Automated Real Time Performance Management for Mobile Networks

Tobias Bandh, Georg Carle
Computer Networks and Internet,
University of Tübingen,
72076 Tübingen, Germany
{bandh,carle}@informatik.uni-tuebingen.de

Henning Sanneck, Lars-Christoph Schmelz
Nokia Siemens Networks
St.-Martin Strasse 76
81541 München
{henning.sanneck, schmelz.lars}@nsn.com

Abstract

Real Time Performance Management (RTPM) is expected to be a key features of future developments of mobile networks. Although it is already requested by operators of mobile networks only basic approaches have been realized.

Implementing access to real time performance data in a mobile network, with a large number of manageable nodes, leads to a huge resource consumption. Not only network bandwidth but also processing resources. Current PM schemes can hardly be extended to support real time performance management.

We present a policy based approach that is capable to deal with challenges imposed by the request for real time performance data. Typical workflows are identified and analysed. Based on these analyses a policy based performance management system is introduced that allows a balanced configuration which does not exceed available resources but still satisfies the information requirements of the human operators.

1 Introduction

Performance management in today's networks is mostly based on offline processing of data. Information on the network performance is gathered in a fixed work cycle: Performance data, generated at the network elements (NEs), is gathered at the end of a 3GPP defined granularity period (GP) [5]. After that the data is copied into a standardized Performance-Measurement-Result-File (PMR-File) [4] which is then transferred to a storage at the Element Management System (EMS), where it is available for further processing.

Operators use an offline tool set to generate measurement reports where raw measurement values are aggregated into more meaningful ones, so called Key Performance Indicators (KPI). Although using standardized counters [6], KPI definitions are vendor and operator specific. However, there are efforts to standardize operator independent KPI definitions [7].

Depending on the information included, these reports are used to conduct different types of analyses:

1. Long Term Analysis: Data collected over a very long period of time (several weeks or months) is analysed to detect long term changes in the Network. For example: increasing load in a certain region of the network.

2. Mid Term Analysis: The data of several days is analysed. For example to check if changes applied to the network have the intended results.

3. Short Term Analysis: These analyses are mainly used to manage the network on special occasions where the operator expects or knows about performance problems. Typical situations when short term analysis is used are the monitoring of the network performance during a fair or sporting event where a load much higher than regular is expected, or on customer complaints.

For two reasons today's Short Term Performance Management does not include real time data. On the one hand the usage of real time data would consume a large share of all available resources and on the other hand 3GPP performance management standards do not define granularity periods less than 5 minutes.

We use the term real time data for data that is generated, transmitted and displayed with a minimal delay. This means that instead of the standard file based performance data transport scheme a stream like transport scheme which ensures a minimal end to end delay is applied.

Today's PM schemes mostly use a minimal GP of 30 minutes to reduce the negative impacts caused by the gathering and transport of data on the network itself and the attached network elements. This results in the fact that even if the data is immediately processed the delay between occurrence and detection of an event can be up to 30 minutes. Real time access to performance data would increase the

possibilities for short time analyses. If the operator had real time access, he could perform sophisticated performance management and possibly detect and solve upcoming problems before there are any noticeable negative effects.

There are several limitations that prohibit a simple reduction of the minimal granularity period down to for example one minute:

1. Processing Power: Due to the large number of available performance counters gathering them into a PMR-File is very time- and processing power consuming. If this had to be done more frequently there would not be enough processing power to finish the task before the next GP ends.All processing of the raw performance data into a human readable format for a large number of network elements is done at EMS level, which consumes a large amount of processing resources at EMS level.

2. Bandwidth: The bandwidth available to transmit the collected data is limited. For example the total management bandwidth for a 3G Radio Network Controller (RNC) is 2 MBit/s, whereof only a fraction is available for performance management data. Therefore it could happen that it is not possible to transfer all data that has been gathered within one granularity period.

3. Human limitations: In this basic scheme, the operator is faced with all incoming data. He has to decide whether the data indicates a normal, critical or upcoming critical situation. Even if short term analyses would only be performed for a small set of network elements it is hard for the operator to monitor and assess all incoming data.

