

An Ontology-driven Approach to support Wireless Network Monitoring for Home Area Networks

Yuqian Song, John Keeney, Owen Conlan

FAME & KDEG

School of Computer Science and Statistics

Trinity College Dublin

Dublin, Ireland

{yuqians, john.keeney, owen.conlan}@scss.tcd.ie

Philip Perry, Adriana Hava

FAME & PEL

School of Computer Science and Informatics

University College Dublin

Dublin, Ireland

{philip.perry, adriana.hava}@ucd.ie

Abstract—Supporting ordinary home users to diagnose and resolve potential problems in their home wireless networks presents a significant opportunity to reduce support costs and increase user satisfaction. However the raw data necessary to identify even simple issues, such as a device being out of wireless range, require specific expertise to interpret. This paper introduces an approach for expert-derived semantic annotation of this raw data in order to allow end users to understand and resolve common network-related problems in real time. This approach demonstrates how semantic web technologies and visualization may be combined to satisfy high-level network monitoring requirements. This semantic approach, with causal reasoning, coupled with semantically-derived visualizations was implemented in a prototype home network monitoring system (HANMS). This paper describes the design decisions, implementation details and evaluation approach of HANMS.

Keywords—home network; network monitoring; ontology; visualization; semantic enhancement

I. INTRODUCTION

In recent years, more and more households are connected to the Internet through broadband access services. In 2009, over 99% of UK households were connected to a broadband enabled exchange [1], while the number of worldwide home broadband consumers is over 300 million [2]. Many of these households have begun to adopt networking technologies to interconnect devices within the home for media sharing, communication, gaming and other applications working together to resulting in a more comfortable, entertaining, and safe home. This flexible Home Area Networking (HAN) approach is increasingly moving towards wireless communication protocols to interconnect networked devices.

However, influenced by the increase of network complexity and management requirements, a key challenge is how to bridge the semantic gap between obscure network domain monitoring data and comprehension capabilities of normal home network users. In a HAN environment, large raw wireless network data logs are generated by wireless devices in real time, but these raw data are uncorrelated and highly heterogeneous. Semantic Web technologies are typically adept at capturing semantic relationships and meaning as specified by

domain experts, and leveraging such semantic encodings to enrich low-level management data. Such domain knowledge, embedded in domain ontologies, is useful not only for data enhancement (semantic uplift), but can also be used for network problem analysis, and has a potential to support high-level network management. This paper presents our Home Area Network Monitoring System (HANMS), which adopts a user-centric approach to combine domain-expert ontologies with user-friendly visual network problem analysis and semantically enhanced data monitoring. We aim to achieve the following outcomes: to discover how to enhance raw network log data with expert defined semantic meaning; to integrate expert knowledge ontologies into wireless network monitoring; to investigate appropriate visual network monitoring approaches for normal HAN users; to design a flexible architecture for expert-driven semantic uplift; to develop and deploy a early stage HANMS prototype in a HAN environment; and to execute an appropriate evaluation to assess the approach.

This research will be described, introduced and discussed in following sections. Relevant projects and research will be briefly reviewed in the following Related Work section. The System Design section will present system goals, investigate ontology-driven network monitoring and introduce a layered HANMS architecture. For HAN deployment, a HANMS prototype was implemented, and is discussed in Section IV. Section V presents a qualitative evaluation using two target groups, normal HAN users and network domain experts, with associated analysis and discussion from different perspectives. Finally, ongoing work is described in the Further Work section.

II. BACKGROUND

This section discusses existing techniques that may assist in bringing more meaningful semantics to complex data and applying semantics to enhance HAN management. Focusing on the complex HAN environment, [14] presents a policy-based federated service management architecture that addresses these concerns for digital home devices participating in end-to-end communications services. Jennings [5] proposed an autonomic network management scheme, which allows a network system to self-govern its behavior within the constraints of the human-specified goals.

