



A programmable web platform for distributed access, analysis, and visualization of data

F. Esquembre ^{a,*}, J. Chacón ^b, J. Saenz ^c, J. Vega ^d, S. Dormido-Canto ^c

^a Department of Mathematics, Universidad de Murcia, Facultad de Matemáticas, Murcia, 30071, Spain

^b Departamento de Arquitectura de Computadores y Automática, Universidad Complutense, Plaza de Ciencias 1, Madrid, 28040, Spain

^c Departamento de Informática y Automática, Universidad Nacional de Educación a Distancia, Juan Del Rosal 16, Madrid, 28040, Spain

^d Laboratorio Nacional de Fusión, CIEMAT, Avenida Complutense 40, Madrid, 28040, Spain

ARTICLE INFO

Keywords:

Collaborative environment
Distributed system
High-performance computing
Machine learning
Nuclear fusion

ABSTRACT

Daily work of Fusion Data Research (FDR) scientists faces three practical challenges: (i) getting access to vast amounts of validated, curated, and (ideally) annotated discharge data, (ii) applying a wide variety of standard, domain-specific, and home-made analysis and visualization software libraries and routines, and (iii) using fast, specialized, and not easy to obtain hardware and software installations. This paper introduces a novel web platform that addresses these three challenges in a federated way. Based on a client-server architecture, the new platform allows for easy use and exchange of curated data, validated analysis and visualization routines, and even networked hardware and software installations among the FDR community. This exchange goes beyond the mere use of a code repository, but facilitates the creation of an actual ready-to-use network of computers which can be used remotely to configure and perform data analysis. The network functions in a federated way, in which each member of the community contributes, using the same web platform, with its data, programming experience, and hardware and software availability. The platform is open source.

1. Introduction

Experimental programs of nuclear fusion devices aim for the characterization of plasma properties. Typically, plasma parameters have to be derived indirectly, which requires the solution of difficult inversion problems. In general, all raw data are stored for off-line analysis, which causes a first important concern in data analysis: the vast amount of data. (As an example, ITER will collect about 2 PBytes of data per day [1].) However, dealing with massive databases is not the only problem in the analysis of discharges. Other issues to be taken into account are: data quality (poor calibration methods, approximate error bars, or insufficient validation), statistical relevance (insufficient and biased databases), image processing (massive quantities of data, difficult interpretation), predictions (complex systems, fast instabilities, and non-linear interactions), diagnostic design and experimental planning (principled design criteria or support to simulations), and adoption of new technologies (FPGAs, GPUs, intelligent data retrieval, and data mining, among others).

The discharge analysis can be seen as made up of two sequential steps. Firstly, automatic data processing (that usually does not require human intervention) is accomplished to carry out data reduction. The objective of this data reduction is to prepare the data for human

inspection. Secondly, specialized software codes are launched to help scientists to obtain results. For the purpose of this article, all the analysis involved in the above description will be referred to as Fusion Data Research (FDR). The daily research activity of FDR scientists faces a number of practical challenges.

In the first place, in addition to having a lot of data, in some cases, these are not open to the public or not downloadable for security or practical reasons. Getting access to experimental data – usually stored at a central, firewall-protected facility – involves obtaining authorization and either a very fast intranet connection or huge hard disk availability. Moreover, data should be validated, curated, and (ideally) annotated. Current experimental data is typically found in raw form, even if it has been already cleaned or processed by other research groups as part of their work. Even more, intermediate research results, produced through the application of Machine Learning (ML) techniques [2–8], may be also of interest to the community, but are typically only accessible to the research group that created it. This limits the FDR scientists' access to provenance-controlled cleaner, annotated data and train/test/validate subsets, that would facilitate both future work and the reproducibility of published research by other members of the community.

* Corresponding author.

E-mail address: fem@um.es (F. Esquembre).

