EuroVis Workshop on Visual Analytics (2018)
C. Tominski and T. von Landesberger (Editors)

A Visual Analytics System for Managing Mobile Network Failures

M. Angelini[1], L. Bardone[2], M. Geymonat[2], M. Mirabelli[2], C. Remondino[3], G. Santucci[1], B. Stabellini[3], P. Tamborrini[3]

[1]University of Rome La Sapienza, Italy
[2]Telecom Italia, Italy
[3]Polytechnic of Turin, Italy

Abstract

Large mobile operators have to quickly react to mobile network failures to ensure service continuity and this task is a complex one, due to the continuous and very fast evolution of mobile networks: from 2G to 3G and onto LTE, each significant milestone in the mobile technology has increased the complexity of networks and services management. Failures must be promptly analyzed and sorted according to different prioritizing objectives, in order to devise suitable fix plans able to mitigate failures impact in terms of money loss or damaged reputation. This paper presents a visual analytics solution for supporting the failure management activities of TIM (Telecom Italia Group), the biggest Italian provider of telecommunications services with over 30M active mobile subscribers. The proposed system has been developed collaboratively by University of Rome "La Sapienza", Polytechnic of Turin, and TIM, analyzing the operators' requirements and viable optimization strategies for prioritizing interventions that rely on statistical data on mobile cells occupation, in order to identify the impact of failures in term of end users' connectivity.
Keywords: Visual Analytics, Mobile network failures, Telecommunications

1. Introduction

Maintenance of mobile networks is a challenging activity: failures must be promptly analyzed and sorted according to different prioritizing objectives, in order to reduce their impact in terms of money loss or damaged reputation. This paper presents a visual analytics solution developed for supporting the decision making processes of the TIM (Telecom Italia Group) operators in charge of detecting and fixing the mobile network failures. The main goal of the system is to provide means for exploring the actual and past states of failures and to monitor and prioritizing interventions based on impact measures. Actual solutions only dealt with technical failure characteristics (e.g., their severity) and fail to prioritize them in a profitable way (e.g., minimizing the number of customers that experiment disservices due to known failures). The paper contributes in supporting failure management by introducing optimization strategies for prioritizing interventions that rely on statistical data on mobile cell occupation. The system allows for getting an overview of cells geographical distribution (about 200.000 cells spread across the whole Italy), their presences values (i.e., the number of connected users), failures distribution and severity, according to different fixing strategies. The paper is structured as follows: Section 2 discusses the TIM scenario, Section 3 describes related proposals, Section 4 presents the implemented prototype, and Section 5 concludes the paper.

2. The TIM Scenario

As a national mobile operator, covering almost 100% of the Italian population and facing strong competition, TIM has to face a continuous and very fast evolution of its operations. From 2G to 3G and onto LTE, each significant milestone in the mobile technology has increased the complexity of network and service management. This paper refers to service assurance within the Radio Access Technologies (BTS, nodeB, eNodeB) and proposes a solution to network failures management during the everyday work of service and network operators. At this level, the massive amount of network elements and the variety of multi-vendor infrastructure equipments create a lot of metrics that could generate alarms such as call drops, resource congestion, access failure; the practical result is that operators are flooded with alarms and they have to deal with them with the following main goals: operators need to have a *geographic overview* of the number of cells, user presences, and alarms together with a *filterable* list of alarms, *prioritized* according to different objective functions, with the main objective of making decisions on how to cope with the current active alarms, planning and prioritizing interventions. The traditional work-flow relies on alarm classification and filtering based on severity metrics (e.g. critical, major, minor), but even the most experienced operator can still find it difficult to know where to start when there are many active alarms that may have different impact on the customer experience. A novel visual analytics solution, linking alarms with traffic data is a key point to improve the efficacy of the whole process.

DOI: 10.2312/eurova.20181108

3. Related Work

The application of visual analytics [KKEM10] [KBB*10] to spatio-temporal data is a well established field of research [AAD*10] [ABM*07] [AMM*07], with several contributions based on tiles [CSK*13] [WF09]. Naboulsi et al. [NFRS16] propose a survey on large mobile traffic analysis: with respect to their classification our solution is positioned in the category *Network Analysis* and subcategories *Aggregate Access Network Traffic* (for the presences data) and *Networking Solutions* (for the fixing strategies).

