Increasing Resilience of ATM Networks using Traffic Monitoring and Automated Anomaly Analysis

Kyle J. S. White, Dimitrios P. Pezaros, Christopher W. Johnson
School of Computing Science
University of Glasgow
Glasgow, G12 8QQ , Scotland
{mail@kylewhite.com, dimitrios.pezaros@glasgow.ac.uk, christopher.johnson@glasgow.ac.uk}

ABSTRACT

Systematic network monitoring can be the cornerstone for the dependable operation of safety-critical distributed systems. In this paper, we present our vision for informed anomaly detection through network monitoring and resilience measurements to increase the operators' visibility of ATM communication networks. We raise the question of how to determine the optimal level of automation in this safety-critical context, and we present a novel passive network monitoring system that can reveal network utilisation trends and traffic patterns in diverse timescales. Using network measurements, we derive resilience metrics and visualisations to enhance the operators' knowledge of the network and traffic behaviour, and allow for network planning and provisioning based on informed what-if analysis.

Categories and Subject Descriptors

C2.3 [**Computer-Communication Networks**]: Network Operations - *network management, network monitoring.*

General Terms

Management, Measurement, Reliability, Security.

Keywords

ATM, Network Monitoring, Automation, Safety

INTRODUCTION

Air Traffic Management (ATM) systems are undergoing significant transformations in both the EU and the US through the Single European Sky ATM Research (SESAR) and the Next Generation Air Transportation System (NextGen), respectively. The transformations will see services coordinated over larger geographical areas and include integration of services such as radar and flight schedules. These new systems are required to cope with the predicted increases in air traffic; more than double from 1999 by 2015 in the EU [1]. As air traffic increases, so will information, as network traffic, on the ATM communications infrastructures making these computer networks more heavily used, meaning any problems or anomalies will have a greater potential impact on services.

To mitigate against these potential disruptions of the ATM service, we propose methods to give IT network operators and engineers higher visibility and greater understanding of the underlying communications infrastructure through effective network monitoring and visualisation software. To complement this real-time perspective we also introduce pre-emptive resilience analysis techniques, based on network measurement, to identify weak aspects of the infrastructure. These techniques introduce simulated network anomalies, so that operators may improve the computing infrastructure to reduce the potential impact of anomalies in an operational context. We believe that such resilience quantification will allow operators to improve reliability and provisioning by identifying weak points. Engineering solutions can then be found to strengthen the network as a whole, furthering our aim to make ATM communication networks as resilient and well-monitored as possible and therefore more resilient than other communications systems. The holistic approach to network resilience [26] ensures a greater overall reliability since individual changes may strengthen one aspect of the system whilst weakening it in others. Our research also fits into the broader idea of resilience for ATM services as a whole, from human factors to technical systems. The initial focus in Hollnagel's concept of resilience engineering include methods “to analyse, measure and monitor the resilience of organisations” [12]. We will introduce such methods for the IT networking systems of the overall ATM service which could be part of a complete resilience evaluation in future.

We present KSWatch, a new passive network monitoring system with unique functionality to playback archived monitored data at any temporal granularity, with the explicit aim to increase operators' visibility and therefore holistic understanding of their network. We also raise the questions of how much automated network-level intervention is a good idea in a safety-critical environment and how effective automated anomaly detection can be achieved through the application of existing generalised, Internet-wide techniques into the specific, less-complex networking environment of ATM systems. Our vision of automated anomaly detection is presented later, where we discuss that an optimal first step could be the real-time

ATACCS’2012, 29-31 May 2012, London, UK.

notification of anomalous trends for operators to consider alongside a policy-based prioritised list of possible causes.

Having visited a major European Area Control Centre (ACC) we believe the aims of this research are the next incremental improvements from the current, initial monitoring solutions and their basic role in detection and diagnosis of issues by network engineers. Muller et al. [21] examined ATM network safety aspects and discusses three hazards: packet delay, loss and corruption also stating there are few network-level safety publications.

The remainder of this paper will review related case studies of failure, discuss appropriate approaches to monitoring in ATM and present our monitoring system. Later sections introduce resilience metrics and question the role of automation, before concluding the paper.

ATM NETWORK FAILURES

There have been previous network failure incidents in ATM and in interconnected systems which provide strong motivation for this research. These incidents have caused major disruptions and their slow detection and diagnosis have aggregated their impact in terms of operational service, financial loss and human factors aspects such as confidence in the technology. While publicised high-impact ATM networking incidents have been infrequent to date, since this is a safety-critical domain it is important to have methods in place which can reduce the impact of problems when they inevitably occur. Our research provides ways to more effectively mitigate issues, through better ways to visualise, review and prepare for anomalies.

Case Studies

There are many examples of network systems failure and technical problems causing disruption in ATM services. Two well-known aviation incidents were directly caused by hardware malfunction of Network Interface Cards (NIC): a major European Airport ATC was shut-down in 2008 and disruption at Los Angeles International Airport (LAX) in 2007. Thankfully, no accidents resulted from these incidents. However, that does not reduce the seriousness of the potential safety implications caused by these IT network level problems. Further motivation is evident from these failures; their slow detection and diagnosis caused enormous disruption and therefore financial loss from loss of revenue to the cost of engineers, in one case prolonged over seven weeks.

The 2002 mid-air collision above Überlingen [9] involved technical changes which brought on system degradation with unmonitored consequences which, while not directly responsible for the accident [15], created a more dangerous environment. Therefore, the technical factors are considered an initial point of failure in the overall incident. While this incident is not directly related to network resilience, the case highlights the importance of monitoring technical systems within ATM. With greater understanding of systems resilience and coupling amongst technical services in place in future, there is a higher likelihood of predicting and more significantly mitigating the impact of technical failures.

