



Carmine Colarusso

Nazionalità: Italiana

ESPERIENZA LAVORATIVA

– Benevento, Italia

Città: Benevento | Paese: Italia

Dottorando presso "Università degli studi del Sannio"

[01/01/2022 – Attuale]

L'attività di ricerca prevede come output la definizione di nuovi modelli architetture, metodi per l'assessment dei livelli di QoS (in particolare scalabilità e prestazioni) dei sistemi software nel continuo edge-cloud, piattaforme sperimentali per il dispiegamento di servizi resilienti e altamente scalabili a supporto delle infrastrutture di una smart city.

Università degli Studi del Sannio – Benevento, Italia

Città: Benevento | Paese: Italia

Assegnista di ricerca

[31/07/2020 – 31/12/2021]

Titolo del Progetto: Piattaforma per attrarre pubblico e diffondere lirica e teatro di prosa con tecnologie immersive - CANTICO

Settore scientifico disciplinare: ING-INF/05 – Sistemi di Elaborazione delle Informazioni.

Titolo della ricerca: Riconoscimento del contesto attraverso l'analisi delle sorgenti informative correlate agli spettacoli

Programma di ricerca: L'attività di ricerca è volta ad analizzare dati da diverse sorgenti informative connesse allo spettacolo in corso, con particolare riferimento alle sorgenti audio, al fine di identificare il contesto corrente da impiegare per la fruizione contestualizzata e personalizzata di contenuti correlati.

ISTRUZIONE E FORMAZIONE

Dottorato di Ricerca in Tecnologie dell'informazione per l'ingegneria

Università degli studi del Sannio [01/01/2022 – Attuale]

Città: Benevento | Paese: Italia | Sito web: <https://www.unisannio.it/it/didattica/corsi-di-studio/dottorato-di-ricerca/dottorato-tecnologie-informazione-ingegneria> | Campi di studio: Tecnologie dell'informazione e della comunicazione (TIC) | Livello EQF: Livello 8 EQF

Laurea Magistrale in Ingegneria Informatica - LM32

Università degli Studi del Sannio [2016 – 26/04/2020]

Città: Benevento | Paese: Italia | Sito web: <https://www.ding.unisannio.it/> | Campi di studio: Ingegneria e studi correlati | Voto finale: 110/110 e lode | Livello EQF: Livello 7 EQF | Tesi: An Application for Traffic Monitoring: Comparison of Microservices and Multi-Tier Approach

iOS Foundation Programm

Università degli Studi del Sannio [12/2018]

Città: Benevento | Paese: Italia | Sito web: <http://iosfoundation.unisannio.it/> | Campi di studio: Informatica

L'iOS Foundation Program, realizzato dall'Università del Sannio in collaborazione con Apple, è un programma accademico per l'offerta di corsi finalizzati alla creazione di App iOS potenzialmente valide per la distribuzione tramite l'App Store di Apple.

Laurea Triennale in Ingegneria Informatica - L8

Università degli Studi del Sannio [2011 – 05/10/2017]

Città: Benevento | Paese: Italia | Sito web: <https://www.ding.unisannio.it/> | Campi di studio: Tecnologie dell'informazione e della comunicazione (TIC) | Voto finale: 97/110 | Livello EQF: Livello 6 EQF | Tesi: Esperienza di sviluppo applicativo nel framework Akka

Diploma Secondario

Liceo Classico "Rosario Livatino" [2012]

Città: San Marco dei Cavoti | Paese: Italia | Livello EQF: Livello 5 EQF

COMPETENZE LINGUISTICHE

Lingua madre: italiano

Altre lingue:

inglese

ASCOLTO C1 LETTURA C1 SCRITTURA C1

PRODUZIONE ORALE C1 INTERAZIONE ORALE C1

francese

ASCOLTO B1 LETTURA B1 SCRITTURA B1

PRODUZIONE ORALE B1 INTERAZIONE ORALE B1

Livelli: A1 e A2: Livello elementare B1 e B2: Livello intermedio C1 e C2: Livello avanzato

COMPETENZE

Sviluppo di sistemi IoT / Edge e cloud / Gestione e manutenzione Server e armadi rack / IaaS PaaS SaaS Development / Web services & Web API / Containerizzazione: Docker, Kubernetes / Conoscenza sistemi operativi(Windows, Linux) / Openstack / OpenTracing (Jaeger Tracing) / Application server: Tomcat, Jboss, Wildfly / Database SQL e NoSQL / Neo4j graphDB / Microsoft Office / Ottima conoscenza dei sistemi di pipeline di CI/ CD in GitHub e GitLab / Git, GitHub, GitLab / Programming languages:C,C++,Java,Python

ONORIFICENZE E RICONOSCIMENTI

Riconoscimenti e premi

CINI Smart Cities University Challenge 2019

La CINI Smart City University Challenge 2019 è la terza edizione dell'University challenge organizzata da CINI Smart Cities e Communities Lab in collaborazione con l'Università degli Studi del Sannio (Benevento).