<https://doi.org/10.1016/j.fusengdes.2023.114049>

Received 1 March 2023; Received in revised form 26 September 2023; Accepted 20 October 2023

Available online 26 October 2023

0920-3796/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

Second, FDR analysis typically use some (actually, many) standard data analysis and visualization software tools, together with domain specific libraries and user-created experimental routines. Both novel and senior scientists need to be fluent in the so many varied software libraries available (which may be written in different programming languages, be updated or expanded every so often, and have a not very up-to-date documentation) to wisely combine standard routines with perhaps their own code to create full or partial data workflows... which may have been built previously by, or be of future interest for other FDR scientists. Lack of availability to reusable and well-tested (even publication-supported) data-processing pipelines, trained ML models, or sets of metadata, sometimes results in the community spending unnecessary time repeating coding or computations.

Finally, the large amount of data to process or the complexity of the algorithms may result in the need for high computational power which requires, perhaps, parallel computation or dedicated hardware (such as GPU or FPGA) [9,10], or the installation of particular software platforms (such as some ML libraries or proprietary software). Reproducing other scientists' work or applying similar techniques to new data would then imply every member of the community having access to every possible such installation. But accessing, installing, and maintaining such frameworks – which can in many cases be shared through a network – can be a very hard, expensive, and time-consuming activity in itself.

These three challenges make research results in FDR laborious to obtain, hard to communicate, and – a key point – particularly difficult to reproduce by other scientists. There is a need for the FDR community to either agree on a common software or platform for the development of the discipline or, better yet, establish a sort of open and *federated* way for scientists to work in this field. A way that simplifies access to annotated data, fine-trained models, and the latest validated analysis software and hardware platforms, offers freedom to choose from the many that exist or to use custom ones and facilitates communication, reproducibility, and replicability of results.

This paper presents IODA (acronym for Input-Output Data Analysis), a novel Web platform created specifically with this need in mind. The platform is based on a client-server architecture with the clients running on any standard Web-enabled device (including PCs, laptops, and tablets) and allowing simple, graphical edition of data workflows, and with a server that offers users access to a whole federated ecosystem providing:

- simplified, secure access to remote, distributed data,
- verified software routines for analysis and visualization, in different programming languages,
- access to network-available specific computing hardware and software platforms,
- the capability to introduce the user's own code for particular analysis, and
- the possibility for members of the community to easily contribute to the ecosystem with new elements.

The IODA platform can serve the FDR community in its need to share access to raw or curated data, publish provenance controlled data-workflows and models, and even share software or hardware resources for members who may need it, in a controlled way.

Other software tools (such as Orange [11], Apache NiFi [12], and Simulink [13], among others) provide users with the capability of designing data workflows in simplified, graphical form. But none of them provides the advanced, federated capabilities offered by IODA.

The client-side user interface of the platform is described in Section 2. The server-side, which provides the real engine of the federated data analysis capabilities is described in Section 3. Finally, Section 4 discusses design decisions and the future implementation of platform components that will address the different needs described above.

Readers and potential users interested to learn in more detail about some technical aspects (installation, adding new servers, security issues, federation management...) can also consult the IODA web site [14].

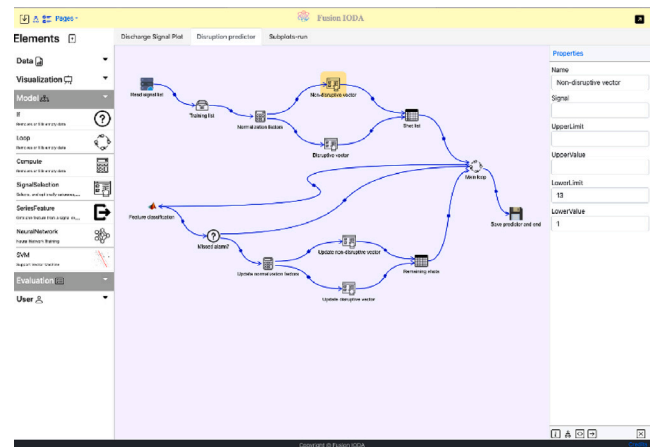


Fig. 1. IODA Web-based graphical user interface running on a tablet and showing a directed graph for data analysis and visualization.

