The Way Forward
Workshop on Future Challenges in Software Engineering

27 June 2023 - Milano, Italy

June 2023

SWForum.eu

Table of Contents

Disclaimer

About SWForum.eu

The SWForum.eu Coordination and Support Action project has received funding from the European Union's Horizon 2020 - Research and Innovation program - under grant agreement no. 957044. The project is coordinated by FUNDACION TECNALIA RESEARCH & INNOVATION "TECNALIA" (ES), supported by its partners from the Politecnico di Milano "POLIMI" (IT), Trust-IT SRL (IT), CONCEPTIVITY (CH) and the Chancellor Masters and Scholars of the University of Oxford "UOXF" (GB). The content of the present report does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of such content.

Executive Summary

The SWForum.eu "Way Forward Workshop on Future Challenges in Software Engineering" took place on 27 June 2023 at the Politecnico di Milano, in the Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB). Its purpose was to bring researchers and practitioners together to assess the current state of software engineering in the context of a set of representative advanced technologies and discuss the resulting challenges that the future will bring in the discipline.

The program committee selected four areas that are considered particularly pertinent at this time.

Security in the Computing Continuum. Security in the computing world is very similar to security in the supply chain world – if one link in the chain or in the process is unsafe or insecure, this can make the whole system or the whole software unsafe or insecure. There have been significant efforts to incorporate "security by design" into software development, but significant challenges likewise remain, including those related to new technologies like machine learning.

Software Engineering and AI. The remarkable worldwide impact of the new Large Language Models is only the latest manifestation of the growing mutual influence that AI and software engineering are exerting on each other. Can current software engineering practices cope with the complexities of AI architectures and applications, including validation (explainability) and security?

Software Engineering for Quantum Computing. Quantum computing is reaching significant and promising advancements and represents one of the ground-breaking initiatives that are expected to change the way we conceive programming today. A crucial issue from the software engineering standpoint is the identification of effective design and programming abstractions that allow people skilled in computer science to take advantage of the enormous power of quantum computing while keeping the complexity of design and programming activities under control and enabling analysis and testing of the developed code.

Sustainable Software Engineering. Sustainability within software is often considered to be how easily or well a particular piece of software can be maintained or developed in the face of changing or developing environments around it. Can the concept be extended to its contribution to urgent challenges such as the United Nations Sustainable Development Goals?

"Software is more pervasive than ever and therefore we need to push together the European research in software engineering. Current technology challenges such as generative AI, Quantum shift paradigm, constrained computing resources in the edge, and cybersecurity for complex software supply chains are highly impacting the way we develop and operate software."

Juncal Alonso, SWForum.eu Project Coordinator.

The workshop began with an introduction to the SWForum.eu project and its objectives in creating a forum for discussion of software engineering challenges and opportunities within the European software community. It closed with a wrap-up session of conclusions among the participants.

The full agenda of the workshop and speaker presentations may be found on the SWForum.eu website.[1] The results of the individual sessions are summarized in the following sections.

1. https://swforum.eu/agenda-swforumeu-way-forward-workshop-future-challenges-software-engineering

