Mauro Andreolini	
Full address	Department of Physics, Computer Science and Mathematics University of Modena and Reggio Emilia Via Campi, 213/A 41125 - Modena
Tel. E-mail Home page	+39 059 2055192 mauro.andreolini@unimore.it https://secloud.ing.unimore.it/people/andreolini/

Personal information and Education

Mauro Andreolini, born in Rome on february 8th 1973, is Associate Professor at the University of Modena and Reggio Emilia (01/INFO-01 group, INFO-01/A area) since May 2026. He joined the Department of Physics, Computer Science and Mathematics in september 2012.

Previous positions

May 2026: Associate Professor (sector INF0-01/A) at the Department of Physical, Computer and Mathematical Sciences of the University of Modena and Reggio Emilia.

December 2023: Attainment of the Italian National Scientific Qualification (ASN) for the position of Associate Professor in Competition Sector 01/B1 (COMPUTER SCIENCE). The Qualification is valid for twelve years starting from 12/12/2023 and expires on 12/12/2035.

December 2023: Attainment of the Italian National Scientific Qualification (ASN) for the position of Associate Professor in Competition Sector 09/H1 (INFORMATION PROCESSING SYSTEMS). The Qualification is valid for twelve years starting from 17/12/2023 and expires on 17/12/2035.

September 2012-May 2026: Researcher (INF/01 area) at the Department of Physics, Computer Science and Mathematics, University of Modena and Reggio Emilia.

January 2005-September 2012: Researcher (INF/01 area) at the Department of Information Engineering, University of Modena and Reggio Emilia.

July 2003-December 2003: Visiting researcher for six months at the IBM T.J. Watson Research Center, Yorktown Heights, NY, USA, hosted by Dr. Philip Yu.

November 2001-November 2004: PhD (XVII cycle) in Automation and Information Engineering at the University of Roma "Tor Vergata".

May 2001-October 2001: Contractor at the Department of Computer Science and Systems, University of Roma "La Sapienza", supervised by Prof. Bruno Ciciani. Topic: *Implementation of a fault-tolerant, high performance Web server system.*

January 2001: Graduated summa cum laude (110/110 laude) at the University of Roma "Tor Vergata" in january 2001. Title of the manuscript: *Meccanismi content-aware per il Web dispatching*, supervisors Prof. Salvatore Tucci and Prof. Michele Colajanni.

Research activity

Mauro Andreolini has published 55 papers in international peer-reviewed conferences and journals in the following areas.

Computer security

Cryptography. In [46] the problem of data integrity in cloud databases is addressed by proposing encrypted Bloom filters that allow users to detect unauthorized modifications to outsourced data. An analytical model is also proposed to optimize storage and network costs according to the database structure and workload. The solution is evaluated through micro-benchmarks and TPC-C, showing improvements over existing alternatives. The contribution combines integrity robustness with reduced operational overhead. In [47] the gap between theoretical models and concrete implementations in encrypted database security is highlighted. In particular, the offline security of the Order-Revealing Encryption scheme originally proposed by Lewi and Wu is analyzed, showing that the scheme, although semantically secure, requires ciphertext ordering for efficiency. It is shown that index construction, ignored in the original paper, may introduce severe information leaks such as the presence of duplicate values, the statistical distribution of values, and the history of DBMS transactions. By analyzing two open-source implementations, vulnerabilities due to the use of standard search trees are identified. Finally, necessary conditions and practical solutions for secure indexing structures are proposed.

Authentication schemes. In [45] a critique of the perimeter security model is proposed, as it is considered inadequate for dynamic Web systems characterized by sudden and continuous addition and removal of users and hosts. The work embraces the "zero trust" approach, which applies controls to every request without relying on the physical location of users and devices. The authors observe that existing zero-trust architectures unrealistically assume the inviolability of some components. They therefore propose a new survivable zero-trust architecture designed for cloud environments. This architecture provides robustness, tolerates compromises, and can recover from software failures and successful attacks.

Side channel. In [51] the possibility of carrying out high-precision timing attacks on embedded USB devices, such as smart cards, using an ordinary PC is studied. In particular, eBPF is adopted to measure the execution times of digital signatures in kernel space, obtaining more accurate data than the user-space measurements proposed in previous work. The authors test the approach on a vulnerable smart card and evaluate the impact of the measurements on the success of a known private-key recovery attack. The results show significant improvements both in the precision of the collected timings and in the probability of compromise.

