Semantic Reasoning for Adaptive Management of Telecommunications Networks

John C. Hoag, *Senior Member, IEEE* and Frederick A. Hayes-Roth, *Senior Member, IEEE*

***Abstract*—Semantics can provide the basis for problem solving in a setting such as telecommunications network management, which maintains multiple simultaneous ontologies for the same set of entities. To avoid overbuilding and to improve quality, network operators desire both a close match between application requirements and allocated resources as well as the capability to perform modifications at will. This paper presents a concept for solving the relevant problem of network resource allocation through semantic reasoning techniques and identifies three specific domains for development of common ontologies**

I. Introduction

DESIGN of wide-area telecommunication networks and internets is a daunting task, often performed once-per-contract, incrementally, or not at all; optimized designs have been rare, and a network operator might simply overbuild capacity. The objective of the system proposed herein is to provide the means for constant iteration of the network design and implementation process – specifically in response to an outage or a change in functional requirements. This approach, which will be new to telecommunications, is not unlike the mediation of wholesale power flows in the U.S. electric transmission grid. In the latter case, a regional transmission operator (RTO) must capture current network topology and model flows such that physical constraints are not exceeded. Investigation of a recent major power grid outage was the impetus for this line of research, and the role of automation in the hastily-prepared setting of electricity deregulation is worthy of further research [1].

Management of telecommunications and electricity grids and other high-consequence domains such as aviation provide potential use cases for an emerging systems architecture known as VIRT, Valued Information at the Right Time, being developed by the second author. The common context for these domains is need for some new system to support realtime replanning of resources subject to changes in a field of underlying data, assumptions, and models. The core of VIRT its ability to monitor conditions of interest in whatever raw form they appear, and to translate them into a common domain ontology for the later consideration [2]. Von Schweber [3] identifies attributes necessary for VIRT as:

- Publish/Subscribe capability,
- Message queuing and transactional messaging,
- Content-based routing,
- Context-aware computing,
- GIS, geo-spatial and mapping,
- Location-based services,
- Temporal data management,
- Active Database,
- Situational Awareness,
- Data, Information and Semantic Interoperability, and
- Service Oriented Architecture.

Further, Von Schweber identifies the following attributes as unique to and sufficient for VIRT:

- Notification and awareness based on:
    - Information “deltas”
    - Prediction and forecasting
    - “Community” entities and attributes
- Extraction of a plan’s critical assumptions and the use of these to drive information filtering and routing

The initial use case for VIRT involved the use of meteorological data by a pilot during a mission, whose plans might change mid-flight. VIRT, aware of the flight plan, would sift through massive realtime data streams, identify significant differences, and forward either distilled information for pilot consideration or alternative plan recommendations.

Architecture of the VIRT engine (i.e., definition of primitives, etc.) is a collaborative effort of the Worldwide Consortium for the Grid (W2COG), whose commercial members are positioned for software development.

VIRT, depicted in Figure 1, incorporates a broad notion of system state that contains information about both goals and conditions, with the capability to obtain sensory information, to provide ontology-based translation on its inputs, to derive the difference between plans and conditions, and to recommend changes to the current plan.

Manuscript received March 30, 2006.

John C. Hoag is with the McClure School of Information and Telecommunication Systems, Ohio University, Athens OH 45701 USA (Phone 740-593-0077; fax 740-593-4889; e-mail hoagj@ohiou.edu).

Frederick A. Hayes-Roth is with the Operational and Information Sciences Department, U.S. Naval Postgraduate School, Monterey, CA 93943 USA (e-mail hayes-roth@nps.edu).