Another effective strategy is to provide tools for non-expert users to understand and manage their network in a visual solution. A novel interactive visualization system [13] was proposed as an approach to monitor the data collected from a home router and control bandwidth usage for family users. Eden [7] is a network management system which uses a drag-and-drop visual representation of network devices and settings for management tasks, which eliminates the barrier to understanding detailed network monitoring and management.

Ongoing network monitoring research points to the appropriateness of visual dashboards for network management tasks by supporting a broad spectrum of information visualization requirements, ranging from raw data presentation to abstract overviews and correlations of information. Dashboards can be classified into three types [8]: strategic, analytical, and operational. Strategic dashboards should provide a quick overview for the decision maker to monitor or forecast. Analytical dashboards present the information along with greater context, and more extensive history. Operational dashboards offer an interactive interface between user and data. A definite synergy exists between dashboards and information visualization for network monitoring/management.

Another problem for the wireless HAN management is how to gather real-time data using appropriate processes and tools. There are many tools that can be used to extract information from the drivers for particular wireless network ports, most notably Wireshark and TShark [18]. These can provide metrics about the packet-level throughput for individual ports and also report on the radio environment, including antenna signal strength and noise levels. The ability to use such metrics for network management functions has been largely addressed in the context of wireless mesh networks [12].

In order to support higher level management by adding domain knowledge to raw monitoring data, a new concept in data enrichment process can be used - Semantic Attributes [6]. Semantic attributes are pieces of semantic encodings created by domain experts to support non-expert user explore an information domain. Semantic attributes encapsulate encoded expert insights for a given domain, supporting personalization and interlinking according to an end-user's preferences. This data abstraction, enhancement, and classification process supports the ability to leverage higher level semantics. OWL ontology [3] provides a series of species, serializations, and syntaxes to represent knowledge which could be understood and reused by the machine and applications on the web.

III. SYSTEM DESIGN

The proposed approach utilizes domain-expert knowledge that is embedded into the system in two processes: the semantic annotation process, and the ontology-driven event detection and analysis process, thereby forming a layered architecture. The core components of HANMS and will be detailed in the following sections.

A. Example Use Case

This section presents a typical HAN that HANMS could operate in. The example HAN has three active devices: an iPad, which can be used for VoIP calls; an Xbox, acting as an IPTV

client; and a laptop PC which is used for web browsing and email. All of these may be connected to a home gateway device via IEEE 802.11 WiFi. Performance degradation may be caused by: (i) wireless interference caused by a radio-controlled toy or microwave oven, (ii) a neighbor initiating some WiFi activity in the same channel – causing interference, or (iii) some hardware may develop a fault, (iv) the iPad or laptop may move out of range of the WiFi basestation, or (v) the PC user may upload a large file which congests the network. Such events will have different impact on the end-user's perceived quality of both the networking service and application-level services. However, the user may not have the expertise to examine device logs to diagnose the cause of problems in real time. Such technical expertise may only be available from highly trained but expensive support staff. If the user could diagnose these easy to resolve problems it would avoid unnecessary contact with support staff and reduce the operating costs for both network provider and HAN users.

B. Semantic Annotation

Semantic Annotation aims to enhance large amounts of collected heterogeneous data with discrete semantic encodings of domain expertise in order to support high-level logic reasoning and bridge the understanding gap between the raw data and non-expert users [6]. In the initial semantic annotation process, the raw data is associated with semantic attributes [6], which encapsulate the expert's subjective insights of a domain. In HANMS, semantic attributes are encoded and identified in an XML-based model. These semantic attributes inform the enhancement of low-level data with semantics, where each contains the following:

- A semantically meaningful concept
- Parameters defined by experts
- Operators related to the particular parameter
- Values related to the particular parameter
- Links to a domain ontology
- Links to domain metadata

Typically semantic meanings are encoded in this model with distinct and indicative concepts from a modeled domain. In the wireless network domain, typical log data collected from wireless devices contains low-level metrics and values including Antenna Noise, Antenna Signal, Throughput, etc. For the non-expert user, these concepts, scales, trends, their values and units, are obscure and meaningless in most situations. By capturing and representing the opinions of a domain expert, the raw data (along with extra information about the source device and time) is correlated with pre-defined metadata. For instance, a piece of data presents as: generated by Device 13, at 13:01:01, with three metrics: Antenna Noise, Antenna Signal and Throughput. Then this piece of data is enriched and annotated according to the definition of semantic attributes.