Anomaly detection and countermeasures. The research activity in anomaly detection, motivated by the goal of effectively countering attackers, has evolved over twenty years from simple attack-recognition heuristics to genuine models based on heuristic artificial intelligence. In [2] HoneySpam is presented, a honeypot-based framework designed to counter spammers. The main idea behind the prototype is to fight spam at the source, rather than during message reception as in most approaches in the literature. The countermeasures range from slowing down the process of harvesting e-mail addresses, to poisoning

e-mail address databases, to tracing offenders through open decoy proxies. In [23] a new attack type, mobility-based evasion, is presented. It exploits device mobility to fragment a malicious payload and evade even advanced intrusion detection systems. The work shows that mobile protocols, designed without adequate security measures, are particularly vulnerable. A cooperative intrusion detection framework is then proposed, which reconstructs and analyzes the distributed information generated during mobility. A prototype on Mobile IPv4, Mobile IPv6 and WiFi demonstrates the effectiveness and practicality of the approach. In [53] DAGA is proposed, an anomaly-detection algorithm for vehicular networks based on n-grams that analyzes only CAN message ID sequences, making it lightweight and suitable for resource-constrained microcontrollers. The framework can generate models with different memory footprints, suitable for hardware platforms with different miniaturization requirements. Tests on three prototypes show that DAGA can outperform the state of the art on more powerful microcontrollers while still remaining operational on very simple devices. The dataset and code were publicly released to support further research. In [52] DOLOS is proposed, an architecture that harmoniously integrates Moving Target Defense and Cyber Deception directly into production systems, overcoming the limitations of existing solutions. DOLOS combines randomization, diversity and redundancy with deception techniques, making it harder for attackers to distinguish real components from fake ones. The experimental evaluation against malware and penetration testers shows a significant slowdown of attacks and better protection against attacks carried out by cyber-criminal groups.

In [36] the weakness of machine-learning-based intrusion detection systems against targeted adversarial attacks is addressed. A new defense methodology is proposed for random-forest models, which are particularly suitable for cyber security. The technique is integrated into a network-traffic detector and tested on millions of flows. The results show that the system is more resistant to attacks and maintains good performance under normal conditions, outperforming existing detectors. In [39] a deep-reinforcement-learning framework is introduced to defend botnet detectors from adversarial attacks. The system automatically generates realistic evasion examples and uses them to harden the models, without degrading their performance in the absence of attacks. The framework is validated on several algorithms and public datasets, showing consistent improvements over existing solutions. The result is a generalizable approach that makes machine-learning-based detection systems more resilient against present and future attacks. In [55] the first dataset designed to evaluate the resilience of botnet detectors against adversarial attacks is introduced. It contains thousands of samples generated automatically through deep-reinforcement-learning techniques, capable of evading state-of-the-art machine/deep-learning detectors. The samples are derived from real botnet flows and include realistic modifications that do not alter the malicious nature of the traffic. The dataset allows researchers to test defenses under realistic conditions and provides useful insights to better understand and counter adversarial attacks. In [37] the threat models used in research on adversarial attacks against intrusion detection systems are criticized, as they are often unrealistic for concrete scenarios. A new model is proposed to describe realistic attacker capabilities and constraints against machine-learning-based Network Intrusion Detection Systems (NIDS). By applying the model to several known attacks, the work shows which attacks are actually feasible and which are not. The contribution helps strengthen defensive systems and guide research toward adversarial attacks that better match real-world conditions.

Automation of cyber operations. In [48, 49] an innovative GNU/Linux distribution, CAINE (Computer Aided INvestigative Environment), is proposed for offline forensic investigations. CAINE's specific novelty lies in its operating environment, which integrates heterogeneous analysis tools and produces a single customizable report.

In [1] a support framework for security assessments is introduced, based on a Prolog expert system that treats assessment activity as theorem proving. The framework can infer new facts, execute actions and dynamically update the knowledge base, overcoming typical limitations of traditional tools. Tests on jeopardy and boot-to-root scenarios show that the approach can identify non-standard objectives, handle isomorphic systems, namely structurally similar systems with the same exploitation path, without reconfiguration, and discover vulnerabilities emerging from the combination of multiple weaknesses.

In [17] an advanced system is proposed to evaluate the performance of trainees participating in cybersecurity exercises in cyber ranges. The proposed solution combines a distributed monitoring architecture, activity modeling through directed graphs, and new graph-based scoring algorithms. This makes it possible to accurately measure trainees' speed and accuracy, while also identifying specific mistakes. Compared with current platforms, the method enables a much more detailed and objective evaluation of the participant's path. In [40] the effectiveness of the proposed system is evaluated on large and heterogeneous exercise scenarios, highlighting some scoring anomalies such as precision losses and distorted scores. A correction to the current model is proposed, decomposing the exercise into local graphs for intermediate challenges and a global graph, thus enabling more accurate evaluations and specific weightings. Tests performed with a Python simulator show that the new approach scales better and produces more reliable scores both locally and globally.

Privacy. In [41] the identification of trajectories of moving users equipped with IoT devices is studied in the presence of privacy countermeasures such as geohashing and K-means clustering. Even with strong generalizations, trajectories often remain unique and therefore easily identifiable, confirming the privacy reduction inherent in the collection of geographic data. The tested techniques also suffer from major usability issues, especially when they drastically reduce spatial resolution. The results confirm that privacy protection in mobility contexts is an extremely complex problem that requires more sophisticated solutions.

