














OPEN
ARTICLE

AiiDA 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance

Sebastiaan P. Huber ^{1,2,9}✉, Spyros Zoupanos ^{1,2,9}, Martin Uhrin ^{1,2}, Leopold Talirz ^{1,2,3}, Leonid Kahle^{1,2}, Rico Häuselmann^{1,2}, Dominik Gresch⁴, Tiziano Müller ⁵, Aliaksandr V. Yakutovich ^{1,2,3}, Casper W. Andersen ^{1,2}, Francisco F. Ramirez^{1,2}, Carl S. Adorf ^{1,2}, Fernando Gargiulo ^{1,2}, Snehal Kumbhar^{1,2}, Elsa Passaro^{1,2}, Conrad Johnston^{1,2}, Andrius Merkys ⁶, Andrea Cepellotti^{1,2}, Nicolas Mounet ^{1,2}, Nicola Marzari ^{1,2}, Boris Kozinsky^{7,8} & Giovanni Pizzi ^{1,2}✉

The ever-growing availability of computing power and the sustained development of advanced computational methods have contributed much to recent scientific progress. These developments present new challenges driven by the sheer amount of calculations and data to manage. Next-generation exascale supercomputers will harden these challenges, such that automated and scalable solutions become crucial. In recent years, we have been developing AiiDA (aiida.net), a robust open-source high-throughput infrastructure addressing the challenges arising from the needs of automated workflow management and data provenance recording. Here, we introduce developments and capabilities required to reach sustained performance, with AiiDA supporting throughputs of tens of thousands processes/hour, while automatically preserving and storing the full data provenance in a relational database making it queryable and traversable, thus enabling high-performance data analytics. AiiDA's workflow language provides advanced automation, error handling features and a flexible plugin model to allow interfacing with external simulation software. The associated plugin registry enables seamless sharing of extensions, empowering a vibrant user community dedicated to making simulations more robust, user-friendly and reproducible.

Introduction

Reproducibility is one of the cornerstones of the scientific method, as it enables the validation and verification of scientific findings^{1–4}. In computational science, for a result to be reproducible, it should be possible to exactly retrace all the data transformations that led to its creation. With the ever-growing availability of computational power, and increasingly complex computational workflows resulting in large amounts of interconnected data, *a posteriori* reconstruction of provenance has become an intractable task. Instead, to guarantee reproducibility, *a priori* provenance should be effortlessly enforced through mechanisms that automatically track data as it is being

¹National Centre for Computational Design and Discovery of Novel Materials (MARVEL), École Polytechnique Fédérale de Lausanne, CH-1015, Lausanne, Switzerland. ²Theory and Simulation of Materials (THEOS), Faculté des Sciences et Techniques de l'Ingénieur, École Polytechnique Fédérale de Lausanne, CH-1015, Lausanne, Switzerland. ³Laboratory of Molecular Simulation (LSMO), Institut des Sciences et Ingénierie Chimiques, École Polytechnique Fédérale de Lausanne (EPFL), Rue de l'Industrie 17, Sion, CH-1951, Valais, Switzerland. ⁴Microsoft Station Q, University of California, Santa Barbara, California, 93106-6105, USA. ⁵Department of Chemistry, University of Zürich, Zürich, Switzerland. ⁶Vilnius University Institute of Biotechnology, Saulėtekio al. 7, LT-10257, Vilnius, Lithuania. ⁷John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, 02138, United States. ⁸Robert Bosch LLC, Research and Technology Center North America, 255 Main St, Cambridge, Massachusetts, 02142, USA. ⁹These authors contributed equally: Sebastiaan P. Huber, Spyros Zoupanos. ✉e-mail: mail@sphuber.net; giovanni.pizzi@epfl.ch

created. Crucially, such a system should not merely store the data itself, but also preserve the explicit connection to the process that generated it, as well as the inputs of the latter and (recursively) their respective provenance.

While essential, automated provenance storage is not the sole requirement that must be satisfied to deliver effective reproducibility. Rather, the data model should be generic enough to be applicable to any computational domain, and the infrastructure should be flexible enough to be interfaced with the diverse range of existing computational software. Moreover, the infrastructure should also provide a system to fully automate complex simulations through the definition of robust workflows. Last, data produced and stored should be easily shareable, such that it can be found, accessed and reused with different tools, following FAIR data principles⁵, and efforts are underway to extend these principles from data to the workflows producing the data⁶. Such concepts of automated workflows, data management and sharing were formalised in the Automation, Data, Environment and Sharing (ADES) model and implemented in the AiiDA informatics infrastructure⁷.

Many computational workflow management systems have been devised over the past decades (s.apache.org/existing-workflow-systems). One popular design choice is to use markup languages to describe the logic of static workflows. While this allows the workflow syntax to be kept simple, it tends to come at the cost of sacrificing the flexibility needed to dynamically change the execution path taken in response to the intermediate results (or errors) encountered. However, such flexibility is crucial in many fields, and in particular in computational materials science, which has been the original driver for the development of AiiDA.

Actually, in computational materials science, multiple tools address the challenge of managing workflows composed of complex simulation codes; Fireworks⁸, AFLOW⁹ and ASE¹⁰ are some of the most popular. In addition to these domain specific workflow managers, there are also GC3Pie¹¹, Signac¹² and Parsl¹³ which are applicable to the domain of computational science and high-performance computing in general. Like AiiDA, all of these are implemented and operated through the Python language, and provide new language constructs to design and automatically execute workflows (in parallel), and store the resulting data. The designs of these workflow managers differ depending on the typical use cases that they target; Parsl aims for many (>10000) short (<100 seconds) concurrent tasks, while Atomate and GC3Pie focus on fewer but computationally more demanding tasks, with Signac somewhere in between. All frameworks store the data produced, but none with a focus on explicitly recording provenance in detail, and, in particular, storing the interconnections between data and the processes that received it as input or produced it as output. A focus on data provenance, as a core principle in AiiDA's design, is a core difference with the other management systems mentioned above.

Early versions of AiiDA have already been successfully used in high-throughput computational studies^{14–18}, making their AiiDA databases publicly available, and the corresponding provenance graphs interactively browsable through, e.g., uploads on the Materials Cloud¹⁹. As a consequence of such uptake of AiiDA, the size and scope of the requirements quickly increased, testing the limits of the high-throughput capabilities of the original architecture. The workflow engine, the component responsible for the automated management of all calculations and workflows, showed its scalability limits under heavy computational loads when trying to manage hundreds of jobs concurrently on different computational resources. AiiDA 1.0, represents the culmination of a complete redesign that not only greatly improves efficiency and aims to deal with the high-throughput loads expected for upcoming exascale computing systems²⁰, but also includes new core features that, amongst others, facilitate querying the provenance graphs and extend core functionalities, while remaining faithful to the principles of the ADES model⁷. Despite fairly radical changes to the code and architecture, a dedicated effort has been undertaken to guarantee that existing databases, generated with earlier versions, can automatically be migrated, guaranteeing and preserving the longevity of existing data. In this paper, we first describe the new architecture of AiiDA 1.0, following with a more in-depth discussion of the new design and features, focusing on the rationale behind those choices and the challenges they are intended to address. The user interface of AiiDA is intentionally not addressed in this paper and as such, it is aimed mainly at plugin developers, power-users and researchers of related fields that may be interested in the presented technical solutions. Prospective users, interested in learning to work with AiiDA, are referred to the extensive online documentation (aiida-core.readthedocs.io/en/latest/), which includes a short tutorial, focused how-to guides, and extensive topical and reference materials.

Architecture Overview

AiiDA aims to provide a framework that enables designing and running complex high-throughput computational workflows with full automatic provenance and built-in support for high-performance computing on remote supercomputers. The architecture, as shown in Fig. 1, is designed with these goals in mind.

One of the core components, the engine, is responsible for running all calculations and workflows that are submitted by the user and will be described in the section “Engine and Workflow Language”. Calculations and workflows can be implemented in the custom language provided by AiiDA's core API which is implemented in Python. A detailed description of the engine's design and implementation, as well as the user interface of the workflow language can be found in ref. ²⁰. In addition, an example of part of the workflow syntax is shown in Supplementary Fig. S2.

Any calculation or workflow that is run by the engine will be automatically recorded in the provenance graph in order to enable the reproducibility of the results. The definition and implementation of the provenance graph is explained in greater detail in section “The Provenance Model”. Besides the workflow language, the ORM also provides the tools to interact with the nodes of the provenance graph and inspect their content. The QueryBuilder is the tool that allows efficient traversal of the provenance graph to select (sets of) nodes of interest and is described in detail in section “Database Abstraction and Querying Language”.

The contents of the provenance graph are stored in a file repository on the local file system and a relational database. The mapping between the database and the Python API is performed by an Object Relational Mapper (ORM): currently the user can choose between the Django (djangoproject.com) or SQLAlchemy (sqlalchemy.org) library.

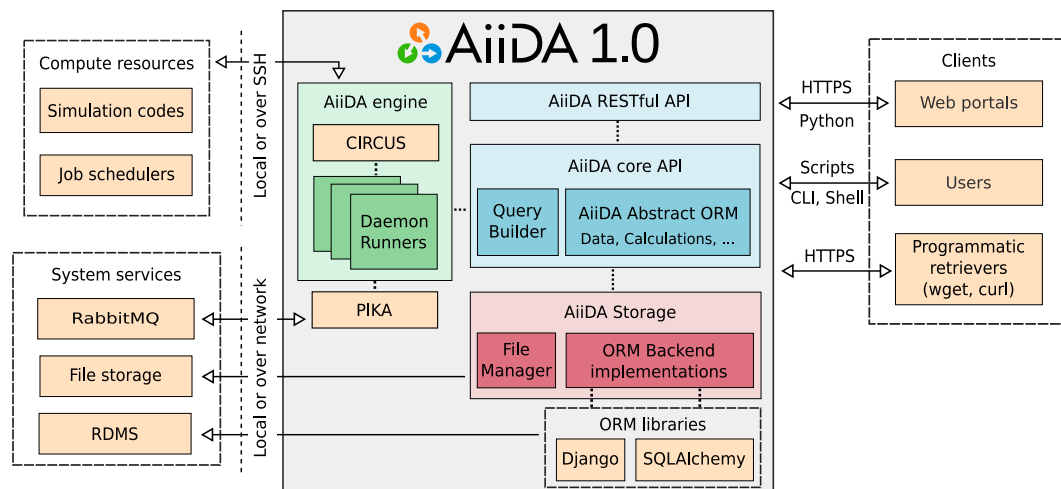


Fig. 1 Schematic overview of the architecture of AiiDA 1.0.

Component	Status	Description
Engine	⊖ Redesigned	Responsible for running all calculations and workflows
Workflow language	+ New	Allows to write robust and modular workflows
Daemon	⊖ Redesigned	Rewritten to be able to scale on demand
Provenance graph	⊖ Redesigned	Core concept extended to include workflows
Query builder	+ New	New tool that allow querying the provenance graph through Python API
REST API	+ New	Provides access to the database over the HTTP protocol
Database	⊕ Optimized	Improved query speed and storage efficiency
Plugin system	+ New	Provides a plug-and-play system to install extensions of AiiDA

Table 1. An overview of the most important additions and improvements in AiiDA 1.0 with respect to the version described in ref. ⁷. Components with the status “Redesigned” have been partially or completely redesigned and reimplemented.

From the outside, users can interact with AiiDA through a command line interface called *verdi*, an interactive Python shell or normal Python scripts. The REST API allows one to query the provenance graph through HTTP calls (see section “The REST API” for more details). AiiDA itself can communicate with computing resources either locally or over SSH to run calculations on those resources and comes with built-in support for most well-known and used job schedulers.