LAX Network Failure

The network problems in LAX lasted approximately 10 hours and caused the US Customs and Border Protection computer systems to have average response times of 2-3 minutes, two orders of magnitude greater than the usual durations of less than 5 seconds [8]. While this system is not within ATM, due to the delays and consequent congestion of people in waiting in terminals, flights were disrupted. The case highlights how issues which are reasonably easy to detect can cause enormous problems when the right methods are not deployed. Up to 17,000 passengers were estimated to have been affected including new arrivals not being allowed to disembark and disruptions to international departures. While the official report found the cause inconclusive, later analysis of the outage confirmed that the initial point of failure, which in turn caused data to overload the system, was caused by a malfunctioning NIC [19]. This case highlights the enormous impact network problems can have, the difficulty in identifying the problem and the length of time taken by manual recovery methods. The recommendations made by The Department of Homeland Security included the introduction of “automatic error notifications” and that staff should be more effective when isolating and resolving outages [8]. Our research will tackle both these aims and go further. The network monitoring system we present gives network operators the opportunity to increase their understanding of the network, perceiving trends and behavioural patterns by replaying data. These trends can be used to simulate failure scenarios, by understanding typical behaviours and then simulating the introduction of an anomaly through resilience measurement techniques which will allow policies and predictions to be made about network vulnerabilities. These contributions could lead towards the automation of detecting vulnerabilities and problems. Another recommendation, which acts as yet more motivation for this research, was the belief by staff that there is “a high risk that a similar outage could occur at other [US] ports of entry”. This shows negligence in understanding the importance of the monitoring and resilience of networked systems, despite their criticality to all ATM services, even within an isolated but interconnected system.

European Airport ATC Network Failure

In 2008, a European airport ATC system experienced a far more disruptive and prolonged networking failure. While this incident had a very similar cause to the LAX network failure, the effects had a greater impact upon ATM services. The root cause of an intermittent faulty NIC and the subsequent error in systems attempting to mitigate the issue led to anomalous manifestations including ATM losing track of planes or associated flight information across several different periods. The problems persisted for seven weeks with sustained durations of normal behaviour. These ongoing problems caused delays, restrictions in air traffic such as the avoidance of aircraft in arrival holding patterns in case a problem occurred, irregularities in information presented on Air Traffic Control Officers' (ATCO) screens, and some periods of

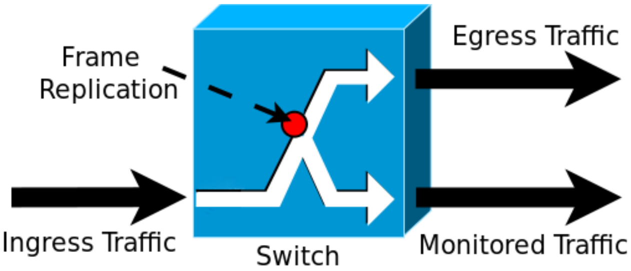

Figure 1: Mirrored port in switch for passive monitoring

complete closure [13]. The recommendations by Thales, the ATM system supplier, and the national aviation authority were “that additional network monitoring be undertaken and that monitoring tools and a passive analyser be installed on the system to aid the early identification of similar malfunctions” and “the possibility of other potential improvements in the network design in order to prevent a re-occurrence” [13]. Adding network monitoring was central to the detection of the faulty NIC. This is similar to the LAX fault recommendation [8] that more network diagnostic tools should be provided.

In each of these situations, Reason's widely cited Swiss Cheese failure model [24] applies since these were not single points of failure, but a series of problems which involved a root technical fault. We argue that in these cases, while network-level problems are unavoidable, aspects of network automation such as monitoring with automated analysis and notification of anomalies, could stop such problems persisting and remaining unknown, and instead be detected very quickly.

These incidents are rare. However, this research provides a method to limit the impact of such faults and many other potential disruptions by increasing the operators' visibility of their network and highlighting weak points through resilience measurement. Networks can be affected by factors beyond hardware malfunctions. Circumstances such as connection failure, natural disasters rendering parts of the infrastructure unavailable and malicious cyber-attacks are all potential threats to safety-critical ATM services. Automated anomaly analysis detection would be applicable for each of these anomalous classifications. With automated anomaly detection the process of detection and remediation could be achievable in real-time. Our first steps in achieving self-management through monitoring and resilience measurements follow.

NETWORK MONITORING

As discussed earlier, successful monitoring of a network provides the operators with many advantages when making management decisions, from shaping traffic policies to the architectural design. Like many enterprise networks, ATM communication networks could be improved given adequate monitoring and an understanding of what these measurements mean. Measurements can provide an understanding of patterns and trends in behaviour. By using knowledge of the traffic behaviour less immediately apparent aspects such as stronger resilience or reliability can be improved, for example by taking the measurements of a busy period and simulating a device or link becoming disconnected. Later we will show resilience measurement examples which can better quantify this abstract metric.

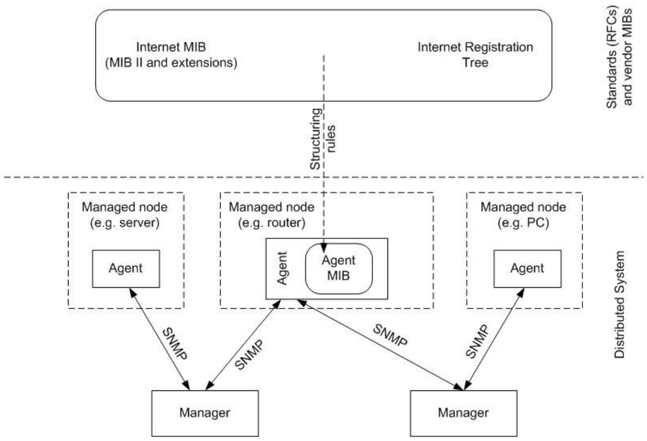

Figure 2: SNMP Internet Architecture Model

Many network monitoring systems provide a basic subset of measurements. However, by using the information we collect more effectively and at an aggregated level we have created a novel automated network monitoring system. Given the unique environment of ATM networks, we believe this bespoke system, introduced in a later section, will be the first step towards a real-time automated anomaly detection system by providing a better view of network behaviour.