Less solutions coped with mobile data: among the major contributions, Van Den Elzen et al. [VDEBH*13] propose a multi-coordinated views environment for traffic analysis, while Angelini et al. [ACF*16] cope with economic decision making support based on mobile data analysis. GANTT-based representations have been used in network analysis (see, e.g., [HHH13] [XGH06]); however, many of these solutions only allow static inspection, partial interactivity, and only focus on temporal aspects, while our interactive solution provides prioritizing strategies and the continuous monitoring of the changes.

For what concerns failures and maintenance analysis, visual analytics proposals range from software maintenance [TEV10] to system maintenance [HSK*10] [JMO*16] in different domains; Matrosov et al. [MHK*15] propose a visual analytics system for exploring key trade-off for London's water supply; differently from our approach, the work uses mostly analytical representations, and focuses more on correlation among different dimensions than spatio-temporal analysis. Janetzko et al. [JSMK14] explore visual methods for anomaly detection, considering only the temporal aspects, while Motamedi et al. [MHA14] propose for the same topic only simple 3d spatial representations.

Finally, regarding automatic planning computation supported by visual inspection and exploration, several approaches exist; among them, Adrienko et al. [AAB08] support evacuation planning considering presences in disaster affected areas, while Kollat et al. [KRM11] use optimization for better network monitoring design: yet again, none of them copes with network failures in mobile communications.

4. The Visual Analytics System

The proposed system is based on a multi-coordinated views paradigm and an overview is presented in figure 1: it is composed of 3 main environments, the geographical analysis environment (A), the temporal analysis environment (B), and the analysis control panel (C). The visual encodings, agreed with TIM operators, reflect the way users deal with the problem: geo-localization of cells and presences and a GANTT like visualization, a visualization the users are familiar with for scheduling interventions.

4.1. The Geographical analysis environment

The main goal of the system is to allow for exploring the actual and past states of failure and monitoring and prioritizing interventions based on the impact a failure has in terms of the number of customers that experiment disservices (see, Section 4.4). The first challenge to overcome was how to visually represent data: TIM cells cardinality is in the range of 200.000 and representing them with a simple plot was discarded as option (requirements asked for a Web based implementation). The design choice was driven by the requirement of being able to identify the zone from which the failure alert was raised and to tie it with characteristics regarding the territory, like the number of cells or the number of presences. Given these requirements, the geographical analysis environment (labeled as A in Figure 1) has been implemented using square tiles of variable size (side of the square ranging from 0.5 to 32 km); each tile is representing the underlying cells and the tile color encodes either the number of cells or the number of presences, allowing for highlight similar areas. Failure markers are superimposed on the map with blue dots. This approach reduced the cardinality to manage from 200.000 elements to around 20.000 at worse, as reported in Table 1, making the application suitable for a Web based implementation. The geographical analysis environment can be customized

Tile size (in Km)	#Number of tiles
32x32	469
8x8	4727
4x4	9990
2x2	14119
1x1	18365
0.5x0.5	21117

Table 1: *Effect of tile size on the number of rendered tiles*

by removing the geographic layer, changing the color-scale (with color blind support) and the opacity of the tiles, searching for particular places through a quick-search function, and using normal and semantic zoom (see, e.g., [PDS*13]): the system allows to change the size of the tile map along preset dimensions; at the same time a normal zoom function is supported to focus the analysis on more specific areas. The two zoom functionalities can be linked by selecting "auto", with the system automatically resizing the tiles based on the actual zoom level. Finally, the system allows to filter data by technology (2G, 3G, LTE) using small multiples maps, and by statistical properties from the menu on the right. In particular it is possible to select the value associated with the tile map (presences or number of cells) and to select subsets of tiles based on uniform intervals (right filter) or box-plot intervals (left filter). The geographical environment supports the operator in exploring the actual state of the network, relating failure alerts to network characteristics (number of cells and presences).