High performance systems

Distributed Web servers. With reference to locally and geographically distributed systems, first- and second-level dispatching algorithms have been proposed. Experimental results showed the effectiveness of the proposed content-aware algorithms compared with traditional policies under realistic load scenarios. The analysis, which confirms some well-known results in load-sharing theory, also raises new questions, such as the usefulness of periodic node-state updates compared with asynchronous overload notifications [24, 28]. The design of innovative local clustering architectures was also pursued. Their scalability and performance proved comparable to those of analogous commercial products [25, 26]. In [20] the main properties of a modern high-performance and reliable e-commerce system, even under load peaks, are characterized. In [4] the main benchmarking tools and methods for Web systems are detailed and their applicability to distributed Web systems is evaluated.

The development of effective dispatching algorithms requires in-depth knowledge of the

internal and external workload characteristics. As regards the external workload, Web workload modeling was carried out by considering the statistical properties of traffic, service dynamism, and geographical effects between clients and servers [M2, C28]. As regards internal workload characterization, a new performance-analysis methodology was proposed. By operating at finer granularity levels, it allows the precise identification of system bottlenecks [R4]. This methodology was later adopted to study the performance sensitivity of a locally distributed system as a function of several hardware parameters, such as memory and connectivity availability [R2, C20]. The preliminary results proved to be of primary importance for studying the statistical properties of internal workload models [R1, C16, C18]. In particular, it was shown that, in realistic workload scenarios, traditional workload representation models are not adequate for the complexity of new Internet-connected distributed systems. Linear and non-linear stochastic workload representation models were therefore proposed, showing accuracy and responsiveness.

An interesting research direction concerned the integration of the load state of a resource with an additional piece of information related to its trend. This information makes it possible to anticipate undesired events and, when applied to load-balancing contexts for Web clusters [C1], proved to significantly improve response times compared with existing algorithms. In [C14], the concept of trend is applied to a linear-regression model for predicting the future load state.

Locally distributed NIDS. In [13] several load redistribution policies are proposed and compared in the context of locally distributed Network Intrusion Detection Systems (NIDS). Experimental results show the effectiveness of some of the proposed solutions in balancing network traffic across several instances of SNORT, a very popular NIDS.

Distributed systems with guaranteed quality of service

The Web has become the preferred interface also for highly critical services, which require privileged treatment compared with others. As a consequence, Web-system architectures have shifted from simple *best effort* systems to structures able to differentiate their behavior according to a given type of user. Research in this context focused on the design and performance evaluation of locally distributed, content-aware Web systems enriched with functionalities for providing services with contractually guaranteed performance levels. These functionalities, well known in QoS theory, were integrated at the application level in order to guarantee end-to-end service quality.

The proposed mechanisms are based on *request classification*, *admission control*, *performance isolation*, and *high resource utilization*. They were integrated into a prototype and the experimental results confirm good stability and robustness with respect to the objectives [6]. Several scheduling algorithms integrating all the above QoS principles were also proposed [5].

In [8] the problem of "graceful degradation" is addressed, namely the gradual and intentional degradation of performance in the presence of a workload volume that saturates a Web server's capacity. To this end, requests arriving at the server are classified and handled in order of importance. In [31] a new admission-control mechanism for Web clusters is proposed, able to guarantee graceful degradation under overload. Unlike other approaches in the literature, characterized by unconditional request rejection under overload, the pro-

posed scheme significantly reduced the number of rejected requests at the price of occasional Service Level Agreement violations.

Operating systems

Budget Fair Queueing (BFQ) is a proportional-share disk scheduler for the Linux kernel. In [54] some BFQ heuristics are presented that significantly reduce response delay for interactive processes and increase overall throughput for a wide range of disks, both rotational and SSD. An extensive experimental analysis shows BFQ's superiority over the official disk scheduler, Completely Fair Queueing (CFQ), in several usage scenarios.

Monitoring of complex systems

Algorithms. Today's computer systems are increasingly characterized by high complexity in terms of hardware/software resources and non-trivial interactions among multiple components. Managing such systems is problematic without appropriate monitoring tools. In this context, research aimed at increasing the scalability of a monitoring system was conducted along two distinct directions: an algorithmic one and an architectural one. In [42] resource-monitoring algorithms were proposed based on continuous sampling performed with standard tools. The main goal is the extraction of a stable and predictable internal representation across different time scales, even in the presence of high variability, dispersion and noise in individual samples. Experimental results show that it is possible to improve the accuracy of predictions provided by standard algorithms, based on linear and autoregressive models, while limiting the resource consumption required to compute the representation. In [28] the possibility of dynamically adapting the resource sampling interval is also studied, with the goal of reducing monitoring management costs while avoiding alteration of the statistical properties of sampled time series. The effectiveness of the proposed solution is demonstrated through several synthetic and real traces.