Highlights from the Event

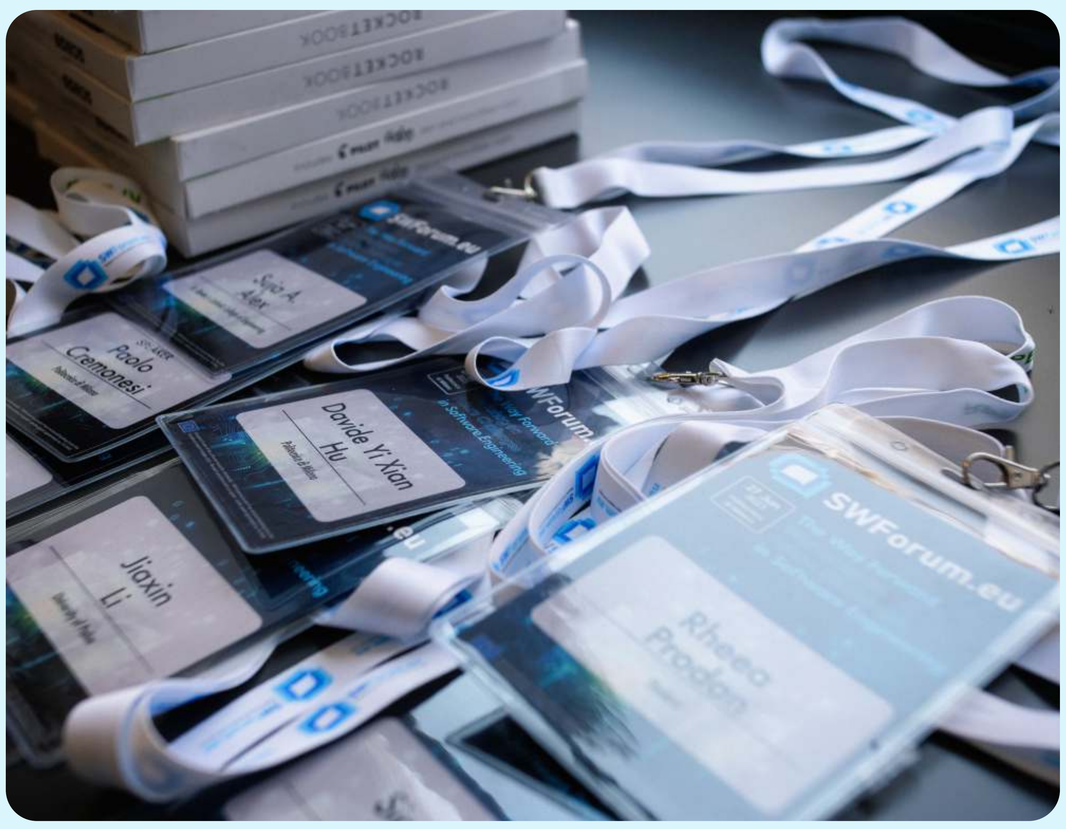

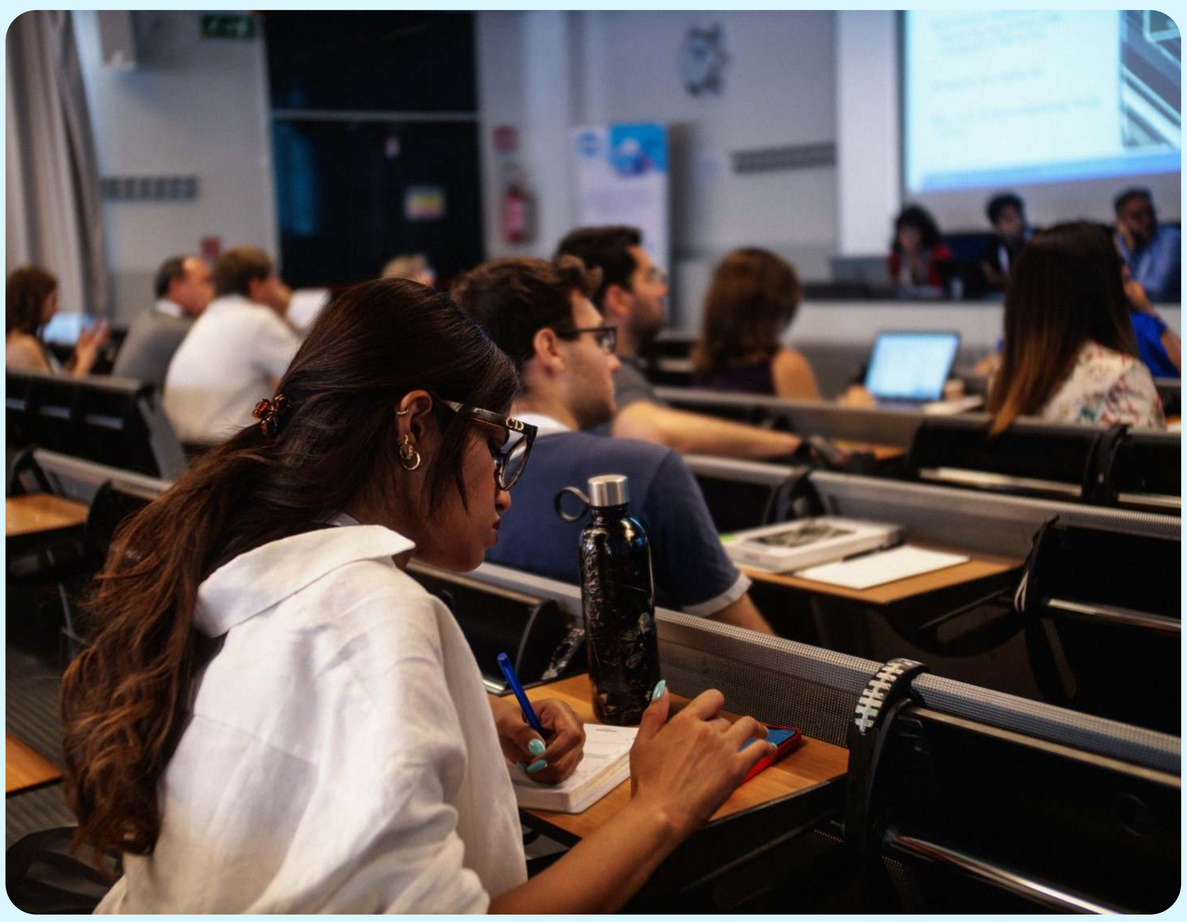

Session 1 09:30 – 10:30

Security in the Computing Continuum

Mark Miller
Conceptivity

CHAIR

The area of cybersecurity is producing new challenges in numerous areas of software engineering. The increasing criticality of cyber-physical systems (e.g., self-driving vehicles, medical equipment, IoT devices) and the growing complexity of the software supply chain is creating challenges for the architecting and development of software systems that interact heavily with the environment and are often connected across the entire cloud-edge continuum. A particularly critical case is power distribution systems, where a successful attack can effectively shut down an entire infrastructure. Legacy power systems are often subsequently connected in unexpected ways, providing originally unforeseen bridges for attackers to access everything downstream in the continuum. New defences are needed, such as frequent changes of interconnections or interactions among system components, to present attackers with moving targets.

Security By Design is a new and compelling paradigm for engineering of secure systems. But in order to become practical in real-world commercial contexts, it needs to incorporate the features that service providers in relevant fields (e.g., telecommunications) are routinely expected offer, such as Service Level Agreements (SLAs). However, features like SLAs for cybersecurity introduce new challenges in software engineering, such as developing meaningful metrics for characterizing levels of security in software systems; such challenges will lead to the development of new and appropriate ways to model systems in order to be able to capture such metrics.

