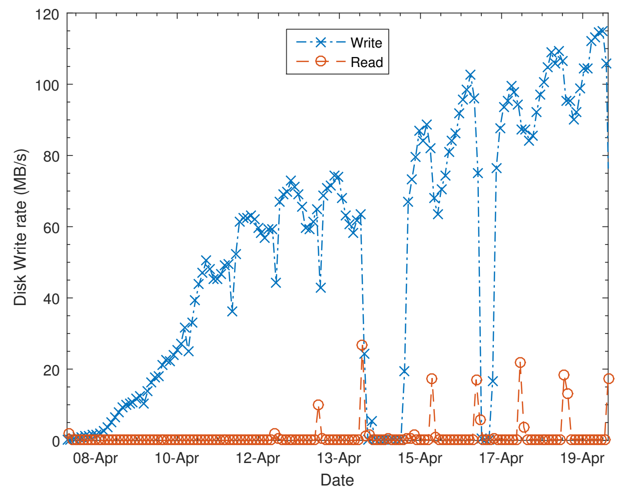

Fig. 2. Disk writes on the database server.

A. Data Collection

Table II shows the main parameters of the measurement executed. The measurement period started at 19:00 hours on April 7 and lasted until 11:00 hours on April 20. The number of VIOLA slaves used was 10, which were all located at the premises of the University of Zurich. The announce interval — the time in which each slave queries trackers of each torrent — was 20 minutes. New torrents were discovered from the Kickass Torrents portal, and only torrents released after the start of the measurement were considered.

In view of these settings, Figure 2 shows the disk read and write rate over the measurement period on the database server. From the beginning of the period the write rate starts to rise due to additional torrents being discovered and measured. This trend continues until the end, with some exceptions. The drops in the write rate can be attributed to the data aggregation job running daily at 10:00 hours. After April 12 these drops coincide with the spikes in the read rate. Due to the growing daily data size, data aggregation takes longer and involves more read operations. On April 14 around 19:00 hours a long query caused a table lock, which lead to the Master's random access memory running full, because events could not be written to the database. This caused a crash of the master which was only discovered and fixed the next day. On April 17 the problem appeared caused by the daily aggregation job. While the root cause for those crashed were analyzed and corrected, for the data set a gap in peer data during these 2 outages remains. Torrent discovery was not affected. During the 6 days before the first outage a total of 7,977,535 unique IP addresses were observed.