2. IODA user interface

The platform is operated by users using any modern Web browser, since the client is programmed using W3C standards: HTML, JavaScript, and CSS [15]. Therefore, virtually any Internet-enabled device can display this Web-based graphical user interface (WebUI), from PCs to tablets and cell phones (but a small screen makes it harder to read). The server authenticates the user and serves the client web page using the HTTPS standard. The WebUI allows users to create one or more projects, each of one of the server-provided dedicated types. The server can customize the types of projects it offers to each user based on user's preferences or authorization level. Projects consist of one or several pages of data analysis and visualization workflows expressed in form of directed graphs (see Fig. 1).

A directed graph is composed of nodes and oriented connections among them. Nodes are displayed in the client's interface shown in Fig. 1 as labeled icons, and represent instances of library elements that the project type offers to the user for its perusal. Each element represents a given piece of code with data algorithms or calls to visualization routines, and has been created or curated by other scientists, experts in the different tasks, and contributed to the platform. Elements are added to a graph and their *properties* are customized by the user using IODA's WebUI. Properties turn into code parameters that help fine-tune the behavior of the element, or represent input or output of the element's code. Users then draw directed arrows that establish connections and send output properties of an origin node to be used as input properties for a target node. This information is sent, during run-time, using standard JSON format [16] and shared temporary data files, allowing for elements written in different programming languages to exchange data. The server will, in run mode, convert these nodes and connections into a valid executable.

Elements are the key building blocks of the ecosystem and the WebUI standardizes the way they are created and used. Projects offer a menu with common, validated elements for most of the tasks required by a project of a given type. For FDR, these include accessing data in remote repositories, data processing or analysis algorithms, trained and well-tested machine learning models, and visualization using specialized libraries. These domain-specific elements are added to the ecosystem in a federated way: members of the FDR community can contribute with new elements according to their programming expertise, access to data, or availability of hardware or software platforms. Besides the elements offered by each project type, users can create new elements using their own algorithms and add them to those used in their graphs. Servers can analyze the code before actually running it to check for unauthorized access or incorrect uses. The process of using,

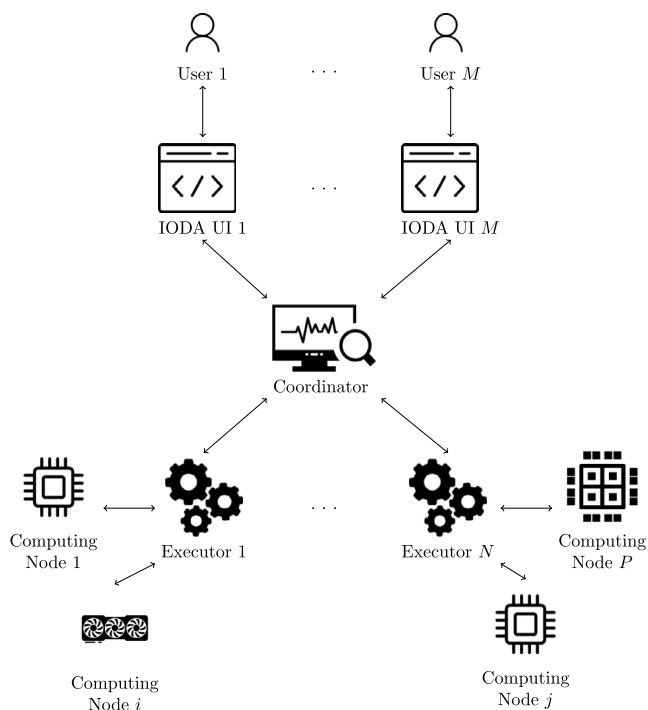


Fig. 2. Big picture of IODA.

creating, and connecting elements to create an effective directed graph is very natural, promotes open source practices (users can inspect the elements' code), provides in-place documentation — which can include links to publications, and makes possible a fast learning curve.

