

A Situation Aware Information Infrastructure (*SAI*²) Framework

Angelos K. Marnerides¹, Dimitrios P. Pezaros², Joemon Jose²,
Andreas U. Mauthe¹, and David Hutchison¹

¹ InfoLab21, School of Computing & Communications, Lancaster University, UK
{angelos.marnerides, a.mauthe, d.hutchison}@lancaster.ac.uk

² School of Computing Science, University of Glasgow, UK
{dimitrios.pezaros, joemon.jose}@glasgow.ac.uk

Abstract. Computer network infrastructures constitute the critical backbone of every socio-economic ICT system. Consequently, they are becoming increasingly mission-critical in our society since they provide always-on services for many everyday applications (e.g., Cloud Data Centres), safety-critical operations (e.g., Air Traffic Control networks), critical manufacturing services (e.g., Utility networks and Industrial Control Systems), and critical real-time services (e.g., Financial Trading Systems). The resilience and ability of such systems to remain operational in the face of threats is therefore paramount; this needs to be done by taking remedial action and intelligently reshaping their resources. At the same time, current communication architectures do not allow for such informed and adaptive provisioning. In this paper, we introduce the concepts, principles and current research activities related to a new Situation Aware Information Infrastructure (*SAI*²) framework being developed for next generation ICT environments.

Key words: Situation awareness, network resilience, security, computer networks

1 Introduction

Undoubtedly, the adequate functionality of today's society relies heavily on the efficiency and performance of services deployed by mission-critical socio-economic ICT systems that operate over networked infrastructures. With the rapid emergence of new areas such as the IoT, it is anticipated that future mission-critical as well as everyday end-user applications will essentially require intelligent and autonomic principles to be adhered to by the underlying networked infrastructures. Such properties will enable future underlying networked infrastructures to deploy the necessary intelligence to dynamically self-protect and self-manage their own operation, hence improve their resilience and resource provisioning.

However, current networked infrastructures do not provide the necessary mechanisms to systematically assess resilience natively through a generic framework, which consequently leads to monolithic solutions targeting only partially

e.g., fault-tolerance, security, or survivability [2, 3]. At the same time, cross-layer resilience schemes tend to insufficiently serve the application layer end-user QoS requirements since they architecturally fail to control and manage their various embedded protocols in a scalable manner [5]. Moreover, situation-awareness under a synergistic use of contextual and operational information has been partially applied in the context of resilience for explicit services and mobile networks but never in the context of mission-critical ICT environments that engage a number of diverse networked infrastructures and services [6].

In general, current communication architectures do not allow for such informed, adaptive and intelligent resource provisioning since there does not exist a generic resilience framework that considers the overall impact of simultaneous challenges manifest in several inter-dependent physical infrastructures. For instance, legacy strategies are bespoke and monolithic (e.g., static firewalls) since they are deployed to protect specific services over specific locations of the infrastructure, against specific and mostly known threats (e.g., signature-based intrusion detection) [6]. Furthermore, network provisioning mechanisms do not incorporate situation awareness or intelligence from the system’s evolution to profile infrastructure-specific behaviour, nor do they harness any local or global context (e.g., prior network attack at another facility) which would aid proactive response to adversarial events. Current anomaly detection practices operate solely on aggregate context-agnostic statistical data over long timescales and are isolated from network control and provisioning algorithms (e.g., routing) [6].

Given the aforementioned limitations, this paper introduces the notion of a Situation Aware Information Infrastructure (SAI^2) framework that is currently an on-going collaborative effort between two UK academic institutions (Lancaster University and the University of Glasgow), driven by pragmatic requirements and input by four industrial partners (The UK National Air Traffic Service - NATS, Solarflare Communications, Jisc, and Airbus Group). The proposed architecture aims to detect and remediate resilience challenges by enabling a deeper understanding of the dynamic evolution of mission-critical ICT systems through harnessing and correlating diverse internal and external network context. In this context of operation, the overarching goal of SAI^2 is to create an adaptive, situation-aware information infrastructure for future mission-critical networked environments. Hence, a range of processes derived out of network resilience, anomaly detection, content-awareness, network instrumentation and measurement, information retrieval, and filtering and semantic processing are merged in order to vertically integrate data, information, measurement and control mechanisms from different layers of the communications stack.

The remainder of this paper is structured as follows: Section 2 introduces the SAI^2 conceptual framework, while Section 3 and Section 4 present two schemes where SAI^2 principles are applied. Finally, Section 5 summarizes and concludes this paper.