Architectures. A preliminary study highlighted the lack of scalability of relational databases as mass-storage support for information related to system monitoring [21]. For this reason, in [16] [C8] an innovative architecture is proposed for monitoring large-scale data centers hosting multi-tenant applications. The proposed approach overcomes the limits of centralized solutions, which cannot scale with the number of resources, and purely hierarchical solutions, which cannot support applications distributed over different data centers. The prototype was subsequently enriched with several monitoring algorithms [30, 27]. In [34] a scalability study of the prototype is detailed, showing its effectiveness.

Cloud Computing

In the context of Cloud systems providing virtualized services, an algorithm was proposed for the allocation and reallocation of virtual machines on the physical nodes of the infrastructure [14, 15]. The algorithm, designed for large-scale systems, is adaptive, threshold-free, and robust with respect to the load applied to the nodes. An important experimental result is the algorithm's ability to significantly reduce the number of live migrations, thereby lowering the management cost of the cloud infrastructure.

In [50] a request-redirection algorithm was studied for Web systems distributed over cloud infrastructures. In previous literature, the server best suited to serve a request is selected by considering at most the server load state and network latency. The proposed algorithm instead tries to predict the redirection cost, also including cloud-infrastructure-specific aspects, and the corresponding response time, choosing redirection only when advantageous. In this sense, the work represents one of the first attempts to integrate load information related to the infrastructure, the user and the network.

Autonomic systems

The research started from the observation that the vast majority of critical Internet-based systems operate on geographically distributed large-scale architectures, typically through the runtime execution of a large number of decision algorithms aimed at solving load-balancing, overload and admission-control, and geographic-redirection problems. The considerable number of components and parameters involved suggested the possibility of adopting models based on "self-* properties" to move toward actual autonomic systems. Research in this context proposed innovative self-inspection and self-decision algorithms and mechanisms in traditional Internet-based systems, with the goal of increasing their scalability and robustness.

The workload aggregation models considered were evaluated through a modular self-inspection support mechanism, generally applicable to Internet-based systems. This mechanism makes it possible to measure the load state of a resource efficiently and robustly across several realistic application scenarios characterized by very different statistical properties [44, 9]. Furthermore, it was shown that the use of these aggregation models increases system scalability and availability while avoiding performance degradation and component overload [34, 13].

In [11] autonomic algorithms for dispatching and request redirection in geographically distributed Web systems were proposed. They demonstrated several advantages over traditional policies: rapid adaptation to sudden load variations, which favors system stability; the use of local information, which minimizes the computational overhead associated with load-state updates among nodes; the absence of ad-hoc configuration parameters; and robustness in the presence of unexpected events.

Peer-to-peer systems

Some studies have shown that peer-to-peer (P2P) traffic is characterized by strong locality in resource accesses, and have proposed the use of caching techniques to limit the impact of P2P traffic on network infrastructures. However, it was observed that disregarding temporal effects in accesses and adopting simplifying assumptions lead to inaccuracies in estimating the fraction of P2P traffic suitable for caching. In [18] an analysis of P2P traffic is proposed with two goals. On one hand, the first analytical model of P2P workload is proposed, highlighting traffic characteristics; on the other hand, the effectiveness of caching solutions is evaluated in light of the results obtained. For each category of resources typically available in a file-sharing network, probability distributions of resource size and popularity, instantaneous traffic, and its temporal evolution with seasonal, weekly and hourly patterns were

provided. The workload model obtained from the analyses was used for more reliable estimates, allowing more accurate evaluation of the effectiveness of file-sharing traffic caching and revising some results previously proposed in the literature.

In [33] a classification of commonly used techniques for characterizing P2P traffic is presented. In particular, the results obtainable through traffic-sample analysis and active probing on a file-sharing overlay network are compared. Some discrepancies between the results obtained with the two methods are highlighted, and it is shown that only a combined use of the two techniques provides a complete view of the traffic patterns associated with P2P applications.

In [32] a new resource-search mechanism based on fuzzy Distributed Hash Tables (DHT) is proposed. The proposed mechanism solves one of the problems associated with using DHTs for search functionalities, namely the impossibility of performing wildcard searches because of the need to use a unique identifier. The properties shown in the experimental evaluation are the following: search flexibility, effectiveness in retrieving resources, efficient use of system resources, and robustness against node failures.

Teaching activity

Mauro Andreolini has taught the following courses at the University of Modena and Reggio Emilia, ordered by academic year.

- Lecturer of the **Vulnerability Research** course (12 hours, 3 ECTS) within the PhD Program in "Computer and Data Science for Technological and Social Innovation (CDS-TSI)" (2023-).
- Lecturer of the **Secure Software Development** course (9 CFU), Master Degree in Computer Science (2018-).
- Lecturer of the introductory module (6 CFU) of the **Operating Systems** course (6+3 CFU), Bachelor Degree in Computer Science (2024-).
- Lecturer of the **Secure Programming** course (6 CFU), Master Degree in Computer Science (2016-2017).
- Lecturer of the **Advanced Web Technologies** course (9 CFU), Master Degree in Computer Science (2009).
- Lecturer of the **Dynamic Languages** course (9 CFU), Bachelor Degree in Computer Science (2008, 2010).
- Lecturer of the **Software Design Methodologies** course (6 CFU), Bachelor Degree in Computer Science (2006-2007).
- Lecturer of the **Computer Applications** course (2 CFU), Bachelor Degree in Geological Sciences (2005-2007).
- Lecturer of the **Operating Systems** course (9 CFU), Bachelor Degree in Computer Science (2004-2023).