Generative AI is acknowledged to have provided new perspectives on software development (see also Session 2), but it is also making contributions to the toolbox of the malicious cybersecurity actor. Phishing attacks account for the largest portion of cyberattacks, and AI is making these kinds of attacks even more efficient and effective to design. Open-source intelligence (OSINT), gathered from the internet and combined with generative AI, makes it possible to personalise phishing message texts. Multimedia forms of generative AI (e.g., image-from-text systems) are making it possible to design and implement credible websites for phishing attacks. Software cybersecurity engineers will need to develop new techniques to recognize and counteract this new category of AI-enabled attacks – including AI-supported countermeasures.

During the session...

Session Agenda

- Cyber-Physical Systems: Attack and Defence – **1. Martin Higgins**, University of Oxford
- Security-by-Design methodologies and security metrics – **2. Valentina Casola**, University of Napoli Federico II
- Usage of AI to automate human-related cyber attacks – **3. Francesco Morano**, CEFRIEL

Session 2 11:00 – 13:00

Software Engineering and Artificial Intelligence

John Favaro
Trust-IT

CHAIR

Software engineering and AI have been closely coupled for decades. AI-enabled assistance for software engineering has been envisioned for more than thirty years, and with the arrival of generative AI, it now seems feasible not only at the coding level, but at all levels ranging from requirements elicitation to operations and maintenance. The size and complexity of the new AI-enabled systems has also opened up new challenges for software engineering support for the construction of these systems.