Passive Monitoring

Network monitoring is a broad topic and covers many techniques. Measuring a network allows operators to gain an understanding of its behaviour and from this perform various tasks from problem diagnosis and anomaly detection to traffic shaping, provisioning and optimisation.

Measurement can be performed passively and actively. Active measurements are performed by adding traffic to the network. The measurements are then made based on the behaviour of the network with this stimulus traffic. Due to the behavioural interference of additional traffic with the operational traffic, its wide-spread use is unsuitable for ATM systems. Contrastingly, passive monitoring is unobtrusive and measurements are taken by observing operational traffic. Different aspects of the network can be measured such as the traffic, performance and structure. Examining each of these gives more measurement options.

Metrics can be derived passively from traffic at various levels, perhaps by using switches with mirrored ports as seen in Figure 1. Mirrored ports or network taps are beneficial for a safety-critical environment as there can be no interference with operational traffic because these methods observe the passing traffic without interaction. The most standardised passive network monitoring technique is Simple Network Management Protocol (SNMP), first defined in 1990 (RFC1157). The majority of devices have this protocol built in and therefore no additional hardware requirements are necessary, unlike mirrored ports or network taps. SNMP-based monitoring is at the highest level where traffic is counted at an aggregated interface level on agents, such as switches or routers, and is then polled by management systems, see Figure 2. The protocol examines Management Information Bases (MIBs) which structure Object IDentifiers (OIDs) in a hierarchy. These OIDs hold the aggregated information about the interfaces on the device and include counters for the number of: packets, bytes and packets with Unicast, Multicast or Broadcast type sent and received, and the number of discarded packets. These measurements show the values between two connected devices. Figure 2 shows managers which poll the MIBs on managed devices using SNMP. Asynchronous trap notifications may also pass from agents to managers e.g. the *sysUpTime* value.

SNMPv3 (RFC3411) is the latest version and is a full Internet standard with additional security over previous versions of the protocol.

Traffic with different sources may pass through common links and devices en route to their possibly distinct destinations. Traffic can be grouped into flows based upon common characteristics, typically based upon its Internet Protocol (IP) source and destination addresses and its transport protocol type. The transport protocol is commonly either the Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) depending upon the traffic requirements of the applications which are sending the communications. These protocols have different properties most noticeably with respect to the time required to establish a connection between two applications and the procedure for dealing with lost packets, e.g. whether or not to retransmit the information. Flow measurements can derive information such as the start and stop times of specific flows, their size and their route. These measurements give more details of the traffic behaviour and can be considered more traffic-centric than SNMP measurements which are device-centric.

The most detailed level of traffic specific monitoring is done at the level of individual packets. This is the finest granularity available and gives metrics on individual packets such as its timing and other information available in the packet's header including number of bytes, protocol, whether the packet is part of a series, its unique number and its type. For high-speed links there is often a requirement for additional hardware support in order to passively measure a large volume of information.

Advanced Technical Methods

Passive measurement relies on observation points. There is much research work involving the combination of network monitoring with network tomography; the process of taking a set of monitoring points in a network and then attempting to infer from these locations what the network behaviour is at other locations in the network. Example applications of tomography include: inferring routing topologies from end-to-end measurements [14], for finding Distributed Denial of Service (DDoS) attack sources [7] or earlier work in estimating TCP packet loss ratios of links within a network [3]. Network tomography and methods which use this technique could be useful for ATM since the number of monitoring points can be reduced whilst retaining the ability to detect issues. If packet loss ratios are observed and improved this may translate to high-level ATM efficiencies. For example, if packets are lost in air-ground-air communication streams, this requires repetition of the communication. Aggregating the impact of improving packet loss rates across every communication would increase the reliability and efficiency noticeably.

Other relevant work has considered the use of highly efficient monitoring at the end-points of a network in order to derive performance [22] through new techniques. The technique used in this example is to piggyback measurement data with operational traffic, in individual packets, and it does not introduce additional packets. It is therefore not active measurement and is more closely aligned with the ideology of passive measurements' unobtrusive approach to measuring operational traffic.

At the level of flow-based measurements there is also considerable work in developing techniques to solve the key high-level networking issues. Flow-based metrics can be used in Intrusion Detection (ID) for classifying network attacks [11] or in anomaly detection and classification [2]. The flow-based anomaly detection work is centred on the observation that a larger fraction of anomalous traffic is concentrated in a small fraction of flows. This concept is similar to that of tomography where there is inference taken from a subset of monitored values, in this case a subset of flows as opposed to a subset of monitored device interfaces, such as the ingress or egress of a network or its end-points. At the most granular level, deep packet inspection is often used for firewalls and security type applications such as intrusion detection. Recent work [25], has evaluated increased efficiency in this computationally intensive process, particularly for high speed networking, in the context of intrusion detection and security for cloud computing. Each of these techniques could bring advantages for ATM whether through an increase in security through ID, anomaly detection using flow metrics or deriving performance efficiently throughout the network using traffic piggyback methods.