Although this is not the first release of AiiDA, the changes with respect to the design, implementation and features are so extensive that AiiDA 1.0 can almost be considered to be a new code, were it not for the fact that old data is guaranteed to be forward compatible. Table 1 gives an overview of the changes in AiiDA 1.0 with respect to earlier versions (a more detailed overview can be found in section A of the Supplementary Information).

Engine and Workflow Language

The engine is the component of AiiDA in charge of automating the execution of calculations and workflows. AiiDA is capable of managing calculations either on the local computer where AiiDA is installed, or on any number of remote resources (see section “Running on external computers: calculation jobs”). In addition, AiiDA provides a workflow language to define the logic to run complex sequences of steps, with potentially nested sub-workflows and calculations.

The engine consists of runners that are executed in parallel as different operating system processes, supervised by a daemon that monitors and relaunches them if they were to die. Each runner can process tasks independently and concurrently, distributing the workload involved in workflow and calculation execution. Task distribution is achieved via a task queue implemented using the AMQP protocol through RabbitMQ (rabbitmq.com), to guarantee reliable scheduling and almost instantaneous reaction to events such as the request for the submission of a new calculation or workflow, or continuing a workflow when the calculations or subworkflows it depends upon are completed. Additional technical details on the implementation of the engine and the use of the event-based task queue can be found in the Supplementary Information.

Thanks to this scalable architecture, the AiiDA engine is able to sustain high-throughput workloads involving tens of thousands of concurrent tasks every hour distributed on multiple computational resources, as we demonstrate in the “Performance” section. Additionally, during execution AiiDA automatically stores all data and actions in the provenance graph (see “The Provenance Model”), including the workflows, the calculations and their inputs and outputs, to provide full traceability.

In this section we define the three core concepts defined by the AiiDA engine (processes, calculations and workflows) and we briefly describe the key aspects of the engine implementation.

Processes in AiiDA. In AiiDA, any entity that handles input data to produce output data, and that is run by the engine, is called a process. Processes come in two flavours: calculations and workflows. These two terms in AiiDA have a more specific meaning than their use in common parlance. In particular, calculations are defined as processes that create new data as output, given certain data as input. A typical case is the execution of a simulation code on a remote computer. In contrast, workflows in AiiDA are solely tasked with the orchestration of subprocesses, calling calculations and/or other workflows in a certain logical sequence. Consequently, workflows are not allowed to generate new data, but can only return existing data, usually created by one of the calculations that they called (either directly or indirectly via a subworkflow).

This distinction is critical in the design of the provenance model of AiiDA, allowing to differentiate the part of the provenance graph that represents exactly how data was generated (the data provenance) from the logical provenance that captures the why of the data creation, i.e., which workflows drove the execution. This distinction is elaborated in more detail in the subsection “Logical and data provenance”.

In routine tasks, users do not interact directly with calculations or workflows, but with specific subclasses that define additional functionalities and are appropriate in different use cases that we describe below.

Process functions: calculation functions and work functions. The simplest way to define a process (a calculation or a workflow) is to add a line (`@calcfunction` or `@workfunction`, respectively) on top of a standard Python function (this line is called a decorator in Python; see the supplementary section “Example of a work chain and calculation function” for an example). Calculation and work functions are collectively referred to as process functions. The logic of the `@calcfunction` and `@workfunction` decorators is defined by AiiDA, and they signal to AiiDA’s engine that any time the function is called, its execution should be recorded in the provenance graph, linking up the inputs passed and the outputs it produced. Furthermore, if a work function calls another process function, this action is also represented in the provenance graph as a connection between the caller and the called process.

This approach of defining processes is very intuitive and powerful, because any standard Python function can be converted into a process function simply by adding a decorator, as long as it accepts AiiDA datatypes as input and returns an AiiDA datatype (or a dictionary of AiiDA datatypes) as output. However, process functions show their limits when running very long workflows that can last for days (e.g., in the case of molecular-dynamics simulations). The reason is that process functions are executed on the local machine and since it is not possible to interrupt the execution of a Python function and resume it later, the local machine cannot be shut down, rebooted or disconnected from the network until the process is finished. In addition, the interface design, through its simplicity, restricts what code can be executed, as it has to be written in Python. Of course one could resort to calling external codes through subprocesses, but this quickly becomes complicated, especially if the external code needs to be run on a remote machine and/or through a job scheduler. For these reasons, AiiDA implements two additional processes: calculation jobs, to manage the execution of external codes through job schedulers, and work chains to implement and orchestrate long-running workflows with the possibility of pausing and restarting between steps.

Running on external computers: calculation jobs. Calculations jobs, implemented by the `CalcJob` process class, are used to manage the execution of external codes, commonly run via a job scheduler and optionally on a remote machine. The `CalcJob` can be adapted for external codes through plugins (see “The Plugin System”), which define how the required raw input files are constructed from the AiiDA datatypes that are provided as input. Once submitted, the engine takes over and performs all the necessary steps to run the calculation to completion. These include uploading the raw input files, submitting the job to the scheduler, querying the job state and waiting for it to finish, and finally retrieving the files to be stored locally. The retrieved files can optionally be parsed into AiiDA datatypes that are registered as outputs of the calculation. The parser logic is also defined through plugins, since this also requires code-specific logic. The explicit parsing into AiiDA datatypes makes the outputs interoperable among different codes and allows for their direct reuse as inputs of new calculations.

To increase the robustness of the engine, we implemented optimised algorithms to automatically reschedule failed tasks through an exponential-backoff mechanism, recovering from common transient errors such as connection issues. If the task fails multiple times consecutively (for example because the remote machine went offline), the process is automatically paused. The process can be resumed effortlessly by the user when the issue is resolved, and the AiiDA engine ensures that no data loss occurs as the calculation resumes from where it was paused.

Additionally, to avoid overloading remote machines with connections, rendering them potentially unresponsive also to other users, AiiDA implements a connection pooling algorithm. Connections to the same computer are grouped and new requests are funnelled within existing open connections, if available. If no connection is open and a new one needs to be created, AiiDA guarantees that a configurable minimal time interval between them is respected. For similar reasons, AiiDA caches the list of jobs in the scheduler queue and respects a minimum (configurable) timeout when refreshing their state.

Interruptible workflows: work chains. As discussed earlier, the limitation of a workfunction is that its execution is atomic and blocking. Therefore, if workfunctions are used to implement complex workflows calling many long-running calculation jobs, intermediate progress within the function body cannot be persisted. As a consequence, if the process is interrupted (for instance to restart the AiiDA daemon or to reboot the machine where AiiDA is running), all progress is lost.

To solve this issue, AiiDA implements the `WorkChain` process class. Work chains allow users to specify units of work (like the processing of data or the submission of subprocesses) as steps. The logical outline of the steps (i.e., their sequence, possibly including `while` and `if/else` logic) is defined in the process specification. By design, the execution of the workflow is managed by the AiiDA engine that gets back control between steps, executing the next one only when all subprocesses have finished. Most importantly, between steps the engine persists the workflow state to the database. Therefore, if the engine is stopped no progress is lost and, upon restarting, the engine will continue the workflow execution from the last persisted step (or the engine will keep waiting for running subprocesses, if any).

In addition to the capability of stopping and continuing execution, work chains have the advantage that inputs and outputs (including information like their type or whether they are required or optional) can be declared in the process specification, together with the logical outline of the steps. Work chains are thus self-documenting because, through simple inspection of the specification, users can determine the interface of the workflow and its main logic. Furthermore, the specification is machine-readable and can be automatically rendered in multiple formats. One use case is the extension provided in AiiDA for the Sphinx ([sphinx-doc.org](https://www.sphinx-doc.org)) engine (used to generate the AiiDA documentation, see “Community building”) that generates human-readable documentation for any work chain, to be displayed, for instance, as a web page.

Since work chains are directly implemented in Python, users have the full power of the programming language at their disposal, as well as direct access to the content of the provenance graph. This allows for the implementation of powerful error-handling mechanisms that can deal with failures of calculations that are run as part of the workflow. AiiDA is a domain and code agnostic infrastructure and therefore does not include error handlers for specific codes. However, various utilities and tools are provided to make implementing these as easy as possible, such as exit codes for calculations and workflows, some of which are already pre-defined by AiiDA, and others that can be defined by the corresponding plugin. Based on these exit codes, error handlers can dynamically decide whether to re-submit calculations with different inputs, adapt postprocessing, etc. Plugins for specific simulation codes can utilise these tools to write the error handling specific to their use case, as has been done for example by the `aiida-quantumespresso` (github.com/aiidateam/aiida-quantumespresso) and `aiida-vasp` (github.com/aiida-vasp/aiida-vasp) plugins for the Quantum ESPRESSO²¹ and VASP²² codes, respectively. The typical failures that are handled range from simple out-of-walltime errors to more complex failures such as problems with the convergence of the self-consistent minimization cycle. These are just two examples of popular computational materials science codes for which AiiDA plugins exist, but many other codes are supported as well, as will be further discussed in “The Plugin System”.

The work chain interface encourages writing modular workflows, with lower-level work chains implemented to solve specific tasks that are often code-dependent, like error handling, restarting or automatic tuning of parameters. Higher-level work chains can then wrap around these, directly exposing all inputs of the lower-level workflows or only a few, with the others being determined by the workflow logic. Such wrapping work chains can implement additional functionality less focused on the technical details of the code execution but rather focusing on the evaluation of scientific quantities of interest. As such, workflows enable the delivery of fully automated turn-key solutions that require only minimal inputs and encode the scientists’ knowledge on how to compute a given result, from code-specific technicalities to the handling of the parameters involved in the scientific models of the simulations.

To conclude, we emphasise that implementing workflows directly in Python is a design choice. Since there is no translation between what the user implements and what the engine executes, the user can leverage the power of the full Python and AiiDA APIs when designing workflows. Additionally, this design directly facilitates debugging using standard Python tools and seamless integration with external libraries for data manipulation.

The Provenance Model

As explained in the previous section, AiiDA’s engine automatically represents the execution of processes, along with their inputs and outputs, as vertices (called nodes in AiiDA) of a directed graph. Graph edges (called links) connect the nodes; process nodes, for instance, have incoming links from their inputs and outgoing links to their outputs. Since output data can in turn be used as input to new processes, extensive graphs are generated. We call them AiiDA provenance graphs because they allow to retrace the exact steps that led to the creation of a piece of data. An example of a simple provenance graph is shown in Fig. 2.

Node types. While all nodes of the graph share a set of common properties, there is a need to define custom properties based on what the node represents. Therefore, the various AiiDA processes (see “Engine and Workflow Language”) as well as data are represented by different node subtypes. This makes it possible to implement functionality specific to each of them and to explicitly target nodes of a certain type when querying the graph (see “Database Abstraction and Querying Language”).

In AiiDA, the `Node` class is the base class to represent any node in the graph. The common properties of any node include the user who created it, the creation and last modification times, an optional computer on which it was run or stored, and a human-readable label and description. `Node` classes are subclassed to build a hierarchy of node types, schematically represented in Fig. 3. In particular, data and process nodes are represented by the `Data` and `ProcessNode` subclasses, respectively.

Various data types are implemented by directly subclassing the `Data` class. AiiDA ships with a few basic data types; for instance, among many others, `Float` to represent a single float value and `Dict` for a dictionary of key-value pairs. Arbitrary new data types can be defined through the plugin system (see “The plugin system”).