I domini applicativi sono:

- Conservazione e fruizione del patrimonio culturale e turistico
- Smart government e open data
- Mobilità
- Smart utilities e gestione delle risorse
- Sicurezza urbana
- Analisi delle Smart City
- Smart living
- Sostenibilità ambientale

Classificazione in graduatoria: Primo

www.consorzio-cini.it/index.php/it/component/attachments/download/946

Data: 20/09/2019

PUBBLICAZIONI

[2020]

[Domain Metric Driven Decomposition of Data-Intensive Applications](#)

Autori: Matteo Camilli, Carmine Colarusso, Barbara Russo, Eugenio Zimeo.

The microservices architectural style is picking up more and more momentum in IT industry for the development of systems as loosely coupled, collaborating services. Companies that undergo the migration of their own applications have aspirations such as increasing maintainability and the scale of operation. Such a process is worthwhile but not easy, since it should ensure atomic improvements to the overall architecture for each migration step. Furthermore, the systematic evaluation of migration steps becomes cumbersome without sensible optimization metrics that take into account performance and scalability under expected operational conditions. Recent lines of research recognize this task as challenging, especially in data-intensive applications where known approaches based, for instance, on Domain Driven Design may not be adequate. In this paper, we introduce an approach to evaluate a migration in an iterative way and recognize whether it represents an improvement in terms of performance and scalability. The approach leverages a Domain Metric-based analysis to quantitatively evaluate alternative architectures. We exemplified the envisioned approach on a data-intensive application case study in the domain of smart mobility. Preliminary results from our controlled experiments show the effectiveness of our approach to support systematic and automated evaluation of migration processes.

Conference, <https://doi.org/10.1109/ISSREW51248.2020.00071>

[2022]

PROMENADE: A big data platform for handling city complex networks with dynamic graphs

Autori: Carmine Colarusso, Antonio De lasio, Angelo Furno, Lorenzo Goglia, Mohammed Amine Merzoug, Eugenio Zimeo.

Continuous data streams, generated by modern sensed cities, open many opportunities and perspectives in terms of developing new innovative services. To exploit this potential, flexible and scalable platforms are needed to ease the design, development, deployment, and operations of new city services. In recent years, several problem-specific platforms have been proposed in different application domains; however, to boost the evolution of smart cities, we claim the need for city-oriented platforms that can be easily customized to address different day-to-day life challenging problems. In this paper, we present the main architectural challenges and solutions proposed for the design of a novel open-source platform (named PROMENADE) characterized by: *i*) a data-driven graph-based modeling support to ensure high generality for addressing disparate problems related to the networked nature of many city infrastructures and systems, *ii*) the dynamic nature of the graph entities updated in real-time from different sources (*e.g.*, IoT/Edge networks, data providers, etc.), and *iii*) high efficiency, scalability and flexibility to easily support new city services. The platform is designed around a general-purpose core that provides a set of built-in standard features such as data ingestion, storage, processing, and visualization exposed as a collection of containerized microservices. A specialization of the platform has been developed for road networks monitoring. It has been deployed in OpenShift/Kubernetes and tested using realistic datasets collected from the city of Lyon, France. The analysis addresses an important problem of big data processing pipelines: the synchronization between data ingestion and processing in order to produce an accurate result in useful time. To this end, we study different approaches for synchronization and show how the end-to-end latency is kept under control by leveraging the scalability of the platform.

[2023]

Actor-driven Decomposition of Microservices through Multi-level Scalability Assessment

Autori: Matteo Camilli, Carmine Colarusso, Barbara Russo, Eugenio Zimeo.

The microservices architectural style has gained widespread acceptance but designing applications according to this style is still challenging. Common difficulties concern finding clear boundaries that guide decomposition while ensuring performance and scalability. With the aim of providing software architects and engineers with a systematic methodology, we introduce a novel actor-driven decomposition strategy to complement the domain-driven design and overcome some of its limitations by reaching a finer modularization yet enforcing performance and scalability improvements. The methodology uses a multi-level scalability assessment framework that supports decision making over iterative steps. At each iteration, architecture alternatives are quantitatively evaluated at multiple granularity levels. The assessment helps architects to understand the extent to which architecture alternatives increase or

decrease performance and scalability. We applied the methodology to drive further decomposition of the core microservices of a real data-intensive smart mobility application and an existing open-source benchmark in the e-commerce domain. The results of an in-depth evaluation show that the approach can effectively support engineers in: (i) decomposing monoliths or coarse grained microservices into more scalable microservices, and (ii) comparing among alternative architectures to guide decision making for their deployment in modern infrastructures that orchestrate lightweight virtualized execution units.

