# A Recommender-System for Telecommunications Network Management Actions

John Keeney
NM lab, LM Ericsson,
Athlone, Co. Westmeath, Ireland
john.keeney@ericsson.com

Sven van der Meer
NM lab, LM Ericsson,
Athlone, Co. Westmeath, Ireland
sven.van.der.meer@ericsson.com

Gabriel Hogan
NM lab, LM Ericsson,
Athlone, Co. Westmeath, Ireland
eeighgn@gmail.com

***Abstract*** **-** Research and products in telecoms network management have long been focused on the automation of processes to keep complex managed networks in an operational state, while being profitable to operate for the network operator. Due to the ever increasing scale and complexity of the problem domain, coupled with specific constraints (legal, regulatory, technological change), a fully automated management approach is virtually impossible. It remains the responsibility of human Network Operations Centre (NOC) operators to oversee and manage the running of the network, where their main task is to respond to huge numbers of messages, errors, warnings and faults constantly flowing from the managed network. In the past, network operators either employed expert systems or became very dependent on expert knowledge of their operational staff. Considering the current explosion in size and complexity of managed telecoms network, it is widely understood and accepted that current manual and semi-automated approaches cannot scale. In this paper we investigate the applicability of recommender systems as an approach to assist NOC operators to correctly respond to indications of incidents in the network they are actively managing.

## I. Introduction

Wireless data traffic has been increasing exponentially in recent years, mainly driven by mobile broadband and smart-phones. As demand increases, cell-size is decreasing, and the number of cells is increasing, with nano- and pico-cells becoming a reality[1]. This results in a rapid increase in the number and type of devices and services to monitor and manage[1] to provide efficient and high quality services to end users. Telecoms operators are already experiencing an explosion in the scale, complexity and cost of managing such networks.

To manage a telecoms network, operators require an Operation Support Systems (OSS). A traditional OSS is usually composed of components, typically categorized as components for managing faults (FM), configurations (CM), accounting, performance (PM), and security. Although these components usually work isolated from each other, more recent integrated approaches (for instance combining fault and performance management functionality) already show that such "management silos" are unsustainable and make end to end monitoring and management of a network a very difficult task.

Currently, monitoring and managing telecoms networks involves significant human interaction, where management personnel are aided by an OSS for semi-automated processing of huge amounts of network management data and configuration management. Despite this, many of the tasks required of human network managers are repetitive and routine, and involve wading through huge amounts of monitoring data.

Given the ongoing explosion in data and complexity in modern telecoms networks, it is also not feasible to continue assigning ever more resources (equipment, staff) to this problem as cost is quickly becoming prohibitive and unsustainable. New technology and approaches are required, quickly, to solve these challenges.

## II. Telecoms Network Management

### A. Monitoring and Root Cause Analysis

Most operational management activities are based on management data related to network equipment, network connections, services running over the network, or environmental conditions (for instance temperature). The first step of processing this data is usually a correlation of incoming data. Correlations document temporal or causal relationship between otherwise unrelated events helping to understand the underlying network condition. Two simple examples illustrate correlations:

a) An alarm event indicating that CPU temperature on a particular base-station processing board has reached a critical level could be correlated to a series of monitoring events that the temperature of the board and other boards in the cabinet, the temperature of the containing cabinet, and the temperature of the room, have all been steadily increasing.

b) Two different radio base stations indicate that a single stationary end user device is generating a huge amount of cell handover traffic as it constantly registers and deregisters ("ping-ponging") between the two cells. This is despite the fact that each individual registration and de-registration might be successful and without error itself and so does not raise any alarms. This can be because either the base stations or the device cannot correctly determine which base station is appropriate.

Associating a cause with a correlation (and thus tagging it as the root cause) is a very complex task that incorporates knowledge about the network topology, in-depth knowledge about the involved technologies and knowledge about the current configuration and context of participating network equipment and environment. Due to its complex and very context-dependent nature, knowledge about correlations and root causes are very difficult to formalize and encode, but quite often such expertise is part of the institutional knowledge of the telecoms network operator and the NOC staff.

### B. Corrective Actions

Given some correlated indications of a non-trivial problem in the managed network, the network manager or NOC operator immediately needs to either fix the root cause of the problem, or mitigate against the

978-3-901882-50-0 

performance-degrading side effects of the problem. There are typically a number of potential actions and the selection of the most appropriate corrective action is based on experience or contextual information. For the examples given above, a correction to the temperature related faults could be related to an air conditioner failure so corrective actions might include: 1) start an auxiliary air conditioning unit or fan; 2) remove load from the cabinet devices to a redundant cabinet and gracefully shut down the affected components to cool; or 3) dispatch an emergency maintenance team to fix or reset the faulty air conditioner. For the second "ping-pong" example: 1) one base station could be instructed to temporarily reject registration requests from the given device; or 2) the transmission strength or alignment of one of the base-station's antennae could be dynamically reconfigured to reduce the radio signal overlap area between cells.