Using Measurements for Network Operations

Measurements allow network operators to gain a better understanding of their network. This means they can have planned rational responses to problems and be aware of the impact of their decisions as opposed to rushed intervention and unknown causes and effects when trouble shooting and resolving issues. All of the monitoring metrics we discussed can be used to significantly enhance anomaly detection, network optimisation and to improve performance, e.g. through monitoring flows it could be found that bottlenecks exist and therefore the architecture could be modified to improve performance or, through deep individual packet inspection, anomalous malicious traffic could be seen and rejected, similar to the actions performed by firewalls. This is true for any network, but ATM communication networks have requirements and advantages in comparison to typical enterprise networks or the Internet generally. ATM traffic is more contained than

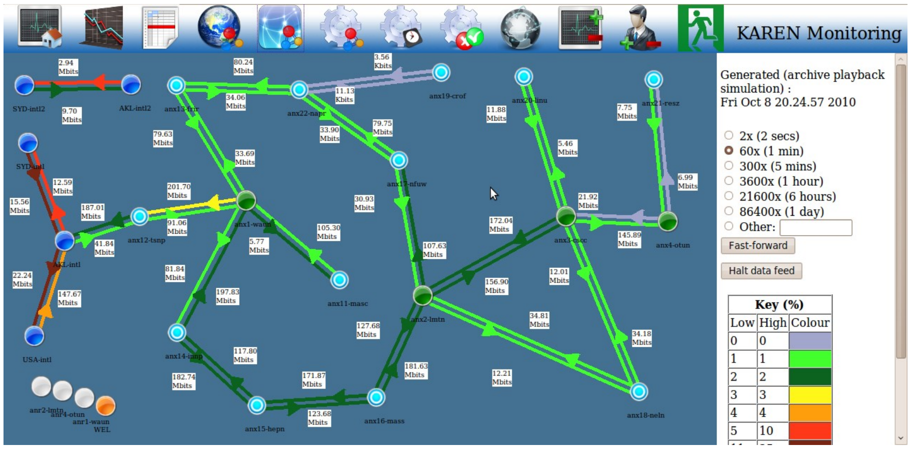

Generated (archive playback simulation) :
Fri Oct 8 20.24.57 2010

○ 2x (2 secs)
● 60x (1 min)
○ 300x (5 mins)
○ 3600x (1 hour)
○ 21600x (6 hours)
○ 86400x (1 day)
○ Other:

Key (%)		
Low	High	Colour
0	0	🟪
1	1	🟩
2	2	🟩
3	3	🟨
4	4	🟧
5	10	🟥

Figure 3: User interface for KSWatch monitoring system

Internet-wide traffic, due to fewer traffic types, lower utilisation, a simpler architecture and a single purpose and set of applications. This reduction in complexity presents opportunities, due to more apparent norms in traffic, in automation which we will discuss later, however from a monitoring perspective this isolated environment allows for highly visible and well understood network behaviour.

The motivation for such an insight and understanding would be strong for any network for generic reasons such as provisioning, optimisation and anomaly detection. However, these benefits are enhanced given the safety-critical domain and therefore the desire for reliability, transparency and understanding in all of its aspects.

We propose automated monitoring, rather than automated resolution of issues. Automatic resolution adds complexity and can lead to unknown problems in the system escalating into run away results which compound each other and make the situation worse. We believe that automated monitoring in combination with automated anomaly detection in real-time with notifications to the operators would be an optimal next step for ensuring the underlying infrastructure of ATM services is well safe-guarded against potential problems.

As we stated earlier, from our discussion with a major European ACC we have seen that some monitoring is present in ATM systems. However, while these systems provide some of the basic measurements we reviewed in this section, the more advanced aspects as well as the visualisation and presentation of this information could be improved in order to better aid the operators and engineers in their network-level tasks.

KSWatch Monitoring Sytem

In this section, we present the design and implementation of KSWatch: an extensible network monitoring and visualisation system that can seamlessly integrate diverse functionality, and can therefore form the basis for a comprehensive network monitoring and management framework. The system eliminates much of the data redundancy found in a typical set-up; it provides for an increased storage and visualisation efficiency, and also offers features like the archive playback of network utilisation. Using this system's unique functionality improves the network operators' understanding of their network, as we have found from our discussions with those who currently use it. KSWatch also gives the operator a clear picture of the network topology, the most highly-used partitions in the network and any bottlenecks, and therefore, vulnerabilities in the infrastructure. These aspects and their associated metrics underpin the resilience techniques we will discuss later.

We are aware of the use of the Zabbix[1] monitoring system in some ATM systems. Anecdotal evidence suggests that engineers use the system to confirm, some time later, the existence of a problem they have already found. This suggests that a far more useful tool for engineers would be the automation of detecting and notifying problems in real-time as opposed to relaying information. However, relayed information is still significantly useful when presented correctly with advanced functionality as we do.

Implementation

While our system[2] uses the common underlying passive monitoring technique of SNMP to gather network data, the functionality it offers is far more useful than that of existing systems. Functionality includes: graphing and tabular statistics of network devices and connections e.g. CPU utilisation or number of lost traffic packets. This system also offers the ability to switch monitoring on and off with a single click, which is beneficial if system resources are needed elsewhere e.g. when performing system tests. The main benefit of this system, is the network *weathermap* tool. The user interface is seen in Figure 3 with options for animated playback displayed on the right side. This tool shows the topology of the network with colour graded connections between devices which equate to the level of link utilisation. This overview of the network can be viewed by the operator using live or archived data for an automated visualisation of the time period specified. This means operators can review the build up to a particular event and analyse busy periods offline. The speed and period of the playback are fully customisable allowing the operator to view, for example, hourly snapshots of the past three months or every minute

[1] Zabbix Monitoring System, http://www.zabbix.com

[2] KSWatch System, http://dcs.gla.ac.uk/~kyle/kswatch.pdf

in the last hour. This information is invaluable when diagnosing problems, considering network provisioning and policy-making. Our specific interest of categorising normal behaviour and vulnerabilities for automated self-management and resilience metrics are also only possible with such highly monitored perspectives.

Application

Using our earlier case study of the European ATC network failure as motivation, here we will explain how the deployment of our system can significantly improve the network staff's understanding of their network. The key aspects which made the NIC failure problem difficult to detect were the intermittency and the lack of sufficient network monitoring.

With our system in place prior to the onset of a problem, such as traffic flooding, the normal behaviour prior to the incident and build up to problematic conditions would be captured and measured. Using our specialised functionality, the data for this period could be replayed and analysed. This replay could be compared with automated visualisations of similar periods such as the weeks prior at the same time or other durations with similar traffic levels. This would give engineers a significantly more informed position to begin their diagnosis than having to work with measurements taken only after the onset of a problem. A self-managed system may be able to go further still and at the point of a problem occurring, be able to detect the exact source of the anomalous traffic. Despite the benefits of hindsight, without detailed monitoring information available from the time of the European ATC failure, it is impossible to state if the deployment of this system would have improved the situation. However, the official recommendations strongly suggests that this research would assist problem detection and diagnosis in this environment.

