

Towards Non-expert Users Monitoring Networks and Services through Semantically Enhanced Visualizations

Owen Conlan, John Keeney, Cormac Hampson, F. Peter Williams
FAME & Knowledge and Data Engineering Group, Trinity College Dublin, Ireland
{Owen.Conlan, John.Keeney, hampsonc, williafp}@cs.tcd.ie

Abstract— Managing networks and services in large virtual enterprises often requires significant expertise to decipher and reconcile information coming from several heterogeneous sources. This presents a significant challenge as the person with the network management expertise is not necessarily the person who has the task or vision of how to maximize the value of these services. Conversely it would be valuable for dedicated network managers to have insight into the meaning of the services running over the network. This paper presents an approach to empowering users to monitor and eventually manage network devices and services that they are not necessarily expert in. This paper details this challenge, presents two enabling technological frameworks and offers two compelling case studies currently being addressed in the Federated, Autonomic Management of End-to-End Communications Services (FAME) project.

I. INTRODUCTION

Due to the continuously growing amount of monitoring data the question of what to show is increasingly becoming as important as how to show it [1]. When considering data from heterogeneous sources, relating the information so that it can be compared and prioritized becomes a significant issue. A semantic layer on top of these streams of data is required to provide a common framework and vocabulary so that comparisons can be drawn. However, a mechanism to dynamically structure the information is also needed to semantically prioritize it, so only the information that is most pertinent to the end user's needs is presented. By introducing appropriate semantics and prioritizations a non-expert user can be offered insights into the underlying information to enable them to monitor and manage underlying networks and services from which the information originates. The work presented in this paper presents an approach to capture and use domain expertise to support visualizations of such monitoring and management information. By applying domain knowledge to the structured exploration of information [2], users are empowered to delve beyond a set of static results to explore a set of meaningful responses.

The Semantic Attribute Reconciliation Architecture (SARA) [3] describes a system that utilizes the subjective expertise of experts to determine what information, from several sources, may be reconciled to highlight significance for the end user. This paper provides an overview of SARA and introduces the Insight Framework and Authoring Tool. Insight facilitates the creation of several visually prioritized information overviews with a sophisticated information exploration paradigm, to sequence and prioritize information for visualization.

This paper presents an overview of SARA and Insight in the

context of two case studies highlighting the challenges from the Federated, Autonomic Management of End-to-End Communications Services (FAME) [4] Irish Strategic Research Cluster. The paper begins by presenting the visual monitoring and management challenges faced in FAME.

II. OVERVIEW AND CHALLENGES

Harmonizing monitoring data requires mechanisms for mapping and federating the large volumes of low-level data produced by monitoring systems into semantically meaningful information and knowledge that may be discretely consumed by all parties involved in producing, delivering or consuming a communications service. Any complex system will have domain experts who are very familiar with how the constituent parts of the system can be managed, and are particularly aware of the end-to-end operating constraints of those constituent parts. Encoding this expertise in a manner that can be utilized by other stakeholders presents a significant challenge [5], and arguably requires significant advances in knowledge representation and engineering. The benefit of enabling such expertise to be encoded, aggregated and interpreted across diverse domains enables other stakeholders of other systems to monitor other parts of the system and relate their operation to their own system. It also enables common knowledge to be shared and reused across multiple domains.

From a practical point of view, there are three major problems with current monitoring and management techniques that this work attempts to solve. The first is the inherent heterogeneity of management information. For example, many systems use proprietary tools or command-line interfaces to configure devices, networks and services, but monitor their performance using the other monitoring tools like SNMP. This requires a skilled human operator that can understand both languages to ensure that an issued command is performed correctly, where this process can require days or weeks of tedious iterative fine-tuning. The second is the potentially large amount of data produced, as seen for example with “alarm storms” that are common in networks under stress. This leads to “information overload”, which complicates an already difficult task of managing networks, which is exacerbated when end-to-end concerns must be understood by different systems using different management philosophies. The third is the presence of stove-pipe based partitioning between end-to-end user quality of experience-focused monitoring and typically technologically- or geographically-bounded network or network element monitoring and configuration functions.