To overcome these limitations two major changes have to be applied to reach the goal of real time performance management:

1. Preprocessing: If the processing steps on the gathered information are already performed at NE level the consumption of as well bandwidth and processing power at EMS level can be reduced.

2. Data Restriction: The second method that can be used to reduce the resource consumption is the restriction of the data gathered. The measurement jobs are configured in a way that they only retrieve the information that is of interest.

To improve such a setting the usage of a policy based RTPM subsystem at EMS level is proposed. It automatically monitors incoming values and independently changes the configuration of the measurement jobs responsible for the data collection. In this approach the operator is only notified in critical situations.

2 Basics

This section shows the basics needed for the presented approach.

2.1 Basic Performance Monitoring Workflow

This section describes the performance management tasks performed by the human operator for monitoring network elements.

1. Monitoring incoming measurement results:
One of the major tasks of a human operator is to monitor incoming measurement results. While there are no critical situations the operators checks the trends of the KPIs in order to assess the future state of the network.

2. Configuration of the network elements:
There are two basic possibilities that are mainly used for NE PM configuration, activation and deactivation of measurement jobs and the adjustment of GPs.
To reduce the load on processing and transmission units only some basic measurement jobs are activated. They provide a basic overview on the networks state, but no detailed information. In case of a detected possibly critical situation, measurement jobs which provide additional information are activated. After the situation is back to normal they are suspended. Even if the performance management system features access to real time data, the operator cannot activate real time data streams with a small GP on all NEs, this would consume all resources available. Therefore only measurement jobs which provide basic information are initialized with a longer granularity period, all other measurement jobs are suspended. The granularity period can be adapted to the state of the network to satisfy the requested level of information.


Figure 1. Basic PM Workflow

2.2 Policy Based Management

Policy based management systems are mainly used within complex, dynamically changing environments. The policy based approach separates the logic from the implementation, which allows the user to install the system and

afterwards specify and change the policies at runtime. Policies can be specified in a very abstract way so that only a few policies are needed to deal with a broad range of events.

The policy framework parses the high level policies and turns them into machine understandable codes, which are then deployed to the policy enforcement point. In case of a policy event, the policy conditions are checked and based on these checks actions are triggered.

The power of policy based management systems can be used to the full extent, if it is possible to formulate as many tasks previously performed manually as possible into non interfering condition-action-pairs. Those can then be transformed into policies and the tasks can be performed automatically.

2.3 Policy Based Performance Management in Mobile Networks

Mobile networks are very complex distributed systems with a large number of attached and interconnected nodes. For performance management the operators have to cope with a large amount of performance data generated by the network elements. Using policy based performance management here aims to introduce the flexibility and the power of the policy based approach. The main intention is to utilize the provided capabilities to improve real time performance management. Especially to reduce the processing time of measurement jobs that provide the performance information.

2.4 RTPM Basics

This section describes a generic RTPM system that satisfies all preconditions.

The generic RTPM system provides measurement jobs which are configurable in two ways. For the basic job specification the operator has to provide the information which counters should be collected and how they should be processed. Which mostly is a formula that uses the counter values to calculate a KPI. After the jobs are up and running they can be adapted to different situations by changing their GP to any operator defined value. This mechanism can even be used to suspend measurement jobs by inserting a infinite value for the GP. If the job should be reactivated the GP is reset to a finite value.

We use a generic RTPM system which consists of three basic domains (c.f. Figure 2):

1. NE-Domain: Contains the network elements.

2. EMS-Domain: Contains the policy based RTPM Subsystem.

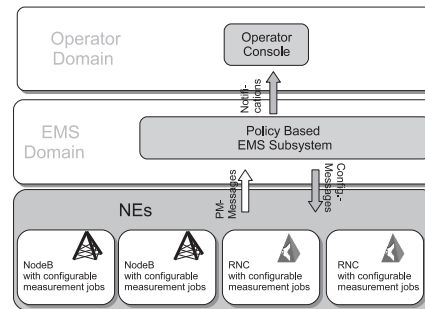3. Operator-Domain: Contains all tools used directly by the operator.