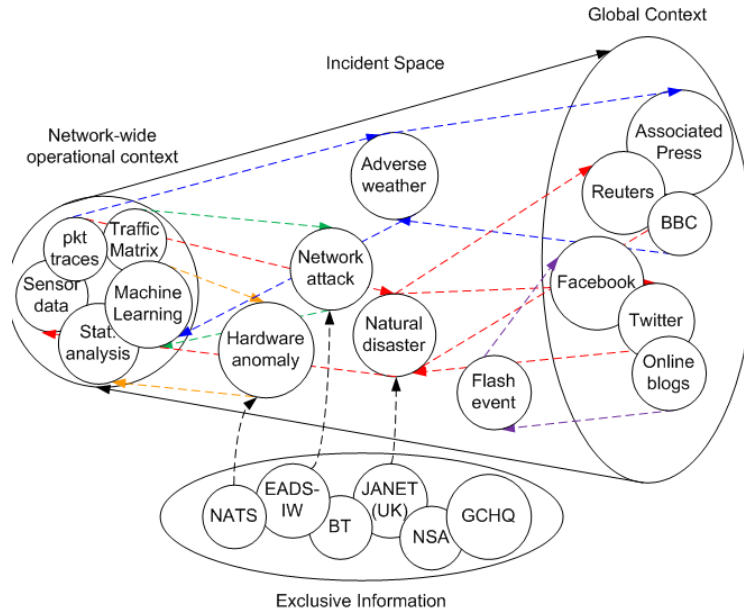


Fig. 1. The SAI^2 Framework

2 SAI^2 Framework

The main hypothesis underpinning the development of SAI^2 is that the legacy static provisioning of networked infrastructures will not be sustainable; this is because it is content and context-agnostic, and therefore cannot adapt at the onset of adversarial events. This is particularly amplified when one considers the growing trend in co-locating critical and commodity services over converged ICT (e.g., Cloud) environments and the consequent centralisation that constitutes the actual and meta-cost of infrastructure failures and outages simply not affordable.

Fig. 1 illustrates the SAI^2 conceptual framework that enables the composition of situation awareness mechanisms from harnessing diverse data sources. In general, the main scientific and technical objectives behind this framework are:

- Development of new statistical techniques derived from machine learning, signal processing and information theory to profile normal network-wide behaviour and detect incidents from aggregating diverse and distributed operational data.
- Compose modelling methods that depend on content analysis in order to adequately map infrastructure-specific context.
- Construction of an an always-on, instrumentation and measurement infrastructure.
- Development of network-wide situation-aware resilience mechanisms.

Based on the above objectives, the on-going activities in SAI^2 have identified a number of use case studies in which situation-aware mechanisms are designed

and developed (e.g., [1, 2, 3, 7]). For the purpose of this paper we restrict ourselves to briefly describing two such studies in the following sections.

3 Arbitrary Packet Matching in Openflow for enabling Situation Aware Applications

In its current form, OpenFlow [4], the *de facto* implementation of Software-Defined Networking (SDN), separates the networks control and data planes allowing a central controller to alter the match-action pipeline using a limited set of fields and actions. Even though SDN can in principle facilitate a programmable control plane capable of monitoring the network operation and of alerting the central controller at the onset of adversarial events, Openflow’s inherent rigidity prevents the rapid introduction of custom data plane functionality that would enable the design of high-speed, in-band packet processing modules for, e.g., custom routing, telemetry, debugging, and security.

We argue that packet matching should be designed independently of any (control) protocol implementation, and allow the control plane to specify the matching process through a set of platform-independent instructions designed to match packets at every layer (rather than protocol version-specific and inflexible hard-coded match fields). Through such instruction set, the execution of the matching could be left as an implementation detail relying on software optimisations (such as Just In Time (JIT) compilation) or hardware acceleration using, e.g., FPGAs or ASICs.

