Via Giovanni Sebastiano Bach 38
Perugia, 06132
075/5853920 – 349/0575782
ingarleo@icloud.com

ALESSIO ARLEO

CURRENT POSITION

PHD student at University of Perugia in Computer and Industrial Engineering

I am studying the problem of 'Big Data' analytics. Given the size of those datasets, the traditional single machine computing approach is obsolete, giving way for new techniques based on parallel programming models. Developing such software, however, is a new challenge to tackle with knowledge and creativity.

My work focuses on Visual Data Analytics for 'Big' datasets using cloud computing, leveraging the computing power of frameworks such as Hadoop and Spark. I am also working on theoretical problems related to the visualisation of big networks.

CONFERENCES

Speaker at the 24th International Symposium on Graph Drawing and Network Visualization

September 19th – 21st 2016, *Athens*, Greece

Graph Drawing is concerned with the geometric representation of graphs and constitutes the algorithmic core of Network Visualization. Graph Drawing and Network Visualization are motivated by applications where it is crucial to visually analyze and interact with relational datasets. The International Symposium on Graph Drawing has been the main annual event in this area for more than twenty years. This year the Steering Committee of GD decided to extend the name of the conference from the "International Symposium on Graph Drawing" to the "International Symposium on Graph Drawing and Network Visualization" in order to better emphasize the dual focus of the conference on combinatorial and algorithmic aspects as well as the design of network visualization systems and interfaces. At this venue I presented my paper "*A distributed Multi-Level force directed algorithm*".

Speaker at the 23rd International Symposium on Graph Drawing and Network Visualization

September 24th – 26th 2015, *California State University*, Los Angeles (USA)

At this venue I presented my paper "*A million edge drawing for a fistful of dollars*".

Speaker at "*The Smart Tourism 2014*" Conference

September 23rd, *Tor Vergata University*, Villa Mondragone (IT)

The topic of the conference was "*Technologies for smart tourism in the Smart Cities era*". During my intervention, I presented the TRART framework and my master thesis. TRART is a project of the IT department of the University of Perugia whose goal is to give tourist an innovative service to enhance their sightseeing experience by providing them custom tours based on their preferences.

EDUCATION

Bertinoro Workshop on Graph Drawing

March 6th – 11th 2016

This is a workshop on themes of Graph Drawing. The idea is to have an informal meeting with a few selected people and work together on both identifying new research directions and on studying some of the questions that may arise. There may be space for presentations but the emphasis will be on working together "pencil-and-paper". Possible outcomes are a new collection of open problems as well as some original research papers

1-st *IEEE* Italy Section Summer School

June 21st – 28th 2015, *University of Perugia*, Perugia (IT)

The school consisted of lectures by international experts coming from industry, academia and other research centres. The School was aimed at providing a basic overview about the modern industrial approaches to research and development, including organization and people management. The School provided a great chance not only for learning new skills but also for networking and sharing experiences.

***BigDat* 2015 Winter School**

January 26th – 30th 2015, *Universitat Rovira i Virgili* , Tarragona (ES)

BigDat aimed at updating about the most recent developments in the fast developing area of big data, which covers a large spectrum of current exciting research, development and innovation with an extraordinary potential for a huge impact on scientific discoveries, medicine, engineering, business models, and society itself.

Computational Social Science Summer School

July 21st – 25th 2014, *University of Catania*, Lipari (IT)

Lipari 2014 school on Computational Social Science examined the intersection between parameters that define a city as smart and the enabling technology that guarantees proper planning and implementation.

Master Degree in Computer and Automation Engineering

2011 - 2014

Thesis title: "*Programming and development of an iOS app for AR navigation*".

Title achieved *cum laude*.

Bachelor Degree in Computer and Electronic Engineering

2008 – 201 1

Thesis title: "*Programming and development of a Vector Graphics Editor based on Java Swing Library*".

PUBLICATIONS

A. Arleo, W. Didimo, G. Liotta, F. Montecchiani

GiViP: A Visual Profiler for Distributed Graph Processing Systems,

To appear at *Graph Drawing 2017*

Analyzing large-scale graphs provides valuable insights in different application scenarios. While many graph processing systems working on top of distributed infrastructures have been proposed to deal with big graphs, the tasks of profiling and debugging their massive computations remain time consuming and error-prone. This paper presents GiViP, a visual profiler for distributed graph processing systems based on a Pregel-like computation model. GiViP captures the huge amount of messages exchanged throughout a computation and provides an in- teractive user interface for the visual analysis of the collected data. We show how to take advantage of GiViP to detect anomalies related to the computation and to the infrastructure, such as slow computing units, anomalous message patterns, and unbalanced graph partitions.

A. Arleo, Oh-Hyun Kwon, Kwan-Liu Ma

GraphRay: Distributed Pathfinder Network Scaling,

To appear at *Large Data Analysis and Visualization 2017*

Pathfinder network scaling is a graph sparsification technique that has been popularly used due to its efficacy of extracting the "important" structure of a graph. However, existing algorithms to compute the pathfinder network (PFNET) of a graph have prohibitively expensive time complexity for large graphs: $O(n^3)$ for the general case and $O(n^2 \log n)$ for a specific parameter setting, PFNET($r = \infty$, $q = n - 1$), which is considered in many applications. In this paper, we introduce the first distributed technique to compute the pathfinder network with the specific parameters ($r = \infty$ and $q = n - 1$) of a large graph with millions of edges. The results of our experiments show our technique is scalable; it efficiently utilizes a parallel distributed computing environment, reducing the running times as more processing units are added.