The hierarchy of `ProcessNode` subclasses reflects the distinction in AiiDA between calculations and workflows, represented by subclasses of `CalculationNode` or `WorkflowNode`, respectively. In practice,

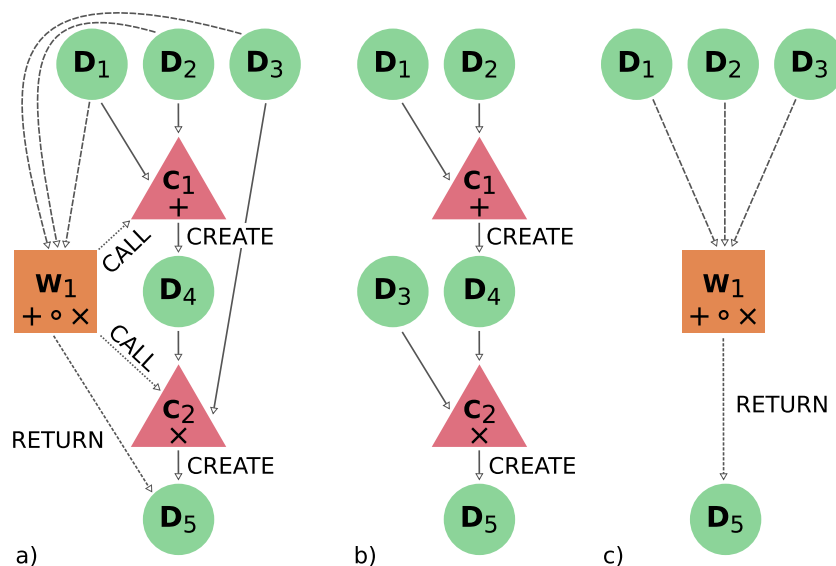


Fig. 2 (a) A schematic provenance graph representing the execution of a workflow W_1 receiving three data nodes D_1 , D_2 and D_3 as input, containing the values x , y and z respectively. W_1 computes the expression $(x + y) \cdot z$ by calling two calculations C_1 (to perform the sum) and C_2 (to perform the product), forwarding the correct inputs to them. C_1 creates the intermediate node D_4 (with the value $x + y$) and C_2 then creates the node D_5 with the final result, that is then also returned by W_1 . While this simplified example is purely for illustrative purposes, it demonstrates that by storing execution information as a graph, the provenance of all data is fully recorded. (b) Data-provenance layer: it includes calculation and data nodes only, showing the exact sequence of steps that led to the creation of the data nodes. (c) Logical-provenance layer: it hides the details of all intermediate results and focuses only on how the workflow produced the final results from a given set of inputs.

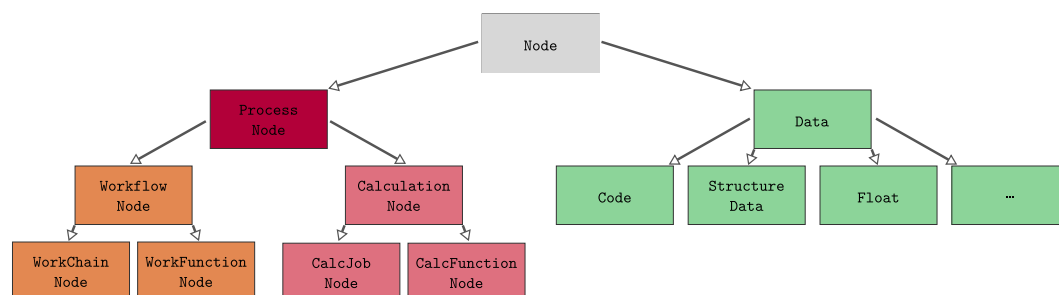


Fig. 3 The hierarchy of the node types in AiiDA. This hierarchy is also mirrored in the Python code, where Python classes are used to represent them, using Python's inheritance model. The different node classes allow to implement custom functionality for each subtype. Additionally, the subclass hierarchy allows to query for specific node types, or a set thereof.

in a provenance graph one finds instances of `CalcJobNode`, `CalcFunctionNode`, `WorkChainNode` or `WorkFunctionNode`, representing executions by the engine of the corresponding process classes (`CalcJob`, `calcfunction`, `WorkChain` or `workfunction`, respectively, as discussed in “The engine”). The intermediate classes in the hierarchy serve mostly as a taxonomic classifier and they are useful when querying the provenance graph (see “Database abstraction and querying language”). For instance, querying for `WorkflowNodes` will match both `WorkChainNodes` as well as `WorkFunctionNodes`.

Link types. All links have a type to indicate the semantic meaning of the relationship. In addition, links have a label that can be used, given a node, to distinguish nodes connected to it with the same link type. For example, labels identify the different input nodes to a process. A summary of all link types in AiiDA is shown in Fig. 4.

Process nodes can have input and output links to data nodes, representing their inputs and outputs. More specifically, in the implementation, input links can be of type `INPUT_CALC` and `INPUT_WORK` depending on the type of the linked process node. Similarly, output links can either be of type `CREATE` or `RETURN` (for calculation and workflow nodes, respectively) explicitly highlighting the difference between calculation and workflow processes described in “Engine and Workflow Language”.

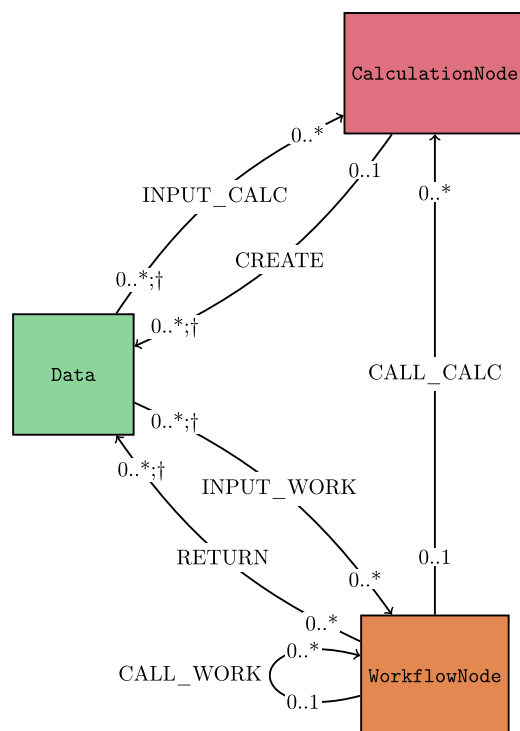


Fig. 4 Link types allowed in the AiiDA provenance graph. Rectangles represent node types and arrows connecting them indicate the direction and the type of each link. The symbols at the start and end of each arrow indicate the cardinality of the corresponding link types: $0..1$ means that at most one node is allowed on that link endpoint for a given node on the opposite endpoint (for instance, a *Data* node can have at most one *CalculationNode* as its creator); $0..*$ means that any number of nodes is possible (for instance, a *CalculationNode* can have an arbitrary number of input *Data* nodes). Additionally, a dagger (†) indicates that link labels must be unique for a given node on the opposite endpoint (for instance, outgoing create links from a *CalculationNode* must have unique labels).

In addition to links between data and process nodes, two additional link types exist between processes, to indicate that a workflow called another process. These links are called *CALL_CALC* and *CALL_WORK*, depending on the type of the called process, and are collectively referred to as call links.

Due to their semantic meaning, link types impose precise validation rules. First, a link type requires specific node types at its two endpoints. In addition, cardinality rules are defined, as illustrated in Fig. 4. In general these dictate that, given a node, any number of nodes can be linked to it by links of a given type. For create and call links there are instead explicit restrictions: a data node can be created by at most one calculation, while it can be returned by multiple workflows; a process can be called by at most one workflow, while a workflow can call an arbitrary number of subprocesses. Finally, uniqueness constraints on the labels of certain link types are enforced: the link labels of input nodes to a given process must be unique, to guarantee that each input can be uniquely identified. The same applies to the output nodes of a process.

Logical and data provenance. Due to the rules defined in the previous sections, AiiDA provenance graphs have useful properties. The subgraph composed exclusively of data and calculation nodes together with the links that connect them forms a Directed Acyclic Graph (DAG). The acyclicity is a direct result of calculations only having outgoing links to data they created: output nodes are always generated as a result of execution and therefore, because of the causality principle, cannot also simultaneously be part of the inputs of the process itself or of any parent process in the graph. We refer to this subgraph as the data provenance, since it is an exact record of the origin of the data. Since a DAG has well defined properties, assumptions can be made by the AiiDA query language (see “Database Abstraction and Querying Language”), for instance to define the concept of an ancestor of a given node as any node that can be reached following the links of the data-provenance DAG in the backward direction (i.e., moving from the node at the head of a link to the one at the tail). Similarly, descendants are defined as nodes that can be reached following DAG links in the forward direction.

We can also define a second subgraph, composed solely of data and workflow nodes and the links connecting them (including *CALL_WORK* links). We call this graph the logical provenance, since it focuses on the logical steps taken by the workflows when processing data and orchestrating processes. Like the data provenance, the logical provenance subgraph also forms a directed graph that, however, is not acyclic. In fact, return links are not bound by the causality principle since they do not signify “creation”, and workflows merely return outputs that have already been created by other calculations. For instance, a workflow selecting one of its inputs based on some criteria will have a return link to that input, which introduces a cycle in the graph.

The two layers of data and logical provenance share data nodes and are interconnected by `CALL_CALC` links. The separation of these two provenance layers has the additional benefit of allowing the granularity for the inspection of the provenance graph, as shown in Fig. 2(b,c), to be selected.

Node properties. In the previous section “Link types”, we described the basic properties that are common to all node types. The various subtypes define the “schema” of the additional properties that fully define and describe the node. AiiDA provides two data stores to persist these properties: a filesystem repository and a relational database (where any JSON-serialisable key-value pair can be saved, see “Database Abstraction and Querying Language”). Properties that are stored in the database are named attributes, which are fully and efficiently queryable. In contrast, properties that do not require querying and/or are very large in size, such as large arrays or raw files, are better stored in the repository so as not to overburden the database. Attributes and the files in the repository are immutable once the node is stored, since together they define the “content” of the node and allowing them to be changed would invalidate the provenance of descendants. Mutable properties are also allowed and are called extras that, like attributes, are stored as key-value pairs in the database. However, in stark contrast to attributes, extras can be added and/or modified at any time. A typical use case is to tag nodes with custom properties that can, for example, be used for more selective querying.

Reproducibility and efficiency. By automatically recording all data transformations through a process together with all the inputs it consumed, the produced graph is in principle fully reproducible. This of course depends on all relevant inputs being stored as well as the code that was ran by the process. The actual source code of process functions is stored by AiiDA in the repository; for all other processes the code is referenced indirectly by storing both the version of AiiDA and of the relevant plugin in the attributes of the process node. The source code of external simulation software is not directly stored, however, although AiiDA provides the user with the possibility to store metadata, such as the version of the software or the libraries and parameters used for compilation. The inclusion of this information is not enforced as it is impossible to validate and it unnecessarily complicates the interface for those users who do not require this level of reproducibility. The obvious solution to improve the provenance of data produced by external codes is containerisation, and development is ongoing to provide integrated support for containerised codes in AiiDA.

Tracking all data transformations in a provenance graph provides another benefit besides enabling reproducibility. The entire provenance graph can serve to implement a “caching” mechanism: if one considers the execution of a process that already has been performed before with identical inputs, the actual execution can be skipped. In this case, since the inputs are identical, we already know what the outputs are going to be as well, and so we can simply take those from the previously executed process instead, avoiding incurring the computational cost once more. This caching mechanism is implemented in AiiDA 1.0 and is explained in detail in the section “Caching”.