Figure 2. Dataflow of the generic RTPM System with 3G NEs

3 System Design

As the basic PM functions of 3G networks are already located at the EMS level it was decided to place the policy framework at the same level. The policy system receives incoming performance measurement values, analyses them and performs the actions predefined by the network operator. There is a broad range of possible actions that can be taken. Information can be gathered, analysed and processed and afterwards be sent to the operator console, or configuration changes at network element level can automatically be triggered. Possible configuration changes are for example the activation or deactivation of measurement jobs, changes of their granularity periods or any other possible configuration of network elements.

The main problem is the definition of the policies. Managing a mobile network is a very complex task, there are a lot of cause and effect interconnections that complicate the specification of simple condition - action pairs. On this account the next section shows an example that is quite simple but demonstrates the power of the policy based approach. We use two different approaches, a basic, EMS centric approach where all decisions are taken at EMS level. The second, distributed approach corresponds to the expected evolution of mobile networks and therefore delegates some parts of the information processing to the network elements.

3.1 Basic, centralized approach

In order to identify automatable tasks, the RTPM workflows were analysed (c.f. Section 2.2.1). These analyses showed, that a major enhancement could be reached by automating two tasks. The monitoring of incoming messages carrying performance information and the application of configuration changes to measurement jobs

Because of the complex interconnections of different causes and effects this system does not trigger any countermeasures but only controls the measurement jobs on a

possibly large number of network elements.

The main intention is to reduce on the one hand the resource consumption as much as possible but on the other hand to increase the information available to the human operator in critical situations.

The basic system is very simple. It can be described in a state-machine-like way with only three states (c.f. Figure 3). It is based on the assumption that critical situations can be detected by defining and monitoring thresholds for KPI values. Changing the message rate corresponds to a state transition.


Figure 3. Finite state machine like representation of the possible states of the system

The different states of the system correspond to the following tasks:

1. Value Prediction: This is the state for non-critical situations. In this state the measurement jobs providing detailed information are deactivated and all active measurement jobs have comparably long granularity periods (e.g. 10 minutes). Incoming KPI values are used for two tasks. They are checked whether the threshold has been crossed. In that case the system directly changes into alarm state. To reduce the risk to miss a critical situation and minimize the time between the occurrence and the detection of critical situations, received KPIs are also used to calculate a prediction of the next values. The prediction is also checked against the defined threshold, in case it is above, the system reduces the granularity period of the affected measurement jobs to a middle value to increase the information available, by increasing the message rate.

2. Threshold Monitoring: After the value prediction indicates an upcoming critical situation and the message rate is increased the system discards the value prediction and enters a pure monitoring mode. Incoming KPI values are checked if they crossed the threshold or not. If it is crossed, the system changes into the alarm state, if the monitored values do not cross the threshold for a certain time the system is reset into value prediction mode.

3. Alarm State: As soon as a basic KPI crosses the threshold all corresponding measurement jobs, that provide detailed information, are activated and the operator is informed about the critical situation. If the KPI values received stay below the threshold for a predefined time, all but the basic measurement jobs are stopped and at the same time the granularity period of the basic measurement job is increased to a middle value.

The policy system can be used to preprocess and analyze the received KPI values before they are forwarded to the operator console. If the operator receives messages which already contain a first analysis, reaction advices and information on possible interconnected problems, the reaction time and the possibility of missconfigurations can be reduced. If causes and effects are completely understood, the policy system can even be used to trigger countermeasures, but this is out of scope of this paper.

3.2 Evolutionary, distributed approach

The presented system uses a centralized approach, where all management logic is located at the element management system. But following the expected evolution of mobile networks, parts of the policy system and large parts of the logic could be distributed to the NEs. Because most of the information from the NodeBs is concentrated at the RNCs the logic is placed at RNC and not at NodeB level. Still, this will result in an additional load reduction at the network links and the policy system at the EMS. Although the presented system already tries to reduce the number of transmitted messages, already in a non critical state the large number of network elements generate a huge number of messages that have to be processed.