Table. 1 shows three candidate semantic attributes for a given device: "Antenna Noise", "Antenna Signal", and "Device Throughput", each with three high-level values "Good", "Moderate", and "Bad". For example, the semantic attribute "Antenna Noise" is given the value "Good" if the low-level antenna noise value is lower than -90dBm. (In the

simplest case a semantic attribute is encoded using a low-level metric, some operator and reference values). This example semantic attribute is also referenced to the antenna noise metric in the metadata and the *AntennaNoise_Good* class in the domain ontology. The “Moderate” and “Bad” parameters are also described in a similar way.

TABLE I. THREE SEMANTIC ATTRIBUTE METRICS

Semantic Attribute	value: ‘Good’	value: ‘Moderate’	value: ‘Bad’
<i>Antenna Noise</i>	<-90dBm	-90dBm to -80dBm	>-80dBm
<i>Antenna Signal</i>	>-74dBm	-74dBm to -86dBm	<-86dBm
<i>Device Throughput</i>	<1 Mbps	1Mbps to 2Mbps	>2Mbps

Semantic attribute are not only applied in the data classification, but also in the data trend description (Table. 2). In the presented prototype, the raw data is inspected every three seconds, and the trend in the three-second-dataset (Semantic Attribute: “Trend”) is associated with corresponding semantic attribute values. The value “Disappear” indicates that the source of the data has vanished, so no new data exists in the dataset. “Gentle up/down” means the change in the dataset is less than 30% and “Sharp up/down” indicates the change is higher than 30%. “Stable” indicates that there is no significant change in the dataset.

TABLE II. SEMANTIC ATTRIBUTES OF DATA TREND

Semantic Attribute	Value: ‘Disappear’	Value: ‘Gentle Down’	Value: ‘Gentle Up’	Value: ‘Sharp Down’	Value: ‘Sharp Up’	Value: ‘Stable’
<i>Trend</i>	No data	<30% Down	<30% Up	>30% Down	>30% Up	<10%

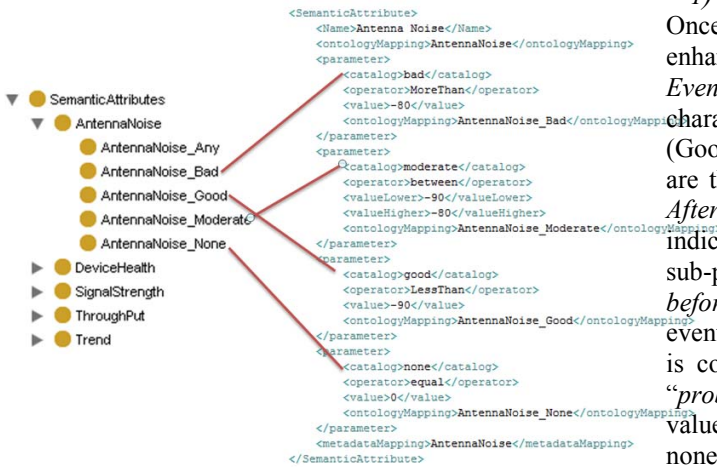


Figure 1. The mapping between Semantic Attributes and domain ontologies

Any parameterization of metric values for deriving values for semantic attributes is subjective and reflects the expert’s perspective. For this reason all semantic attribute encoding rules are configurable at runtime, but based on domain expert input, each semantic attribute also has default values. Because not all domain experts are familiar with XML encoding, it is difficult to hand-craft semantic attributes. This work presents visual ‘expert widgets’ to allow non-technical users to create and adjust the semantic attribute values in real time. This

means that HANMS gets accurate human-created semantic attributes without the cost of large amounts of manual effort.