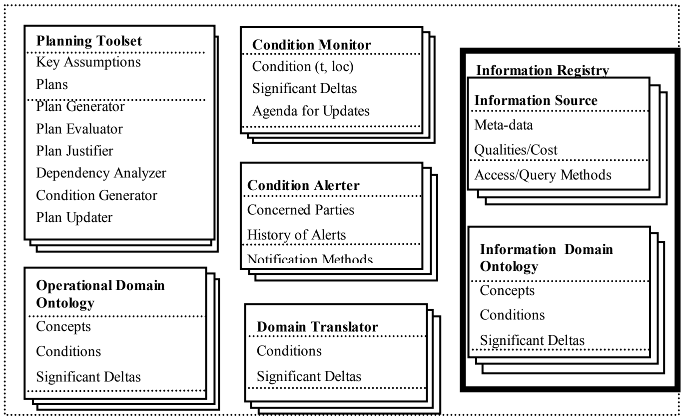

Fig. 1. Graphical representation of VIRT. The initial use case for VIRT was delivery of aviation weather enroute, hence the special and temporal reference

While most details in Figure 1 are somewhat self-explanatory, the planning toolset merits special attention. The VIRT process can be viewed as always alert and able to execute a procedure when a state-changing event is detected. The basis for logical reasoning within VIRT is both the current plan and the justification for selection of that plan with respect to assumptions. Within VIRT, a set of semantic relationships rules would be developed that concern the stability of the assumptions and the dependency of the plan upon them. In the second author's use case, assume that a flight mission (i.e., plan) requires the use of certain navigation instruments, hence the selection of a certain airplane. If the instruments or the assigned aircraft are unavailable, the mission planner must be alerted – and perhaps VIRT can recommend an alternative plane, another destination, a delay for better (e.g., VFR) weather, etc.

The organization of the remainder of this paper is as follows: an introduction the problems of telecommunications network management; discussion of semantic reasoning and identification of three specific subject domains for ontology development within telecommunications; and an overview of the proposed system in operation. The paper concludes with an assessment of project status and plans for future work.

II. Telecommunications Network Management

There are two competing interpretations of the task called network management, but no configuration of these systems combines to automate the sequence of steps of designing and implementing large networks (cf. wide area networks, WANs, and internets). These two competing genres of systems are termed Operations Support Systems (OSS), which fulfill administrative and business transactions, and Network Monitoring Systems (NMS), which detect and respond to events such as failures. A network operator (e.g., carrier/provider or enterprise customer) may have hundreds of OSS function points and a network topology consisting of thousands of monitoring points.

The purpose of adaptive network management is to acquire substitute facilities to fulfill communications requirements, which is essentially to build a bridge between OSS and NMS environments. While the primary setting shall be fault recovery, the capability to fulfill "normal" requests to add, modify, or delete services is a modest extension. Central to this capability is an inclusive definition of system state:

- The set of all application functional requirements per endpoint, aggregated only for similar quality of service (i.e., bandwidth, delay, etc.)
- The set of all network elements in provider(s) inventory and their availability to be deployed
- The set of all provisioned (i.e, assigned) resources, which is a cross-reference of the previous two items
- The set of generic templates for configuring network monitoring all classes of devices and elements

OSS capabilities include applications for order processing, provisioning, service activation, testing, maintenance and repair, billing, and customer information. These functions are central to telecommunications service providers and are increasingly integrated, distributed, and accessible to competitor and reseller firms. While interactive, OSS systems as defined neither monitor the performance of network elements nor exhibit autonomy of control [4].

There are two fundamental types of SNMP standards-based Network Management: systems that poll elements status, and those that are triggered by exception events. Status information may be filtered or correlated in context of nearby devices or exhibited as part of a larger structure (i.e., topology). Interpretation, filtering, or diagnosis of event signatures is feasible but is beyond the scope of baseline commercial software, and the maximum extent of this functionality is to interoperate with the corresponding OSS to generate, say, a trouble ticket based for a confirmed fault [5].

Both kinds of network management systems operate largely in an open loop or in one closed only through actions of a human operator. Each system contains capabilities, respectively, to control or to monitor complex, distributed information flows and their infrastructure. No complete integration is available between the systems for the benefit of either enterprise users or providers.

A small number of academic and commercial collaborative projects are seeking to advance network management technology. Where open source exists, it has been acquired by this project.