Database Abstraction and Querying Language

When running automated high-throughput research projects with AiiDA, very large provenance graphs containing millions of nodes or more are easily generated (as for instance in the study of ref. ¹⁴). Tools that can efficiently query such graphs become essential to perform data analysis. In order to allow the implementation of a performant tool, AiiDA uses a relational database to store the provenance graph with its links and nodes including most of their properties (see “Node properties”). In addition, we have optimised AiiDA’s database schema and indexes to ensure that typical graph queries are very efficient, as discussed in section “Performance”.

Direct database queries must be expressed in its native language. AiiDA’s current database solution PostgreSQL ([postgresql.org](https://www.postgresql.org)) is based on the SQL language, which, while known for its efficiency, requires (in the context of the AiiDA’s provenance graph) writing queries that are long and cumbersome even for database experts. Furthermore, the exact query structure depends on the specific implementation choices of AiiDA, that could change between versions to improve efficiency (for instance if the schema is improved or the SQL backend is replaced by a graph database). AiiDA does not directly write the SQL statements itself but relies on object-relational mapping (ORM) libraries like Django and SQLAlchemy as an intermediate layer to express the queries in Python. However, these queries still depend on the specific database schema and have essentially the same complexity level as the native ones. It is therefore crucial to provide a tool that abstracts this process and makes writing queries not just as simple as possible but also independent of the ORM and database implementation.

The query builder is the tool in AiiDA that satisfies these criteria: it allows users to express any query on the provenance graph using a familiar Python syntax that is automatically translated into an optimised SQL query and executed. To illustrate the concept of the query builder, we describe in Fig. 5 a simple provenance graph and how a query is mapped onto it. Essentially, this is the problem of subgraph isomorphism, which, given two graphs G and H , consists of finding all the subgraphs within G that are isomorphic to H ²³. Here G is the entire AiiDA provenance graph and H is the subgraph represented by the query as expressed through the query builder.

The schematic query shown in Fig. 5(b) represents the search for all crystal structures (`StructureData`) used as input to a calculation (`CalculationNode`) that created a `Dict` node as output. The query encodes filters on the link directions, link types (`INPUT`, `CREATE`) and node types (`StructureData`, `CalculationNode`, `Dict`). On top of these constraints, additional query specifications are available which are not shown in Fig. 5 but are explained in detail in the “Query builder syntax example” section in the Supplementary Information. In particular, filters can be set on the node properties, for example the structure must contain a given chemical element or the output must have a value within a certain range. Moreover, the user can specify a list of projections, i.e., which subset of properties of the matched nodes should be returned. Once the query is fully defined and executed by the user, the query builder returns all the subgraphs embedded in the AiiDA provenance graph that match the query constraints. Figure 5 shows all four embedded subgraphs that

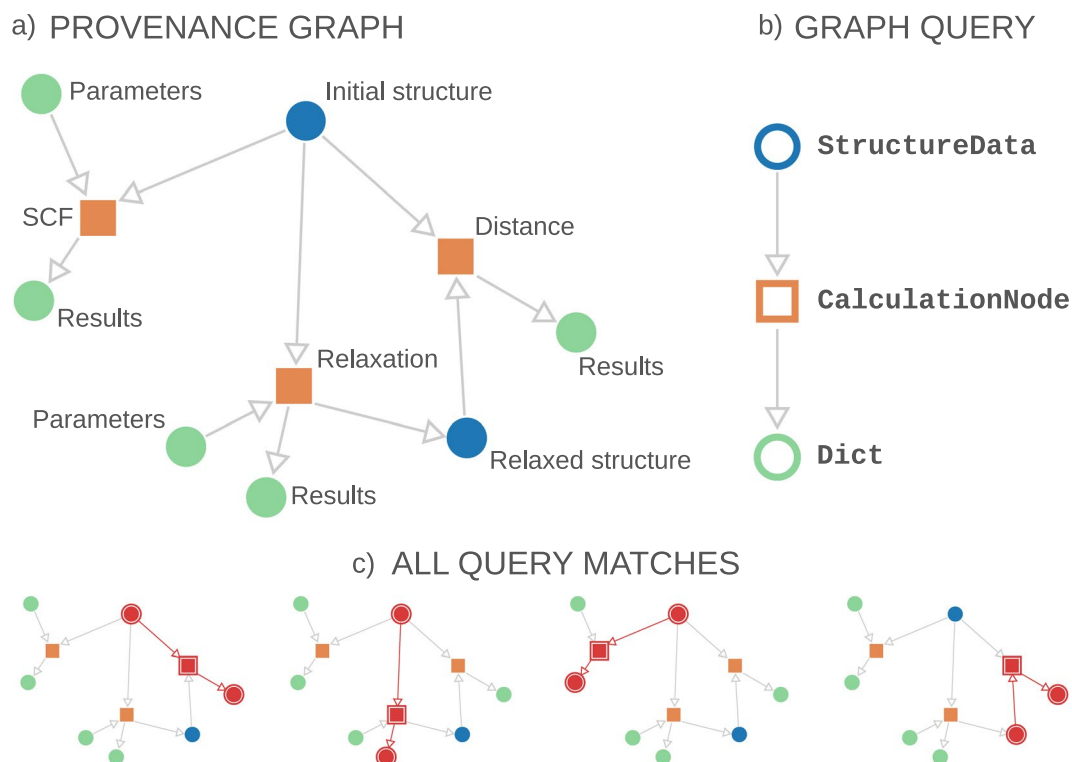


Fig. 5 (a) Schematic of an AiiDA graph that could result from a materials science simulation: as described by the labels, a Density Functional Theory self-consistent field (SCF) calculation and a geometry relaxation of a crystal structure, and a calculation of the “distance” between the initial and final structure. Orange squares represent nodes of type `CalculationNode`, circles represent `Data` nodes: blue for crystal structures (of type `StructureData`) and green for nodes of type `Dict` (dictionaries of key–value pairs with input parameters or parsed results). (b) Representation of a graph query searching a `StructureData` node that was an input of a `CalculationNode` that created a `Dict` node as output. Labels on the right represent the filter on the node type applied while querying. (c) The four subgraphs that match the query embedded in the entire provenance graph, where the matching nodes and links are colored in red and highlighted by a surrounding border.

match the query of this particular example. Finally, the query builder converts the results into Python objects, that can directly be used with common data processing libraries like SciPy ([scipy.org](https://www.scipy.org)) and pandas (pandas.pydata.org) for further data analysis.

The reason for the existence of two ORM backend implementations (Django and SQLAlchemy) is mostly historical. The original implementation used only Django, and as a result AiiDA’s API was tightly coupled to this library. When the query builder was first introduced, interaction with the database was implemented instead via SQLAlchemy, as it provided a richer feature set allowing for more general queries. Since SQLAlchemy at the time already provided support for JSONB (unlike Django), it was also used as an alternative ORM backend implementation to benefit from the significant improvements in database performance (see “Performance”). To achieve this, extensive work was performed to decouple the backend ORM from Django-specific constructs, and to create a new layer of backend-independent AiiDA frontend classes (those directly exposed to users to interact with nodes). Moreover, all functionalities of AiiDA (command line tools, query builder, import/export functionality, ...) were critically revised and updated so as not to rely anymore on backend-specific logic but to use general backend-independent interfaces instead. As a consequence, in the current version of AiiDA the user interface does not depend anymore on the chosen backend. Additionally, implementing new ORM backends is now greatly streamlined and simplified. Django also started supporting JSONB recently and, therefore, we have updated the Django interface in AiiDA to use JSONB fields for attributes and extras instead of our original solution based on entity-attribute-value tables, increasing database performance also for the Django backend (see “Performance”).

The REST API

The query builder is the tool of choice for querying data directly via the Python interface, which in turn is the preferred approach when one has direct access to the machine running AiiDA. If, instead, data needs to be made available to users without full access to the machine, for example over the web, a solution is required that can serve the data in a secure way. In addition, such a solution should not just be able to serve the database contents as a whole, but it should provide the necessary functionality to query for specific data.

A widespread approach to share data over the web is through a REST API ([w3.org/2001/sw/wiki/REST](https://www.w3.org/2001/sw/wiki/REST)), a stateless protocol that allows to query and retrieve data via HTTP requests ([w3.org/Protocols](https://www.w3.org/Protocols)). REST APIs are not only generally adopted on the web, but they have also become widespread in specific scientific disciplines and

domains. For instance, in the materials science community the OPTIMADE (optimade.org) consortium has been formed to define a common REST interface for querying material property databases. Many of the major materials databases have implemented (or are planning to implement) the OPTIMADE interface like AFLOW⁹, the Crystallography Open Database²⁴, TCOD²⁵, the High-Throughput Toolkit (httk.openmaterialsdb.se), Materials Cloud¹⁹, MPDS²⁶, Materials Project²⁷, NOMAD²⁸, OQMD²⁹, and AiiDA, among others.

AiiDA implements a REST API server that can be launched directly from the command line or deployed via scalable web servers like Apache (httpd.apache.org). API endpoints are available to access data associated with the AiiDA graph, including the list of nodes, their properties (like attributes, extras and files in the repository) as well as the graph information (incoming and outgoing links for a given node). Nodes are identified by their UUID (ietf.org/rfc/rfc4122.txt) to ensure that resources are uniquely identifiable even if data is shared and then made available by a different server. Custom queries can be performed by specifying filters in the query string in order to narrow the matched subset of nodes. The web server is implemented using the flask (palletsprojects.com/p/flask) web framework and a number of flask plugins (including flask-sqlalchemy to manage sessions to the database and flask-restful to handle requests using the REST approach). Once a request is received, this is translated into a database query using AiiDA's query builder, which is then executed. The results of the query are then mapped to the format defined by the REST API and serialised into a JSON response that is returned to the web client. Results are paginated to facilitate downloads of large amounts of data without overloading the server. In addition to endpoints common to all node types, additional endpoints are available that provide functionality specific to only certain node subtypes. For example, it is possible to directly obtain raw inputs and outputs of a calculation, or to download the content of data nodes in a specific format other than AiiDA's internal representation.

In conclusion, AiiDA's REST API endpoints provide a complete set of features to interact with AiiDA programmatically and in a secure way from the web. An emblematic example of its use is provided by the Explore section of the Materials Cloud¹⁹ portal, where AiiDA databases are made available as interactive web pages and the provenance graph can be browsed via a graphical user interface. Indeed, all data needed to display the provenance graph and the node contents on Materials Cloud Explore is obtained through AiiDA's REST API.

Performance

In this section, we examine the performance of the database and the engine. Where applicable, a comparison is drawn with earlier versions of AiiDA⁷, to illustrate how the new design has improved performance.

The database schema of AiiDA 1.0 includes many improvements and optimisations with respect to the schema published in ref. ⁷. Here we highlight two changes that have had the greatest impact on storage and query efficiency. First, in section “EAV replaced by JSONB” we describe how the schema of node attributes has been changed from a custom entity-attribute-value solution to a native JSON binary (JSONB) format. Second, section “On-the-fly transitive closure” details how the transitive-closure, originally a statically generated table, has been replaced with one that is generated on the fly in memory.

EAV replaced by JSONB. As described in “The provenance model”, the attributes of a node are stored in a relational database. The exact schema for these attributes depends on the node type and cannot be statically defined, which is in direct conflict with the *modus operandi* of relational databases where schemas are rigorously defined *a priori*. This limitation was originally overcome by implementing an extended entity-attribute-value (EAV) table that allowed storing arbitrarily nested serialisable attributes in a relational database⁷. While a successful solution, it comes with an increased storage cost and significant overhead in the (de-)serialisation of data, reducing the querying efficiency.