In HANMS, semantic attributes are mapped to corresponding domain ontologies. In the mapping example (Fig. 1), the parameters of the semantic concept models (r.h.s) are linked to related domain ontology classes (l.h.s). In this example “AntennaNoise_Any” indicates that the antenna noise metric could be any value except none, which is used in the ontology reasoning process for event detection and analysis. These mappings enable high-level data annotation supporting ontology-based logic reasoning and endow more complex semantic meanings to real-time data from heterogeneous resources, which is a foundation of the ontology-driven event detection and analysis process.

C. Ontology Driven Event Detection and Analysis

HANMS enhances the raw data with an ontology-based semantic annotation process by mapping semantic attributes and domain ontologies, and the enriched streams of log/performance data are observed and an event detection and analysis process is based on expert defined domain ontologies.

In domain ontologies, network problems are defined as Events, which could be a significant change and may affect other devices or services of the HAN. In an extended domain ontology (Fig. 3), the *Event* class has a subclass “*SuddenProb*” which models network problems happening within a short time period. The *Event* properties *afterPresents* and *beforePresents* indicate the data trend within a dataset after or before event instances. Another property, *hasReason* links *Event* instances to instances of the *Reason* class, which represents possible reasons to cause this event.

1) Event Detection

Once every second, HANMS observes the semantically enhanced data and events, and checks for new instances of *Event* to detect if there is a new event occurred. The characteristics of the data streams, e.g. trends or data status (Good, Moderate, and Bad), before and after the current second are the fundamental criterion for the event detection process. *AfterPresents* (if available) and *beforePresents* are used to indicate data trends and characteristics, both of which have two sub-properties: *afterPresents_essential*, *afterPresents_optional*, *beforePresents_essential* and *beforePresents_optional*. If the event meets all restrictions in essential properties, this detection is considered to be high reliability, encoded with the value “*probably*”; if it only agrees with some essential restrictions the value is “*maybe*”; where the value “*maybe not*” indicates that none of the restrictions are satisfied. E.g., for the “*DeviceDisappear*” event presented in Fig. 3, the dataset should have following characteristics: the antenna noise value disappears resulting in an *AntennaNoiseDisappear* event; the signal strength value disappears resulting in a *SignalStrengthDisappear* event; and the throughput value disappears resulting in a *ThroughPutDisappear* event. Meanwhile, the data trend could present with any trend before the event. Therefore for a given unidentified event, *iff* the semantic attribute representing throughput for the given device has the value *ThroughPutDisappear* after the event and

ThroughPut_any before the event, the HANMS reasoning process infers the event as a high-level *DeviceDisappear* event.




Figure 2. The visualized domain ontologies

2) Event Analysis

If an event is detected, the next step is to judge its reason. This event analysis process uses the data characteristics in a period (9 seconds) before and after the detected event. The semantic property “hasReason” links instances of the *Event* class to instances of the *Reason* class. The *Reason* class properties *causedByData* and *causedByTrend* have essential and optional sub-properties and are evaluated and presented according to three probability levels, “probably_causedBy”, “maybe_causedBy” and “maybe_not_causedBy”. As encoded by the domain expert in the domain ontology, the “*DeviceDisappear*” event has three potential reasons: manual or automatic device disconnection, interference, and out of range. For each potential reason, we analyze annotated semantic attributes of the antenna noise, signal strength and throughput before and after this event and based on the data characteristics and trends, HANMS can infer a probable cause for a given newly identified and reasoned event.