Our system is currently running on the Kiwi Advanced Research and Education Network (KAREN)[3]. This deployment has shown this system is fully functional and has shown the operators ways to improve the efficiency of their traffic policies. These improvements have since been made and monitored. The system continues to run and the functionality is being regularly extended due to the flexible architecture and data management scheme.

KSWatch provides operators with perspectives which highlight trends and patterns in network behaviour over time. By examining these trends, the network and the services running on the network are better understood. When considering resilience, these trends are the starting point from which problematic scenarios can be simulated. Standard operational data will be present when failures occur and therefore it is important to know the impact a range of failures could have on operational traffic before they occur so that resilience can be pre-emptively improved with careful and adequate consideration.

RESILIENCE

With monitoring in place such as that provided by KSWatch, it is possible for network engineers to recreate events, simulate potential threats and model the resilience of the infrastructure. Network resilience is central to the primary aim of the ATM service: to maintain and provide reliable services resisting challenges which may occur.

Since the Communications Network Services (CNS) and network data are a fundamental part of the ATM service operation, it is vital that there is a greater understanding of their resilience. This in turn impacts upon the resilience of the ATM service overall. There are many considerations with resilience, discussed in depth by Fry [10] and Cholda [6], which include the idea that resilience requirements for different services may differ despite the fact that these services run on the same networks. Using flow-based monitoring, services can be distinguished and the individual resilience of an ATM service calculated.

Network resilience can be quantified over many factors including: link utilisation, packet loss and connectivity. These example factors link with the SNMP monitoring OIDs of: bytes sent and received and discarded packets. This means operational data can be used in resilience simulation to provide comparable measurements of the impact of problems. ATM networks have a high level of redundancy in their infrastructures and often have complete back up network architectures. This makes connectivity metrics more suitable for future research. While there may be bottlenecks or weak points at the connections to the redundant architectures it is likely these will be specific to each ATM operation. However, there are common parameters which will affect the network resilience of any ATM system regardless of the topology. Since ATM is a real-time service, any degradation in the network which affects an application's timely delivery of information is unacceptable. Therefore, we have considered which resilience metrics would be most valuable to model initially and to understand their impact. These are the amount of disruption in the network, (represented by the available capacity on each link) and the amount of packet loss.

We have run simulations based on an ATM communications network experiencing a NIC network card failure similar to that encountered by the European ATC centre, by extracting technical details from the analytical reports. These simulations will be used as a demonstration of how to produce example resilience envelopes which would typically be performed prior to an incident in order to understand its potential impact.

Resilience (Metric) Envelopes

Resilience envelopes were developed by the ResiliNets[4] networking project at The University of Kansas and Lancaster University. These graphs show the upper and lower bounds of a specific resilience metric plotted against predetermined failures in the network. It is therefore important when considering resilience issues to contemplate all potential threats as well as the metrics

[3] KAREN, http://www.karen.net.nz

[4] ResiliNets Wiki, https://wiki.ittc.ku.edu/resilinets

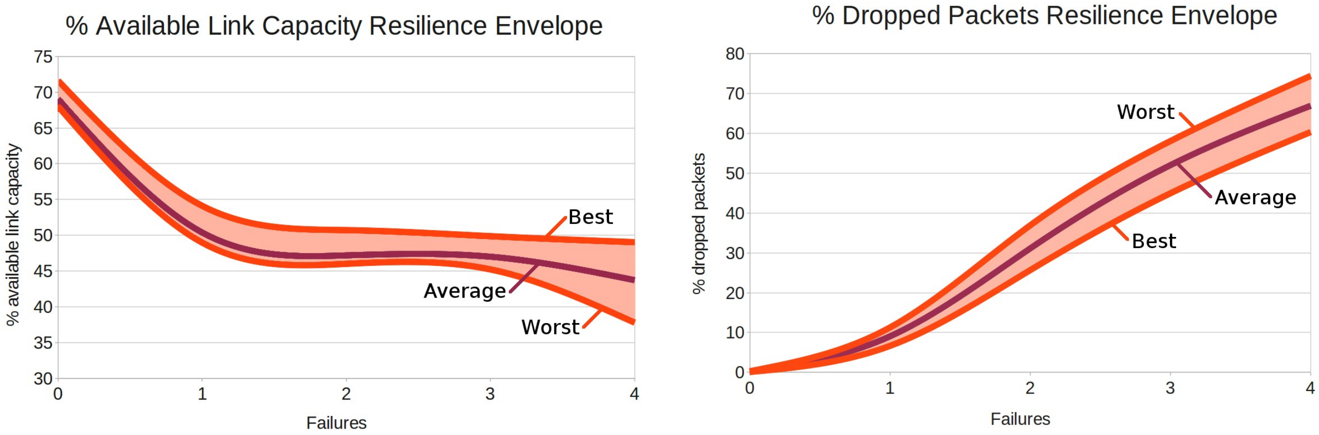

Figure 4: Resilience metrics for utilisation by NIC failure

which could affect reliable performance in the network. The information visualised can then be used to make improvements to the network in anticipation of potential issues in order to mitigate their impact. The effects of these improvements can then be measured by re-modelling the resilience envelopes and observing whether the envelope range has improved. Changes for one resilience aspect may also affect others, perhaps through modifying the topology. We therefore propose that an implementation of this research would involve an automated *sandbox* test-bed environment where simulated changes could be rerun automatically against all resilience parameters with notifications of any other measurements which were degraded as a result of the change.