In this evolutionary approach the tasks: "Value Prediction" and "Threshold Monitoring" will be located at the network elements. The measurement jobs at the NEs will not emit any messages until the critical threshold is crossed. As soon as this happens an alarm message will be sent to the policy system at EMS level. Now it can be decided whether to request high rate real time information from the measurement job at network element and possibly activate more measurement jobs or not. Such a decision can be based on the load situation at the EMS, link load or any other operator defined criteria.

It is obvious that not all of the logic can be shifted from the EMS to the NEs. To be able to deal with problems affecting multiple NodeBs assigned to different RNCs.

4 Evaluation and Results

The main focus of the evaluation is on the reliability in detecting critical situations, the number of messages transmitted and their impact on the network links. In particular

it is examined if it is possible to discover all critical situations with a minimal average delay between occurrence and detection. Finally it will be estimated how many concurrent measurement jobs in which configuration can be supported by a present day third generation mobile network.

4.1 Evaluation Setting

All evaluations are conducted with an identical data basis. The data originates from a real network[10], the information is stored in a file containing all counter values needed to calculate the "BranchAdditionFailure-Rate" as basic KPI and the "Intra NodeB-", "Intra RNC-" and "InterRNC" KPIs as so called Sub-KPIs. Sub-KPIs incorporate only a subset of the counters used for the basic KPI calculation, for this reason they provide more specific information. For the evaluation a PM data tracefile holding the values of a 24 hours measurement cycle with a resolution of one minute was needed. As there are no such files available, the file was generated by interpolating another file holding the values with a 30 minutes resolution. Critical situations are defined as periods where the KPI value exceeds an operator defined threshold. We used a static threshold of 20%. This means that seven critical situations with duration between 15 and 24 minutes are detectable. Figure 4 shows the changes of the basic KPI within 24 hours. The following
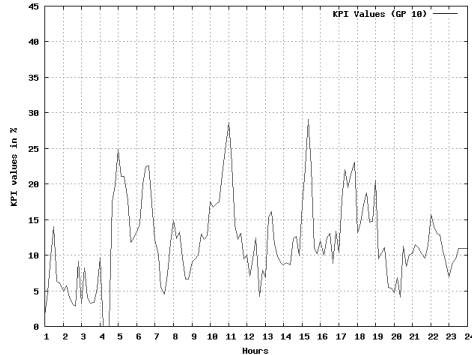

Figure 4. Changes of the basic KPI within 24 hours

PM schemes are compared:

1. Conventional PM: Constant 30 minutes GP, transmission of the full counterset, no matter how many of the counters are needed.

2. RTPM1: Fixed one minute GP, four active measurement jobs, constantly calculating and sending all KPI and SubKPI values.

3. RTPM2: Variable GP, controlled by the policy based EMS subsystem (possible values are 1, 5 and 10 min-

utes). Dynamic activation and suspension of measurement jobs.

4. RTPM3: This scheme is basically identical with RTPM2, but most of the logic is shifted to the network elements (Compare Section 3. 3.2). The system is configured to activate high rate messages from all four measurement jobs whenever the threshold is crossed.

The comparison of the RTPM-Schemes is based on the following values measured:

1. Total number of messages emitted within 24 hours

2. Comparison of the number of messages sent per hour (min, max, avg)

3. Comparison of the delay between occurrence and detection (min, max, avg)

4. Number of detected critical events.

5. Number of predicted but not occurred threshold crossings.

6. Impacts on communication links

4.2 Results

The following subsection shows the results of the measurements conducted.

4.2.1 Number of Messages Exchanged

It is obvious, that compared to today's PM-Schemes, the number of messages generated by the RTPM-Schemes with shorter granularity periods and per job message transmission will be increased. In the scenario with a fixed one minute GP and all four measurement jobs active, a single network element will emit 5760 messages within 24 hours. But already by applying the first scheme with a variable GP (RTPM2) the number of messages drops to a total of 947. RTPM3 does not result in a comparably strong reduction. Using this configuration 757 messages are transmitted within 24 hours. Table 1 shows the message statistics for the schemes with variable granularity period. The numbers in brackets give the total number of messages for all four measurement jobs. Looking at the message rate statistics shows that there is only a minor difference of the maximal number of messages. This can be explained by the fact that when in Alarm state both schemes emit the same number of messages. The more network element centric scheme (RTPM3) suppresses only all low and middle rate messages. If there was a critical situation with a duration greater than 60 minutes the maximal number of messages would be 60 (240) for both schemes.