4.2. Temporal analysis environment

Failures are not only characterized by their spatial distribution, but even by temporal aspects; for this reason the system supports a coordinated failures analysis environment based on a temporal representation. Failures are characterized by several dimensions; among others, the most important are their severity (critical, major, minor), affected technology (2G, 3G, 4G), status (open, solved) and starting and ending time of the relative alert. By existing design, data collection happens every fifteen minutes. This constraint, paired with the requirements of being able to identify still open alerts and to prioritize the interventions, led the design choice of a "GANTT-like"

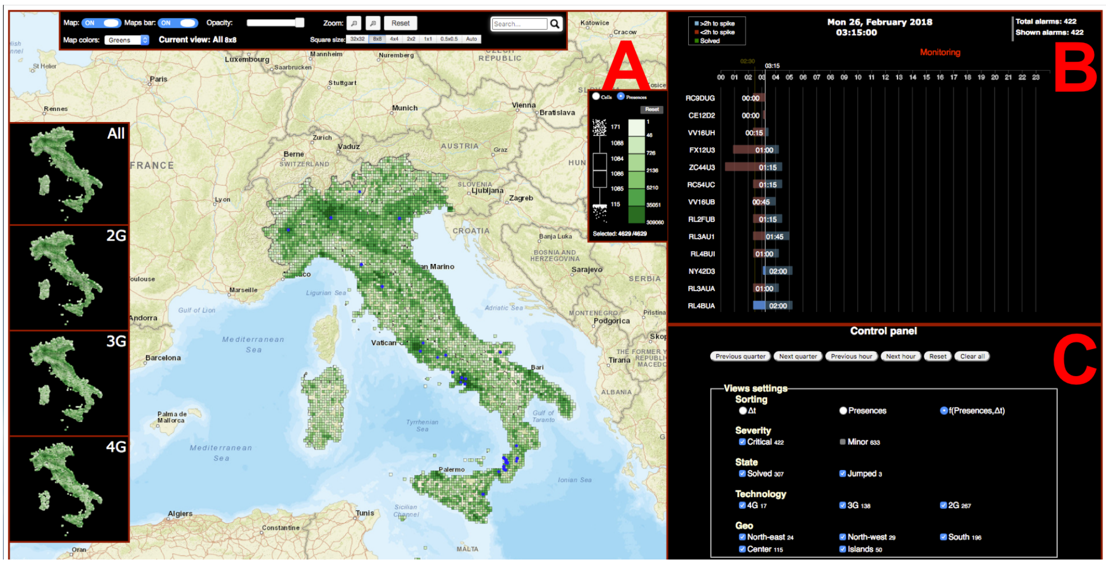

Figure 1: *Overview of the visual analytics system for failures analysis that includes three main views: the geographical analysis environment (A), for identification of zones with failures, the temporal analysis environment (B), for estimating the failures impact and prioritize interventions, and the analysis control panel (C), to parametrize the analysis.*

temporal visualization (see Figure 1 B). Failures are reported vertically, sorted by a criteria discussed in Section 4.4; X axis reports the 96 time-frames, each of 15 minutes, that compose a single day (reporting only 24 hours is dictated by historical data analysis showing that the large majority of failures are fixed within 24 hours). Each failure is represented by an horizontal bar whose length represents the distance between alert detection time and the time instant in which the impact on the service, depending on the affected cell, would be maximum if the failure is not fixed; in other words, the length encodes the amount of time available for fixing the failure without incurring the worst damage. When this time interval is under 2 hours the bar will be colored in red with blinking effect, otherwise it will be colored in blue. When a failure is solved, the corresponding bar will be colored in green. If a failure is not solved in time, the system will automatically compute the next maximum impact time and will reorder the failure accordingly. To keep track that a particular failure has already produced a maximum negative impact, an orange marker will be drawn on the bar at the time in which the impact occurred. This view allows to monitor the actual failures state and to inspect the proposed fixing order according to the chosen criteria, and to browse past failures to evaluate how good the fixing strategy worked, or how many failures were not fixed in time, or which was the impact of these failures on the quality of service. The potential impact and the resulting fixing strategy are modeled using the analytical component described in section 4.4.