He has also supervised or co-supervised 80 Bachelor and Master theses. He is the scientific tutor of two PhD students in the PhD Program in "Computer and Data Science for Technological and Social Innovation (CDS-TSI)" at the University of Modena and Reggio Emilia. He has been a member of the PhD Faculty Board of the PhD Program in "Computer and Data Science for Technological and Social Innovation (CDS-TSI)" since 2023.

Third Mission activities

Mauro Andreolini has carried out the following teaching activities within the Third Mission.

- Lecturer of the advanced training course **Advanced Penetration Tester** (40 hours) (2025).
- Plenary lecture "Hacker, questo illustre sconosciuto!" at the event "A tu per tu con la Scienza" (2024-).

- Lecturer of the advanced training course **Penetration Tester** (40 hours) (2022-2023).
- Lecturer of the **Exploitation** teaching module (24 hours) of the "Cyber Academy" advanced training course (2018).
- Lecturer of the **Secure Programming** teaching module (24 hours) of the "Cyber Academy" advanced training course (2016-2018).
- Lecturer of the **Vulnerability Assessment and Penetration Testing** teaching module (24 hours) of the "Cyber Academy" advanced training course (2016-2018).
- Lecturer of the **Operating Systems** teaching module (24 hours) of the "Cyber Academy" advanced training course (2016-2017).
- Lecturer of the **Vulnerability Assessment** teaching module (24 hours) of the CyberDefense Master at the Armed Forces School of Telecommunications in Chiavari (2014-2018).
- Lecturer of the **Operating Systems** teaching module (24 hours) of the CyberDefense Master at the Armed Forces School of Telecommunications in Chiavari (2013-2017).
- Lecturer of the **Operating Systems** and **Penetration Testing** teaching modules of the second-level Master in "Sicurezza dei Sistemi Informatici: Normative e Tecniche Avanzate di Protezione" (2010-2014).

Mauro Andreolini has carried out the following public engagement activities within the Third Mission.

- Coordinator of the "CyberChallenge.IT" booth at the "Notte della Ricerca" event (2023-).

Duties to the scientific community

Scientific committees

Mauro Andreolini has been a member of the scientific committee for the following international conferences:

- International IEEE Symposium on Network Computing and Applications (IEEE NCA, 2014-2024).
- International Workshop on Emerging Technologies for Next-generation GRID (ETNGRID, 2006-2014).
- IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems (IEEE DPDNS, 2013-2013).
- IEEE International Symposium on Parallel and Distributed Processing with Applications (IEEE ISPA 2012).

Organizing committees

Mauro Andreolini has been a member of the organizing committee for the following international conferences:

- Financial Chair of the International IEEE Symposium on Network Computing and Applications (NCA 2023).
- Program Chair of the International IEEE Symposium on Network Computing and Applications (NCA 2022).
- Member of the organizing committee of the IFIP WG7.3 International Symposium on Computer Performance Modeling, Measurement and Evaluation (PERFORMANCE 2002).

Reviewing activity

Mauro Andreolini has been a reviewer for the following international journals: IEEE Transactions on Parallel and Distributed Systems, ACM Performance Evaluation Review, Computer Journal, IEEE Computer Networks, Journal of Systems and Software, Journal of Network computing and Applications, Int'l Journal of Web Engineering and Technology, Pervasive computing. Mauro Andreolini has been a reviewer for the following international conferences: World Wide Web Int'l Conference (2002, 2003, 2004, 2005, 2006, 2007, 2009), AAA-Idea (2006, 2007), Europar (2007), DISC (2005), Globe-comm (2007), HotP2P (2004, 2005, 2006, 2007), IEEE MASCOTS (2004, 2005), MP2P (2005), Perf 2002, PE-WASUN(2005), ACM SIGMETRICS 2007, UIC 2007, WTAS 2006, COOPIS (2007), NCA (2008-2025).

Scientific responsibilities

Mauro Andreolini has been Scientific Coordinator for the following research projects for the University of Modena and Reggio Emilia:

2017-2021 PNRM A2016.099bis "UNAVOX";

2016-2017 PNRM E.F. 2015 A2013.060 "Smart Environment Area Firing RangE (SEAFIRE)";

2015-2021 PNMR E.F. 2014 A2012.154 "Digital Trunk Communication in Hostile Environment (DTCHE)";

2008-2010 Local research-unit leader for UNIMORE in the PRIN AUTOSEC "Autonomic Security" project.

Awards and recognitions

Within his research activity Dr. Mauro Andreolini has received the following international awards and recognitions.