As storing semi-structured data is a common requirement for many applications, PostgreSQL added support for a native JSON and JSONB datatype as of v9.2 (postgresql.org/docs/current/static/release-9-2.html) and v9.4 (postgresql.org/docs/current/static/release-9-4.html), respectively, which is an efficient storage and indexing format (postgresql.org/docs/current/static/datatype-json.html). In AiiDA 1.0, the custom EAV implementation for node attributes has been replaced with the native JSONB provided by PostgreSQL, which yields significant improvements in both storage cost and query efficiency.

The replacement of EAV by JSONB significantly reduces storage cost in two ways: (a) the data itself is stored more compactly as it is reduced from an entire table to a single column and (b) database indexes can be removed while still providing a superior query efficiency. Figure 6(a) shows the space occupied when storing 10000 crystal structures, comparing the size of the raw files on disk, and with their content stored in the EAV and JSONB schema. In the case of raw files, the XSF format (xcrysden.org/doc/XSF.html) was used since it contains only the information that is absolutely necessary.

This benchmark was performed on a PostgreSQL 9.6.8 database using the ORM backends as implemented in AiiDA 1.0. When comparing the EAV format to the JSONB format, a decrease in storage space of almost two orders of magnitude becomes apparent. The space gains of the new format do not only apply to the occupied space on disk, but also to the amount of data transferred when querying JSON fields, as shown in Table 2. This effect is, however, only a part of the increase in query efficiency thanks to the JSONB schema.

Using the JSONB-based format also carries significant speed benefits. These mainly come from the more compact JSONB-based schema with respect to the EAV schema, as described in the previous section. This results in (a) less transferred data from the database to AiiDA, and (b) a reduced cost of deserialising the raw query result into Python objects.

Figure 6(b) shows benchmarks carried out with PostgreSQL 10.10 on an AiiDA database generated for a research paper¹⁴ which contains 7318371 nodes. The benchmarks were carried out on a subset of 300000 crystal-structure data nodes on a machine with an Intel i7-5960X CPU with 64GB of RAM. Three different kind of attributes were queried: *cell*, *kinds* and *sites*. The *cell* is a 3 × 3 array of floats, *kinds* contain information on

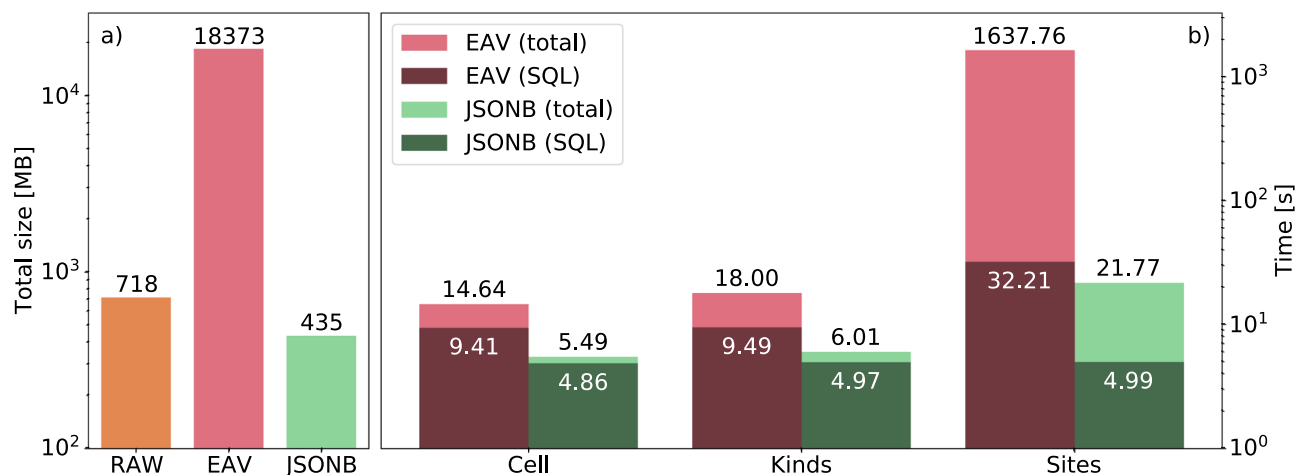


Fig. 6 Comparison in a log scale of the space requirements and time to solution when querying data with the two AiiDA ORM backends. (a) Space needed to store 10000 structure data objects as raw text files, using the existing EAV-based schema and the new JSON-based schema. The reduced space requirements of the JSON-based schema with respect to the raw text files are due to, among other things, white-space removal. The JSONB schema reduces the required space by a factor of 1.5 compared to the raw file size and a factor of 25 compared to the EAV-based schema. (b) Time for three different queries that return attributes of different size for the same set of nodes. The benchmarks are run on a cold database, meaning that the database caches are emptied before each query. We indicate separately the database query time (SQL time) and the total query time which includes also the construction of the Python objects in memory. The total query time of the *site* attributes in the JSONB format is 75 times smaller compared to the equivalent query in the EAV format. The SQL time for the same query is roughly 6.5 times smaller for the JSONB version of SQL query compared to the EAV version of the query.

Attribute	EAV		JSONB	
	Rows	Size (MB)	Rows	Size (MB)
Cell	255,671	71	19,667	8
Kinds	537,303	150	19,667	11
Sites	15,352,241	4412	19,667	201

Table 2. Result size and number of rows of attribute queries presented in Fig. 6(b) on a database table of 300000 crystal structures.

atomic species (therefore, typically there is one kind per chemical element contained in the structure), while there is a *site* per atom, explaining the increase in result sizes as shown at Table 2.

Due to the specific format of the EAV schema, more rows need to be retrieved for every crystal structure data node. The effect of the different result size is visible both in the SQL time (reflecting the time to perform the query and to get the result from the database) and in the total amount of time spent which includes the deserialisation of raw query results into Python objects. As shown in Fig. 6(b), the total query time of the *site* attributes in the JSONB format is 75 times smaller than the equivalent query in the EAV format. The SQL time for the same query is roughly 6.5 times smaller for the JSONB version of SQL query compared to the EAV version of the query. The increased final speedup at the Python level is due to the fact that in the EAV based schema there is the overhead of serialising the attributes at the Python level, which is largely avoided in a JSONB-based schema.

On-the-fly transitive closure. Very often, when querying the provenance graph one is only interested in the neighbours directly adjacent to a certain node. However, some use cases require to traverse the graph taking multiple hops in the data provenance to find a specific ancestor or descendant of a given node. To make the queries for ancestors and descendants at arbitrary distance as efficient as possible, early versions of AiiDA computed and stored the transitive closure (TC) of the graph (i.e., the list of all available paths between any pair of nodes) in a separate database table. Storing these paths in a dedicated database table with appropriate indexes allowed AiiDA to query for ancestors and descendants with time complexity $\mathcal{O}(1)$ independent of the graph topology and the number of hops.

However, the typical size of the TC table is significant even for moderately sized provenance graphs, and quickly has an adverse effect on the general performance of the database. For example, a subset of just one million nodes from the database generated in ref. ¹⁴ has 226 million rows in the TC table, corresponding to 200 GB on disk. In addition to the raw disk storage cost, the time needed to store a new link also increases, as the TC is updated with automatic triggers at each update of the links table. This becomes more expensive as the table grows

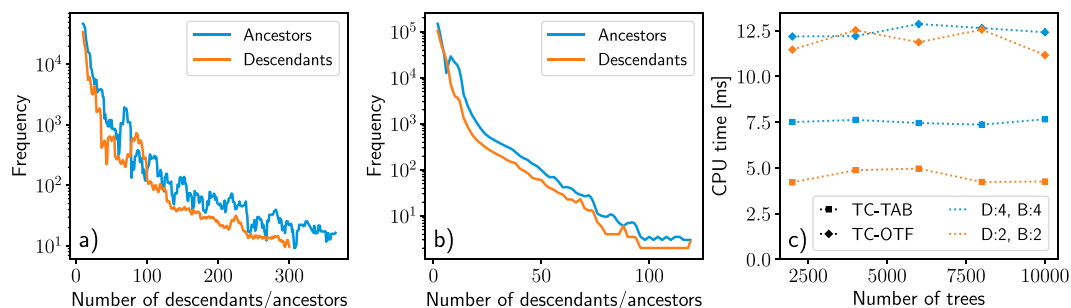


Fig. 7 Analysis of a sample of one million nodes of the AiiDA graph published in ref. ¹⁴. **(a)** Frequencies of the number of ancestors and descendants of all nodes. **(b)** Frequencies of the number of hops, i.e., the distance to reach the farthest ancestor/descendant. **(c)** Required CPU time when querying for all descendants of 50 top-level nodes in a graph that consists of a number of binary trees of breadth B and depth D using the transitive closure on-the-fly (TC-OTF, diamonds) or the explicitly tabulated transitive closure (TC-TAB, squares).

because table indexes need to be updated as well. AiiDA 1.0 replaces the TC explicitly stored in a table with one that is computed lazily, or on-the-fly (OTF), whenever ancestors or descendants of a node are queried for. This is implemented in the query builder via SQL common table expressions to recursively traverse the DAG. The OTF method greatly reduces the time required to store new links and does not require any disk space for storing the TC, albeit at the cost of slightly slower queries. However, the impact on the efficiency of the recursive queries for AiiDA provenance graphs is minimal, since the typical graph topology is relatively shallow and often composed of (almost) disjoint components. This can be seen in Fig. 7(a,b) that shows frequency graphs capturing the topology of a subgraph of the database of ref. ¹⁴, composed of one million nodes. Indeed, the vast majority of nodes only have a handful of ancestors and descendants and these can be reached in a relatively small number of hops.

To compare the performance of the explicit and lazy implementations of the TC, we performed benchmarks on multiple graphs with topologies comparable to typical AiiDA provenance graphs. Each graph consists of N binary trees with a depth and breadth (the number of downward branches and the number of outward going edges from each vertex, respectively) of 2 or 4. The benchmark records the total time it takes to query for all descendants of 50 top-level nodes using either the explicit TC table (TC-TAB) or the on-the-fly TC (TC-OTF).

Figure 3(c) clearly shows that in both cases the number of trees does not affect the query efficiency. Moreover, as the depth and breadth of the graph increases, the TC-TAB query time increases. In contrast, for the TC-OTF, the topology of the graph has little impact on the query time. Note that this holds for these particular topologies, which match that of typical AiiDA provenance graphs, but is not necessarily the case for more complex graph topologies. Finally, as expected, the TC-TAB is faster than the TC-OTF, albeit by just a factor of two. We deem this increased cost more than acceptable, given the considerable savings in storage space provided by the TC-OTF, the performance independence from the graph topology and the faster storage of new links. For these reasons, all recent versions of AiiDA implement only the TC-OTF.

Event versus polling-based engine. To evaluate the performance of the event-based engine of AiiDA 1.0 compared to the polling-based one of earlier versions, we consider an example work chain that performs simple and fast arithmetic operations. The work chain first computes the sum of two inputs by submitting a `CalcJob` and then performs another addition using a `calcfunction`. For the `CalcJob`, the `ArithmeticAddCalculation` implementation is used, which wraps a simple Bash script that sums two integers. Each work chain execution then corresponds to the execution of three processes (top work chain, a calculation job and a calculation function) and is representative of typical use cases. For each benchmark, 400 work chains are submitted to the daemon and the rate of submission and process completion is recorded, as shown in Fig. 8. These benchmarks were performed on a machine with an Intel Xeon E5-2623 v3 CPU with 64GB of RAM. Beside the results obtained for the old and new engine using optimised parameters (number of workers, transport intervals, ...), for a fair comparison and to highlight the effect of different engine types, in Fig. 8(a) we also show the results for the new engine with some artificial constraints. In particular, we run the new engine with four workers only (which is roughly comparable with the old engine, with four independent tasks for submitting, checking queued jobs, retrieving files, and processing work chain steps) rather than twelve. Additionally we set a minimal interval between connections of 5 seconds in the new daemon to simulate the polling behaviour (with default polling time of 5 seconds) of the old daemon, despite all calculation jobs being run on the local host where an interval of zero is optimal.