Personalized visualization of monitoring data, after it has been normalized and homogenized, will be key to including non-technical service consumers and provider-consumers in the end-to-end service quality of experience-based control loop [6]. Presenting and comprehending complex systems and their constituent data in a visual manner makes it much easier for the viewer to elicit knowledgeable conclusions from the information presented [7][8]. This is also essential to empower non-technical service consumers or suppliers to understand complex tasks and be able to abstract and contextualize possibly unfamiliar information from a viewpoint that makes sense to them [8][9]. Human interactions with complex systems, particularly for monitoring, management or tuning, tend to follow a pattern of negotiation [10], therefore the governance of such complex systems must be a dynamic and converging two-way process. However, the relevance of some information or mechanism is dependent on the context of the user and it must be contextualized for the user [9] in a way that the user appreciates its relevance. This has led to the popularity of personalized “dashboard” type monitoring applications for visualizing key performance indicators of managed systems [11].

III. SEMANTICS AND PRIORITIZATION FOR VISUALIZATION

When presented with a large volume of data it is easy for end users to become disoriented, especially when they are not looking for one specific thing. Through exploration a user can refine their needs whilst getting an impression of the information space. However, the paths of exploration need to be laid out using semantics that are accessible and familiar to the user. SARA [3] provides mechanisms for semantically reconciling heterogeneous information sources, such as databases and web services, into a coexistence model (see Fig. 1), i.e. SARA allows the schemas of each source to coexist and it performs the mappings as needed. SARA allows domain experts (even those with no computer coding experience) to encode domain concepts and their accompanying rules using metadata from the sources registered to SARA. It also provides a means to apply expert oriented subjective analysis and trends, in the form of semantic attributes, across this data to bring additional value to the end user. Semantic attributes [6] are discrete encodings of domain expertise that can be combined together and personalized to support user exploration of an information domain. They often act as abstractions and simplifications from the raw data that are intended to make it more accessible for the ordinary, non-expert user. Given the subjective nature of these semantic attributes, end users, or specifically the client tools they are using, have the opportunity to choose and balance between different semantic attributes. This enables a novice user to access an information space in accordance with one or several experts' perspectives [3]. SARA provides an authoring tool to enable the domain experts to create and combine semantic attributes. One of the key successes of SARA is that these semantic attributes may be further personalized to the user's specific needs by adjusting the parameters.

Insight is a middleware framework that sits between SARA and the client visualization tool that aims to support exploration of semantically enhanced visualizations (fig.1). It provides the capability to prioritize semantic and numerical information into several quantized views. These quantizations offer visualization clients layered access to the information, where the labeling of the layers can also introduce meaningful semantics. For example, all network transfers below a certain speed may be grouped into layers named ‘Slow’. The top layer in each view contains the most pertinent information as determined by a domain expert. Several layers may be selected by the client, thus showing a depth of information, e.g. all ‘Slow’ and some ‘Medium’ layers. The key benefit of Insight is the ability to manage and prioritize this depth in order to manage the end user’s attention and mitigate cognitive overload, i.e. being presented with more information than they can process. Insight also offers an authoring environment (Fig. 2) that is designed to enable the domain expert to define a set of reusable overviews of the data source using sequenced rules that combine filtering and data reordering techniques (i.e. to define views, and rules for view membership according to the values and combinations of semantic attributes and other numerical information values).

Semantic differentiation is accomplished by assigning responses from systems like SARA to a layered display hierarchy known as the Quantized Semantic Sort (QSS). Subjective comparators derived from domain semantics inform a structured set of assignment rules which populate each view. Through creating several views designed to expose a sequence, a narrative across the underlying information may be exposed. This enables the domain expert to highlight trends and features of the information that may be pertinent to the end user. Exploration is supported by manipulation of the current view. This is done in part through adjusting the parameters of the semantic attributes from SARA, and reconciling user interaction with a narrative. This behavior model is known as elasticity. Elasticity enables the end user to interact with an overview and alter the prioritization of the information, effectively enabling them to ask 'what if' questions of Insight. For example, the end user may alter the parameters on the ‘Slow/Medium’ view to see transfers that they consider to be ‘Very Slow’. Essentially, they would only be presented with