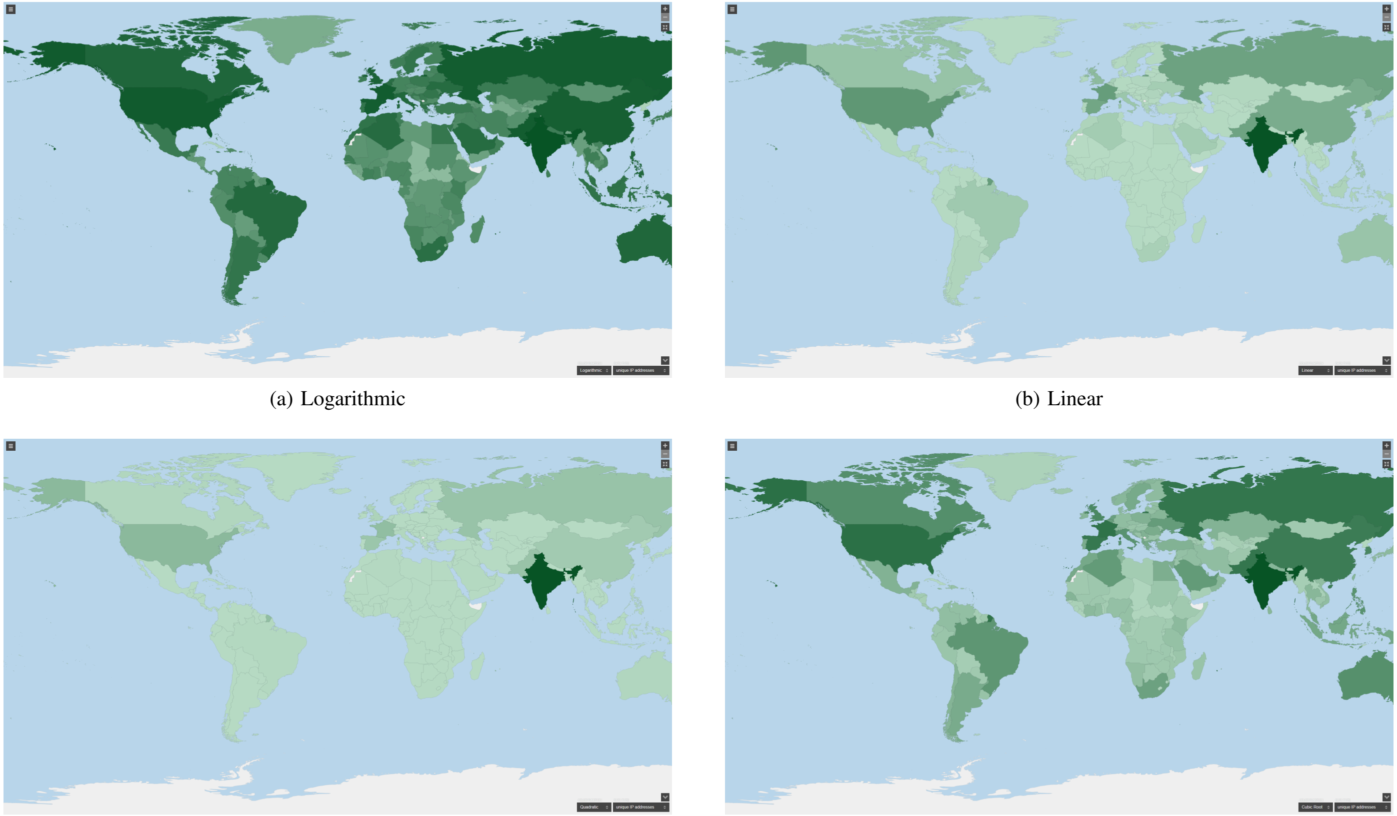

(a) Logarithmic

(b) Linear

(c) Quadratic

(d) Cubic root

Fig. 3. The same distribution of all unique IP addresses on April 13, 2015 with different color functions.

B. Evaluation of GeoChart

The data collected was aggregated on a daily basis to reflect unique IPs per country for all torrents and for the top 5 largest swarms individually. A respective interactive demo version of GeoChart based on the data set collected is available under [20]. Different data types were defined, one for IP addresses in all torrents and one for each torrent among the top 5 on a day during the measurement period, to allow users to select and display them. Figure 3 shows the data of April 13 in four different views using different color functions. The relative number of IP addresses for the top 3 countries are 13.5% India, 8% USA, and 7.4% Spain. The logarithmic function in Figure 3a shows a more details among the countries with little percentage, *e.g.*, inside Africa. Therefore, this function is suitable in case of very few high values and many small values being available that are close to each other. With the linear function in Figure 3b the finest detail is lost, but it is easier to distinguish the top countries, *e.g.*, India and USA. This function suits well for evenly distributed values. The quadratic function in Figure 3c emphasizes large values and, hence, should be used for data sets with few low values and many large values. The cubic root function 3d is similar to the logarithmic function keeping the fine details inside Africa, but also showing a clear difference between India and China. The square root function is similar to cubic root, preferring brighter colors though.

The examples of different color functions, shown next to each other in Figure 3, address he fact that there are significant

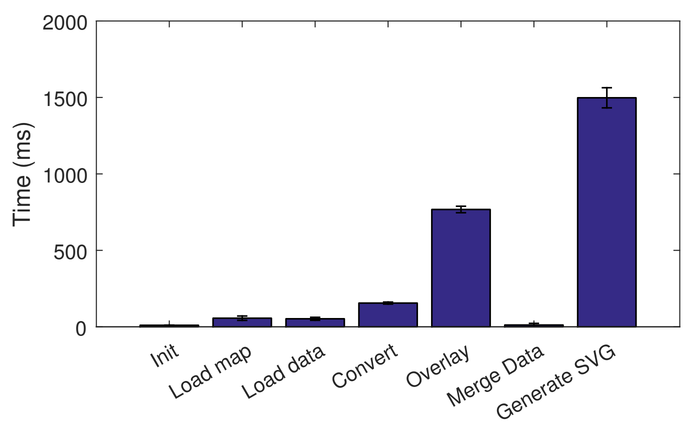

Fig. 4. GeoChart performance measurements.

differences in visual appearance between different functions. Therefore, possible studies can influence the conclusions drawn based on choropleth maps greatly by choosing the color function that suits their goals. An interactive representation of choropleth maps, as provided by GeoChart, is more transparent than a fixed map.