4.3. Coordination and parametrization

Geographical and failure environments are coordinated: the operator can either filter the failures selecting subsets of tiles in the geographical environment, or select failures in the failure environment, having them highlighted in the geographical view. This coordination allows for comprehensive spatio-temporal explorations starting from the more relevant operator's perspective (geographical vs prioritized failure list). Filtering the failures (using the severity, geographical zone, type of connectivity, etc..) or the selection of the fixing strategy (taking into consideration time, impact, and coverage), is possible through the analysis control panel (see Figure 1 C). All together they let to visually explore the incidence of actual failures on the relevant characteristics (e.g., impacted presences, impacted areas) and the temporal dimension, like temporal distances from disservice peaks, and ordering of failures to support different fixing strategies. As an example of practical use, the operator can inspect the list of alerts sorted by the remaining time to the disservice peaks (alerts that have little time for being fixed) comparing it to the alerts sorted by impact on users (alerts that, if not fixed, will produce a disservice on a large number of users). Such data, together with the geographic location of alerts allow the operator to devise a fix schedule.

4.4. The Analytical Component

The main goal of the TIM operators is to monitor the current list of failures, quickly identifying the more relevant ones to prioritize interventions. Together with the users we have designed and implemented different fixing strategies that rely on statistical data about the number of cellular phones (i.e., presences) connected to a cell c within the time frames of 15 minutes q of every day. In particular, for each cell a 7 elements vector is maintained, one element week-day that contains 96 values (we have 96 fifteen minutes intervals per day) representing the weekly average of presences for each time interval in that day. We denote this statistical information with $PR(c, quarter_i, wd)$, while $MaxPR(c, wd)$ returns the maximum presences for the day. We use the term *maximum peak* at time t of cell c, to denote the time interval in the next 24 hours that exhibits the maximum number of connections. Using such an information the system prioritizes the active failures according to different strategies, described in the following.

- Time distance from the maximum peak. The system sorts all the failures according to *time distance* from the maximum peak. The idea is to prioritize interventions to fix failures that in a short time are going to impair a cell during its maximum usage.
- Affected presences. Maximum peaks are local maximums; if the goal is to minimize the impact on presences, disregarding the time, failures are sorted on the *number of presences* corresponding to the maximum peaks. In this way operators can focus on the failures that, within the next 24 hours, will produce the maximum impact.
- Time and presences. It is easy to spot situations in which local maximums will lead the second strategy to fail (e.g., the operator is focusing on a failure that will have a very high impact on presences in 10 hours, while three other failures will produce a greater cumulated impact in 3 hours). To deal with this issue, a third strategy has been defined, using a customizable function f(time, presences) that combines both time and daily peak presences:
$$f(time, presences, c) = \frac{p_1 * maxNumberOfPresences}{p_2 * timeToMaxPeak + 15} \quad (1)$$
and allows for tuning their relevance with two coefficients, p_1 and p_2, ranging in $(0,1]$.
- Local recovery capability. This strategy is under development and is the most challenging strategy required by the operators; it tries to forecast the recovery capability of the cellular network using working cells close to the failure point. Each cell has a nominal range that is affected by the topography (flat land, mountain, city), technology (2G, 3G, 4G), elevation and, for the sake of simplicity, that is represented by the radius of the covered area, $r(c) = f(topography_c, technology_c, elevation_c)$. We compute the residual presences, $RS(c, quarter, wd)$, as:
$$RS(c,q,wd) - \sum_i maxPR(a_i,wd) - PR(a_i,q,wd) \quad (2)$$
where $Distance(a_i,c) \leq r(c)$. The formula estimates the capability of the cells close to c to bear the c presences that cannot be accommodated due to the failure. According to the TIM operators this is a very relevant strategy: it allows to postpone the fix of cells that have a low *RS* and a critical maximum peak because the impact on customers is mitigated by the surrounding cells.