While some corrective actions are standardized and others are described in best practice documents, most network operators will have their own individual set of actions for a given set of network characteristics. This can be due to a wide range of considerations, e.g.: specific local legal and technical constraints; different network topologies; different usage patterns at different times in different parts of the network; the operators' business-level preferences; different internal management processes and tools; different QoS commitments for different customers; external considerations e.g. weather, location, holidays, events, etc..

## III. Our Approach

Our research investigates the applicability of advanced analytics to analyze and understand network management data (events) and the use of recommender systems to support management decisions and network planning activities. In particular, we are investigating:

***Intelligent data collection*** – current tools collect huge amounts of monitoring data, which is then warehoused for off-line inspection (if any). Applying intelligent filtering/selection of the data at collection time will help to collect more useful information, thereby supporting more analysis at, or just after collection time.

***Cross-layer network-wide correlation of monitoring data*** – A key difficulty with current monitoring approaches is mapping from low-level monitoring data and events to service-level performance indicators. New approaches investigate (deductive) data mining to perform deterministic service-level diagnoses and causal analysis of both stored and streaming monitoring data.

***Predict occurrence of serious errors before they occur*** – Usually, once a serious error occurs, the managed system will have already degraded or failed. Combining (deductive) data mining techniques with (inductive) predictive analytics can identify temporal and causal signatures of errors or avoidable degradations, which can then be used to predict the occurrence of such events.

***Proactively recommend management actions*** – Current telecommunication network management actions are usually triggered by human operators, even where the most reasonable management action for a given situation could be inferred from past actions in similar situations. Combined with a prediction system, a recommender system could identify previously successful corrective actions before or as the need for such actions becomes apparent to the human operator. This exploits the Recommender System approach, whereby a number of recommended and ranked suggested solutions are presented to a human user. This approach observes the response actions of the user to seed a set of recommended solutions, and where the use (or not) of recommendations is used to inform subsequent recommendations.

## IV. Recommending Management Actions

Before we can proactively recommend a management action, we must first receive live or predictive information about the state of the managed network and its associated managed objects. Using existing state and context a Recommender System can select appropriate corrective actions, interact with the human operators for feedback and validation, and learn from past successful recommendation. Recall the examples given in section 2:

We could predict with some confidence that there was a general air-conditioning problem at the base-station location because all temperature sensors in that location or cabinet have been increasing over a period of days. Whenever this pattern occurred and was alerted before, the NOC agent performed one of a small set of actions. Based on these previous actions, the system can recommend to again proactively schedule maintenance on the air-conditioning systems and cabinet fans at the site. By being proactive, the maintenance can be scheduled at a time when the maintenance personnel are available or when the base station is not busy, so potential traffic disruption will be avoided. If the NOC agent rejects or adjusts the recommended response, then that information will be available for the next time similar patterns occur.

Given a ping-pong situation, we predict that the given user device will continue to "ping-pong" between base stations. Whenever this pattern was detected and alerted before, the NOC agent selected from a set of responses to mitigate the issue. The system can therefore rank a list of recommendations: 1) that the particular handset be temporarily prioritized on one base station; 2) temporarily reconfigure the coverage area of the base stations; 3) proactively schedule a reconfiguration or replanning of the base stations' radio parameters. Again, depending on which action is chosen, the ranking may change next time the pattern occurs.

A recommender system is a software system that recommends items to a user based on a prediction that the recommended items will be useful to the user based on their current task or state (context). There are a number of different approaches to implement this prediction algorithm, but most are based on selecting items, or similar items, that the user (or a similar user) in a similar context "liked" previously. In this case, we model the user's context (task or state) to be based on the characteristics of the problem they are trying to solve. We model the "like" operation to be the use or refinement of a recommended action by the user (or a similar user). The recommendation algorithm then boils down to comparing and ranking context states in a context space, where each