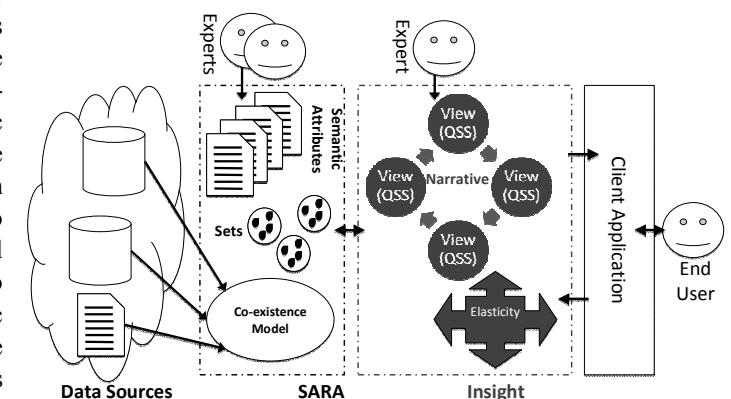


Figure 1. Supporting Personalized Information Visualizations

the slowest of the ‘Slow’ layers. The elastic part of elasticity enables this user-altered perspective to settle back to the original view. Through this transition the user may see key differences between the view and the perspective they just created. In this way the user can easily derive a tacit understanding of the data by dynamically and graphically manipulating the rules that create semantic attributes and assign values to layers in the QSS view.

Fig. 2 shows the Insight Authoring Tool interface and it highlights the views and their quantization. In this screenshot a new view is composed by sequencing assignment rules that filter input from specific layers in order to populate selected ranges of named layers. (In the screenshot below Insight is used to sort and prioritize query results about a large set of movies according to their length and grossing values. This forms part of an IPTV movie-streaming scenario within the FAME project.) The authoring tool provides live access to the underlying data so the domain expert can see the current output of the views and quantizations they are creating (on the right of the screenshot). Color-coding is important as it provides a way of tracing what data is being quantized into what layers. If several layers have the same label they are also assigned the same color.

Continuing the network monitoring example above, the top layers may show the slowest connections being composed into a new view that highlights ‘Unacceptable performance’. In this way a novice user can benefit from the insights encoded by the expert. This encoding is performed by semantic selection and the application of thresholds. Several elements of the underlying data may be combined into a single view. For example, ‘Unacceptable performance’ may be seen as a function of transfer speed, latency and service performance.

IV. CASE STUDY – VISUAL NETWORK SERVICE MONITORING

Managing networks in medium and large organizations is a highly specialist job requiring detailed knowledge of both network topologies and actively managing the flow of different types of data across those networks. To date the management of such networks has typically maintained a very data centric

view, with little consideration given to the actual tasks the network users are performing. Specifically, there is rarely any semantic view offered to the network administrator that highlights the nature or relative importance of the tasks being performed on the network. For example, if there is a particularly important videoconference being held between the Chief Operating Officer (COO) and the board of directors the network manager may not be aware of it.

Through using the Semantic Attribute Reconciliation Architecture (SARA) and the Insight Framework it is possible to provide the network administrator with a continuously cycling view of the health of both the network and the tasks (manifested as services) being performed on it. SARA is responsible consolidating the heterogeneous sources into a canonical model of the network and service health. This is typically performed by transforming information into a common ontological representation, but for services that dynamically produce data, this is done by wrapping the service interface to transform the data on the fly. These steps in themselves are not particularly innovative.

The next step is to apply subjective expert perspectives in the form of rule-based Semantic Attributes [3] to this canonical model. As data is processed by these rules it is assigned to different sets that grow and shrink dynamically. Continuing the example, one expert perspective may state that any communication service originating from any member of an organization’s senior management should be added to an ‘Important Communications’ set. Similarly there may be another perspective that states that any communication with poor quality should be added to the ‘Low Quality’ set. Both of these semantic attributes have subjective aspects. The definition of what warrants important communication may vary strongly between different contexts, as may the definition of poor quality. The personalizable nature of semantic attributes enables the network administrator to tailor these expert perspectives to better represent their needs. SARA offers a query interface that enables other systems to capitalize and tailor such dynamic semantic associations.