The performance of GeoChart.js depends mainly on the capability of the rendering device, *e.g.*, laptop, and the accuracy of the map data. The time consumption of the rendering process has been measured with $n = 19$ separate measurements on a consumer notebook (Asus ZenBook UX32VD) with Google

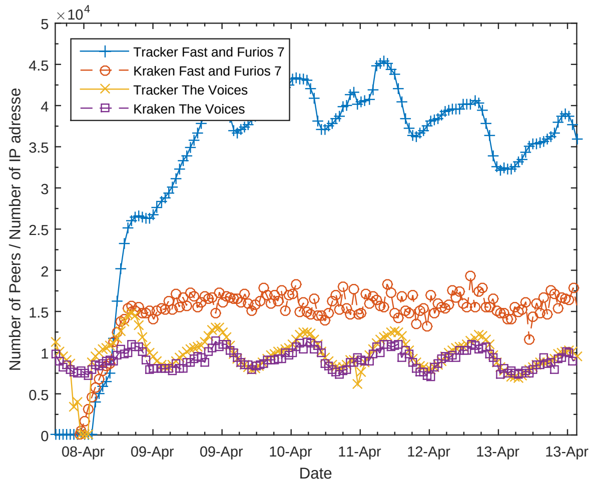

Fig. 5. Comparison of VIOLA discovered IP addresses versus Peers Reported by Trackers.

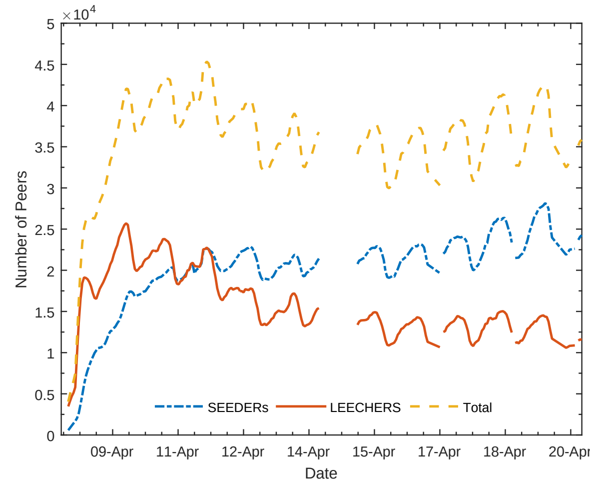

Fig. 6. Swarm composition of "Fast and Furios 7", the largest swarm measured.

Chrome 42.0.2311.135 m. The measurement was categorized into 7 steps (cf. Figure 4), whereas the computing of the SVG elements took the most amount of time (on average 1498 ms with a standard deviation of 66 ms).

Overall, the process of rendering the complete map took an average of 2,553 ms with a standard deviation of 86 ms. To decrease the rendering time of the SVG elements, the TopoJSON map can be replaced by a map with a lower resolution, which will lead to fewer zooming-in levels being readable and fewer calculations. Thus, this decrease of accuracy of the map on higher zooming levels cannot be avoided, but saves time.

C. VIOLA Accuracy

Collecting all IP addresses for a swarm is still a great challenge, because trackers will always reply with a random set of IP addresses from their pool, which is a variant of the coupon collector problem [29]. Issuing more queries to the tracker does not scale linearly, *i.e.* querying a tracker 10 times more will not yield 10 times more unique IP addresses, because of the random peer selection. Therefore, the larger a swarm is, the more difficult it becomes to discover all IP addresses. Although, it might not be possible to discover a complete swarm, it needs to be investigated which part of a swarm can be discovered.