- **Best paper award** for Mauro Andreolini, Sara Casolari, Stefania Tosi, "A hierarchical architecture for on-line control of private cloud-based systems", *Proc. of 10th World Wide Web Internet Conference*, Timisoara, Romania, October 2010.
- **Best paper award** for Mauro Andreolini, Sara Casolari, Michele Colajanni, "Self-inspection mechanisms for the support of autonomic decisions in Internet-based systems", *Proc. of 3rd International Conference on Autonomic and Autonomous Systems*, Athens, Greece, June 2007.
- **Candidate best paper award** for Mauro Andreolini, Marcello Pietri, Stefania Tosi, Andrea Balboni, "Monitoring Large Cloud-Based Systems", *Proc. of the International Conference on Cloud Computing and Services Science*, Barcelona, Spain, 3-5 April 2014.

Organizational activities

- Director of the CyberChallenge.IT project for the University of Modena and Reggio Emilia (2023-).

Research products

- Original designer of the CAINE (Computer Aided INvestigative Environment) GNU/Linux distribution for computer forensics.
- Contributor of a modification to the Linux BFQ (Budget Fair Queueing) disk scheduler aimed at improving throughput in interleaved workload scenarios.

Modena, May 21, 2026

Bibliography

References

[1] Mauro Andreolini, Andrea Artioli, Luca Ferretti, Mirco Marchetti, Michele Colajanni, Claudia Righi, et al. A framework for automating security assessments with deductive reasoning. In *CEUR WORKSHOP PROCEEDINGS*, volume 3488. CEUR-WS, 2023.

[2] Mauro Andreolini, Alessandro Bulgarelli, Michele Colajanni, and Francesca Mazzoni. Honeyspam: Honeypots fighting spam at the source. *SRUTI*, 5:11–11, 2005.

[3] Mauro Andreolini, Claudia Canali, and Riccardo Lancellotti. Impact of request dispatching granularity in geographically distributed web systems. In *Sixth IEEE International Symposium on Network Computing and Applications (NCA 2007)*, pages 45–52. IEEE, 2007.

[4] Mauro Andreolini, Valeria Cardellini, and Michele Colajanni. Benchmarking models and tools for distributed web-server systems. In *IFIP International Symposium on Computer Performance Modeling, Measurement and Evaluation*, pages 208–235. Springer, 2002.

[5] Mauro Andreolini, Emiliano Casalicchio, Michele Colajanni, and Marco Mambelli. A cluster-based web system providing differentiated and guaranteed services. *Cluster Computing*, 7(1):7–19, 2004.

[6] Mauro Andreolini, Emiliano Casalicchio, Michele Colajanni, Marco Mambelli, et al. Qos-aware switching policies for a locally distributed web system. In *Proc. of the 11th Int'l World Wide Web Conf.* Honolulu, Hawaii, May, 2002.

[7] Mauro Andreolini and Sara Casolari. Load prediction models in web-based systems. In *Proceedings of the 1st international conference on Performance evaluation methodolgies and tools*, pages 27–es, 2006.

[8] Mauro Andreolini, Sara Casolari, and Michele Colajanni. A distributed architecture for gracefully degradable web-based services. In *Fifth IEEE International Symposium on Network Computing and Applications (NCA'06)*, pages 235–238. IEEE, 2006.

[9] Mauro Andreolini, Sara Casolari, and Michele Colajanni. Self-inspection mechanisms for the support of autonomic decisions in internet-based systems. In *Third International Conference on Autonomic and Autonomous Systems (ICAS'07)*, pages 53–53. IEEE, 2007.

[10] Mauro Andreolini, Sara Casolari, and Michele Colajanni. Trend-based load balancer for a multi-tier distributed system. In *2007 15th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, pages 288–294. IEEE, 2007.

[11] Mauro Andreolini, Sara Casolari, and Michele Colajanni. Autonomic request management algorithms for geographically distributed internet-based systems. In *2008 Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 171–180. IEEE, 2008.

[12] Mauro Andreolini, Sara Casolari, and Michele Colajanni. Models and framework for supporting runtime decisions in web-based systems. *ACM Transactions on the Web (TWEB)*, 2(3):1–43, 2008.

[13] Mauro Andreolini, Sara Casolari, Michele Colajanni, and Mirco Marchetti. Dynamic load balancing for network intrusion detection systems based on distributed architectures. In *Sixth IEEE International Symposium on Network Computing and Applications (NCA 2007)*, pages 153–160. IEEE, 2007.

[14] Mauro Andreolini, Sara Casolari, Michele Colajanni, and Michele Messori. Dynamic load management of virtual machines in cloud architectures. In *International Conference on Cloud Computing*, pages 201–214. Springer, 2009.

[15] Mauro Andreolini, Sara Casolari, Marcello Pietri, Stefania Tosi, et al. Resilient and adaptive networked systems., 2014.

[16] Mauro Andreolini, Sara Casolari, Stefania Tosi, et al. A hierarchical architecture for online control of private cloud based systems. In *Proc. of 10th World Wide Web Internet Conference (WWWCONF2010), Timisoara, Romania*, 2010.