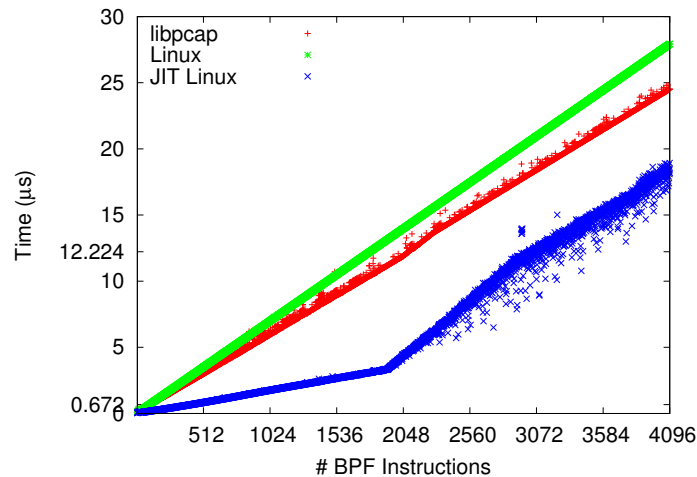


Fig. 2. Performance of different BPF engines

We have developed an arbitrary matching framework for OpenFlow switches based on the Berkeley Packet Filter (BPF) packet matching instruction set, and

have defined a new OpenFlow eXtended Match field (OXM) to match packets using BPF at the switches [1]. We show that our proposed match scheme reduces the number of flow entries, allows matching on fields and protocols not supported by current OpenFlow, and mitigates the use of the SDN controller for the classification of packets unmatchable by the switch. Using diverse prototype and compiler implementations on a software switch platform, we have demonstrated the feasibility of BPF matching for line-rate processing, as shown in Fig. 2. We have further developed example programs that can be executed as part of this BPF-based arbitrary packet matching engine that include network telemetry (e.g., real-time packet size and inter-arrival time distributions computation), and lightweight (e.g., EWMA) anomaly detection that can form the basis (infrastructure-support) for natively resilient and self-managed networked environments. Such functionality is not possible using today’s SDN/Openflow technology since packet matching rules are restricted to a protocol version-specific set, and do not allow for, e.g., arbitrary functionality, or even port range matching to be implemented.

4 Situation Awareness in the SmartGrid

A part of the activities within *SAI*² has placed a strong emphasis on composing solutions for the adequate profiling of power consumption in SmartGrids and further relating such measurements with diverse sources of information. In particular, we have developed and introduced short term power load forecasting schemes that rely on Deep Learning Neural Network models [8] and enable the correlation of power consumption with external data feeds such as weather conditions and basic human-oriented behavioral aspects (e.g. holidays, work hours etc.). Hence, it was possible to provide a deeper understanding on how, why and when the power load demand was distributed and in parallel how it will be distributed in immediate time periods (e.g. next day or next week forecast).

As already mentioned, we employ Deep Neural Network models and in particular we exploit the principles of Feed-foward Deep Neural Networks (FF-DNN) and Recurrent Deep Neural Networks (R-DNN) in order to predict short term electricity load. Achieving higher accuracy in forecasts requires to include all the related features that affect the overall electricity consumption that go beyond the raw power measurements but also consider external data sources that can also provide an insight regarding the status (i.e. current state of a given situation) of the SmartGrid. We accomplish the latter by employing the methodology depicted in Fig. 3 that considers power measurements as well as external datasets gathered by 6 US states in New England ¹.

The methodology depicted in Fig. 3 initially treats the gathered data on the time and frequency domain independently and subsequently frequency domain components are transformed back to the time domain. The resulted time-frequency (TF) features efficiently capture dominant effects i.e. weather, time,

¹ ISO New England Dataset: <http://www.iso-ne.com/isoexpress/web/reports/load-and-demand>