4.5. Evaluation

TIM operators used the system depicted on Figure 1 in an informal test environment, comparing it with the previously used solution, constituted by an environment that lists the failures in order of arrival and reports as a table their various characteristics. From the comparison they appreciated the easier way of coping with the monitoring aspects and get interested in enriching both the geographical environment (with values derived from failures) and the analytical engine (suggesting improvements regarding the historical analysis and subsequent failures sorting criteria). They were even able, by using our system, to spot areas of the covered territory where the data collection was acting strangely (e.g. not reporting the total number of expected cells) and correcting their data collection system. On the other hand, they underlined difficulties on exploring and reading data presented on the analysis control panel, Figure 1 *C*, highlighting the need for a better visual communication of control parameters. Additionally, they asked specifically to add mechanism for easily projecting the user knowledge in the system (e.g., the capability to search for specific territorial entities, like cities, or regions).

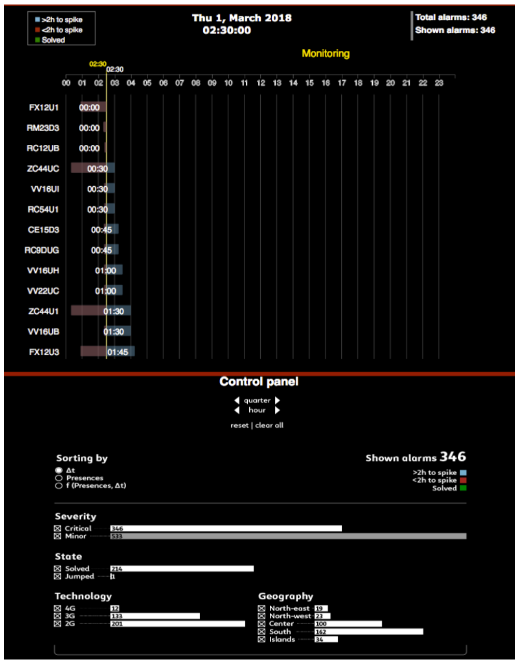

Figure 2: *Redesigned control panel, allowing for quicker comprehension of the actual parametrization and filtering capabilities*

5. Conclusion & Future Work

The paper presented a visual analytics solution for supporting the failure management activities of TIM, solution that has been developed analyzing the TIM operators' requirements and implementing optimization strategies for prioritizing interventions. Starting from the evaluation results, we are redesigning the control panel on the right side. The main goal of the operation is to respect the requirements of clarity, precision, and efficiency, working with affordance to offer to users an intuitive way to interact with parameters and memorize information, see Figure 2. Moreover, we plan to allow users to interact with the optimization strategies tweaking the formulas parameters, like the cell nominal range (they were skeptical about having the possibility of defining optimization functions from scratch). Moreover we are expanding the system capabilities toward classification of similar failures based on suitable similarity metrics and we are collecting statistics about the effectiveness of fixing activities in order to take them into account in the fixing strategy computation.

References

[AAB08] ANDRIENKO G., ANDRIENKO N., BARTLING U.: Interactive visual interfaces for evacuation planning. In *Proceedings of the working conference on Advanced visual interfaces* (2008), ACM, pp. 472–473. 2

[AAD*10] ANDRIENKO G., ANDRIENKO N., DEMSAR U., DRANSCH D., DYKES J., FABRIKANT S. I., JERN M., KRAAK M.-J., SCHUMANN H., TOMINSKI C.: Space, time and visual analytics. *International Journal of Geographical Information Science 24*, 10 (2010), 1577–1600. 2

[ABM*07] AIGNER W., BERTONE A., MIKSCH S., TOMINSKI C., SCHUMANN H.: Towards a conceptual framework for visual analytics of time and time-oriented data. In *Proceedings of the 39th conference on Winter simulation: 40 years! The best is yet to come* (2007), IEEE Press, pp. 721–729. 2

[ACF*16] ANGELINI M., CORRIERO R., FRANCESCHI F., GEYMONAT M., MIRABELLI M., REMONDINO C., SANTUCCI G., STABELLINI B.: A visual analytics system for mobile telecommunication marketing analysis. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2016), Eurographics Association, pp. 7–11. 2

[AMM*07] AIGNER W., MIKSCH S., MÜLLER W., SCHUMANN H., TOMINSKI C.: Visualizing time-oriented data – a systematic view. *Computers & Graphics 31*, 3 (2007), 401–409. 2