Figure 5 shows a comparison between the number of peers reported from trackers and the number of unique IP addresses discovered by VIOLA for 2 torrents. The data from VIOLA as from Trackers is aggregated to hourly intervals. The first torrent is "Fast and Furious 7 2015 HD-TS XVID AC3 HQ Hive-CM8" (fast7), which is the largest torrent measured with an average of 34684 and a maximum of 45337 peers as reported by trackers. The second largest torrent is "The Voices 2014 TRUEFRENCH BRRip Xvid-BLUB avi" (voices) with an average of 8150 and a maximum of 14808 peers reported by trackers. For voices, VIOLA discovered almost the complete swarm except for the peak in the beginning. At some points VIOLA reported even more IP addresses (08:00 hours on April 8 and 03:00 hours on April 13). This is explained by tracker failures and the fact that some peers only use DHT trackers that cannot collect swarm statistics. For fast7 VIOLA can only cover about $1/3$ of the swarm due to the large size.

Considering that voices is the 16th largest out of 5,068 torrents monitored in that period, this observation reveals, that more than 99% of those torrents measured were accurately accounted for by VIOLA. By adding more slaves or decreasing the interval even more coverage can be achieved.

D. Data Set Analysis

For a high level overview of the measurement period the top 9 largest swarms covered by these measurements are displayed in Figure 7 according to their size over time as reported by trackers. Note that Figure 7 is split into two graphs, where for Figure 7a the y-axis' scale is twice the one of Figure 7b. The two gaps in the data, on April 14 and 17, were caused by system failures (cf. Section V-A). Besides the top 9 largest torrents shared during that period the plots show also the evolution of these torrents. All swarms exhibit fast growth in the beginning of their lifetime, as seen in flash crowds. However, the swarm size does not shrink rapidly over time as typical flash crowds do. Swarms follow a regular diurnal pattern. Interestingly, fast7 shrinks to its minimum on April 15 and reaches another peak on April 19. This shows that swarm sizes, *i.e.* content popularity, does not necessarily shrink with the age of the swarm.

Figure 6 provides a detailed insight into the composition of the largest swarm fast7, as reported by trackers. A swarm consists of seeders — peers that have the complete file — and leechers — peers that are still downloading the file. It took three days after release of the torrent until the number of seeders and leechers broke even. The amount of seeders is constantly increasing, while the number of leechers decreases after the initial peak. This means that leechers become seeders and do not immediately leave the system after they completed their download. Furthermore, the total number of peers increases again after April 15. Peers show an altruistic behavior and free riding is not a problem in this case.

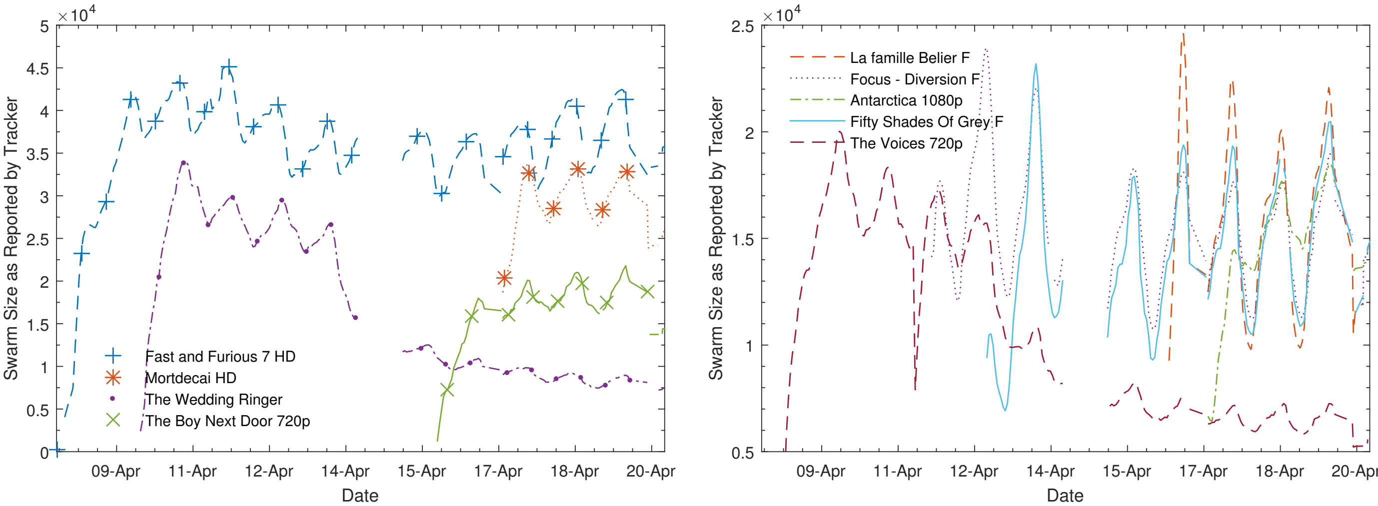

(a) Rank 1 to 4

(b) Rank 5 to 9

Fig. 7. The 9 largest swarms ranked after average swarm size.