D. System Architecture

Home network challenges are compounded by the need to provide a system for normal HAN users to understand and observe how wireless HAN devices work in real time and to perform a dynamic problem analysis using the expert knowledge embedded in defined semantic attributes and

domain ontologies. For these reasons HANMS is required to realize the following goals:

- extensible architecture with high compatibility and usability
- support heterogeneous input data with semantic attribute and domain ontology consuming capabilities
- enhance and annotate the raw data with semantic meanings
- domain ontology driven network intelligent problem detection and analysis
- user-friendly visual interface for non-expert users
- appropriate visual widgets for domain experts to define/adjust semantic attributes and analyze actual problems

Based on these goals, HANMS uses a layered architecture (Fig.4) with three distinct layers: Data Preparation Layer, Semantic Processing Layer, and the Visualization Layer. HANMS was designed using Flex, a popular RIA technology, to realize the Visualization Layer and a server-client structure for the whole system.

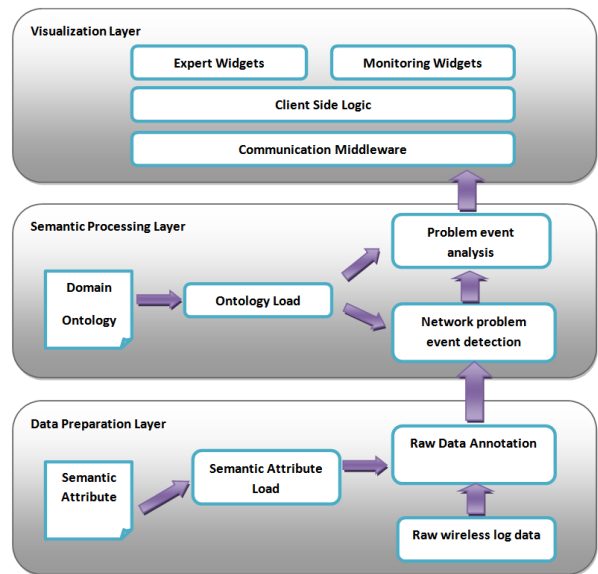


Figure 3. HANMS Architecture

1) Data Preparation Layer

The heterogeneous input data collected from the home wireless devices are associated with related Semantic Attributes, which are defined in XML format by the domain expert (with visual support) and loaded into the system. In this process, the raw data are annotated with the pre-defined semantic meanings in real time and prepared in a semantic meaningful structured data stream for the Semantic Processing Layer.

2) Semantic Processing Layer

This layer loads additional, expert-defined domain ontologies, and classifies them according to upper-level Event Ontologies and Reason Ontologies. The Event Ontology describes the features of a network problem event. The Reason Ontology is used to infer the reason causing a network event. Based on these new network problem events are detected from the real-time structured data stream and causal reasons are judged and analyzed. The real-time annotated data, detected events and related reasons are then pushed to the Visualization Layer.

3) Visualization Layer

Several user-friendly widgets were designed for the non-expert users to understand and monitor the HAN according to the data objects from the layer below, through a communication middleware. As an RIA system, the client side logic can be executed in this layer to realize complex interaction behaviors. Expert widgets are also provided in this layer to assist the domain expert to define/adjust the Semantic Attributes and review network problem events from the data point of view. It is particularly noteworthy that the data objects received from the semantic processing layer are independent of any particular visualization widget, so the visualization layer can embed additional expertise-driven logic to select or personalize the most appropriate presentation widget for a given combination of data/user. This separation of domain-specific expertise from visualization-specific expertise improves on the current approach of embedding domain reasoning associated assumptions, in the presentation layer.

IV. SYSTEM IMPLEMENTATION

By using the selected technologies and approaches to meet the goals mentioned above, a prototype has been developed and deployed locally. Raw wireless log data was collected from a physical wireless monitoring test-bed and loaded into a simulation engine, which feeds HANMS by generating several wireless device log data streams seeded by the data collected from real wireless devices. The expert defined semantic attributes are stored in an XML database, eXist [9]. The server-side HANMS logic, which imports semantic attribute definitions and domain ontologies, was developed in J2EE and deployed locally on a Tomcat server. The HANMS client-side (implemented in Flex) visual components communicates through Blazeds middleware [10].