Once a graph is created, the client can automatically check the correctness of the nodes and connections, and offers a simulated step-by-step execution, so that the user can see what nodes will be executed and in which order. Finally, a single click sends all the information to the server for the actual execution. The server then takes charge and serves back, as the graph is executed, all the information that is produced, including which executor (see below) run which node, the hardware or software platform used, the time elapsed by each node, and any intermediate or final outcome produced by the execution. This information is provided in form of HTML pages that the client displays to the user. The user can then replay the execution, inspecting inputs and outputs of each step, which facilitates the detailed inspection, detection of bottlenecks or errors, and download of intermediate and final results of the whole workflow.

3. Server-side architecture

The server-side architecture, depicted in Fig. 2, is based on three conceptually different components: one *coordinator*, that serves directly the users and initiates and monitors all server-side tasks, one or more *executors* that execute the directed graphs, and one or many *computing nodes* that perform the actual computing. The three components may be running on a single machine (for small tasks) or escalate to a cloud of federated computers connected by an intranet or, in some cases, the Internet. Although quite sophisticated, the client WebUI makes the complexity of the platform transparent to the user.

Users are given, by the coordinator server machine, a profile and a disk quota where the platform stores their projects, user-created private elements, and all results created by the executions of their graphs. Users can also upload their data to this disk for use in their analysis.

Once the user requests the execution of a graph from the WebUI, the client creates a deep copy of the task, packaging the directed graph with any user-uploaded data, current definitions of the elements, and

any other information relevant for the execution. (Since this information could later be changed by their authors, this deep copy ensures reproducibility.) The life cycle of the task begins when this package is sent to the coordinator, which is responsible of (1) finding an executor to process the task, (2) monitoring the running state, and (3) serving back to the client all intermediate and final output results.

The executor relieves the coordinator from dividing the task into steps, finding computing machines for each of them, and driving the execution of the graph, passing information and data to and from computing nodes as the execution proceeds. The coordinator server can hence serve several clients simultaneously and these tasks will not affect the performance of the execution of a graph. How the coordinator chooses one executor from the possibly many available depends on its capabilities or availability. It is important to remark that the executors may be heterogeneous machines and therefore provide different computing possibilities. Some executors may have access to computing nodes with installed specific software platforms or dedicated hardware (GPU or FPGA), or have local (fast) access to a data repository. Others may be standard computational nodes or grids of them, suitable for parallel computation. Since elements may include in their definition run-time requirements or recommendations, the execution may demand a given configuration from the executor. After the executor is assigned, the coordinator prepares an execution environment, extracting the information of the task in a distributed file system which the executor can access. The coordinator then requests from the executor the start of the graph execution.

When requested, the executor proceeds to process the nodes in turn, following connections as indicated by the directed graph. The executor then uses the asynchronous task queue manager *Celery* [17] to communicate with its available computing nodes to prepare (schedule) and synchronize the execution, and passing along any intermediate data that would be required. Each node of the directed graph is mapped to a Celery task that is, in turn, provided by one of the computing nodes. The data exchange between nodes is done using different backends, from passed JSON structures for small variables, to access details from a temporary shared file system or database for larger data. This shared file system is provided as a disk space reserved by the coordinator exclusively for the execution of the graph. This disk space is then used by the executor and computing nodes in turn as a common space where to read and write intermediate data needed for the execution. The executor coordinates this access so that only computing nodes involved in a particular graph execution access this data, and only when they need to. Once the graph execution is completed, access to the file system is closed.

The computing nodes are the server-side components that perform the actual processing. They receive the information on a given node processing code, customized properties, and inputs received and create and execute a computer process. They then send back the node's output and run-time information to the executor so that the workflow proceeds. The term computing nodes apply not only to number-crunching machines. Some nodes may act as *data providers*, smart repositories of experimental data both storing raw numbers and offering data curation services to the users. These data-providers access their organization data, which is *not* shared directly with the rest of the federation, and optionally preprocess the data to respond to the user's request, expressed via the provided element's properties. This ensures that the platform respects the access policy of each member organization to its data. Others nodes offer direct access to acceleration hardware or specialized software [18].