1. The open-source *OpenOSS* Initiative comes from the Telemanagement Forum's reference architecture for Next Generation Operations Support Systems (NGOSS). Its prototype incorporated fault management (FM) and trouble ticketing (TT) functions using a service-oriented architecture [6,7]; this project's academic partner is Southampton.
2. The *OSS Through Java Initiative* (OSS/J) is a consortium of commercial OSS vendors implementing TMF's

NGOSS architecture, which has defined publicly-available interfaces using WSDL and an XML schema [8].

3. *OpenNMS* is also open-source with over 5000 downloads of the latest revision. It has functions for service polling, data collection, and event management, and can acquire over 200,000 datapoints from 20,000 elements in five minutes – performance similar to commercial products and limited only by the pace of serial SQL posting [9].
4. The *KING* Next Generation Network Architecture uses centralized intelligence to provide adaptive, automatic, and autonomic network resource allocation by means of admission controls and policy-based routing [10]. This project is conducted at Karlsruhe.

III. SEMANTIC REASONING AND ONTOLOGIES

The most general contribution of this paper as well as the essential enabling technology for adaptive network management is the use of semantic reasoning in three specific domains: applications requirements, location and WAN services. Each of these domains merits development of a specific new ontology, but currently different process stakeholders use different names for the same entity. With ontologies in place, the VIRT function may reduce structured queries iterating within a loop in search of least-cost feasible solutions.

Semantic reasoning can enable the following kinds of transactions:

1. A customer requests service to support some new application like group conferencing among selected endpoints.
2. An engine translates this into specific technical specifications and considers aggregating similar traffic at common endpoints.
3. The engine repeatedly polls its service providers for feasible infrastructure that is available where needed.
4. The engine consummates transactions and activates services; it also informs the NMS of new elements to monitor.

A fuller example is presented shortly.

Each ontology is being developed as a hierarchical set of metadata using Protégé, a panel of which is displayed next:

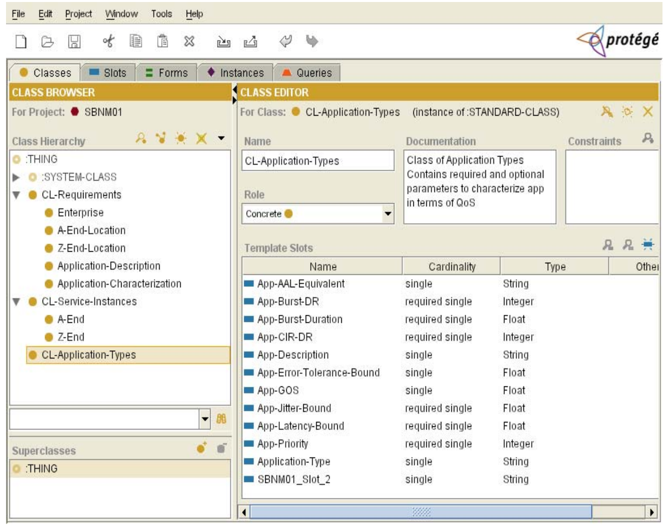

Fig. 2. Protégé screen capture showing Application Requirements Class.

The first semantic relationship involves translation of application requirements. We define applications broadly, to include all varieties of voice, data, and video purposes – realtime and not, reliable or not, interactive or not, and so on; and are geographically located. We seek equivalent descriptions to characterize applications with their inherent service quality (QoS) parameters such as latency, jitter, instantaneous and burst datarate, instantaneous and cumulative error rate, and so on. Moreover, designers are empowered to aggregate similar data streams between endpoints. For example, a call center might describe its needs, in its busy hour, to receive forty calls of ten minutes duration; Erlang techniques render a solution of nine traditional trunks or a 216kbps path for VoIP trunking under 8:1 compression.

The second equivalency, location, is more straightforward. The same location may be identified by street address, latitude and longitude, or telco "Vertical and Horizontal" (V&H) coordinates. GIS can enable co-located service endpoints to be revealed regardless of the original coordinate system.