[CSK*13] CHENG D., SCHRETLEN P., KRONENFELD N., BOZOWSKY N., WRIGHT W.: Tile based visual analytics for twitter big data exploratory analysis. In *Big Data, 2013 IEEE International Conference on* (2013), IEEE, pp. 2–4. 2

[HHH13] HAO L., HEALEY C. G., HUTCHINSON S. E.: Flexible web visualization for alert-based network security analytics. In *Proceedings of the Tenth Workshop on Visualization for Cyber Security* (2013), ACM, pp. 33–40. 2

[HSK*10] HAO M. C., SHARMA R. K., KEIM D. A., DAYAL U., PATEL C., VENNELAKANTI R.: Application of visual analytics for thermal state management in large data centres. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 1895–1904. 2

[JMO*16] JÄGER A., MITTELSTÄDT S., OELKE D., SANDER S., PLATZ A., BOUWMAN G., KEIM D.: Lessons on combining topology and geography – visual analytics for electrical outage management. In *Proceedings of the EuroVis Workshop on Visual Analytics* (2016), Eurographics Association, pp. 1–5. 2

[JSMK14] JANETZKO H., STOFFEL F., MITTELSTÄDT S., KEIM D. A.: Anomaly detection for visual analytics of power consumption data. *Computers & Graphics 38* (2014), 27–37. 2

[KBB*10] KEIM D. A., BAK P., BERTINI E., OELKE D., SPRETKE D., ZIEGLER H.: Advanced visual analytics interfaces. In *Proceedings of the International Conference on Advanced Visual Interfaces* (2010), ACM, pp. 3–10. 2

[KKEM10] KEIM D. A., KOHLHAMMER J., ELLIS G., MANSMANN F.: *Mastering the information age-solving problems with visual analytics*. Florian Mansmann, 2010. 2

[KRM11] KOLLAT J. B., REED P. M., MAXWELL R.: Many-objective groundwater monitoring network design using bias-aware ensemble kalman filtering, evolutionary optimization, and visual analytics. *Water Resources Research 47*, 2 (2011). 2

[MHA14] MOTAMEDI A., HAMMAD A., ASEN Y.: Knowledge-assisted bim-based visual analytics for failure root cause detection in facilities management. *Automation in Construction 43* (2014), 73–83. 2

[MHK*15] MATROSOV E. S., HUSKOVA I., KASPRZYK J. R., HAROU J. J., LAMBERT C., REED P. M.: Many-objective optimization and visual analytics reveal key trade-offs for london's water supply. *Journal of Hydrology 531* (2015), 1040–1053. 2

[NFRS16] NABOULSI D., FIORE M., RIBOT S., STANICA R.: Large-scale mobile traffic analysis: a survey. *IEEE Communications Surveys & Tutorials 18*, 1 (2016), 124–161. 2

[PDS*13] PITTAPPILLY T. B., DEUTSCH R., SOEGIONO O. W., WAGGONER N. R., KUEHNLE H., KUSHNER M. H., CARR W. D., LUENGEN R. N., KWIATKOWSKI P. J., BARLOW A. G., ET AL.: Semantic zoom, Mar. 14 2013. US Patent App. 13/228,707. 2

[TEV10] TELEA A., ERSOY O., VOINEA L.: Visual analytics in software maintenance: Challenges and opportunities. *In Proceedings of EuroVAST, Eurographics* (2010), 65–70. 2

[VDEBH*13] VAN DEN ELZEN S., BLAAS J., HOLTEN D., BUENEN J.-K., VAN WIJK J. J., SPOUSTA R., MIAO A., SALA S., CHAN S., ET AL.: Exploration and analysis of massive mobile phone data: A layered visual analytics approach. In *Proceedings of the 3rd International Conference on the Analysis of Mobile Phone Datasets (NetMob'13). Boston, MA, USA* (2013). 2

[WF09] WILKINSON L., FRIENDLY M.: The history of the cluster heat map. *The American Statistician 63*, 2 (2009), 179–184. 2

[XGH06] XIAO L., GERTH J., HANRAHAN P.: Enhancing visual analysis of network traffic using a knowledge representation. In *IEEE Symposium On Visual Analytics Science And Technology* (2006), IEEE, pp. 107–114. 2