The resilience envelope concept is a component of the larger ResiliNets strategy of D^2R^2+DR (defend, detect, remediate, recover, diagnose, refine) which is discussed more comprehensively along with principles and approaches in these works [26, 27].

The core idea is to have real-time network resilience through successful iterations of each of these high-level concepts. Our research aims to progress the defence and detection elements of this strategy in ATM systems.

Experimental Setup

We have used the popular network research software ns-3[5] to simulate an ATM communications network, with an accurate example topology, and traffic flow between applications. We have chosen to base our simulation on the case study regarding network card failure since it is widely known and is the best recorded example of the impact of such a problem. This technique can be used for any set of anomalous events. In our experiment we tracked the remaining free capacity on links and the number of packets dropped, averaged across the network as a whole. This reverts back to the concept of considering resilience of the network in a holistic way. However, an implemented test suite of these measurements could see simulations and resilience envelopes produced on a per-link or per-flow basis depending on the visibility desired.

These chosen metrics were mapped against the introduction of failures. We simulated the hardware failure of a NIC flooding event by introducing a UDP application which produces packets without considering the establishment of connections or retransmission. To show the escalation of resilience, we increase the number of failures over various simulations or in the simulated context, an increase in the volume of anomalous traffic. All the experiments were repeated and the results averaged across numerous simulations. There is variance introduced in the results from the random start and stop time of normal traffic applications. This means in some simulations, more normal traffic will be present in the network when the failure occurs causing a greater impact. Similarly, the best case may be when the failures occurred at the period of quietest normal activity due to the random timing of events across the simulations.

Simulated ATM Resilience Envelopes

In Figure 4 we see the resilience envelopes produced from our simulations. Assumptions regarding network traffic utilisation were made when creating these measurements since no archived monitoring traffic was available from the time of the incident. However, the benefits of using this measurement concept to pre-emptively validate and gain reassurances in the resilience of an ATM network are apparent from these examples.

It can be seen that there is a clear trend between the number of failures and the network disruption for both metrics. A consideration when examining the above envelopes, particularly the dropped packets example, is that for Voice over Internet Protocol (VoIP) applications, dropped packets in the stream of between 5 - 10 % affects the quality significantly [20]. With a single failure on our basic simulation this range is reached even in the best case scenario. Clearly such a failure would cause enormous disruption in this real-time safety-critical environment. For other networked communications, retransmission can be achieved without as much disruption. This is because when there are delays or portions missing in audio or video, humans will find it difficult to interpret the meaning as it is heard or viewed. However, in an ATM context any delays through packet loss and their consequent retransmission will be apparent since the

[5] NS-3: Network Simulator, http://www.nsnam.org

service operates in real-time. Data which is neither audio nor video can tolerate more delay. Given a disruptive scenario, these will still suffer significant delays. The failures we model here are infrequent as are other potential anomalies. However, the entire ATM service operation is totally reliant on the power systems and network level operation to deliver information and so additional safe-guards at this granularity are of significant benefit. The back up systems in place for power supply are automated and highly sophisticated, to deliver continuous power from a point of failure and whilst transitioning among different power sources. The network infrastructures also have high redundancy with often numerous physical connections to avoid loss of services. However, currently the process of detecting anomalous network behaviour is less sophisticated than what we present here. We believe the application of this research would increase the overall resilience of ATM services. The left envelope in Figure 4 shows the relationship between the available link capacity in the network and the number of failures, or quantity of anomalous traffic in our simulations.

This means that a very severe single flooding incident could have the impact shown as though it were four less intensive failures occurring simultaneously. Assuming approximately 30% utilisation across the links in the network under normal circumstances, the introduction of a failure grossly reduces the available link capacity from approximately 70% down to on average nearly 50%. This reduction is significant since this pushes the network towards a state where additional resources would be required in order to ensure reliability of services with an increase in traffic. This is because high utilisation does not allow new services to initiate with a guarantee of reliable performance. Existing services could also require additional available capacity, for example if a nearby airport was closed due to weather and a significant level of traffic was redirected to the affected ATM service. There is therefore a need to know the available capacity at any given time which can be provided by KSWatch monitoring. These examples are uncommon yet possible and with an ever-increasing demand for air travel and expansion of air traffic predicted, necessary additional resilience can be found by ensuring the underlying fundamental CNS are highly visible and well understood. As a basis for future work, we have abstractly considered the classification of ATM CNS specific threats, similar to the abstract taxonomy by Plonka [23].

ATM Potential Network Threats

We have considered a subset of all potential network challenges which are applicable on a typical ATM communications architecture. As we mentioned in our motivation, there is a diverse range of potential threats from low-level issues such as hardware faults to high-level events e.g. natural disasters. These threats also span a wide range at a lower technical level. While some will impact in a very noticeable manner such as a link failing, others are likely to be less apparent such as a hardware malfunction. The traffic types will vary as well from potentially corrupt information from a single source through a software fault to enormous crippling volumes from different sources via a DDoS.

Future work will consider the taxonomy of these more closely and with a greater emphasis on their impact in an ATM context at a technical level, along with potential detection techniques from the observation of their identifying characteristics. Considerable related work exists on this topic but often they examine a minimal number of technical parameters when creating the classifications. Such work is also not specialised to ATM systems and therefore is unable to exploit the benefits we discussed earlier.

With such a variety of potential anomalies the concept of automated detection is difficult. However, there is a strong desire for this level of self-management in networks since problems can be noticed more rapidly. In the next section we will raise the open question of what would be the optimal level of self-management in ATM communication networks and consider how to approach any level of automation in this context.

CONSIDERING AUTOMATION

Adding automation to a system adds to its complexity. This is not desirable for interconnected, safety-critical systems since the number of potential points of failure increases. Therefore complete network self-management, where the network automatically looks after its own configuration, optimisation, protection and healing is not yet viable. As we have suggested however, some elements of automation could prove invaluable such as automated anomaly detection. The level of self-management in this case would be the detection and notification of what is considered anomalous traffic for operators and engineers to consider. Of course false-positives and false-negatives may arise and users of any implementation would have to be aware of such possibilities and not be fully reliant on such a system. It would be optimal to have such a system aid the Network Operations Centre (NOC) team in their activities and possibly provide an early warning for any anomalies. With any system automation, high visibility and a well understood operation are vital before trusting its reliability.