[17] Mauro Andreolini, Vincenzo Giuseppe Colacino, Michele Colajanni, and Mirco Marchetti. A framework for the evaluation of trainee performance in cyber range exercises. *Mobile Networks and Applications*, 25(1):236–247, 2020.

[18] Mauro Andreolini, Michele Colajanni, and Riccardo Lancellotti. Peer-to-peer workload characterization: techniques and open issues. In *2004 International Workshop on Hot Topics in Peer-to-Peer Systems*, pages 66–71. IEEE, 2004.

[19] Mauro Andreolini, Michele Colajanni, and Riccardo Lancellotti. Impact of memory technology trends on performance of web systems. In *International Conference on Next Generation Web Services Practices (NWeSP'05)*, pages 7–pp. IEEE, 2005.

[20] Mauro Andreolini, Michele Colajanni, and Riccardo Lancellotti. Web system reliability and performance. In *Web Engineering*, pages 181–218. Springer, 2006.

[21] Mauro Andreolini, Michele Colajanni, and Riccardo Lancellotti. Assessing the overhead and scalability of system monitors for large data centers. In *Proceedings of the First International Workshop on Cloud Computing Platforms*, pages 1–7, 2011.

[22] Mauro Andreolini, Michele Colajanni, Riccardo Lancellotti, and Francesca Mazzoni. Fine grain performance evaluation of e-commerce sites. *ACM Sigmetrics Performance Evaluation Review*, 32(3):14–23, 2004.

[23] Mauro Andreolini, Michele Colajanni, and Mirco Marchetti. A collaborative framework for intrusion detection in mobile networks. *Information Sciences*, 321:179–192, 2015.

[24] Mauro Andreolini, Michele Colajanni, and Ruggero Morselli. Performance study of dispatching algorithms in multi-tier web architectures. *ACM SIGMETRICS Performance Evaluation Review*, 30(2):10–20, 2002.

[25] Mauro Andreolini, Michele Colajanni, and Marcello Nuccio. Kernel-based web switches providing content-aware routing. In *Second IEEE International Symposium on Network Computing and Applications, 2003. NCA 2003.*, pages 25–32. IEEE, 2003.

[26] Mauro Andreolini, Michele Colajanni, Marcello Nuccio, et al. Scalability of content-aware server switches for cluster-based web information systems. In *The Twelfth International World Wide Web Conference-WWW 2003: proceedings of the tracks education, global community, industrial practice and experience, web services*, pages 200–208. World Wide Web Conference Committee., 2003.

[27] Mauro Andreolini, Michele Colajanni, and Marcello Pietri. A scalable architecture for real-time monitoring of large information systems. In *2012 Second Symposium on Network Cloud Computing and Applications*, pages 143–150. IEEE, 2012.

[28] Mauro Andreolini, Michele Colajanni, Marcello Pietri, and Stefania Tosi. Real-time adaptive algorithm for resource monitoring. In *Proceedings of the 9th International Conference on Network and Service Management (CNSM 2013)*, pages 67–74. IEEE, 2013.

[29] Mauro Andreolini, Michele Colajanni, Marcello Pietri, and Stefania Tosi. Adaptive, scalable and reliable monitoring of big data on clouds. *Journal of Parallel and Distributed Computing*, 79:67–79, 2015.

[30] Mauro Andreolini, Michele Colajanni, and Stefania Tosi. A software architecture for the analysis of large sets of data streams in cloud infrastructures. In *2011 IEEE 11th International Conference on Computer and Information Technology*, pages 389–394. IEEE, 2011.

[31] Mauro Andreolini, Michele Colajanni, and Paolo Valente. Design and testing of scalable web-based systems with performance constraints. In *2005 Workshop on Techniques, Methodologies and Tools for Performance Evaluation of Complex Systems (FIRB-PERF'05)*, pages 15–25. IEEE, 2005.

[32] Mauro Andreolini and Riccardo Lancellotti. A flexible and robust lookup algorithm for p2p systems. In *2009 IEEE International Symposium on Parallel & Distributed Processing*, pages 1–8. IEEE, 2009.

[33] Mauro Andreolini, Riccardo Lancellotti, and Philip S Yu. Analysis of peer-to-peer systems: workload characterization and effects on traffic cacheability. In *The IEEE Computer Society's 12th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems, 2004.(MASCOTS 2004). Proceedings.*, pages 95–104. IEEE, 2004.

[34] Mauro Andreolini, Marcello Pietri, Stefania Tosi, Andrea Balboni, et al. Monitoring large cloud-based systems. In *CLOSER 2014-Proceedings of the 4th International Conference on Cloud Computing and Services Science*, pages 341–351. SciTePress, 2014.

[35] Mauro Andreolini, Marcello Pietri, Stefania Tosi, and Riccardo Lancellotti. A scalable monitor for large systems. In *International Conference on Cloud Computing and Services Science*, pages 100–116. Springer, 2014.