Figure 7b shows, among others, three French movies — marked with an F in the legend —, which show a very similar pattern, although with different amplitudes. This observation shows the strong locality of these swarms. Since the time-wise behavior is almost equal they must be downloaded in the same time zone. This explanation is supported by the fact that these movies are in French language — indicated in the description of the torrent — and, therefore, mainly downloaded by users located in France. A look to the choropleth map produced by GeoChart, shown in Figure 8 with data from April 13, confirms these assumptions. The ranking proves that more than 75% of all peers were located in France, Belgium, and Algeria, which all share a common time zone. These results indicate that there is a clear relation between the pattern of swarm size over time and the locality of a swarm.

The maximal swarm size of "Fifty Shades of Grey F" from Figure 7a (approximately 28,000) and the number of unique IP addresses shown in Figure 8 for the same swarm (approximately 80,000) show a significant discrepancy. This difference is explained by the different meaning of the values. A tracker calculates the swarm size at any time, while GeoChart shows all unique IP addresses measured during the whole day. Therefore, there must be a high churn during a day, which leads to a swarm being completely replaced almost three times a day.

Figure 9 depicts the number of unique IP addresses measured per continent over the 24 hours of April 13. Although, India had the most unique IP Addresses on this day, Europe in total had more. The time zone patterns are clearly visible, even for those continents with few IP addresses, *e.g.*, Oceania (OC) and South America (SA). North America (NA) and SA are very much in sync with their peak at 04:00 hours, followed by Asia (AS) and Europe (EU). Europe, spanning 3 hours in time difference, has the narrowest peak while Asia, spanning 9 hours, has a very smooth peak. NA and the other continents with even fewer peers show smooth transitions as well.

Comparable measurements were undertaken by [24] in November 2008 and 2010. The data set used in [24] was collected from the viewpoint of individual peers and does not allow for the identification of the content shared. However, the continental diurnal patterns can be compared and show similar behavior. A notable difference is the order of continents with regard to the number of peers they contribute. While Europe was already the largest contributor, AS has overtaken NA. This trend was already visible form 2008 to 2010, were NA and AS were on the same level. The growth of AS compared to NA is influenced by several factors: file sharing prosecution in NA, emerging video streaming services in NA, or improved internet connectivity in AS. The VIOLA dataset, compared to [24], offers also details about content (cf. Table I), which allows to investigate individual swarms.

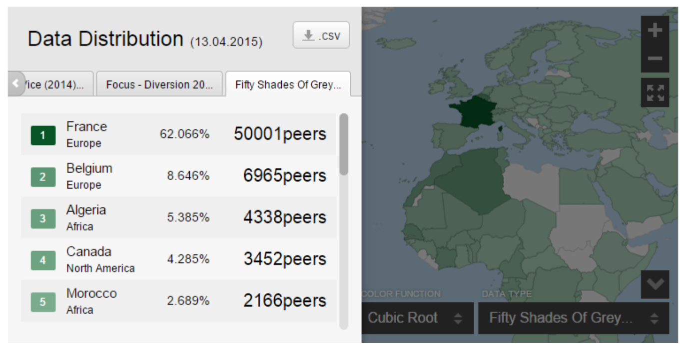

Fig. 8. Unique peer distribution on April 13 with cubic root for "Fifty Shades of Grey F".