To present the semantically meaningful data, a set of Visualization Widgets have been applied into three focused perspectives to expose the system functionality (See figures 4, 5, 6).

A. Focus1: home area network monitoring

The HANMS network monitoring panel (Fig. 4) contains two widgets to express the HANs diverse perspectives using semantically meaningful real-time data. In this scenario, there

are three linked HANs, HAN1 in blue, HAN2 in green and HAN3 in red. The radar chart widget is used to present the general health of the whole network. The equi-angular spokes indicate the number of devices, the number of external services, the number of internal services, the aggregate health level, and the total throughput of the different HANs. In this widget, the bigger area means the corresponding HAN is more active. Agraph widget is used to describe the topology and state of individual devices in the different HANs. The icon of each device presents the device type, as laptop, xBox, iPad, etc., and the icon color shows the device health level, green (good), yellow (moderate), red (bad) or grey (disconnected). If an event is detected, there will be a red alert icon linked to the device where the event happened.

B. Focus on the event detection and analysis

If a detected event is clicked, the HANMS event panel (Fig. 5) will emerge to present the event analysis process. In this scenario, according to the domain ontology-based reasoning, the inferred “*Device Disappear*” event is annotated with three reasons, each with an associated probability: i.e.: “*probably*” caused by “*Disconnect*”, “*maybe not*” caused by “*OutOfRange*”, and “*maybe not*” caused by “*interference*”.

Based on the analysis result, the “*Device Disappear*” event is most probably caused by a sudden physical disconnection of the device. The line charts shows what happened within the period of 9 seconds before and after the event, and the background color indicates the semantic attribute thresholds. In this case, the antenna noise, antenna signal, and throughput suddenly go to zero, so the system inferred that the event was caused by “*Disconnect*”. By presenting the *real* data, along with the correlation thresholds, it is much easier for the non-expert user to understand the analysis process.

C. Focus on the expert’s point of view

HANMS also provides suitable widgets for a domain expert (or advanced user) to adjust the semantic attribute thresholds and personalize the event analysis process. A temporal visual interface allows the expert to navigate back previous intervals to manually evaluate whether the event is detected and analyzed correctly by HANMS.

The HANMS expert panel (Fig. 6) adjusting/redefining the value of semantic attributes. Any change is reflected to the