Table 1. Result Table: Message rate statistics

	RTPM2	RTPM3
Min /hour	6	0
Max /hour	47 (173)	43 (161)
Avg /hour	16	8

4.2.2 Detection Reliability

One of the main intentions in using real time data is to enhance the detection reliability and reduce the detection delay for critical situations. Figure 5 shows that both objectives are met. The policy controlled schemes detect all seven critical situations with an average delay of 3.57 minutes, while the conventional scheme misses three critical situations and has an average detection delay of about 25.5 minutes. Because of a very simple prediction algorithm, which only uses the average changes of the last five measurements to calculate a possible next value, three predicted critical situations did not become critical.


Figure 5. Delay between the start and the detection of a critical situation

4.2.3 Impacts on Communication Links

In order to estimate the impacts on the communication links, variations of the following scenario were examined.

The question is which configurations can be used to monitor a RNC domain, using the policy controlled RTPM schemes. As a basic scenario a RNC with a 2 MBit/s O&M link to the EMS is assumed. Only 33% of the O&M link are assumed to be available for performance management (Furtherly referred to as PM Link). This limits the maximal amount of data per hour to 300 MB. The RNC domain comprises 1000 cells with 10 basic KPIs that are to be monitored. Each of them has three additional SubKPIs. A KPI message is assumed to be 500 Bytes long.

Four different variations of the scenario are examined:

1. Fixed one minute GP (RTPM1): Already the monitoring of the basic KPIs on all cells will use exactly 100% of the PM-Link capacity. SubKPI values cannot be requested without suspending other measurement jobs.

2. Policy Based PM with changing GP (RTPM2): Using the results shown in Table 1 shows that monitoring only the basic KPIs a maximal number of 47 messages per hour has to be transmitted. This results in a PM-Link usage of $\approx 79\%$, which allows the operator to request some SubKPI information without suspending the basic monitoring.

3. Identical with Scenario 2 but with a minimal GP of 10 seconds (instead of 1 minute): This scenario generates ≈ 1335 MB of data or a PM-Link usage of 445%. In addition to the overall traffic, the data generated within one GP cannot be transmitted within the next GP. So the monitoring is restricted to the basic KPI $< 25\%$ of all cells. No SubKPI measurement jobs can be activated.

4. Identical with Scenario 2 but the additional three SubKPIs are activated as soon as the GP is reduced to 1 minute: This scenario generates a total of ≈ 895 MB of data ($\approx 298\%$ PM-Link usage). If this scheme should be applied it may not be activated on more than $\approx 33\%$ of the cells.

These examples are based on a worst case assumptions, where all cells behave identically, all critical situations appear at the same time in all cells and last exactly the same time. For real networks this is not a very probable assumption. Mostly only clusters of cells show this behavior. That means that in real networks with an average load level the operator can use these PM schemes on a larger number of network elements and activate the calculation of more KPIs and SubKPIs. The operator has to specify the policies in a way to achieve an optimal distribution between resource consumption and information availability.

5 Related Work

Although there was no directly related work in the area of policy based performance management for mobile networks, this work has been influenced by a multitude of projects and techniques. In the area of mobile networks and their management an essential source of information are the standards provided by the 3GPP and the TMF [4, 5, 6, 7] to name only a few. In addition to these fundamental information there is a broad range of information on network and real time network management [8, 15, 13]. As already

mentioned policy based management systems tend to be developed to a state of the art solution for managing complex networks. Therefore there are already a large number of research papers on this topic [16, 17, 9, 12, 11].

To be able to choose a suitable policy framework that could be used as a basic for the presented system, several frameworks were examined. From all policy frameworks examined ([14, 3, 2, 1]) Ponder2 was chosen because it satisfied all requirements.

6 Conclusion

Real time access to the results of the measurement jobs enhances performance management. Critical situations can be detected before they happen and the operator is provided with exactly the information he needs to counteract to the upcoming critical situations.