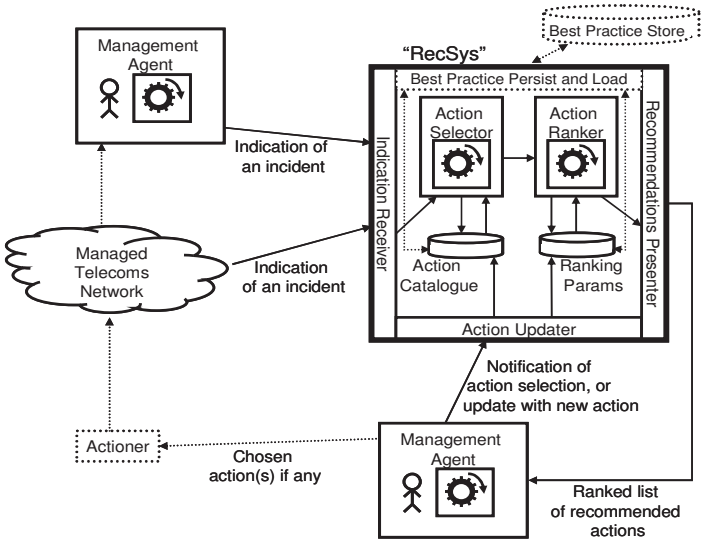


**Figure 1: Recommender System for Network Management**

context state (network problem) has one or more candidate items (management actions) associated with it. Based on the ranked similar context states, a ranked list of associated items can be extracted, thus forming a ranked list of recommended management actions in response to a network problem. Figure 1 shows such a recommender system for telecoms network management.

When an incident occurs, or is predicted by the advanced analytics system, an incident report (e.g. trouble ticket) is forwarded to the recommender system. Incident reports may originate from the managed network or from one or more (human or automated) management agents. The parameters of incident indication are used to select from a list of previously applied or suggested actions, drawn for the action catalogue by an action selector. The action list is then sorted and ranked by the action ranker based on ranking parameters derived from guidelines of best practice and/or the historical adoption of similar actions in response to similar past incident indications. A ranked list of suggested actions may then be further manipulated, e.g. select only the highest ranked suggestions. Because more than one recommendation is presented, but the recommendations are ranked, the management agent can quickly see alternative recommendations to a given incident, and the degree to which those recommendations are deemed suitable. The manager then retains the authority to re-rank the list, select one or more of the recommendations, ignore all recommendations, or explicitly select, define or refine a different action. This step provides feedback to the recommendation system to adjust or extend its action catalogue and refine its ranking parameters. Using this continuous feedback process, the recommender system learns and evolves. This feedback mechanism provides a way to deal with the "cold start" problem inherent in all recommender systems. This approach also provides a mechanism to evolve as the network, context, institutional knowledge, or business priorities for the network evolves.

Recall the use cases presented in section 2, where many times a single corrective action would not be appropriate. For the first air-conditioning case, one network operator might have an emergency call-out contract with an air-conditioning team nearby, while another may have suitable redundancies in place to suspend the affected components. Yet another operator may already have a maintenance team onsite dealing with another issue. For the second case: an operator might decide that the problematic device should be restricted from registering to one of the base stations while another operator might consider if the device user is a high-priority customer where such actions may not be suitable.

Another key feature of this system is the ability to load or persist the internal state of the system, in particular the action catalogue and ranking parameters. This has a number of advantages: a solution set of best practice recommendations can bootstrap the recommender system, thus alleviating the "cold-start" problem. Additional newer best practice solutions and procedures can be pushed by vendors to customers as new technological advances occur. This approach also supports network providers to dynamically and continuously encode their own institutional knowledge, which can otherwise be lost, or never fully developed, due to high staff turnover in some NOC environments.

## V. RELATED WORK

Most existing solutions for human-oriented telecoms network management use static workflow approaches to respond to incidents in the network. These workflows are difficult to change as the behavior or characteristics of the network evolves. Workflows have to be written by hand based on existing 'best-practice', and are manually targeted to individual telecommunication networks. Within specific networks, there is typically one 'best' response to a given problem which cannot learn from other solutions employed in other telecommunication networks. These approaches invariably select only one statically defined candidate management action.

Automated and semi-automated telecoms network management decision processes are currently realized using similar static workflows, codebooks or run-books implemented as rules, or in extended policy rules. If such rules are accompanied by an inference engine and a knowledge base, this approach can be termed an "expert system" for decisions. Telecoms network management based on or using expert systems rely on knowledge (rules) and inference using logic (for instance propositional logic, 1st order logic, fuzzy logic) to select a management action. Statically encoded policy-based systems, workflow systems, generic rule-based systems, expert systems, etc., are designed to emulate the decision making of human experts, but invariably select only one statically defined candidate management action. In traditional policy-, rule-, expert system-based management, this presented action may also be directly and immediately applied to the managed system without the intervention of a human operator. When this automated approach is taken the human operator no longer has fine-grained dynamic control to exploit their expertise to optimize the management of the network.

The use of policy-based, rule-based, workflow-based or expert system-based systems for telecoms network management has several other serious disadvantages. It

requires a careful selection of the rule language and a (constant) maintenance of the defined and used rules, including contradictions and conflicts in the maintained rule set [7]. Even where the rules are auto-generated, their maintenance becomes a major overhead [8]. The logic employed in the inference engine constrains the ability to infer solutions from the rule base (or knowledge base). For instance, fuzzy logic can be used to interactively infer towards a defined optimum, while 1st order logic can be used to establish sound and complete statements. Such systems are also brittle in terms of their ability to cope with unknown cases without existing rules or knowledge.