Figure 4. HANMS network monitoring panel

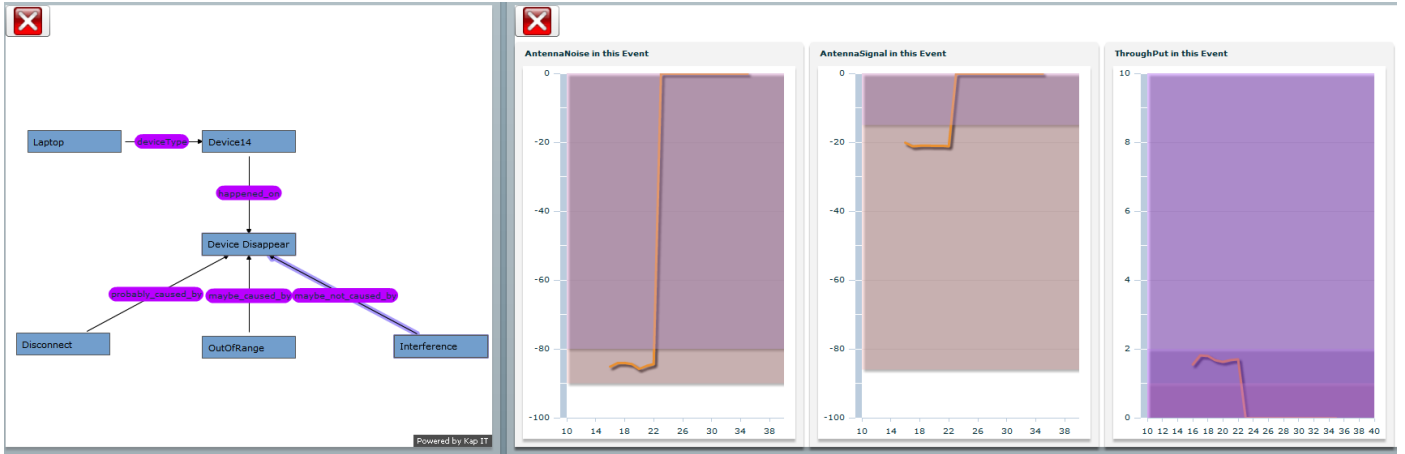


Figure 5. HANMS event panel

monitoring panel and the event panel immediately.

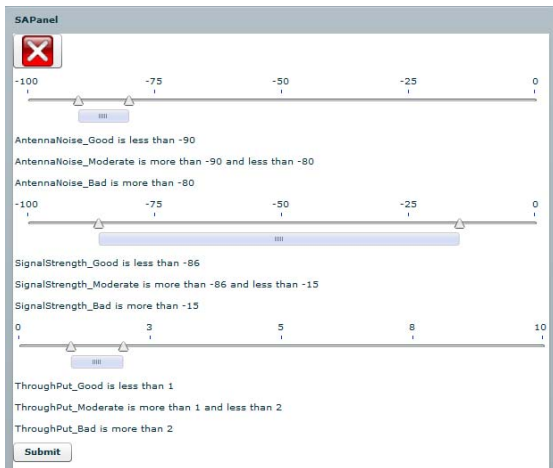


Figure 6. HANMS expert panel

V. SYSTEM EVALUATION

The HANMS prototype was evaluated to gather feedback about its operation, whether it matches defined goals, and user perception. Evaluation participants were divided into two groups, normal user and expert, based on networking knowledge, and they were invited to use this system for three simulated different tasks. They then completed a feedback questionnaire on the functionality, effectiveness, efficiency, usability, user-interface, and limitations of HANMS.

The questions were presented in a SUS-based [11] questionnaire, followed by statements which participants can rate on a Likert scale: Strongly Agree, Agree, Disagree, and Strongly Disagree, for qualitative analysis. Participants were encouraged to provide qualitative feedback on the tasks to capture feedback not collected for quantitative analysis. The tasks expose three common problems in HAN usage scenarios.

To date, 19 volunteers have participated this evaluation, and 4 of them were domain experts. The answers are ranked to express the feedback result from negative to positive.

According to the result table (Table. 3), the majority of feedback is positive, especially in the usability and UI design. Most participants agree that HANMS is usable by a non-expert user. More than 70% of participants believe HANMS could detect and analyze the simulated network event effectively. As

a result of the test tasks, 50 out of 57 HANMS observations were correct, which indicates that even in its current prototype state HANMS is suitable for different tasks in different scenarios. 16 of the 19 volunteers thought the existing functions of HANMS were sufficient. Moreover, nearly all participants agree that this network monitoring tool is more efficient than the traditional ones. Experts could also get appropriate assistance in HANMS. However, the evaluation results reflect some limitations: HANMS is not sufficient to show multiple complex events; the historical data should be associated with the monitoring panel; and the probability of reasons could be quantified and presented in a more intuitive way.

TABLE III. HANMS EVALUATION RESULTS

Rank	Strongly Disagree	Disagree	Agree	Strongly Agree
Usability	2	2	20	33
UI Design	0	7	38	31
Effectiveness	4	13	21	19
Functionality	3	3	24	6
Efficiency	0	0	11	8
Expert support	0	4	6	6

With the results from the evaluations above, we can conclude this HANMS prototype has already achieved most of the requirements and the addressed limitations could be improved in the future development. This prototype will be setup in a real home scenario to evaluate the commercial potential for HAN usage.