The problem that there is no chance to manually generate an optimal configuration for thousands of cells due to the continuous changing conditions, can be solved by the presented scheme. A policy based approach can be used to assure a balanced configuration of the measurement jobs, an optimal resource usage and to overcome the human limitations.

Defining policies that ensure that the information requirements of the human operators are satisfied and include the limitations of the transmission and processing capabilities is a non-trivial task. Therefore it is advantageous that policies can be defined offline and on a very abstract level. As soon as a problem is fully understood, policies can be defined to match a class of similar problems.

Another major advantage of the policy based approach is the possibility of dynamic adaptation. In contrast to hard coded logic, policies can easily be changed at runtime without the need to recompile and restart the complete system, even if the conditions within a mobile network are not expected to change that frequently. For example: the operator has defined a policy set that works fine in 99% of all situations. For exceptional situations as for example a fair, an additional policy set for the concerned cells can be defined and inserted into the system, while the rest of the system stays untouched and will continue working.

Future work can enhance the system by introducing dynamic thresholds that can reduce the number of false positives or negatives. This could be done by using prediction algorithms that consider seasonal trends or learn from operator response on warnings emitted by the system. To achieve these enhancements the possibility of dynamic adaptation of the policy system will be used to a greater extent as it is done in the presented approach.

References

[1] IBM pmac - http://www.alphaworks.ibm.com/tech/pmac. 02.02.2007.

[2] JBossRules - http://labs.jboss.com/portal/jbossrules/docs/. 02.02.2007.

[3] Jessrules - http://herzberg.ca.sandia.gov/jess/docs/70/. 02.02.2007.

[4] 3GPP. 3GPP TS 32.104 V4.0.0. Technical Specification Release 4, 3rd Generation Partnership Projekt; Technical Specification Group Services and System Aspects; Telecommunication Management; 3G Performance Management(PM)(Release4),032001., 03 2001.

[5] 3GPP. 3GPP TS 32.401 V6.4.1. Technical Specification Concepts and Requirements - Version 6, 3rd Generation Partnership Projekt, Techinal Specification Group Services and System Aspects; Telecommunication Management; Performance Management(PM);,022005., 02 2005.

[6] 3GPP. 3GPP TS 32.403 V7.1.0. Technical Specification Performance measurements - UMTS and combined UMTS/GSM, 3rd Generation Partnership Projekt; Technical Specification Group Services and System Aspects; Telecommunication Management; Performance Management (PM);,122005., 12 2005.

[7] T. Forum. Cooperative OSS Project (CO OP) High Level Architecture . Technical Specification Ver. 1.0, Telemanagement Forum - CO OP.

[8] H.-G. Hegering. *Integriertes Management vernetzter Systeme*. dpunkt.verlag, 1999. ISBN 3-932588-16-9.

[9] S. Iacono, F. Arneodo, K. Cardoso, M. G. Genet, and D. Zeghlache. Policy Based Management for Next Generation Mobile Networks. *Wireless Communications and Networking*, 2:1350– 1354, 2003.

[10] S. C. M. N. S. Implementation. RNC Performance Data Tracefile. Tracefile, 03 2006.

[11] E. Lupu and M. Sloman. Conflicts in Policy-Based Distributed System Management. *IEEE Transactions on Software Engineering, Special Issue on Inconsistency Management*, 25(6):852–869, 1999.

[12] E. Lupu and M. Sloman. Security and Management Policy Specification. *Special Issue on Policy-Based Networking*, 16(2), 2002.

[13] D. Nikitopoulos, B. Pagurek, and T. White. Real-time WLAN Monitoring in a 4G Multiplatform Environment. Technical report.

[14] ponder2.net. Ponder2 - http://ponder2.net. 02.02.2007.

[15] P. G. s, P. Magnusson, J. Oom, and N. Storm. Real-time performance monitoring and optimization of cellular systems. *Ericsson Review*, (1), 2001.

[16] J. Strassner. *Policy-based network management : solutions for the next generation*. Morgan Kaufmann, 2004. ISBN 1-55860-859-1.

[17] D. Verma. Simplifying Network Administration Using Policy-Based Management. *IEEE Network*, 16(2):20–26, 2002.