Input and output data required and produced, respectively, by each computing node is passed among graph elements via elements' properties links. Each element must provide output properties for any information required by further processing of the data by other elements (such as time stamps, metadata, or results of any preprocessing done). Since elements may run in different programming languages, data produced by these languages must be converted to a set of standardized property

types used to declare each element property type (this standardization also helps IODA's WebGUI check that links connect compatible data types). Computing nodes automatically run ancillary code, before and after running the code defined by the element, which cares for this conversion, in a way transparent to the element's designer. Each implementation of a computing node, using a given programming language, cares for converting these types into the programming language types that can be operated equivalently. Current data types include booleans, integer and double-precision numbers, strings, one and two-dimensional matrices, structures, and others...providing support for typical data-exchange needs. This list can be extended to accommodate more types in the future, if it becomes necessary.

4. Design decisions and status

The initial requirements for the IODA platform included the need to:

- Provide secure, remote access to large amounts of data and computing power hosted in institutional data servers.
- Allow users to operate the client part in any modern device (computers, tablets, even smartphones).
- Be intuitive to use and provide immediate on-line help, making it easy to use, even for sophisticated data analysis.
- Run in a distributed environment, profiting from remote hosting and computing power.
- Allow integration of heterogeneous data and user provided data analysis codes in different programming languages.
- Support the building of a community of scientists which contribute and use the platform in a federated form.

The decision to use a client-server architecture which communicates using Internet standards was an obvious one. It allows separating the implementation and maintenance of the client and server sides, using complete different technologies. The use of W3C standard software on the client side allows nearly universal deployment with no installation, permits using common tools for modern, easy to apprehend user interfaces, and offers resorting to the Internet for extended documentation and links to publications.

The decision to configure the user experience in form of projects with pages of directed graph for data analysis was based on the daily evidence that scientist very frequently use graphs when they explain to each other their data analysis workflows [19]. In fact, several other applications use directed graphs to build a complete algorithm dedicated to one or more of the following jobs:

- Big data analysis: Usually allow users to access, modify, filter, and visualize data [20]. Most of these tools include machine learning and data visualization capabilities (KNIME, Orchest, Easy Data Transform, Orange).
- Handling data streams: The main purpose of these tools is to build data pipelines using a graphical interface. The pipelines are created to automate a data flow between systems, where different data sources, functions, and data sinks can be used [21]. Some examples are Node-Red, Storm, Apache NiFi or StreamSets.
- Visual programming and hardware interaction: Allow the users to focus on the analysis and the data exploration instead of coding. Widgets, nodes, and pipes allow the creation of workflows as easy-to-understand directed graphs [22]. A few examples of this type of software are LabVIEW, Simulink, or Orange.

Though some of these tools satisfy some of the requirements for the platform, none covered all of them.

On the server-side the well-know Python-based Django server [23] was used to implement the coordinator. The platform for sharing disk space is currently based on good-old NFS, while support for other solutions is open. Executors are also developed in Python, implementing

both communication with the Django coordinator and the computation nodes using Celery. Celery was chosen as task queue manager because (i) it is fast, relatively simple, and very flexible, (ii) has a large supporting community and is in active development, and (iii) it is interoperable with practically every technology that could be needed by the IODA platform. The message broker RabbitMQ (RabbitMQ 2023) was chosen as the underlying communication mechanism for Celery, because of its wide range of uses, open source nature, and embedded security mechanism.

The implementation of each computation node is open. Although the project includes several prototypes (in Python) that each member of the community can adapt, any contributor can design its own computation node using different platforms, as long as it complies with the Celery, NFS, and data type conversion requirements. An overall design principle of IODA is to encapsulate the computing in individual tasks, each with only one specific purpose, that will be exposed to the executors using the Celery backend. The rationale for this decision is to allow for an extensive list of individual reusable tasks that can be useful for different users, to discourage the dependency on specific machines, and to facilitate the contribution by members of the community using their own available platforms and skills. For those with less experience, easily customizable, ready to use instances of computing nodes are distributed using Docker containers [24]. The complete platform is open source.