Figure 8(a) shows that the submission rate of the old engine is slightly faster compared to the new engine, because the procedure was significantly simpler with no communication with remote daemon workers. Nevertheless, the total completion time of the new engine (even in the constrained configuration) is shorter, with the optimised new engine completing all processes three times faster than the old one in this simple example. Additionally, for the old engine all work chains complete towards the end of the time window at roughly the same time, in a few discontinuous jumps because of the polling-based architecture. In contrast, the completion rate of the event-based engine is much smoother (beside being faster in general) because of the continuous operating character of the new engine, with processes managed concurrently and executed immediately, without waiting for the next polling interval.

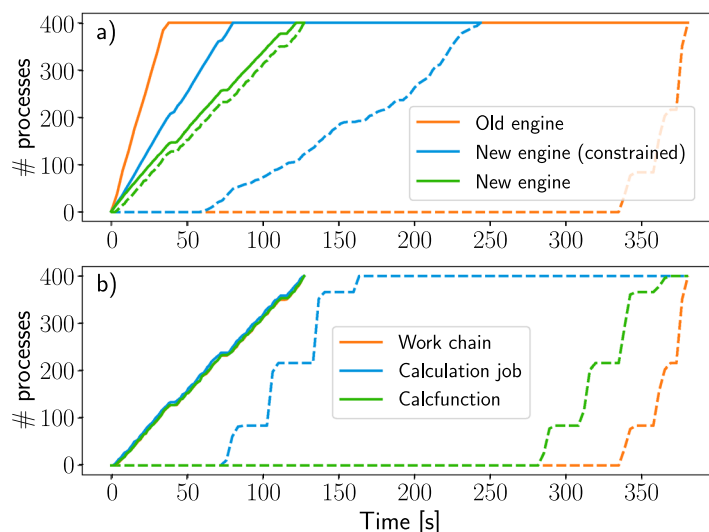


Fig. 8 Process submission and completion rates for the old and new engine. (a) Number of submitted (solid lines) and completed (dashed lines) processes over time for the new engine (both with optimised parameters and with artificial constraints, see text) and the old engine. The submission of the old engine is slightly faster, but despite this the completion rate of the new engine is clearly higher, even under constrained conditions. (b) Number of completed processes for the old (dashed lines) and new (solid lines) engine, decomposed in the separate (sub)processes. The polling-based nature of the old engine is clearly reflected in the stepwise behaviour of the completion rate with processes being finalised in batches. In contrast, the curves for the new engine, due to its event-based design, are smooth and closely packed together, indicating processes being executed in a continuous fashion.

The concurrency of the new daemon is highlighted even more in Fig. 8(b) where we compare the completion time for the new (optimised) daemon and old one, showing independently the completion for the top work chain and each of the two subprocesses. The reason of the large delay in the old engine is because, even though the workflow only runs two subprocesses, the internal logic consists of multiple steps, where only one is processed per polling interval. In contrast, the new engine executes all workflow steps in quick succession without interruption.

We stress that the efficiency improvements of the new engine are even larger in real high-throughput situations, since the daemon is never idle between polling intervals. Most importantly, the new daemon is scalable and the number of daemon workers can be increased dynamically to distribute heavy work loads. This effect is made visible in Fig. 8(a), where the optimised new engine (with 12 workers and without connection delay) completes all processes in half the time required by the constrained one. The effective throughput of the new engine for this experiment, which was run on a modest work station, amounts to roughly 35000 processes per hour. Due to the scalable design of the new engine, this rate can be easily increased by running more daemon runners on a more powerful machine.

Caching. The storing of complete data provenance as described in “The Provenance Model” does not only guarantee the reproducibility of results, it can also reduce the unnecessary repetition of calculations. If the engine is asked to launch a calculation, it can first check in the database if a calculation with the exact same inputs has already been performed. In that case, the engine can simply reuse the outputs of the completed calculation saving computational resources. This mechanism is referred to as caching in AiiDA and users can activate it for all calculation types or only for specific ones.

To rapidly find identical calculations in the database, one needs an efficient method to determine whether two nodes are equivalent. For this purpose, AiiDA uses the concept of hashing, where the whole content of a node is mapped onto a single short hexadecimal string. In AiiDA we employ the cryptographic BLAKE2b algorithm (blake2.net), which has a relatively low computational cost combined with an overwhelmingly unlikely probability of hash collisions. The latter property means that any two nodes with the same hash can be assumed to be identical. The content of a node that is included in the computation of its hash consists of the immutable node attributes and the file repository contents. In addition, for a calculation node the hashes of all its inputs are also included, such that looking for calculations with identical inputs can be done merely by looking at the hash of the calculation itself.

As soon as a node is stored and it becomes immutable, its hash is computed and stored as a node property, making it queryable. When the engine is asked to launch a new calculation, it first computes its hash and searches the database for an existing node with the same hash. If an identical calculation is found and caching is enabled, the engine simply clones the output nodes of the existing calculation and links them to the new calculation. This saves valuable computational resources and results in the same provenance graph as if the calculation had actually been run. Nevertheless, specific extra properties are added to indicate which calculation was used as the cache source, making it possible to identify cached calculations (mostly for debugging purposes).

The concept of caching is especially powerful when developing and running complex workflows consisting of many calculations. If any calculation fails, the workflow that launched it fails as well. If the error can be resolved (e.g., because it was due to a bug in the workflow), the workflow can be fixed and simply rerun from scratch: thanks to the caching mechanism, it will effectively continue from where it previously failed without repeating successful calculations.

The Plugin System

The AiiDA ontology and provenance graph are designed to be employed in any field of computational science, and therefore they are intentionally domain agnostic. To enable users to extend core functionalities to suit the needs of a specific discipline, AiiDA provides a flexible plugin system to add custom data types, to interface external simulation codes with specific input generators and parsers, to implement custom workflows, and more. These extensions are registered upon installation through entry points, as explained in “Architecture Overview”, which allows them to be developed and installed completely independently of the AiiDA code base. To promote sharing of plugins, AiiDA provides an online registry (see “Registry”) where plugin packages can be registered, discovered and downloaded.

Architecture. The plugin system builds on top of the `setuptools` project (setuptools.readthedocs.io), currently the *de-facto* standard tool for bundling and installing Python packages. `Setuptools` provides a feature called “entry points”, which are handles to specific Python resources (e.g., a class or function) of a package. A Python package can define these entry points (categorised in entry point groups) such that, once installed, the corresponding resources are registered and become automatically discoverable and usable by any other package. AiiDA leverages the entry point system by defining several groups specific to the type of resources that a plugin can extend, for example `aiida.data` and `aiida.workflows` for new data types and workflows, respectively (ten groups are currently defined, to extend support to new codes, parsers, job schedulers, transport protocols to connect to remote computers, external databases, etc.). When AiiDA plugin packages register plugins in the appropriate entry point groups, AiiDA automatically discovers them and makes the functionality available to users, integrating it in its core infrastructure. Plugin packages are encouraged to namespace their entry points by prefixing them with the package name to avoid overlap with entry points of other packages.

Installing a plugin package can be performed with a single command of the `pip` (pypi.org/project/pip) Python package manager if the package is published on PyPI (pypi.org). We emphasise, however, that publication on PyPI is not required. Local or private packages can be installed just as easily and they can also register plugin entry points, giving developers full freedom on how to maintain and distribute their plugin packages.

Registry. The AiiDA plugin registry (aiida.team.github.io/aiida-registry/) is an online resource with a list of all known plugin packages. This centralised overview makes plugins easily discoverable by users of the AiiDA community and encourages code sharing and reuse. Authors can register their plugin packages through a pull request to the registry repository providing minimal information like package name, development status, links to the code repository and documentation, as well as a URI pointing to a JSON file maintained by the plugin developers. The latter contains additional information in `setuptools` format, such as the name and description of the project, the authors list and the entry points that are provided. Based on the information provided by the registered packages, the registry website is automatically built through continuous deployment. The plugin system in combination with the registry provides a powerful and effective tool to allow users and developers to easily create and share extensions of AiiDA’s core functionality.

Community Building

Science is a collective enterprise and scientific reproducibility can only be realised through a concerted effort of the scientific community. Making a software tool like AiiDA available is an important first step towards improving the reproducibility of computational science, but it is not enough. One needs to facilitate the uptake of the tool by the community, such that the data produced by it become interoperable and reusable. We have undertaken various approaches to build a community around AiiDA, such as guaranteeing the quality and robustness of the code, ensuring the longevity of data produced with AiiDA, and actively promoting knowledge transfer by online and offline training of users and developers through tutorials and workshops. Additionally, we are now member of NumFOCUS (numfocus.org), an organisation that promotes open science and open research data.

Code quality, testing and continuous integration. AiiDA’s source code is hosted on GitHub (github.com) and all code development contributions, both from internal and external contributors, go through its pull request system. These pull requests facilitate the review and improvement of the suggested changes before they are accepted into the main code base. On top of this quality control performed by peers, we have enabled an automatic continuous integration testing system³⁰. Using GitHub actions (github.com/features/actions), each code commit triggers the running of a suite of over 1000 unit and integration tests that verify that the changes do not break existing functionality. These measures are crucially important to guarantee the longevity of the project as the codebase and number of contributors keep growing. The GitHub repository not only serves as a code hosting platform, but also facilitates discussion, interaction and collaboration between all contributors through its issue and milestone tracker, and the project boards.

Data longevity. With data longevity we refer here to the possibility of accessing data produced with earlier code versions from newer versions. This is an important aspect of any data infrastructure and it is particularly critical to AiiDA. Indeed, AiiDA aims at improving the reproducibility of computational science simulations and thus data should be accessible also years after it has been produced. However, AiiDA also focuses on being

performant for high-throughput workloads (see “Performance”), which often requires changes to existing data layouts.

To guarantee the longevity of data, AiiDA can automatically migrate data across code versions without the need of human intervention. Particular care has been devoted to implement robust database schema and data migrations for both database backends (Django and SQLAlchemy, see “Database Abstraction and Querying Language”) with 42 migrations per backend implemented in v1.0. These migrations ensure that early AiiDA users can seamlessly migrate their databases generated years ago to the most recent version without losing any of their data or provenance. Additionally, to guarantee that all migrations are correct and thus avoid the potentially irreversible corruption of user data, the integrity of each migration is verified by automated individual unit tests as described previously. While the development and maintenance effort for these migrations is significant, they are essential for data longevity and with that, for the uptake of the platform by users, by giving them confidence that their data will be accessible in the future. Similar migrations and the corresponding tests have been implemented for AiiDA archive files that contain exported (parts of) AiiDA databases. Since these archive files are not concerned with the database schema but merely with the data itself, the database migrations cannot be reused, but instead separate migrations in Python had to be developed.

A final contribution to AiiDA’s data longevity is the compatibility of AiiDA 1.0 with both Python 2 and Python 3. Despite Python 2 reaching its end-of-life in January 2020, AiiDA 1.0 supported this version for another 6 months, thereby providing a grace period for users and developers to upgrade their code and plugins to Python 3.