VI. Summary and Conclusions

VIOLA distributed measurement system, which had been prototyped and tested on the BT network. The resulting data set covers the dimensions: content, peers, and time, which is especially unique for BT data sets due to the integration of time. VIOLA was active during a 13 days period in April 2015 and measured over 4,000 swarms and in the first 6 days almost 8 million. unique IP addresses were observed. The data set was also aggregated and used in conjunction with GeoChart

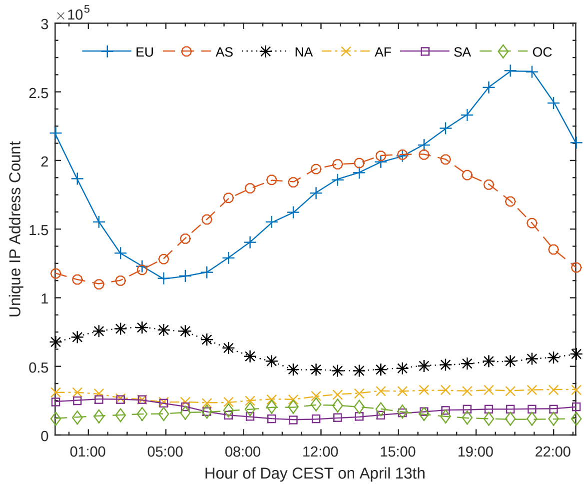

Fig. 9. Diurnal pattern of unique peers per continent on April 13.

to visualize the total unique IP addresses of one day and the top 5 largest torrents per day. For an easy handling GeoChart offers a Web-based interactive choropleth map with selectable color functions and data types.

The initial evaluations of VIOLA show that a large part of the BT network can be monitored accurately with this approach. Based on the VIOLA data set, it can be confirmed that BT swarms follow a distinct diurnal pattern, which varies between continents. Even for worldwide downloaded torrents this observation holds true. Thus, VIOLA makes it possible to analyze individual swarms over time *and* location, which led to the observation that locality of swarms is reflected in their daily peaks. The data shows that popularity of a content is not necessarily continuously shrinking after the peak popularity is reached. Furthermore, the data suggests as well that peers are altruistic and continue seeding a torrent after they finished the download, contrasting the assumption of the selfish user.

The interactive choropleth maps produced by GeoChart have provided a valuable range of interpretations in the analysis of country referenced data. Features, such as color function selection, zooming, and data type selection, enable the users to customize quickly a choropleth map to their needs. Thus, using different color functions allows for more transparency in those choropleth maps.

Acknowledgments

This work was supported partially by the SmartenIT and the FLAMINGO projects funded by the EU FP7 Program under Contract No. FP7-2012-ICT-317846 and No. FP7-2012-ICT-318488, respectively.

References

[1] C. Adamson, “Architecture of a Highly Scalable NIO-Based Server,” https://community.oracle.com/docs/DOC-983601, [Online, accessed 2015-05-05].

[2] G. Andrienko, N. Andrienko, A. Savinov, “Choropleth Maps: Classification Revisited,” The IEEE International Conference on Communications (2001), Helsinki, Finnland, June 2001.

[3] Azureus Software, Inc., “Distributed Hash Table,” https://wiki.vuze.com/w/Distributed_hash_table, [Online, accessed 2015-05-06].

[4] “BitTorrent,” http://www.bittorrent.com/, [Online, accessed 2015-08-26].

[5] D. Choffnes, F. Bustamante, “Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems,” in *ACM SIGCOMM Computer Communication Review*, Vol. 38, No. 4, October 2008.

[6] CISCO, “Cisco Visual Networking Index: Forecast and Methodology, 2012-2017,” http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/, [Online, accessed 2016-02-15].

[7] D. D. Clark, S. Bauer, W. Lehr, K. Claffy, A. D. Dhamdhere, B. Huffaker, M. Luckie, “Measurement and Analysis of Internet Interconnection and Congestion,” The 42nd Research Conference on Communication, Information, and Internet Policy (TPRC 2014), Arlington, VA, USA, September 2014.

[8] C. Decker, R. Eidenbenz, R. Wattenhofer, “Exploring and Improving BitTorrent Topologies,” IEEE Thirteenth International Conference on Peer-to-Peer Computing (P2P 2013), Trento, Italy, September 2013.

[9] Global Internet Maps, “DNSSEC Deployment Stage (ccTLD),” http://www.internetsociety.org/map/global-internet-report/#dnssec-deployment-stage-cctld, [Online, accessed 2016-01-15].

[10] F. Hecht, T. Bocek, B. Stiller, “B-Tracker: Improving Load Balancing and Efficiency in Distributed P2P Trackers,” IEEE International Conference on Peer-to-Peer Computing (P2P 2011), Kyoto, Japan, August/September 2011.