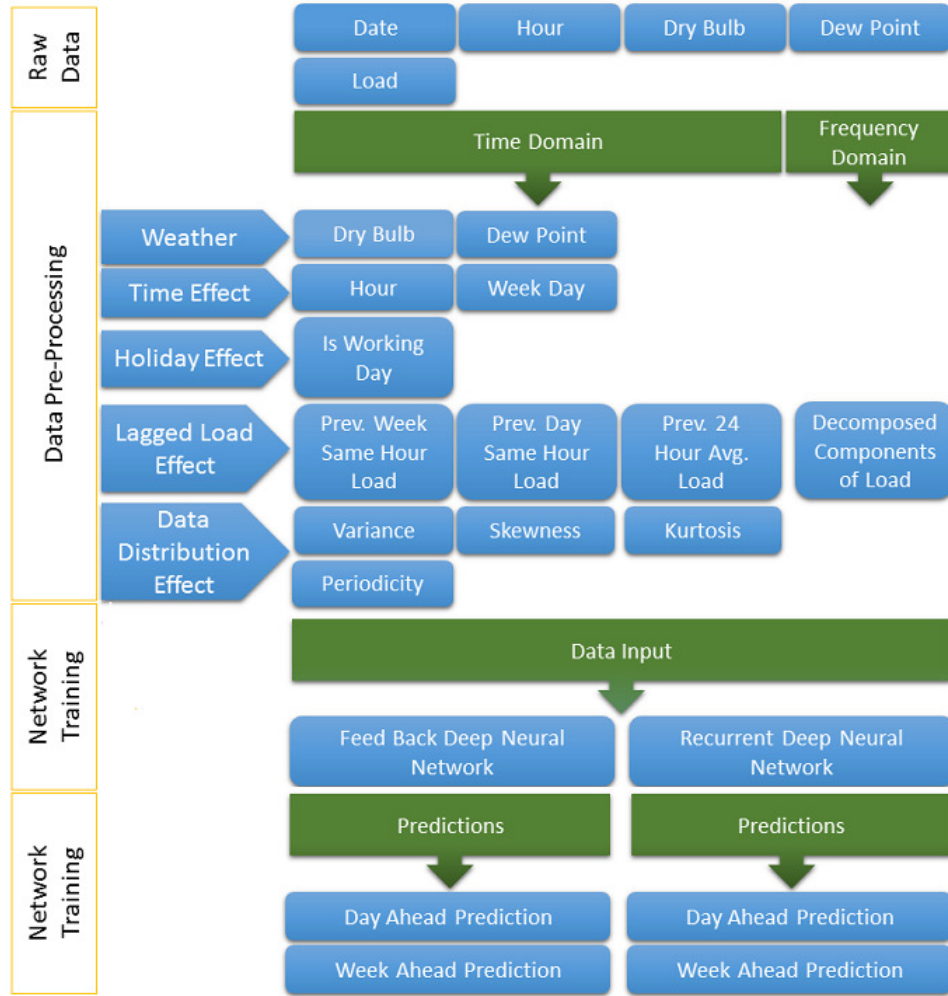


Fig. 3. Visual Description of Load Prediction Methodology

working and non-working days, lagged load and data distribution effects, thus providing an insight on the current and potential state of the SmartGrid. Nonetheless, due to the constantly changing environment, electricity consumption patterns of domestic as well as commercial users carry complex characteristics. Hence, these characteristics are analyzed in time and frequency domain where we compare the prediction performance of the utilized deep neural network models on the basis of the Root Mean Square Error (RMSE), Mean Average Error (MAE) and Mean Absolute Percentage Error (MAPE) error metrics. As we present in our recent work in [8], the results obtained with the presented methodology indicate least MAPE errors as compared to other existing models (e.g., [9]).




Fig. 4. MAPE error comparison using Time domain features.

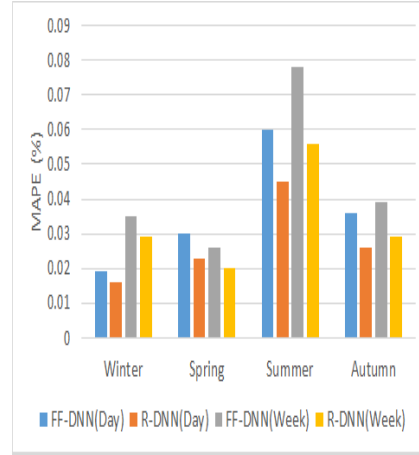


Fig. 5. MAPE error comparison using joint Time-Frequency domain features.

Thus, they can be definitely considered for future situation aware mechanisms that ideally can be instrumented in modern SmartGrid control centres.

Apart from identifying the correlations between the non-power related features with the actual power consumption distribution, this study has also demonstrated the usefulness of considering a joint time-frequency (TF) representation of data features. In fact, we have compared the prediction errors in the scenario where strictly time domain features were used with the scenario where join TF features were used.

Fig. 4 and Fig. 4 present the comparison conducted between the prediction errors obtained when strictly using time domain features with those achieved with the use of TF data metafeatures of the original raw datasets. The comparison was performed for predictions generated for the "next day" and "next week" power consumption from data sampled randomly in the various seasons (i.e. winter, autumn, spring and summer) of the year 2011 throughout all the 6 US states in New England, USA. As evidenced, the MAPE metric resulted to hold much smaller values in the scenario of using TF features rather than those obtained when time domain features were used. Thus, the prediction accuracy was much more improved.