[36] Giovanni Apruzzese, Mauro Andreolini, Michele Colajanni, and Mirco Marchetti. Hardening random forest cyber detectors against adversarial attacks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 4(4):427–439, 2020.

[37] Giovanni Apruzzese, Mauro Andreolini, Luca Ferretti, Mirco Marchetti, and Michele Colajanni. Modeling realistic adversarial attacks against network intrusion detection systems. *Digital Threats: Research and Practice (DTRAP)*, 3(3):1–19, 2022.

[38] Giovanni Apruzzese, Mauro Andreolini, Mirco Marchetti, Vincenzo Giuseppe Colacino, and Giacomo Russo. Appcon: Mitigating evasion attacks to ml cyber detectors. *Symmetry*, 12(4):653, 2020.

[39] Giovanni Apruzzese, Mauro Andreolini, Mirco Marchetti, Andrea Venturi, and Michele Colajanni. Deep reinforcement adversarial learning against botnet evasion attacks. *IEEE Transactions on Network and Service Management*, 17(4):1975–1987, 2020.

[40] Andrea Artioli, Mauro Andreolini, Luca Ferretti, Mirco Marchetti, et al. Evaluating trainees in large cyber exercises. In *CEUR WORKSHOP PROCEEDINGS*, volume 3731. CEUR-WS, 2024.

[41] Andrea Artioli, Luca Bedogni, and Mauro Andreolini. Effects of geohashing and k-means clustering on uniqueness in a mobility dataset. In *2024 IEEE/ACM Symposium on Edge Computing (SEC)*, pages 389–394. IEEE, 2024.

[42] Sara Casolari, Mauro Andreolini, and Michele Colajanni. Runtime prediction models for web-based system resources. In *2008 IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems*, pages 1–8. IEEE, 2008.

[43] Michele Colajanni, Mauro Adreolini, and Valeria Cardellini. Benchmarking of locally and geographically distributed web-server systems. *World Wide Web, Budapest*, 2003.

[44] Michele Colajanni, Mauro Andreolini, Riccardo Lancellotti, et al. Open issues in self-inspection and self-decision mechanisms for supporting complex and heterogeneous information systems. *Proceedings of. Int'l SELF-STAR*, page 13, 2004.

[45] Luca Ferretti, Federico Magnanini, Mauro Andreolini, and Michele Colajanni. Survivable zero trust for cloud computing environments. *Computers & Security*, 110:102419, 2021.

[46] Luca Ferretti, Mirco Marchetti, Mauro Andreolini, and Michele Colajanni. A symmetric cryptographic scheme for data integrity verification in cloud databases. *Information Sciences*, 422:497–515, 2018.

[47] Luca Ferretti, Mattia Trabucco, Mauro Andreolini, Mirco Marchetti, et al. How (not) to index order revealing encrypted databases. In *CEUR WORKSHOP PROCEEDINGS*, volume 3488. CEUR-WS, 2023.

[48] G Giustini, Mauro Andreolini, Michele Colajanni, et al. Caine: A new open-source live distribution for digital forensics. In *First Workshop on Open Source software for computer and network forensics (OSSCONF2008)*, pages 51–61. Springer, 2008.

[49] Giancarlo Giustini, Mauro Andreolini, and Michele Colajanni. Open source live distributions for computer forensics. In *Open source software for digital forensics*, pages 69–82. Springer, 2009.

[50] Riccardo Lancellotti, Mauro Andreolini, Claudia Canali, and Michele Colajanni. Dynamic request management algorithms for web-based services in cloud computing. In *2011 IEEE 35th Annual Computer Software and Applications Conference*, pages 401–406. IEEE, 2011.

[51] Gianmarco Lusvardi, Luca Ferretti, Mauro Andreolini, et al. Timing side-channel attacks on usb devices using ebpf. In *CEUR WORKSHOP PROCEEDINGS*, volume 3731. CEUR-WS, 2024.

[52] Giulio Pagnotta, Fabio De Gaspari, Dorjan Hitaj, Mauro Andreolini, Michele Colajanni, and Luigi V Mancini. Dolos: A novel architecture for moving target defense. *IEEE Transactions on Information Forensics and Security*, 18:5890–5905, 2023.

[53] Dario Stabili, Luca Ferretti, Mauro Andreolini, and Mirco Marchetti. Daga: Detecting attacks to in-vehicle networks via n-gram analysis. *IEEE Transactions on Vehicular Technology*, 71(11):11540–11554, 2022.

[54] Paolo Valente and Mauro Andreolini. Improving application responsiveness with the bfq disk i/o scheduler. In *Proceedings of the 5th Annual International Systems and Storage Conference*, pages 1–12, 2012.

[55] Andrea Venturi, Giovanni Apruzzese, Mauro Andreolini, Michele Colajanni, and Mirco Marchetti. Drelab-deep reinforcement learning adversarial botnet: A benchmark dataset for adversarial attacks against botnet intrusion detection systems. *Data in Brief*, 34:106631, 2021.