A. Arleo, W. Didimo, G. Liotta, F. Montecchiani

A Distributed Multilevel Force-directed Algorithm,

Presented at *Graph Drawing 2016*

The wide availability of powerful and inexpensive cloud computing services naturally motivates the study of distributed graph layout algorithms, able to scale to very large graphs. Indeed, companies are increasingly relying on the use of PaaS (Platform as a Service) infrastructures to process their big data, thus saving money to buy and maintain complex and expensive hardware. So far, only few examples of basic force-directed algorithms that work in a distributed environment have been described. Instead, the design of a distributed multilevel force-directed algorithm is a much more challenging task, not yet addressed. We present the first multilevel force-directed algorithm based on a distributed vertex-centric paradigm, and its implementation on Giraph, a popular platform for distributed graph algorithms. Experiments show the effectiveness and the scalability of the approach. Using an inexpensive cloud computing service of Amazon, we draw graphs with ten million edges in about 60 minutes.

Alessio Arleo, Carla Binucci, Emilio Di Giacomo, William Evans, Luca Grilli, Giuseppe Liotta, Henk Meijer, Fabrizio Montecchiani, Sue Whitesides And Stephen Wismath

Visibility Representations of Boxes in 2.5 Dimensions,

accepted at *Graph Drawing 2016*

We initiate the study of 2.5D box visibility representations (2.5D-BR) where vertices are mapped to 3D boxes having the bottom face in the plane $z = 0$ and edges are unobstructed lines of sight parallel to the x-axis and y-axis. We prove that: (i) Every complete bipartite graph admits a 2.5D-BR; (ii) The complete graph K_n admits a 2.5D-BR if and only if $n \leq 19$; (iii) Every graph with pathwidth at most 7 admits a 2.5D-BR, which can be computed in linear time. We then turn our attention to 2.5D grid box representations (2.5D-GBR) which are 2.5D-BRs such that the bottom face of every box is a unit square at integer coordinates. We show that an n-vertex graph that admits a 2.5D-GBR has at most $4n - 6\sqrt{n}$ edges and this bound is tight. Finally, we prove that deciding whether a given graph G admits a 2.5D-GBR with a given footprint is NP-complete. The footprint of a 2.5D-BR Gamma is the set of bottom faces of the boxes in Gamma.

A. Arleo, W. Didimo, G. Liotta, F. Montecchiani

A Distributed Force-Directed Algorithm on Giraph: Design and Experiments,

Published on *Information Sciences*

In this paper we study the problem of designing a distributed graph visualization algorithm for large graphs. The algorithm must be simple to implement and the computing infrastructure must not require major hardware or software investments. We design, implement, and experiment a force-directed algorithm in Giraph, a popular open source framework for distributed computing, based on a vertex-centric design paradigm. The algorithm is tested both on real and artificial graphs with up to million edges, by using a rather inexpensive PaaS (Platform as a Service) infrastructure of Amazon. The experiments show the scalability and effectiveness of our technique when compared to a centralized implementation of the same force-directed model. We show that graphs with about one million edges can be drawn in less than 8 minutes, by spending about 1\$ per drawing in the cloud computing infrastructure.

A. Arleo, W. Didimo, G. Liotta, F. Montecchiani

A Million Edge Drawing for a Fistful of Dollars,

accepted at *Graph Drawing 2015*

In this paper we study the problem of designing a graph drawing algorithm for large graphs. The algorithm must be simple to implement and the computing infrastructure should not require major hardware or software investments. We describe a simple implementation of a spring embedder in Giraph, a vertex-centric open source framework for distributed computing. The algorithm is tested on real small-world, scale-free graphs of up to 1 million edges by using a cheap PaaS (Platform as a Service) infrastructure of Amazon. We can afford drawing graphs with about one million edges in about 20 minutes, by spending less than 2 USD per drawing for the cloud computing infrastructure.

A. Arleo, F. De Luca, G. Liotta,F. Montecchiani, I. Tollis,

***GraphBook*: Making graph Paging Real,**

in *Graph Drawing*, Springer, 2014

GraphBook is a new and innovative way to explore and understand complex networks on mobile devices. By using specific algorithms, it allows the user to explore and comprehend the network topology, even on a small-sized display.

EXPERIENCE

Visiting PHD student

January 5th – April 05th 2017, *University of California*, Davis (CA)

During this experience I further expanded my knowledge of distributed computing and algorithms. I applied this computing paradigm to visualization problems studied at UC Davis, at prof. Kwan-Liu Ma's VIDI Lab.

Tutor at University of Perugia (Terni branch)

04/2017 – 06/2017

I assisted the professor in the teaching "Computer Science fundamentals" at University of Perugia branch in Terni.

Contract Job at University of Perugia

07/2014 – 09/2014

Development of a graphical user interface for the visualisation of geo-tagged relational data. The task is part of the project: “Development of models and algorithms to represent information about the driving style on a large user database”.

Contract Job at University of Perugia

06/2014 – 09/2014

Feasibility study for developing a tourist and travel information website for Umbria Regional Administration integrating innovative services for tourists.

Internship at University of Perugia

01/2013 – 04/2013

Development of a Web Client to interact with the server hosting *TRART* algorithm software implementation, using *JavaScript* and *Ajax*.

Web Developer at Nuova Dimensione Medica (NDM)

09/2012 – 03/2013

Development of an e-learning platform.

I hereby authorize the processing of the personal data contained in this CV in compliance with the Italian Personal Data Protection Code (Legislative Decree no. 196 of 30 June 2003).

Perugia, August 21st 2017

(Alessio Arleo)