[2024]

A distributed tracing pipeline for improving locality awareness of microservices applications

Autori: Carmine Colarusso, Assunta De Caro, Ida Falco, Lorenzo Goglia, Eugenio Zimeo.

The microservices architectural style aims at improving software maintenance and scalability by decomposing applications into independently deployable components. A common criticism about this style is the risk of increasing response times due to communication, especially with very granular entities. Locality-aware placement of microservices onto the underlying hardware can contribute to keeping response times low. However, the complex graphs of invocations originating from users' calls largely depend on the specific workload (e.g., the length of an invocation chain could depend on the input parameters). Therefore, many existing approaches are not suitable for modern infrastructures where application components can be dynamically redeployed to take into account user expectations. This paper contributes to overcoming the limitations of static or off-line techniques by presenting a big data pipeline to dynamically collect tracing data from running applications that are used to identify a given number k of microservices groups whose deployment allows keeping low the response times of the most critical operations under a defined workload. The results, obtained in different working conditions and with different infrastructure configurations, are presented and discussed to draw the main considerations about the general problem of defining boundary, granularity, and optimal placement of microservices on the underlying execution environment. In particular, they show that knowing how a specific workload impacts the constituent microservices of an application, helps achieve better performance, by effectively lowering response time (e.g., up to a 61% reduction), through the exploitation of locality-driven clustering strategies for deploying groups of services.

[2024]

Towards business continuity with Edge-Cloud continuum

Autori: Carmine Colarusso, Ida Falco, Eugenio Zimeo.

With the large volume of data produced by IoT applications, a pure Cloud model could not satisfy application requirements. Edge-Cloud continuum is a viable solution since it enables data processing closer to the source, leveraging the Cloud for global and large-size computations. This paper presents an Edge-Cloud continuum framework for adapting Cloud-based IoT applications to reduce latency and ensure business continuity even in case of network interruptions. The framework is built to be highly configurable to ensure a seamless transition to Edge-Cloud, guaranteeing minimal effort. The approach has been evaluated using an IoT application that simulates an intensive tracking activity from multiple retail shops geographically distributed. The experimental results show the framework's effectiveness in ensuring business continuity and lower response times even in case of network failure.

<https://ieeexplore.ieee.org/abstract/document/10743024>

[2024]

A greedy data-anchored placement of microservices in federated clouds

Autori: Carmine Colarusso, Ida Falco, Eugenio Zimeo.

In a multiple cloud environment, the placement of execution environments is crucial and may be subject to data constraints. Data could be anchored to some environments due to regulatory compliance, data sovereignty issues, or performance optimization. Consequently, applications and microservices must be designed to operate efficiently with these constraints. This enforces specific placement strategies to obtain good performances and scalability. This paper proposes a technique to enforce a constrained data-centric deployment placement in a federated multi-cloud

environment. The algorithm analyzes a graph model of microservices interaction, considering communication with data storage and adopting “anchors” for implementing a data-centric placement strategy. The results show how a better placement based on data position in multiple clouds improves performance in terms of overall system response time. This also allows microservices to offload near data sources for multi-cloud environments, improving overall system performance without violating data movement constraints.

Link: <https://ieeexplore.ieee.org/document/10838374>

COMPETENZE ORGANIZZATIVE

Competenze organizzative e gestionali

Buona capacità di gestione dei progetti e dei gruppi, capacità di leadership e senso dell'organizzazione, ottenuti grazie all'esperienza universitaria. In particolare i numerosi progetti in gruppo hanno fornito anche empatia e flessibilità, capacità di organizzazione e gestione del tempo, pianificazione e gestione dei progetti, rispetto delle tempistiche, buona resistenza allo stress.

COMPETENZE COMUNICATIVE E INTERPERSONALI

Competenze comunicative e interpersonali.

Spirito di gruppo (derivante da una serie di progetti portati a termine in ambito universitario)

Buona capacità di comunicazione

Autorizzo il trattamento dei miei dati personali presenti nel CV ai sensi dell'art. 13 d. lgs. 30 giugno 2003 n. 196 - “Codice in materia di protezione dei dati personali” e dell'art. 13 GDPR 679/16 - “Regolamento europeo sulla protezione dei dati personali”.

12/2024