The final equivalency concerns provider WAN services and represents an important skill set for a network analyst. This list should exhaustively state alternative mechanisms to achieve the same QoS objective. Our trend is to replace physical resources such as a 1.544 synchronous T-1 circuit with a virtual work-alike acceptable to that application. For voice services, DS3 and OC1 services are equivalent, as could an Ethernet WAN segment at the specified datarate.

IV. State-Based Network Management

In the adaptive management scheme, agents at local nodes such as enterprise customer endpoints and provider edges must maintain their own local system state, defined previously in this paper. This architecture embeds VIRT in a state engine node whose objective is to fulfill the current set of telecommunications requirements by acquiring provider resources and monitoring them.

Interestingly, our craft has operated on the premises that applications should not know anything about the underlying network, and vice versa; also, that any incorporation of state information would limit scalability, i.e., growth.

Development of the adaptive network management concept benefits from the study of grids. The notion of an information grid not only implies meshed interconnection, it also infers interoperability among collaborative systems. The current network management paradigm, polled and driven by events, has been organized around centralized intelligence, perhaps organized in a hierarchy of element managers. The proposed approach borrows from grids to incorporate [11]

- Insightful messaging between nodes
- Local state maintained within each node
- Local resources, functionality, and self-awareness
- Ability to coordinate passage of transit traffic

The syntax and semantics of inter-node communication are arguably the biggest challenges in development of an information grid in any domain. The key enabler for this project is a service-oriented architecture and a semantic web created for the telecommunications management domain, which implies common metadata and a common ontology accepted by participants [12]. In either the management or control planes, the ability for autonomic devices to carry out their full spectrum of duties without operator intervention is the long term vision of this project.

Whereas in the current telecommunications paradigm, information is pulled into the network manager for action, this research embraces "smart push" between the NMS and the state engine. One motivation for this is the mitigation of the problem of massive (and duplicated) data sent to managers. The volume of messages increases all forms of latency: congestion delays, queuing delays, processing delays in the host operating system, and processing delays in the management application itself. Event correlation and filtering, in this design, remain delegated to the element manager software, but responsibility for topology, state, and action will reside in the state engine model.

The following drawing represents the information flows between components within the architecture proposed by our research. Nodes will be discussed below, and information flows are labeled by their protocol (e.g., XML or SQL) and an identifier that will be used to describe the use cases.

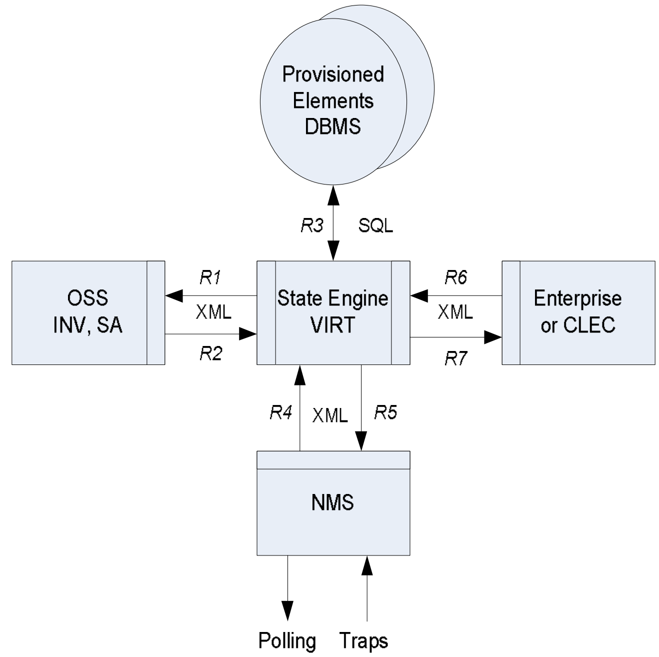

Fig. 3. Graphical representation of State-Based Network Management

In Figure 3,

- The central engine maintains state as defined previously and executes the VIRT process – to monitor conditions in order to maintain state, and to obtain resources necessary to fulfill requirements.
- The rightmost box represents an agent acting on behalf of the customer. In the network management grid, it manages information about users, applications, and locations and must transact with the central engine in terms of aggregated demand and service quality.
- The upper node represents the database local to the central engine. It is currently depicted as a SQL relational database, but we anticipate the availability of native SOA-complaint software.
- The leftmost box represents the collective OSS/BSS capability of one or more providers. The scope of these interfaces is intentionally limited in our subsequent examples.