One such challenge is the integration of considerations of trustworthiness and ethical values in AI-enabled systems, if they are to be used in critical, industrial-grade applications. In Value-Sensitive Design (VSD), end users and technology developers work together from the very beginning of the engineering cycle to ensure this integration, while generating a statistical model and other general evaluation mechanisms to understand and quantify the satisfaction of ethical and trustworthiness criteria.

A related challenge is the integration of economic considerations in AI-enabled systems, of which one very well-known example is binary classification. Binary classification is often used to separate “good” from “bad” categories in arbitrary domains. However, mistaken classifications (e.g., false positives or false negatives) can have high economic costs – which could be the opposite depending on the domain. By integrating domain-specific (economic) value-based considerations into these AI-enabled systems, their real-world usefulness can be greatly enhanced.

If software engineering can hope to ameliorate the construction of machine learning-based systems, a plausible research agenda must be formulated. Challenges currently being identified for possible inclusion in a research agenda include improving the deployment and operation of federated learning systems in a self-adaptive manner, which in practice involves addressing numerous issues, such as the sharing of private data among the clients and the optimisation of workload and resource allocation. Another example is the automatic detection of unknown domains and adaptation of neural networks to them.

Yet even if these challenges are addressed, a major challenge remains of finding a good paradigm for the integration of software engineering and AI development processes. In an experience with “MLOps”, a phenomenon of two separate teams (ML and DevOps) was observed that did not successfully integrate at the cultural level, each demonstrating little knowledge of or interest in the other culture. This led to costly development errors, such as a drastic underestimation of the amount of support infrastructure needed for MLOps, with consequent rise in technical debt.

Yet another challenge concerning the integration of AI and software engineering is the increasing importance of AI in applications that span the computing continuum. This leads to the need for new software engineering approaches to managing both applications and AI-related activities across the continuum. For example, for the validation and improvement of ML models used on the near and far Edge (e.g., for robotic and automated driving applications), techniques are needed to optimally move data from devices to the cloud to train, validate, and improve AI models to be used in further missions. The concept of a “meta” operating system that can transparently span the continuum is one way of helping to simplify the management of application development and operations in the continuum, including AI-enabled applications.

Another way to help to simplify AI-related development in the continuum is the provision of an advanced framework that fills known technology gaps, by providing solutions to efficiently develop and deploy AI enterprise applications across the computing continuum. The framework contains not only strategically developed tools and components that cover those known gaps, but also a design workflow that pulls together those key tools into a coherent methodology.

During the session...

Session Agenda

- Evaluating and validating hybrid AI algorithms and asserting their adherence to an ethical and legal overaching approach to trustworthiness – **1. Nuria Quintana Fernandez**, TECNALIA
- Software Engineering for Machine Learning, some first experiences – **2. Luciano Baresi**, Politecnico di Milano
- MLOps the hard way: a journey on integrating AI & software engineering – **3. Michele Ciavotta**, University of Milano
- On the evaluation of binary classifiers for Software Engineering - **4. Luigi Lavazza**, Universita degli Studi dell’Insubria
- Development of an IoT2cloud Operating System (ICOS) – **5. Marina Giordanino**, Stellantis
- AI-SPRINT: Design and Runtime Framework for Accelarating the Development of AI Applications in the Computing Continuum – **6. Francesco Lattari**, Politecnico di Milano

Session 3 14:10 – 15:30

Software Engineering for Quantum Computing

Elisabetta Di Nitto
POLIMI

CHAIR

Quantum computing is rapidly leaving the laboratory with the promise of becoming industrially relevant technology. Software engineers are attempting to create a new discipline of “quantum software engineering” for providing new methods, frameworks, and programming languages for quantum applications. A kind of roadmap and call for action has been elaborated in the Talavera Manifesto for Quantum Software Engineering and Programming, which attempts to define and present a set of fundamental principles. Yet survey results indicate that much current research concentrates on software testing, proposals for solutions, and “philosophical” papers rather than more traditional software engineering topics like processes and professional practice.