Default installation files, in form of Docker-enabled packages, for coordinator, executor, and computing node servers are provided. General hardware requirements for the three servers are limited to running the Docker application. The coordinator is the central piece of a IODA federation and requires customization to set the IP address and user/password security of the RabbitMQ communication system, set NFS' file system access permissions, and add new users' profiles (using standard Django administrative tools). Once done, the standard Docker *compose* command will automatically download and install all required software and run the server. Executor and computing node servers need to be customized to comply with the access security provided by the coordinator administrator and are then installed and run similarly. These servers communicate with the coordinator to register themselves in the federation at start-up. Administrators of computing nodes can, optionally, customize the distribution to add extra libraries that only run in their nodes; for instance, to provide access to their data, or run specialized software or hardware, which needs to be installed separately. (As an example, a Matlab computing node needs to install a valid licensed copy of Matlab.) These procedures are described in detail in the IODA web site [14].

The IODA platform is currently in its beta-testing period and is being used to host a small experimental community. This community is building a wide variety of commonly used data analysis and visualization elements, and installing executors and computing nodes using different approaches and offering a wide range of services. The goal of this final phase is to illustrate with examples the benefits of the contribution of federated nodes with diverse computing resources or access to data repositories. But, the most important objective is to show how a large group of data analysis scientists, such as the FDR community, can use the IODA platform to share their resources and know-how.

The IODA platform can also be of interest for communities in other scientific domains which need to share access to data and run different types of code to study it. The flexibility and openness of the process of adapting and creating specialized graph elements, using routines particular of each domain, facilitates adoption by scientists in other fields looking for a similar federated platform.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgments

This work was supported in part by the Spanish Ministry of Science and Innovation under Projects PID2019-108377RB-C31, PID2019-108377RB-C32, PID2022-137680OB-C31, and PID2022-137680OB-C32.