VI. FUTURE WORK AND CONCLUSION

This work will aims to extend ongoing work towards the abstraction of complex management tasks so that the non-expert user can understand how to iteratively view, control and compose the high-level monitoring information and high-level management tasks. This would support the manipulation and control of the managed system from a high-level abstraction of user goals to low-level concrete control actions [4].

According to the deployment and evaluation result, this HANMS prototype could assist ordinary end users to understand, diagnose and resolve potential problems in their home wireless networks, which presents a significant opportunity to reduce support costs and increase user satisfaction. This prototype also proved that it is possible to

combine semantic technologies and visualization technologies to support network monitoring/management.

ACKNOWLEDGEMENT

This work was funded by Science Foundation Ireland via grant 08/SRC/I1403 — Federated, Autonomic Management of End-to-End Communications Services (FAME).

REFERENCES

- [1] D. Emma, "Broadband Access", Key Issues for the New Parliament 2010, UK Parliament, <http://www.parliament.uk/documents/commons/lib/research/key%20issues/Key%20Issues%20Broadband%20access.pdf>
- [2] "Broadband Internet Statistics", Internet World Stats, 2007, <http://www.internetworldstats.com/dsl.htm>
- [3] "OWL Web Ontology Language Guide", W3C Recommendation, February 2004, <http://www.w3.org/TR/owl-guide/>
- [4] J. Keeney, O. Conlan, D. O'Sullivan, D. Lewis, V. Wade, "Towards the Visualisation of Collaborative Policy Decomposition", Proc. 9th IEEE Workshop on Policies for Distributed Systems and Networks (POLICY 2008), Palisades, NY, USA, 2-4 June 2008.
- [5] B. Jennings, S. Der Meer, S. Balasubramaniam, D. Botvich, M. Foghlu, W. Donnelly, and J. Strassner, "Towards autonomic management of communications networks," IEEE Communications Magazine, vol. 45, Oct. 2007, pp. 112-121.
- [6] C. Hampson and O. Conlan, "Supporting Personalized Information Exploration through Subjective Expert-created Semantic Attributes," 2009 IEEE International Conference on Semantic Computing, IEEE, 2009, p. 384-389.
- [7] J. Yang, W.K. Edwards, and D. Haslem, "Eden: supporting home network management through interactive visual tools," Proc. 23rd annual ACM symposium on User interface software and technology (UIST '10). ACM, pp 109-118, 3-6 Oct. 2010
- [8] G. E. Krasner, S. T. Pope, "A cookbook for using the model-view controller user interface paradigm in Smalltalk-80". J. Object Oriented Program. 1, 3 (Aug. 1988), 26-49.
- [9] eXist Open Source Native XML database, <http://exist.sourceforge.net/>
- [10] "BlazeDS online document", Adobe Open Source, <http://opensource.adobe.com/wiki/display/blazeds/BlazeDS/>
- [11] I. Brook, "SUS: a "quick and dirty" usability scale". in P. W. Jordan, B. Thomas, B. A. Weerdmeester, & A. L. McClelland. Usability Evaluation in Industry. London: Taylor and Francis, 1996. <http://www.usabilitynet.org/trump/documents/Suschapt.doc>.
- [12] F. Jan, B. Mathieu, D-E, Meddour, "A monitoring tool for Wireless Multi-hop Mesh Networks," Network Operations and Management Symposium, NOMS 2008. IEEE , pp.587-601, 7-11 April 2008
- [13] M. Chetty, R. Banks, R. Harper, T. Regan, A. Sellen, T. Karagiannis, P. Key, and C. Gkantsidis, "Who's Hogging The Bandwidth?: The Consequences Of Revealing The Invisible In The Home," *Technology*, 2010.
- [14] R. Brennan, D. Lewis, J. Keeney, Z. Etzioni, K. Feeney, D. O'Sullivan, J. Lozano, and B. Jennings, "Policy-based integration of multiprovider digital home services," *IEEE Network*, vol. 23, Nov. 2009, pp. 50-56.