One area that has garnered considerable interest is quantum optimization, due to its potential to achieve significant improvements over classical approaches, with vast applications ranging from transportation (cf. the famous Traveling Salesman problem) to finance and energy. There is particular interest in hybrid quantum-classical solvers, due to its promise of both performance improvements and dealing with current limited quantum resources. But different research groups are approaching the problem from different perspectives, and still need to form a united community to guide research and avoid fragmentation.

Optimization serves as a good example of the conceptual challenges that quantum computing can present to software engineers used to traditional approaches. While the quantum gate model is closer to the traditional digital computing model, quantum annealing may be seen as an analogue version which, while capable of achieving highly performant results (for example, in optimization problems like the Traveling Salesman), requires programmers to understand the nonintuitive quantum mechanics underlying its behaviour.

During the session...

Session Agenda

- The Quantum Frontier of Software Engineering and its Adoption – **1. Dario Di Nucci**, University of Salerno
- The Role of Formal Probabilistic Models in Quantum Computing – **2. Vlado Stankovski**, University of Ljubljana
- Quantum Optimisation: current trends and open opportunities – **3. Eneko Osaba**, TECNALIA
- Quantum Annealing as the analog version of Quantum Computing – **4. Paolo Cremonesi**, Politecnico di Milano

Session 4 16:00 - 17:30

Sustainable Software Engineering

David Wallom
University of Oxford

CHAIR

“Sustainable software engineering” is a concept that may be usefully viewed from different perspectives. For example, it may be viewed from the perspective of traditional best practices in development and maintenance in order to remain robust and valid – sustainable – over time. But it also may be viewed from a perspective linked more closely to current preoccupation with societal challenges such as climate change and the UN Sustainable Development Goals in general. Big data and its associated processing intensity is known to be a source of potential problems in this regard.

A fertile ground for seeking progress in this area Is presented by the processing of graph structures, which have become a data representation of choice for many important analytics applications (including those related to the Sustainable Development Goals mentioned above). The efficient processing of massive graphs, as measured by appropriate metrics, can make a concrete contribution to sustainability goals.

Such metrics can include “volume” (for the size of graphs), “velocity” (for dealing also with dynamically changing topologies), and a novel metric, defined explicitly to measure capability of sustainable processing at scale, known as “viridescence”. These can help not only to construct effective graph processing optimizers, but also graph processing “greenifiers”, with an eye toward reducing the harmful effects of the underlying computing activities.

During the session...