Insight uses this interface to assemble several overviews, where the information included is semantically prioritized towards the user’s preferences and display device capabilities. For example, an overview may show the health of all internal VOIP and video communications within the organization. The same view may automatically highlight all scheduled, i.e. those listed in calendaring system, voice or video calls with three or more participants.

Insight is responsible for sequencing several such overviews and managing the transitions between

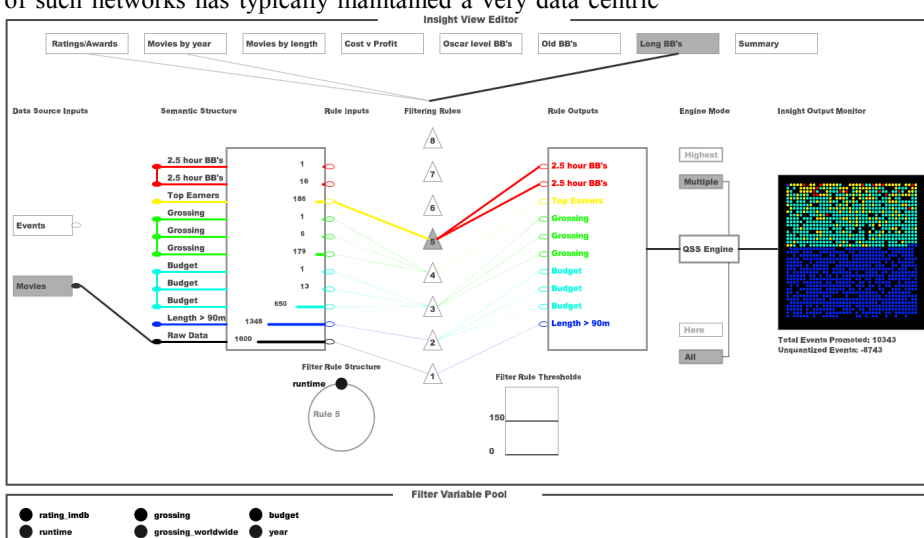


Figure 2. Insight Authoring Tool

them. Indeed the transitions are important for highlighting outlier information that may otherwise go unnoticed. In this example, the COO's call with the board of directors is prioritized on this overview as it is an 'Important Call'. This is semantic information that the network administrator would not otherwise be able to determine by just looking a quality of service metrics such as jitter and latency. In this instance the Quantized Semantic Sort (QSS) has determined that this call is of higher priority and should be displayed prominently. This means that other calls may not be displayed unless there is a specific problem, or some other important reason to do so. The network administrator may decide to explore further and alters the display depth (thus increasing how many items are shown). When he releases this visualization it will elastically settle back to its default view enabling him to see the distance between the COO's call and those that are deemed 'Low Quality'.

Through the meaningful narrative sequencing of several *overviews* the visualizations supported by Insight and SARA can tell a story of the health of a network and services offered over it, in a manner that is customizable and tailored towards the capabilities and priorities of the manager.

V. CASE STUDY – HOME-USER HAN WIRELESS MONITORING

Consider a user trying to diagnose a problem in a home area network with 2 wireless access points (APs). Specifically, a problem manifests as poor streamed media quality. Expert-derived simplifying semantics can be used to highlight and visualize the potential causes to the non-expert user. In this case the non-expert user as someone with limited computing experience, but wants to diagnose their problem uses the SARA and Insight tools, pre-informed with expert guidance.

In this HAN use case the user does not understand values for "Signal-to-noise-ratio=64", "Signal-Strength=-86dBm" or "UDP-Frame-Retransmission-Ratio=2.46%", but would understand higher-level derived values, such as "Bandwidth loss from interference vs. weak signal from poor antenna placement". For example a value of signal noise less than -90dBm is good. -90dBm to -80dBm, if consistently measured, indicates a poor connection to the antenna or a mediocre chipset. A sudden move to this range from good indicates external interference that the customer should rectify. Consistent values above -80dBm show a system failure. Transients to this value indicate strong external interference that must be rectified (for example, a microwave oven).