[11] T. Hoßfeld, D. Hock, S. Oechsner, F. Lehrieder, Z. Despotovic, W. Kellerer, M. Michel, “Measurement of BitTorrent Swarms and their AS Topologies,” University of Würzburg, Tech. Rep. 464, Nov. 2009.

[12] T. Hoßfeld, F. Lehrieder, D. Hock, S. Oechsner, Z. Despotovic, W. Kellerer, M. Michel, “Characterization of BitTorrent Swarms and their Distribution in the Internet,” *Computer Networks*, Vol. 55, No. 5, April 2011.

[13] K. Jünemann, P. Andelfinger, “JKad - A Java Implementation of the Kademlia Protocol, Designed for Academic Research,” https://dsn.tm.kit.edu/english/misc_3148.php, 2014, [Online, accessed 2015-05-05].

[14] N. Jurran, “Hinter den Kulissen von Netflix,” http://heise.de/-2304264, August 2014, [Online, accessed 2015-12-23].

[15] S. Kaune, R. C. Rumin, G. Tyson, A. Mauthe, C. Guerrero, R. Steinmetz, “Unraveling BitTorrent’s File Unavailability: Measurements and Analysis,” IEEE Tenth International Conference on Peer-to-Peer Computing (P2P 2010), Delft, The Netherlands, August 2010.

[16] kickasstorrents.to, http://kickasstorrents.to/, [Online, accessed 2015-05-06].

[17] M. Kryczka, R. Cuevas, C. Guerrero, A. Azcorra, A. Cuevas, “Measuring the BitTorrent Ecosystem: Techniques, Tips, and Tricks,” *IEEE Communications Magazine*, Vol. 49, No. 9, September 2011.

[18] A. Lareida, T. Bocek, M. Waldburger, B. Stiller, “RB-Tracker: A Fully Distributed, Replicating, Network-, and Topology-aware P2P CDN,” 2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013), Ghent, Belgium, May 2013.

[19] A. Lareida, “VIOLA - Video Consumption in Overlay Networks,” http://www.csg.uzh.ch/research/viola.html, [Online, accessed 2016-01-15].

[20] A. Lareida, S. Schrepfer, “Configurable Example of Geochart.js,” http://geochart.csg.uzh.ch, [Online, accessed 2015-12-23].

[21] A. Loewenstern, “DHT Protocol,” http://www.bittorrent.org/beps/bep_0005.html, [Online, accessed 2016-01-15].

[22] MaxMind, Inc., https://www.maxmind.com/, May 2015, [Online, accessed 2015-05-05].

[23] Natural Earth, “Natural earth data,” http://www.naturalearthdata.com/, [Online, accessed 2016-01-15].

[24] J. S. Otto, M. A. Sánchez, D. R. Choffnes, F. E. Bustamante, G. Siganos, “On Blind Mice and the Elephant: Understanding the Network Impact of a Large Distributed System,” ACM SIGCOMM 2011, Toronto, Ontario, Canada, 2011.

[25] M. Petazzoni, “Ttorrent, a Java Implementation of the BitTorrent Protocol,” https://github.com/mpetazzoni/ttorrent, [Online, accessed 2015-05-05].

[26] R. C. Rumin, N. Laoutaris, X. Yang, G. Siganos, P. Rodriguez, “Deep Diving into BitTorrent Locality,” IEEE INFOCOM 2011, Shanghai, China, April 2011.

[27] Sandvine Inc. UCL, "Global Internet Phenomena Report1H," http://www.sandvine.com/downloads/documents/Phenomena_1H_2013/Sandvine_Global_Internet_Phenomena_Report_1H_2013.pdf, [Online, accessed 2013-05-30].

[28] D. M. Theobald, "Understanding Topology and Shapefiles," http://www.esri.com/news/arcuser/0401/topo.html, [Online, accessed 2014-11-10].

[29] W. Xu, A. K. Tang, "A Generalized Coupon Collector Problem," *Journal of Applied Probability*, Vol. 48, No. 4, December 2011.

[30] C. Zhang, P. Dhungel, D. Wu, K. W. Ross, "Unraveling the BitTorrent Ecosystem," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 22, No. 7, July 2011.