Interoperability. Through its plugin system, AiiDA forms a natural interoperability interface for various simulation codes that is centred around the reproducibility provided by the provenance graph. Any external code that has a command line interface can in principle be run with AiiDA. Other interoperability requirements involve importing and converting data and interfacing to existing libraries. Indeed, we try not to reimplement features provided by existing robust codes, but rather we provide interfaces to them. Specifically, AiiDA comes with built-in support for various widespread computational materials science libraries for data analysis and conversion such as ASE¹⁰, pymatgen³¹, spglib³² and seekpath³³. Additionally, tools are provided to directly interface with existing databases, such as ICSD³⁴, COD²⁴, TCOD³⁵ and OQMD²⁹, and import the data into an AiiDA database. Finally, while AiiDA is a fully independent tool, it integrates seamlessly with the Materials Cloud dissemination platform¹⁹ that provides, among other things, a web application for browsing AiiDA provenance graphs interactively. Powered by the AiiDA REST API, the provenance browser can be used both for AiiDA databases uploaded to the Materials Cloud as well as for those residing on the computer of the AiiDA user.

Outreach. In addition to the aspects outlined earlier, we actively pursued various activities to strengthen the user and developer community of AiiDA. Notably, we organised a significant number of events (see aaida.net/events), with 19 tutorials, schools and workshops over the 2017–2020 period. These targeted both new users, introducing the code and the concepts of provenance and reproducible workflows, as well as advanced developers (with yearly coding weeks, and events aiming to provide direct support to AiiDA plugin developers).

To broaden the accessibility of the educational material presented during these events beyond people attending in person, the tutorial texts are made available online (aaida-tutorials.readthedocs.io) and live recordings of presentations are uploaded to the Materials Cloud¹⁹ Learn section (materialscloud.org/learn/). The tutorial texts are distributed together with virtual machines, based on the Quantum Mobile¹⁹, which contain preinstalled and preconfigured versions of AiiDA, its plugins, and the simulation codes and data, such that no setup is required to follow the tutorial. With this innovative approach, the barrier to access and learn the code is significantly lowered, and interested new users can quickly try out AiiDA and understand if it suits their research needs. In addition to the interactive tutorials, the code comes with extensive online documentation (aaida-core.readthedocs.io) and a mailing list is operated by the core developers to provide direct user support (aaida.net/mailling-list).

Thanks to these efforts, AiiDA has seen a rise in its use and adoption for research projects^{14–18,36–43} (see also aaida.net/science for a more complete overview of recently published work that has employed AiiDA). The community of plugin developers has also grown substantially and as of May 2020 the plugin registry (see “The Plugin System”) hosts 49 plugin packages, including for BigDFT⁴⁴, CASTEP⁴⁵, CP2K⁴⁶, CRYSTAL⁴⁷, FLEUR (flapw.de), Gaussian⁴⁸, GULP⁴⁹, Phonopy⁵⁰, Quantum ESPRESSO²¹, Raspa⁵¹, Siesta⁵², VASP²², Wannier90⁵³, Yambo⁵⁴ and many more. This developer community has also been strengthened by the targeted events that we have organised, for example to help developers migrate their plugins to AiiDA 1.0, to extend support to Python 3, and to design common workflows and interfaces for codes implementing similar methods.

Finally, as AiiDA grows, we have now standardised the approach to provide concrete ideas for improvement and further extensions of AiiDA by the community at large. Inspired by the concept of the Python Enhancement Proposals (PEPs) ([Python.org/dev/peps/](https://python.org/dev/peps/)), we have implemented a repository to host new AiiDA Enhancement Proposals (AEPs) (github.com/aidata/AEP) and a standardised protocol to provide new suggestions. Discussions on each suggested AEP are facilitated by the GitHub platform and remain available also after approval, allowing to go back and review the reasoning that justified specific design decisions. AEPs provide a way to extend the discussions on the future roadmap of the code to all interested users, beyond the pool of core developers.

Conclusions

We have presented AiiDA 1.0, a Python open-source, high-throughput and scalable infrastructure for computational science with a strong focus on automated data provenance. We have highlighted the new design of AiiDA 1.0 that makes it scalable to exascale high-throughput computational loads. AiiDA 1.0 can easily sustain tens of thousands of processes per hour and dispatch them on a broad range of computing resources, from local computers to large high-performance supercomputers. Key to achieving this goal has been the redesign of the workflow

engine from a polling-based to an event-based paradigm. The engine can be scaled on demand to an arbitrary number of workers that operate independently, where communication with and among them is made possible by the RabbitMQ message broker.

In addition to the engine, the design and implementation of the provenance graph have been significantly improved. The concept of workflows is now fully integrated into the provenance graph, including now also the logical steps in the data provenance. The implementation of the provenance graph has been optimised and made more efficient by migrating the storage of node properties to native JSONB column types and by computing the transitive closure on the fly, instead of storing it in a static table. The `QueryBuilder` is a powerful tool that allows users to inspect their provenance graph and extract information, without having to be able to write SQL queries and that provides a simple Python syntax that is independent of the ORM backend used for the implementation. The REST API provides yet another way of extracting data from an AiiDA provenance graph, which is especially useful when direct access to the machine running the instance is not available. Crucially, despite the profound improvements and changes in the code compared to early versions, existing data and their provenance can be automatically migrated, and are therefore guaranteed to remain compatible with the present version of AiiDA.

Finally, the plugin system makes AiiDA a flexible tool interoperable with any simulation software. It is not just limited to calculation plugins; users can create their own data types, command line interface extensions, and workflows, and share any of those. The plugin registry allows developers to register their plugin packages and other users to discover them, which directly fosters a lively community of developers, that has grown beyond the original field of application in materials science, supporting now over 90 codes spanning many fields of research ranging from mechanical engineering, to chemistry and physics. Thanks to advanced tools available within the Python ecosystem, installing plugin packages can be performed with a single command, automatically registering the included plugins with AiiDA. Through these tools, AiiDA provides a platform and community geared towards making computational science more transparent, user-friendly, and ultimately fully reproducible, in full compliance with the FAIR principles⁵.

Data availability

The data used to create Figs. 6, 7 and Table 2 for the analysis of the database performance come from ref. ¹⁴. The raw data is available on the Materials Cloud Archive⁵⁵.

The data used to create Fig. 8, along with the scripts that were used for the creation and analysis of the data, are available on the Materials Cloud Archive⁵⁶.

Code availability

The source code of AiiDA is released under the MIT open-source license and is made available on GitHub (github.com/aidata/aiiida-core). It is also distributed as an installable package through the Python Package Index (pypi.org/project/aiiida-core).

Received: 24 March 2020; Accepted: 19 August 2020;

Published online: 08 September 2020

References

- Ioannidis, J. P. A. *et al.* Repeatability of published microarray gene expression analyses. *Nat. Genet.* **41**, 149–155, <https://doi.org/10.1038/ng.295> (2009).
- Peng, R. D. Reproducible research in computational science. *Sci.* **334**, 1226–1227, <https://doi.org/10.1126/science.1213847> (2011).
- Stoddart, C. Is there a reproducibility crisis in science? *Nat.*, <https://doi.org/10.1038/d41586-019-00067-3> (2016).
- Allison, D. B., Brown, A. W., George, B. J. & Kaiser, K. A. Reproducibility: A tragedy of errors. *Nat.* **530**, 27–29, <https://doi.org/10.1038/530027a> (2016).
- Wilkinson, M. D. *et al.* The FAIR guiding principles for scientific data management and stewardship. *Sci. Data* **3**, <https://doi.org/10.1038/sdata.2016.18> (2016).
- Goble, C. *et al.* FAIR computational workflows. *Data Intell.* **2**, 108–121, https://doi.org/10.1162/dint_a_00033 (2020).
- Pizzi, G., Cepellotti, A., Sabatini, R., Marzari, N. & Kozinsky, B. AiiDA: automated interactive infrastructure and database for computational science. *Comput. Mater. Sci.* **111**, 218–230, <https://doi.org/10.1016/j.commatsci.2015.09.013> (2016).
- Jain, A. *et al.* FireWorks: a dynamic workflow system designed for high-throughput applications. *Concurr. Comput. Pract. Exp.* **27**, 5037–5059, <https://doi.org/10.1002/cpe.3505> (2015).
- Curtarolo, S. *et al.* AFLOW: An automatic framework for high-throughput materials discovery. *Comput. Mater. Sci.* **58**, 218–226, <https://doi.org/10.1016/j.commatsci.2012.02.005> (2012).
- Larsen, A. H. *et al.* The atomic simulation environment—a python library for working with atoms. *J. Physics: Condens. Matter* **29**, 273002, <https://doi.org/10.1088/1361-648x/aa680e> (2017).
- Maffioletti, S. & Murri, R. GC3pie: A python framework for high-throughput computing. In *Proceedings of EGI Community Forum 2012/EMI Second Technical Conference — PoS(EGICF12-EMITC2)*, <https://doi.org/10.22323/1.162.0143> (Sissa Medialab, 2012).
- Adorf, C. S., Dodd, P. M., Ramasubramani, V. & Glotzer, S. C. Simple data and workflow management with the signac framework. *Comput. Mater. Sci.* **146**, 220–229, <https://doi.org/10.1016/j.commatsci.2018.01.035> (2018).
- Babuji, Y. *et al.* Parsl. In *Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing — HPDC 2019*, <https://doi.org/10.1145/3307681.3325400> (ACM Press, 2019).
- Mounet, N. *et al.* Two-dimensional materials from high-throughput computational exfoliation of experimentally known compounds. *Nat. Nanotechnol.* **13**, 246–252, <https://doi.org/10.1038/s41565-017-0035-5> (2018).
- Kahle, L., Marcolongo, A. & Marzari, N. High-throughput computational screening for solid-state Li-ion conductors. *Energy & Environ. Sci.* **13**, 928–948, <https://doi.org/10.1039/c9ee02457c> (2020).
- Mercado, R. *et al.* In silico design of 2d and 3d covalent organic frameworks for methane storage applications. *Chem. Mater.* **30**, 5069–5086, <https://doi.org/10.1021/acs.chemmater.8b01425> (2018).
- Prandini, G., Marrazzo, A., Castelli, I. E., Mounet, N. & Marzari, N. Precision and efficiency in solid-state pseudopotential calculations. *npj Comput. Mater.* **4**, <https://doi.org/10.1038/s41524-018-0127-2> (2018).
- Vitale, V. *et al.* Automated high-throughput Wannierisation. *npj. Comput. Mater.* **6**, 66, <https://doi.org/10.1038/s41524-020-0312-y> (2020).