These are features that an expert could easily identify and codify in rules. Its importance may be subjective, but it gives the non-expert user an ability to inspect and diagnose problems, such as two APs placed too close together. This is a problem they could not otherwise identify. With the proliferation of networked home devices this is a real issue.

VI. CURRENT STATUS AND FUTURE WORK

Prototypes of the SARA and Insight systems currently exist and are being continually developed and enhanced. Ongoing work is focusing on extending the current systems to better support asynchronous event-based data in addition to query or

stream based data. A particular focus is the incorporation of more semantic queries, temporal and causal relations, and the use of symptom databases to perform real-correlation of events and faults at runtime[12]. Another body of work is building up a repository of composable customizable visualization widgets to consume the information from SARA and Insight and to support the manager in monitoring and managing their networks and services. It is important that such visualization widgets can remain loosely coupled from the underlying managed system to avoid rigid and uncustomizable visualizations.

VII. CONCLUSION

This paper has highlighted, through case studies, how the SARA and Insight frameworks are being applied in the FAME project to support non-expert users in monitoring and managing networks and services. It advocates the encoding of expert knowledge to enable discrete and meaningful visualizations of complex information from heterogeneous sources. These visualizations not only offer a means to monitor the network and services, but they also can be used to manipulate the underlying system in a causally connected manner [17] thereby managing the underlying resources. This work is part of the FAME [4] Strategic Research Cluster funded by Science Foundation Ireland and forms the basis of one of the key scenarios of this five year program.

REFERENCES

- [1] Hauser, H., "Towards New Grounds in Visualization", ACM SIGGRAPH Computer Graphics, Issue 2, May 2005, pp5-8
- [2] Shneiderman, B., "Inventing Discovery Tools: Combining information visualization with data mining". Information Visualization, 2002, issue 1.
- [3] Hampson, C., Conlan, O., "Supporting Personalised Information Exploration through Subjective Expert-created Semantic Attributes", IEEE International Conference on Semantic Computing (ICSC), Berkeley, CA, USA, Sept 2009.
- [4] Federated, Autonomic Management of End-to-end communication services (FAME) Strategic Research Cluster (SRC), Available online at <http://www.fame.ie>
- [5] Strassner, J., "Policy-Based Network Management: Solutions for the Next Generation", in the Morgan Kaufmann Series in Networking, 2003
- [6] Hampson, C. "Semantically Holistic and Personalized Views Across Heterogeneous Information Sources". International Workshop on Semantic Media Adaptation and Personalization, Dec 2007, London.
- [7] González, V. and Kobsa, A. "Benefits of Information Visualization Systems for Administrative Data Analysts". Conference on Information Visualization (IV 2003). London, UK, Jul 2003.
- [8] Simon, H. A., & Larkin, J. H. "Why a diagram is (sometimes) worth ten thousand words", Cognitive Science, vol. 11, 1987.
- [9] Novak, J., "Helping Knowledge Cross Boundaries: Using Knowledge Visualization to Support Cross-Community Sensemaking", Conference on System Sciences (HICSS-40), Hawaii, Jan 2007.
- [10] Maglio, P. Campbell S. Kandogan E. "On the Need for Negotiation in Policy-based Interaction", Conference on Autonomic Computing (ICAC 2005), Seattle, WA, Jun 2005
- [11] Palpanas, T., Chowdhary, P., Mihaila, G.A., Pinel, F., "Integrated model-driven dashboard development" Journal of Information Systems Frontiers, 9(2-3). Jul 2007
- [12] Holub, V., Parsons, T., O'Sullivan, P., Murphy J. "Run-time correlation engine for system monitoring and testing". International Conference on Autonomic Computing, (ICAC), Barcelona, Spain, Jun 2009.
- [13] Keeney, J., Lewis, D., Wade, V., "Towards the use of Policy Decomposition for Composite Service Management by Non-expert Endusers" International Workshop on Business-driven IT Management (BDIM 2010) at NOMS 2010, Osaka, Japan, Apr 2010.