As an extension of expert-systems, a number of AI-based techniques have also been used for Telecoms Network Management. For example the systems presented in [2][3][4][5] uses a case-based reasoning (CBR) approach on network fault management trouble tickets to select previous incident trouble tickets (and their encoded solutions) based on the degree of similarity between the given reported incident and previous incidents. While these CBR approaches progressively refine similarities between trouble tickets, or similarities between solutions, so these cases can be retrieved when similar new case occurs. However, they do not automatically learn to rank proposed solutions to problems based on uptake by the manager.

Several exiting systems also use AI techniques such as CBR, artificial neural networks, Bayesian belief networks, etc, for fault filtering, root cause analysis, failure prediction and correlation for telecoms network management [6][8][9], often as part of an autonomic system [10][11]. While some of these techniques can progressively learn and identify causes for network errors, they do not in general present solutions. For those systems that can provide a solution: they generally require a tight binding/encoding between the problem and solution; or they do not provide multiple solutions; or they do not rank those solutions based on uptake by the management agent.

## VI. CONCLUSIONS

There is no known existing or proposed state of the art telecoms network management system using a Recommender System. This Recommender System approach has the advantage that the network management system will learn the human expert's approaches, preferences and knowledge over time and evolve the recommendations to exploit this acquired knowledge. An important distinction of this recommender system approach from some state of the art systems is that the expert human manager retains the ultimate decision to ignore or select from the recommendations presented. Thus, the human manager maintains overall control of the managed network, but with the assistance of the Recommender System. This maintains management oversight and control over the network while still allowing and supporting special cases and exceptions.

Most existing recommender systems focus recommendations based on feedback from a large number of users, however the system presented here focuses on a relatively small number of users, with clear roles or tasks, but with a large number of interactions with the system. Unlike any state of the art telecoms network management system (proposed or actual) this system is designed to present ranked lists of recommended actions, with the ranking partly based on historical uptake of the recommended actions. Yet another important difference of this system is that the action selection mechanism and ranking mechanism are updated when recommendations are used, ignored or extended thus allowing older selections to decay as the managed network evolves.

Because more than one recommended solution is presented, but the recommendations are ranked, the manager quickly sees alternative recommendations to a given incident, and the degree of suitability of those recommendations. This will reduce uncertainty and errors in the management of the telecoms network.

The proposed use of a recommender system approach will support faster resolution of network issues as candidate solutions are presented to the network manager (NOC operator) rather than a list of incidents. Encodings of 'best practice' can be kept up to date by exporting the recommender system state and may be shared between network deployments, perhaps even between network operators. As the system learns, tuning 'best practice' for a given network deployment and behavior characteristics, the expertise required by the NOC personnel for day-to-day operation and maintenance of the network is reduced, thus reducing cost, improving management throughput, and freeing up time and resources for the managers to concentrate on more strategic management issues.

## VII. REFERENCES


[1] Ericsson: "More than 50 billion connected devices", http://www.ericsson.com/res/docs/whitepapers/wp-50-billions.pdf (Ericsson Whitepaper: February 2011)

[2] Lewis, L.: "A Case-Based Reasoning Approach to the Management of Faults in Communications Networks". IEEE, Conference on Computer Communications. 1993.

[3] Melchiors, C., Tarouco, L, Althoff, K-D., "Fault Management in Computer Networks Using Case-Based Reasoning: DUMBO System" in Case-Based Reasoning Research and Development, LNCS1650, 1999.

[4] Melchiors, C., Tarouco, L.M.R., "Troubleshooting network faults using past experience," IEEE/IFIP NOMS, 2000

[5] Penido, G., Nogueira, J.M., Machado, C., "An automatic fault diagnosis and correction system for telecommunications management," IFIP/IEEE IM, 1999.

[6] Catal, C., Diri, B., "A systematic review of software fault prediction studies", Expert Systems with Applications, 36(4), May 2009

[7] Lupu, E.C., Sloman, M. "Conflicts in policy-based distributed systems management", IEEE TOSE, 25(6), 1999

[8] Sterrit, R., "Facing fault management as it is, aiming for what you would like it to be" Soft-Ware, LNCS 2311, 2002

[9] Covo, A.A.; Moruzzi, T.M.; Peterson, E.D. "AI-assisted telecommunications network management". Global Telecommunications Conference, (GLOBECOM), 1989

[10] "Autonomic Network Management Principles: From Concepts to Applications", Agoulmine., N., (ed), Academic Press, 2010

[11] Samaan, N., Karmouch, A., "Towards Autonomic Network Management: an Analysis of Current and Future Research Directions," IEEE Commun. Surveys Tuts., 11(3), 2009