Given desirable automation in the detection and protection aspects of ATM systems, the next question is how to achieve such a vision. Considerable networked systems research has investigated anomaly detection on Next Generation Networks (NGN) and the Internet. Lim [18] provides a good overview, with a security perspective, of the taxonomy of different detection techniques which shows the variance in existing approaches but highlights that research in this area is still ongoing. Flow sampling is an emerging technique for anomaly detection and increasing efficiency is being derived [2], albeit in this case, on specific, well-defined anomalies. Improved visualisations of network behaviour for engineers to gain better perspectives, such as those in our monitoring system, are also a widely accepted concept. Algorithms have been proposed to better visualise traffic data for the purpose of understanding anomalies by Celenk [5]. This is an anomaly estimation technique, but it highlights existing work in this area which attempts to reduce the obscurity of the network, particularly in terms of problem diagnosis. A different approach [16] is the use of histogram-based anomaly detection whereby normal flows are classified and modelled according to their features. Appropriately, this work highlights the LAX incident as motivation. However, the common problem of defining normal traffic is present with this solution since it is based on anomaly-free training data, which cannot be guaranteed and it is aided by administrators labelling known anomalous traffic.

Characteristics of ATM Networks and the Internet

These ideas are relevant since the statistical characteristics of Internet traffic are not so disperse from those of ATM CNS. As a starting point, the hardware involved is standard in terms of switches, routers and devices. The physical cabling infrastructure is also common with either Ethernet or fibre optic links. Topologies differ amongst networks, however the components and general architectures and their behavioural characteristics are known such as trees, rings or interconnected meshes. Ethernet forwarding and Media Access Control (MAC) addressing are the same underlying low-level method.

The protocols are also shared from the standard SNMP management agents to the transport protocols such as TCP or UDP and the ubiquitous Internet Protocol (IP). Therefore, the only significant difference which remains between the types of networked system is the traffic behaviour, which is defined by the applications which are running. By examining these more closely we see that ATM services, whilst sharing common characteristics with Internet-wide networks, are far less complex. Fewer services, with distinct roles, such as network traffic being relayed from radar sites, will be more predictable than other enterprise traffic which may include competing, simultaneous applications such as video streaming, peer-to-peer services, Virtual Private Network (VPN) traffic, viruses and active measurements. This environment is far less predictable than ATM systems which could even have pre-determinable elements with correct modelling. While ATM network traffic will experience bursts, for example in the case of additional flights arriving due to a near neighbour closing, these bursts will be of similar traffic type, are likely to be proportional to aircraft, and always have an upper limit: either the maximum number of aircraft which can be safely managed in the ATM's domain or the maximum capacity able to be scheduled to arrive and depart at a given airport. Neither of these limits are likely to exceed the utilisation of the network. The anomalies which can affect both types of networked system are also similar. Our case studies of faulty NICs are essentially instances of network flooding. While “normal” Internet traffic is more complex behaviourally, a flooding event will still show the same characteristics, perhaps less noticeably, as a network flooding event in an ATM networked system. It is therefore reasonable to consider equivalent techniques for automated anomaly detection on Internet-wide networks since they are of a statistically similar nature in terms of hardware, protocols and anomalous traffic characteristics. The differences are apparent in the applications and services on ATM systems which will make the process of automatically detecting anomalies more viable than in other networks.

It is a commonly held belief that SNMP monitoring alone is too coarse-grained for effective classification for anomaly detection on Internet-wide networks [4]. Recent work has progressed SNMP-based anomaly detection [17]. In each of these cases the results are at best a reduction in false-positives or provide only reasonable accuracy which supports the claim that SNMP is too coarse-grained for such traffic behaviour. However, this conclusion may not hold true in an ATM context with different application characteristics making it worth exploring. If future work proves inconclusive we can adapt our methodology to consider some of the notable work based around flow measurements as we previously discussed.

Motivation for automated anomaly detection was reviewed earlier. An additional benefit of this would present itself with regards to network security. While ATM systems are currently kept relatively isolated from external sources such as the Internet, wireless or portable devices, it is possible that this will change in the future with an increasing demand on flexible networking solutions and the adoption of aggregated resources or outsourced services such as data centres for data back up or externally administered Cloud email systems. By introducing effective network monitoring and gaining an in-depth understanding of network resilience through the systems we present in this paper, as well as automated anomaly detection, this would more easily allow for the future integration of such external services. This is important for different reasons since it is likely in terms of quality, cost and security that external solutions are better than bespoke alternatives. It is also a motivational concern that currently, should unauthorised network access be achieved, a great deal of undetected malicious activity could be performed. However, since ATM networks are currently isolated, this allows us to better model and understand the components involved prior to the introduction of further complexities, including more traffic types, protocols and if considering wireless access the increase in dropped or corrupt network packets.

Another possible automation could be with regards to network failure. As mentioned earlier, ATM power supplies are highly resilient and automatically transition amongst sources in the event of failure. The network infrastructures also have high redundancy with often numerous physical connections to avoid loss of services. In contrast, the current process of detecting anomalous network behaviour and opting to switch infrastructure is not automated. This is adequate for when it is known there is a fault, but not for detecting the occurrence of faults.

We suggest that this may not be sufficient for the increased volumes of network traffic which will occur with the forecasted levels of air traffic. A potential automation could be that when anomalous behaviour occurred an analysis could be performed. If this problem was determined hazardous to operations, as specified in a control policy, the traced area of the network where the problem was detected could be isolated or bypassed allowing operations to continue unaffected. This would be in line with current ATM service policies which outline delivering fully functioning services as opposed to fully functioning systems.