19. Talirz, L. et al. Materials cloud, a platform for open computational science. *Sci. Data*. <https://doi.org/10.1038/s41597-020-00637-5> (2020).
20. Uhrin, M., Huber, S. P., Yu, J., Marzari, N. & Pizzi, G. Workflows in AiiDA: Engineering a high-throughput, event-based engine for robust and modular computational workflows. Preprint at <https://arxiv.org/abs/2007.10312> (2020).
21. Giannozzi, P. et al. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *J. Physics: Condens. Matter* **21**, 395502, <https://doi.org/10.1088/0953-8984/21/39/395502> (2009).
22. Kresse, G. & Furthmüller, J. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169–11186, <https://doi.org/10.1103/physrevb.54.11169> (1996).
23. Ullmann, J. R. An algorithm for subgraph isomorphism. *J. ACM (JACM)* **23**, 31–42, <https://doi.org/10.1145/321921.321925> (1976).
24. Gražulis, S. et al. Crystallography open database (COD): an open-access collection of crystal structures and platform for world-wide collaboration. *Nucleic Acids Res.* **40**, D420–D427, <https://doi.org/10.1093/nar/gkr900> (2011).
25. Gražulis, S. et al. Launching the theoretical crystallography open database. *Acta Crystallogr. Sect. A Foundations Adv.* **70**, C1736–C1736, <https://doi.org/10.1107/s2053273314082631> (2014).
26. Blokhin, E. & Villars, P. The PAULING FILE project and materials platform for data science: From big data toward materials genome. In *Handbook of Materials Modeling*, 1–26, https://doi.org/10.1007/978-3-319-42913-7_62-1 (Springer International Publishing, 2018).
27. Jain, A. et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002, <https://doi.org/10.1063/1.4812323> (2013).
28. Draxl, C. & Scheffler, M. The NOMAD laboratory: from data sharing to artificial intelligence. *J. Physics: Mater.* **2**, 036001, <https://doi.org/10.1088/2515-7639/ab13bb> (2019).
29. Kirklín, S. et al. The open quantum materials database (OQMD): assessing the accuracy of DFT formation energies. *npj Comput. Mater.* **1**, <https://doi.org/10.1038/npjcompumats.2015.10> (2015).
30. Duvall, P., Matyas, S. M. & Glover, A. *Continuous Integration: Improving Software Quality and Reducing Risk (The Addison-Wesley Signature Series)* (Addison-Wesley Professional, 2007).
31. Ong, S. P. et al. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Comput. Mater. Sci.* **68**, 314–319, <https://doi.org/10.1016/j.commatsci.2012.10.028> (2013).
32. Togo, A. & Tanaka, I. Spglib: a software library for crystal symmetry search. Preprint at <https://arxiv.org/abs/1808.01590> (2018).
33. Hinuma, Y., Pizzi, G., Kumagai, Y., Oba, F. & Tanaka, I. Band structure diagram paths based on crystallography. *Comput. Mater. Sci.* **128**, 140–184, <https://doi.org/10.1016/j.commatsci.2016.10.015> (2016).
34. Belsky, A., Hellenbrandt, M., Karen, V. L. & Luksch, P. New developments in the inorganic crystal structure database (ICSD): accessibility in support of materials research and design. *Acta Crystallogr. Sect. B Struct. Sci.* **58**, 364–369, <https://doi.org/10.1107/s0108768102006948> (2002).
35. Merkys, A. et al. A posteriori metadata from automated provenance tracking: integration of AiiDA and TCOD. *J. Cheminformatics* **9**, 56–67, <https://doi.org/10.1186/s13321-017-0242-y> (2017).
36. Gröning, O. et al. Engineering of robust topological quantum phases in graphene nanoribbons. *Nat.* **560**, 209–213, <https://doi.org/10.1038/s41586-018-0375-9> (2018).
37. Atambo, M. O. et al. Electronic and optical properties of doped TiO₂ by many-body perturbation theory. *Phys. Rev. Mater.* **3**, <https://doi.org/10.1103/physrevmater.3.045401> (2019).
38. Wang, S. et al. On-surface synthesis and characterization of individual polyacetylene chains. *Nat. Chem.* **11**, 924–930, <https://doi.org/10.1038/s41557-019-0316-8> (2019).
39. Mishra, S. et al. Topological frustration induces unconventional magnetism in a nanographene. *Nat. Nanotechnol.* **15**, 22–28, <https://doi.org/10.1038/s41565-019-0577-9> (2019).
40. Li, W. et al. Interface engineered room-temperature ferromagnetic insulating state in ultrathin manganite films. *Adv. Sci.* **7**, 1901606, <https://doi.org/10.1002/adv.201901606> (2019).
41. Abbott, D. F. et al. Design and synthesis of Ir/Ru pyrochlore catalysts for the oxygen evolution reaction based on their bulk thermodynamic properties. *ACS Appl. Mater. & Interfaces* **11**, 37748–37760, <https://doi.org/10.1021/acsami.9b13220> (2019).
42. Mateo, L. M. et al. On-surface synthesis and characterization of triply fused porphyrin–graphene nanoribbon hybrids. *Angewandte Chemie Int. Ed.* **59**, 1334–1339, <https://doi.org/10.1002/anie.201913024> (2020).
43. Stamminger, A. R., Ziebarth, B., Mrovec, M., Hammerschmidt, T. & Drautz, R. Fast diffusion mechanism in Li₄P₂S₆ via a concerted process of interstitial Li ions. *RSC Adv.* **10**, 10715–10722, <https://doi.org/10.1039/d0ra00932f> (2020).
44. Mohr, S. et al. Accurate and efficient linear scaling DFT calculations with universal applicability. *Phys. Chem. Chem. Phys.* **17**, 31360–31370, <https://doi.org/10.1039/c5cp00437c> (2015).
45. Clark, S. J. et al. First principles methods using CASTEP. *Zeitschrift für Kristallographie - Cryst. Mater.* **220**, <https://doi.org/10.1524/zkri.220.5.567.65075> (2005).
46. Hutter, J., Iannuzzi, M., Schiffmann, F. & VandeVondele, J. cp2k: atomistic simulations of condensed matter systems. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **4**, 15–25, <https://doi.org/10.1002/wcms.1159> (2013).
47. Dovesi, R. et al. Quantum-mechanical condensed matter simulations with CRYSTAL. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **8**, e1360, <https://doi.org/10.1002/wcms.1360> (2018).
48. Frisch, M. J. et al. Gaussian–16 Revision C.01 (2016). Gaussian Inc. Wallingford CT.
49. Gale, J. D. GULP: A computer program for the symmetry-adapted simulation of solids. *J. Chem. Soc. Faraday Transactions* **93**, 629–637, <https://doi.org/10.1039/a606455h> (1997).
50. Togo, A. & Tanaka, I. First principles phonon calculations in materials science. *Scripta Materialia* **108**, 1–5, <https://doi.org/10.1016/j.scriptamat.2015.07.021> (2015).
51. Dubbeldam, D., Calero, S., Ellis, D. E. & Snurr, R. Q. RASPA: molecular simulation software for adsorption and diffusion in flexible nanoporous materials. *Mol. Simul.* **42**, 81–101, <https://doi.org/10.1080/08927022.2015.1010082> (2015).
52. Soler, J. M. et al. The SIESTA method for ab initio order-n materials simulation. *J. Physics: Condens. Matter* **14**, 2745–2779, <https://doi.org/10.1088/0953-8984/14/11/302> (2002).
53. Pizzi, G. et al. Wannier90 as a community code: new features and applications. *J. Physics: Condens. Matter* **32**, 165902, <https://doi.org/10.1088/1361-648x/ab51ff> (2020).
54. Sangalli, D. et al. Many-body perturbation theory calculations using the yambo code. *J. Physics: Condens. Matter* **31**, 325902, <https://doi.org/10.1088/1361-648x/ab15d0> (2019).
55. Mounet, N. et al. Two-dimensional materials from high-throughput computational exfoliation of experimentally known compounds. *Materials Cloud* <https://doi.org/10.24435/materialscloud:2017.0008/v3> (2018).
56. Huber, S. P. et al. AiiDA 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance. *Materials Cloud* <https://doi.org/10.24435/materialscloud:2020.0027/V1> (2020).

Acknowledgements

This work is supported by the MARVEL National Centre of Competence in Research funded by the Swiss National Science Foundation (grant agreement ID 51NF40-182892), the European Centre of Excellence MaX “Materials design at the Exascale” (grant no. 824143) and by the Swiss Platform for Advanced Scientific Computing (PASC). Additional support was provided by the “MaGic” project of the European Research Council

(grant agreement ID 666983), the swissuniversities P-5 “Materials Cloud” project (grant agreement ID 182-008), the “MARKETPLACE” H2020 project (grant agreement ID 760173), the “INTERSECT” H2020 project (grant agreement ID 814487), the “NFFA” H2020 project (grant agreement ID 654360) and the “EMMC” H2020 project (grant agreement ID 723867). We would like to thank the following people for code contributions, bug fixes, improvements of the documentation, and useful discussions and suggestions: Oscar D. Arbeláez-Echeverri, Michael Atambo, Valentin Bersier, Marco Borelli, Jocelyn Boullier, Jens Bröder, Ivano E. Castelli, Keija Cui, Vladimir Dikan, Marco Dorigo, Y.-W. Fang, Espen Flage-Larsen, Marco Gibertini, Daniel Hollas, Eric Hontz, Jianxing Huang, Christoph Koch, Ian Lee, Daniel Marchand, Antimo Marrazzo, Simon Pintarelli, Gianluca Prandini, Philipp Rübmann, Philippe Schwaller, Ole Schütt, Christopher Sewell, Andreas Stamminger, Atsushi Togo, Daniele Tomerini, Nicola Varini, Jason Yu, Austin Zadoks, Bonan Zhu and Mario Zic.

Author contributions

The initial contributions to AiiDA are listed in ref. ⁷; here we list the developments and contributions since then. G.P. and N.Ma. supervised and coordinated the project. M.U., G.P. and S.P.H. designed and M.U. and S.P.H. implemented the new engine, with contributions from D.G. in parts of the user interface. B.K. designed the use of process functions to track the provenance of Python functions, and G.P., N.Mo., A.M., M.U. and S.P.H. implemented it. R.H., M.U. and S.P.H. designed and implemented the new daemon. B.K. formulated the extension of the provenance graph definition to include link types. S.P.H. and M.U. designed the current provenance graph formalising the two provenance levels, incorporating the extension of B.K., and implemented it with contributions in both design and implementation by L.T., G.P., A.V.Y. and S.Z. B.K. developed an initial scheme for JSONB node attributes and their queries via the SQLAlchemy ORM. S.Z. oversaw and implemented the abstraction of the backend ORM and the addition of SQLAlchemy support, with contributions by A.C., L.K. and F.G. S.Z. and S.P.H. implemented the use of JSONB fields also with the Django backend. M.U. designed the front-end/back-end ORM separation and S.P.H., M.U., G.P., L.T., S.K., S.Z. and C.J. implemented it. L.K. and G.P. implemented the on-the-fly transitive closure. L.K. designed and implemented the QueryBuilder. L.K. designed the graph traversal tool and implemented it with F.F.R. T.M. and R.H. designed and implemented the plugin system. R.H. redesigned the command line tool and implemented it with S.P.H., S.Z., A.V.Y. and L.T. T.M. added Python 3 support. B.K. suggested the idea of caching of calculations. D.G. designed and implemented the caching mechanism. S.Z., G.P., C.W.A. and A.V.Y. improved the export and import functionality. F.G., S.K. and E.P. implemented the REST API and A.M. designed and implemented the tools to import crystal structures from external databases and to export them to external databases. C.S.A. and L.T. have setup and contributed to the AiiDA Enhancement Proposal (AEP) system. S.P.H., S.Z., G.P., M.U., L.T., L.K., R.H., N.Mo., A.V.Y., C.W.A., F.F.R., C.S.A., C.J., A.M., A.C., F.G., S.K. and E.P. are or have been part of the AiiDA core developers team that maintains the software. N.Mo., A.M. and A.C. integrated various external libraries for data transformation and implemented many command line utilities to export and visualise data. N.Mo., G.P., F.G., S.Z., S.P.H., M.U., L.T., L.K., A.V.Y., R.H., C.W.A., A.M., A.C., S.K. and E.P. contributed to the various tutorials that have been organised. All authors have read and approved the final version of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41597-020-00638-4>.

Correspondence and requests for materials should be addressed to S.P.H. or G.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020