In general, the contribution behind this work lies with the utilisation of a time-frequency (TF) feature selection procedure from the actual "raw" dataset that aids the regression procedure initiated by the use of Deep Neural Networks (DNNs). We show that the introduced scheme may adequately learn hidden patterns and accurately determine the short-term load consumption forecast by utilising a range of heterogeneous sources of input that relate not necessarily with the measurement of load itself but also with other parameters such as the effects of weather, time, holidays, lagged electricity load and its distribution over

the period. Overall, our generated outcomes reveal that the synergistic use of TF feature analysis with DNNs enables to obtain higher accuracy by capturing dominant factors that affect electricity consumption patterns and can surely contribute significantly in the context of situation awareness for the recently introduced SmartGrid.

5 Conclusions

Situation awareness is acknowledged as a critical property in which next generation, mission-critical ICT environments are required to possess in order to confront the dynamic and unpredictable behaviour of their use as well as to maintain their overall resilience. In order to achieve situation awareness that will essentially offer societal and economical benefits, it is quite important to devise mechanisms that are not strictly dependent on the explicit properties of a given system but rather consider external and diverse sources of information. Therefore, in this paper we have briefly presented the main principles and some of the on-going activities conducted within the *SAI*² project. We have introduced the *SAI*² framework and initially described an SDN-based scheme that can aid towards the dynamic deployment of applications that can serve situation awareness. In addition, we have presented a use case study that aims on strengthening situation awareness in the SmartGrid scenario by exploiting the regression capabilities of Deep Neural Networks by harnessing diverse sources of information. We argue that the work reported herein can set concrete foundations towards the refinement of situation aware mechanisms and further establish a strong basis for future design of resilience mechanisms for next generation mission-critical ICT systems.

Acknowledgments

The work has been supported in part by the UK Engineering and Physical Sciences Research Council (EPSRC) grants EP/N033957/1, EP/L026015/1, and EP/L005255/1. The authors are grateful to Simon Jouet, Long Chen, Noor Shirazi, Steven Simpson and Ghulam Mohi Ud Din for their valuable input.

References

1. S. Jouet, R. Cziva, and D.P. Pazaros, "Arbitrary Packet Matching in OpenFlow", 16th IEEE International Conference on High Performance Switching and Routing (IEEE HPSR), Budapest, Hungary, 1-4 Jul 2015.
2. A. K. Marnerides, A. Bhandari, H. Murthy and A. U. Mauthe, "A multilevel resilience framework for unified networked environments," 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM), Ottawa, ON, 2015.

3. M. A. M., Ariffin, A. K. Marnerides and A. U. Mauthe, "Multi-level Resilience in Networked Environments: Concepts & Principles", to appear in IEEE CCNC 2017, Las Vegas, NV, USA, 2017
4. N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling Innovation in Campus Networks", ACM SIGCOMM, March 2008.
5. J. Sterbenz, D. Hutchison, E. Cetinkaya, A. Jabbar, J. Rohrer, M. Schoeller, P. Smith: "Redundancy, Diversity, and Connectivity to Achieve Multilevel Network Resilience, Survivability and Disruption Tolerance", Journal Telecommunications Systems, vol. 56, no. 1, May 2014
6. A.K. Marnerides, A. Schaeffer-Filho, A. Mauthe, "Traffic anomaly diagnosis in Internet backbone networks: A survey", Computer Networks, Volume 73, 14 November 2014, Pages 224-243, ISSN 1389-1286, <http://dx.doi.org/10.1016/j.comnet.2014.08.007>.
7. M. R. Watson, N. u. h. Shirazi, A. K. Marnerides, A. Mauthe and D. Hutchison, "Malware Detection in Cloud Computing Infrastructures," in IEEE Transactions on Dependable and Secure Computing, vol. 13, no. 2, pp. 192-205, March-April 1 2016.
8. G. M. Ud Din, and A. K. Marnerides, "Short Term Power Load Forecasting Using Deep Neural Networks", to appear in IEEE International Conference on Computing, Networking and Communications (ICNC) 2017, Silicon Valley, USA, Jan. 2017
9. S. Fan and R. J. Hyndman, "Short-term load forecasting based on a semi-parametric additive model," *IEEE Trans. Power Syst.*, vol. 27, no. 1, pp. 134-141, 2012.