The bottom node contains the NMS for monitoring and event correlation. This icon represents use of SNMP as well as perhaps a subtree of other element managers. Northbound messaging via XML is expected to be filtered and terse, passing information only about faults with high confidence.

V. Conclusions and Future Work

Much of the motivation for this project comes from anticipated applications involving defense and homeland security. Combining these novel applications with the agile nature of warfighting and coalition-building, expectations for future networks are immense.

The U.S. Defense Information Systems Agency (DISA) has engaged the Net-Centric Enterprise Services (NCES) pilot project to demonstrate wide-ranging applications such as [14]:

- Session Management
- Text Messaging within a web conference
- Whiteboarding
- Application Sharing
- Audio (within a conference session)
- Presence and awareness (of others attending)
- File Transfer
- Voting and Polling
- Video, shared live or playback
- Recording
- Instant Messaging
- Chat rooms
- Virtual Spaces (for asnynchronous collaboration)

Manual methods will be neither effective nor efficient to deploy and monitor these kinds of services in the desired settings. Adaptive network management can be an excellent means to implement and manage these services.

This paper identified semantic reasoning and ontologies as the foundation of adaptive network management and provided an overview of their development.

The next phase of this project is to obtain, when available, a database product that can perform semantic reasoning based on native XML schemas.

Another task within the VIRT project team is to formalize and validate the software architecture.

References

[1] Hoag, John, "Telecommunications Systems Failure: The Electricity Blackout of August 14, 2003," presented at the 12th International Conference on Communication Systems Modeling, 2004.

[2] Hayes-Roth, Richard, "Model-Based Communication Networks and VIRT," Presented to the 10th International Command and Control Research and Technology Symposium, 2005.

[3] Von Schweber, Erick, "VIRT: Issues the VIRT Working Group Must Address," presentation notes dated February 19, 2006.

[4] Terplan, Kornel, OSS Essentials: Support System Solutions for Service Providers, Wiley, 2001.

[5] Keana, Francis, and Andrew Exton, "HP Openview Service Activator 4.1 Software Data Sheet," May, 2005

[6] Gallen, C., and J. Reeve, "Investigating the Feasibility of Open Development of Operations Support Systems," in Proc. of the 9th IFIP/IEEE Int. Conf. on Integrated Network Management, 2005.

[7] British Telecommunications plc, "Open OSS: A New and Complementary Approach to Drive the Implementation of OSS Standards," http://www.tmforum.com,, accessed January 12, 2006.

[8] "OSS Through Java Initiative, The Only Game in Town," http://www.ossj.org/learning/whitepapers.shtml, accessed Jan. 12, 2006.

[9] Balog, Tarus, "Enterprise-Wide Network Management with OpenNMS," O'ReillySysAdmin, http://www.oreillynet.com/pub/a/sysadmin/2005/09/08/opennms.html, accessed January 12, 2006.

[10] Walter, Uwe, and Joachim Charzinski, "Architecture of a Network Control Server for autonomous and efficient operation of Next Generation Networks," presented at the 13th International Conference on Telecommunications Systems Modeling, 2005.

[11] Foster, Ian, "What is the Grid? A Three Point Checklist," Grid Today, October 2002, http:/www-fp.mcs.anl.gov/~foster/Articles, accessed January 12, 2006.

[12] Hayes-Roth, Richard, "Model-Based Comunication Networks: Filtering Information by Value to Improve Collaborative Decision Making," Naval Postgraduate School, Unpublished Manuscript, June 23, 2004.

[13] Hayes-Roth, 2005, *ibid*.

[14] Defense Information Systems Agency, "Net-Centric Enterprise Services Collaboration Service Request for Quotation, Statement of Objectives," Attachment 2, December 16, 2005.