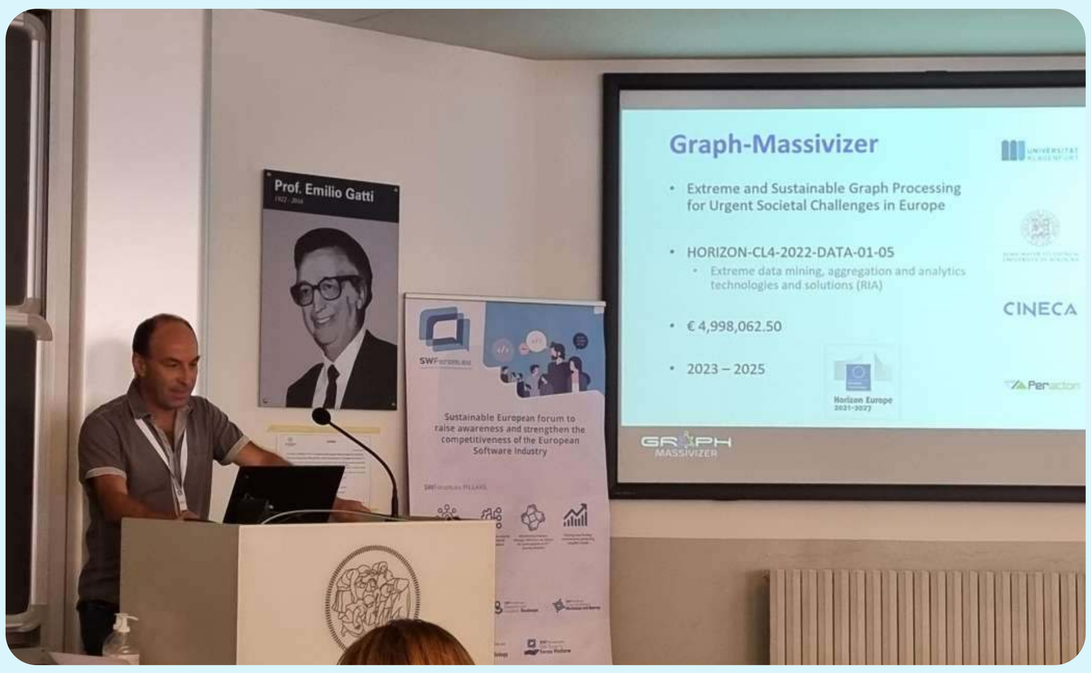

Session Agenda

- Extreme and Sustainable Graph Processing for Urgent Challenges in Europe – **1. Radu Prodan**, University of Klagenfurt
- Commonalities, Themes and Take-Home Messages – **2. David Wallom**, University of Oxford

Final Discussion

A pair of major themes emerged from the final discussion of the workshop among all participants, with implications for the way forward in software engineering. These themes were characterized by their appearance transversally across all of the sessions.

The first theme concerned **uncertainty about the true applications** of the new technologies. Almost without exception, in each of the sessions the participants observed significant levels of uncertainty about where the new, cutting-edge technologies will find their true use cases, even after a number of years of research and development. This was especially emphasized in the quantum software engineering session; but it was also highlighted in the sessions on artificial intelligence and cybersecurity technologies (as well as noting that sustainable software engineering is a very new area in any case). This observation has also been made in other new fields such as blockchain and distributed ledge technologies. A challenge going forward will be to teach students in software engineering to be able to seek out and recognize the appropriate applications for the new technologies they are being taught.

The second theme emerged early in the session on AI, where two very **siloed communities** in software engineering and machine learning were described, with the attendant difficulties and failures in collaboration. The theme emerged again in the quantum session, where the difficulty of integrating communities of researchers, physicists, and practitioners was cited, because each was concerned with its own perspective, each speaking about problems in the language of its own community. The new technologies discussed in this workshop are so advanced that they are leading to a phenomenon of deep specialization in one area to the detriment of other areas. How will integrated development approaches – DevSecOps, MLOps, QuantumOps, “EcoOps” – be possible when there is no shared base of cross-cutting expertise? This will be a challenge once again for educators, and was a motivator for the conception of the cross-fertilization workshops of SWForum.eu.

One conclusion agreed by all participants was that we have not reached “the end of software engineering” as a core discipline. We will need a full cycle of renewal and adaptation of consolidated practices going forward, to give rise to a whole new generation of software engineering paradigms, models, and best practices that will help us build the systems that will accommodate both the current technologies and those yet to come.

Resources

Explore the SWForum.eu insights and finding from

“The Way Forward: Workshop on Future Challenges in Software Engineering.”

Download the speakers' presentations on zenodo

and read the post-workshop news piece **here** to gain valuable resources to effectively address future challenges in software engineering.

Consortium

CONCEPTIVITY 360 SECURITY

UNIVERSITY OF OXFORD

POLITECNICO MILANO 1863

tecnalia Inspiring Business

Trust-IT Services communicating to markets

Thank you!

European forum of the software research community - SWForum.eu project has received funding from the European Union's Horizon 2020 - Research and Innovation program - under grant agreement no. 957044. The content of this report does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of such content.