CONCLUSION

In this paper, we have highlighted the advantages of and motivations for effective monitoring in ATM communication networks, and have presented KSWatch monitoring which gives network operators a new perspective of their network which may provide an increased understanding of network behaviour. From passive measurements recorded by KSWatch, trends and patterns can be observed and animated over time. These measurements can be used to calculate resilience metrics by simulating potential network anomalies at various intensities. We have shown how resilience can be quantified and visualised in order for operators to pre-emptively mitigate the impact of potential threats, thereby increasing the reliability of their networks and, as a direct result, the overall resilience of the ATM service. Finally, we have offered informed views on the debate of how much network automation and self-management would be appropriate in this safety-critical environment.

REFERENCES

1. Air Traffic Action Group (ATAG), European Air Traffic Forecast 1985–2015, Geneva, Switzerland, 1999.
2. Androulidakis, G. et al., "Network anomaly detection and classification via opportunistic sampling," *Network, IEEE*, vol.23, no.1, pp.6-12, Feb. 2009
3. Benko, P.; Veres, A.; , "A passive method for estimating end-to-end TCP packet loss," *Global Telecommunications Conference IEEE* , vol.3, no., pp. 2609-2613 vol.3, 17-21 Nov. 2002
4. Callado, A.; Kamienski, C.; Szabo, G.; Gero, B.; Kelner, J.; Fernandes, S.; Sadok, D.; , "A Survey on Internet Traffic Identification," *Communications Surveys & Tutorials, IEEE*, vol.11, no.3, pp.37-52, 2009
5. Celenk, M.; Conley, T.; Willis, J.; Graham, J.; , "Predictive Network Anomaly Detection and Visualization," *Information Forensics and Security, IEEE*, vol.5, no.2, pp.288-299, June 2010
6. Cholda, P. et al., "A Survey of Resilience Differentiation Frameworks in Communication Networks," IEEE Commun. Surveys & Tutorials, vol. 9, no. 4, 2007
7. Demir, O.; Khan, B.; , "Finding DDoS attack sources: Searchlight localization algorithm for network tomography," IWCMC, pp.418-423, July 2011
8. Department of Homeland Security, Office of Inspector General , "Lessons Learned from the August 11, 2007, Network Outage at Los Angeles International Airport (Redacted)" August 2007
9. EUROCONTROL, Johnson, C., Final Report: "Review of the BFU Uberlingen Accident Report", Dec 2004
10. Fry, M.; Fischer, M.; Karaliopoulos, M.; Smith, P.; Hutchison, D.; , "Challenge identification for network resilience," Next Generation Internet, pp.1-8, June 2010
11. Hancock, D.L.; Lamont, G.B.; , "Multi agent system for network attack classification using flow-based intrusion detection," Evolutionary Computation (CEC), 2011, vol., no., pp.1535-1542, 5-8 June 2011
12. Hollnagel E.; Woods D. D., Leveson N., "Resilience Engineering: Concepts and Precepts", Ashgate Publishing, p.6, 2006
13. IAA, "Report of the Irish Aviation Authority into the ATM System Malfunction at Dublin Airport", Sep 2008
14. Jian N.; Haiyong X.; Tatikonda, S.; Yang, Y.R.;, "Network Routing Topology Inference from End-to-End Measurements" *INFOCOM*, pp.36-40, Apr 2008
15. Johnson C.W., "Linate and Uberlingen: Understanding the Role that Public Policy Plays in the Failure of Air Traffic Management Systems", Proceedings of the ENEA International Workshop on Complex Networks and Infrastructure Protection, 2006
16. Kind, A. and Stoecklin, M.P. and Dimitropoulos, X. "Histogram-based traffic anomaly detection" Network and Service Management, IEEE, June 2009
17. Lee D.C.; Byungjoo P.; Kim E.K.; Lee J. J.; "Fast traffic anomalies detection using SNMP MIB correlation analysis," *Advanced Communication Technology*, vol.01, pp.166-170, Feb. 2009
18. Lim, S.Y.; Jones, A.; , "Network Anomaly Detection System: The State of Art of Network Behaviour Analysis," *Convergence and Hybrid Information Technology*, pp.459-465, 28-30 Aug. 2008
19. Los Angeles Times, "Customs acts fast to shore up systems", Quote by Jennifer Connors, Chief Customs and Border Protection Officer, Sep. 4, 2007
20. Mansfield, K. C.; Antonakos, J. L., "Computer Networking from LANs to WANs: Hardware, Software, and Security", Cengage, pp.501, 2010
21. Muller, J.; Perruchoud, G.; Borel, H.; Bomme, P.; , "What does safety mean for networks?," *Systems Safety 2009.* IET, pp.1-6, Oct. 2009
22. Pezaros, D.P.; Hoerdt, M.; Hutchison, D.; , "Low-Overhead End-to-End Performance Measurement for Next Generation Networks," Network and Service Management, IEEE, vol.8, no.1, pp.1-14, Mar. 2011
23. Plonka, D.; Barford, P., "Network anomaly confirmation, diagnosis and remediation," *Communication, Control and Computing, Allerton* pp.128-135, Oct. 2009
24. Reason, J. (1990) "Human Error". Cambridge: University Press, Cambridge.

25. Smallwood, D.; Vance, A., "Intrusion analysis with deep packet inspection: Increasing efficiency of packet based investigations," *Cloud and Service Computing*, pp.342-347, Dec. 2011

26. Smith, P.; Hutchison, D.; Sterbenz, J.P.G.; Schöller, M.; Fessi, A.; Karaliopoulos, M.; Lac, C.; Plattner, B.; ,"Network Resilience: A Systematic Approach," *Communications Magazine, IEEE*, vol.49, no.7, pp.88-97, July 2011

27. Sterbenz, J.P.G. et al., “Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines,” Elsevier Computer Networks, Special Issue on Resilient and Survivable Networks, vol. 54, no. 8, pp. 1243–42, June 2010