References

- [1] ITER, ITER web page, 2023, URL: <http://www.iter.org>. (Accessed 28 February 2023).
- [2] G. Farias, S. Dormido-Canto, J. Vega, I. Martinez, L. Alfaro, F. Martinez, Adaboost classification of TJ-II Thomson Scattering Images, *Fus. Eng. Des.* 123 (2017) 759–763, <http://dx.doi.org/10.1016/j.fusengdes.2017.05.042f>, 29th Symposium on Fusion Technology (SOFT), Prague, CZECH REPUBLIC, SEP 05-09, 2016.
- [3] J. Vega, F. Hernandez, S. Dormido-Canto, A. Isayama, E. Joffrin, G. Matsunaga, T. Suzuki, Assessment of linear disruption predictors using JT-60U data, *Fus. Eng. Des.* 146 (A, SI) (2019) 1291–1294, <http://dx.doi.org/10.1016/j.fusengdes.2019.02.061>, 30th Symposium on Fusion Technology (SOFT), Messina, ITALY, SEP 16-21, 2018.
- [4] D. Humphreys, A. Kupresanin, M. Boyer, J. Canik, C. Chang, E. Cyr, R. Granetz, J. Hittinger, E. Kolemen, E. Lawrence, V. Pascucci, A. Patra, D. Schissel, Advancing fusion with machine learning research needs workshop report, *J. Fus. Energy* 39 (2020) 123–155, <http://dx.doi.org/10.1007/s10894-020-00258-1>.
- [5] G. Farias, E. Fabregas, S. Dormido-Canto, J. Vega, S. Vergara, Automatic recognition of anomalous patterns in discharges by recurrent neural networks, *Fusion Eng. Des.* 154 (2020) <http://dx.doi.org/10.1016/j.fusengdes.2020.111495>.
- [6] J. Vega, A. Murari, S. Dormido-Canto, G.A. Ratta, M. Gelfusa, J. Contributors, Disruption prediction with artificial intelligence techniques in tokamak plasmas, *Nat. Phys.* 18 (7) (2022) 741–750, <http://dx.doi.org/10.1038/s41567-022-01602-2>.
- [7] J. Degraeve, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. Casas, C. Donner, L. Fritz, C. Galperti, A. Huber, J. Keeling, M. Tsimpanoukelli, J. Kay, A. Merle, J.-M. Moret, M. Riedmiller, Magnetic control of tokamak plasmas through deep reinforcement learning, *Nature* 602 (2022) 414–419, <http://dx.doi.org/10.1038/s41586-021-04301-9>.
- [8] W. Zheng, F. Xue, C. Shen, Y. Zhong, X. Ai, Z. Chen, Y. Ding, M. Zhang, Z. Yang, N. Wang, Z. Zhang, J. Dong, C. Tang, Y. Pan, Overview of machine learning applications in fusion plasma experiments on J-TEXT tokamak, *Plasma Sci. Technol.* 24 (12) (2022) 124003, <http://dx.doi.org/10.1088/2058-6272/ac9e46>.
- [9] M. Astrain, M. Ruiz, A. Carpeno, S. Esquembri, D. Rivilla, Development of deep learning applications in FPGA-based fusion diagnostics using IRIO-OpenCL and NDS, *Fusion Eng. Des.* 168 (2021) 112393, <http://dx.doi.org/10.1016/j.fusengdes.2021.112393>.
- [10] M. Ruiz, J. Nieto, V. Costa, T. Craciunescu, E. Peluso, J. Vega, A. Murari, Acceleration of an algorithm based on the maximum likelihood bolometric tomography for the determination of uncertainties in the radiation emission on JET using heterogeneous platforms, *Appl. Sci.* 12 (2022) 6798, <http://dx.doi.org/10.3390/app12136798>.
- [11] University of Ljubljana, Orange data mining web page, 2023, URL: <https://www.orangedatamining.com>. (Accessed 28 February 2023).
- [12] The Apache Software Foundation, Apache NiFi web page, 2023, URL: <https://nifi.apache.com>. (Accessed 28 February 2023).
- [13] Mathworks, Simulink web page, 2023, URL: <https://www.mathworks.com/products/simulink.html>. (Accessed 28 February 2023).
- [14] IODA, IODA web page, 2023, URL: <https://t.um.es/ioda>. (Accessed 20 September 2023).
- [15] W3C, World wide web consortium web page, 2023, URL: <https://www.w3.org>. (Accessed 28 February 2023).
- [16] ECMA International, JSON standard Web Page, 2023, URL: <https://www.ecma-international.org/publications-and-standards/standards/ecma-404>. (Accessed 28 February 2023).
- [17] V.C. Gomes, G.R. Queiroz, K.R. Ferreira, An overview of platforms for big earth observation data management and analysis, *Remote Sens.* 12 (8) (2020) 1253.
- [18] S. Aluru, N. Jammula, A review of hardware acceleration for computational genomics, *IEEE Des. Test* 31 (1) (2013) 19–30.
- [19] J.J. Thomas, K.A. Cook, Illuminating the path: The research and development agenda for visual analytics, in: *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, 2005, pp. 74–97.
- [20] S. Shankar, A. Kini, D.J. DeWitt, J. Naughton, Integrating databases and workflow systems, *SIGMOD Rec.* 34 (3) (2005) 5–11, <http://dx.doi.org/10.1145/1084805.1084808>.
- [21] F. Daniel, M. Matera, Mashups concepts, models and architectures, in: *Mashups Concepts, Models and Architectures*, Springer, Berlin, Heidelberg, 2014, pp. 137–181, http://dx.doi.org/10.1007/978-3-642-55049-2_6.
- [22] J.P. Morrison, *Flow-Based Programming: A New Approach to Application Development*, CreateSpace, 2010.
- [23] Django, Django project web page, 2023, URL: <https://www.djangoproject.com>. (Accessed 28 February 2023).
- [24] Docker, Docker web page, 2023, URL: <https://www.docker.com>. (Accessed 28 February 2023).