



This is a repository copy of *The environmental impact, carbon emissions and sustainability of computing in the ATLAS experiment*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/235581/>

Version: Published Version

Article:

Aad, G. orcid.org/0000-0002-6665-4934, Aakvaag, E. orcid.org/0000-0001-7616-1554, Abbott, B. orcid.org/0000-0002-5888-2734 et al. (2935 more authors) (2025) The environmental impact, carbon emissions and sustainability of computing in the ATLAS experiment. The European Physical Journal C, 85 (12). 1397. ISSN: 1434-6044

<https://doi.org/10.1140/epjc/s10052-025-14976-3>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>



The environmental impact, carbon emissions and sustainability of computing in the ATLAS experiment

ATLAS Collaboration*

CERN, 1211 Geneva 23, Switzerland

Received: 14 May 2025 / Accepted: 22 October 2025
© CERN for the benefit of the ATLAS Collaboration 2025

Abstract ATLAS, a general-purpose experiment at the Large Hadron Collider (LHC), makes use of a large internationally-distributed computing infrastructure, including over 10^6 TB of managed data on disk and tape and almost one million simultaneously running CPU cores. Upgrades for the High-Luminosity LHC (HL-LHC) will increase the required computing resources by a factor of 3–4 by the beginning of the 2030s, and by an order of magnitude before the conclusion of data taking at the beginning of the 2040s. These resources are spread over around 100 computing sites worldwide. Efforts are underway within the experiment to evaluate and mitigate various aspects of the environmental impact of the sites, with the additional long-term goal of making recommendations to the sites that will significantly reduce the total expected environmental impact in the HL-LHC era. These efforts take several forms: building awareness in the experiment community, adjusting aspects of the computing policy, and modifications of data center configurations, either in ways that take advantage of particular features of ATLAS workloads or in generic ways that reduce the environmental impact of the computing resources. This paper describes the ongoing investigations and approaches that have already provided useful and actionable outcomes.

1 Introduction

As the world continues to experience the increasing effects of global warming [1], the sustainability, environmental impact, and carbon footprint of various aspects of our society have come under scrutiny. A series of efforts have been undertaken at CERN to characterize and reduce the energy, environmental, and carbon footprints [2,3] of the organization's activities. Computing is one component of the larger issue of sustainability in high-energy physics (HEP), and some aspects of this issue have been addressed by the Worldwide LHC Computing Grid (WLCG), the collaboration that encompasses

most of the computing resources used by the LHC experiments today [4].

The ATLAS experiment [5,6] utilizes an extensive distributed computing system detailed in Ref. [7]; the planned expansion of these resources over the next 5 years in preparation for the High-Luminosity LHC (HL-LHC) upgrade is detailed in Ref. [8]. In short, the experiment typically uses around 600 000–700 000 cores of compute, with peaks of more than one million cores. Over 10^6 TB of storage is in use, comprising over 400 PB of disk and 600 PB of tape.¹ These resources are divided among about 100 distributed computing centers around the world, including WLCG sites, high-performance computing (HPC) centers, volunteer computing, and various other resources.

By 2030, when the HL-LHC upgrades are complete, the average required computing capacity of the experiment will be around 1.5–3 million cores; that number is expected to more than double by the end of the experiment's life in 2041. Similarly, the disk storage needs are set to exceed $1\text{--}2 \times 10^6$ TB in 2031 and the tape needs to exceed 2×10^6 TB, with both doubling again by 2041. This growth of about an order of magnitude over 15 years implies a significant increase in the carbon footprint of the experiment's computing effort unless mitigated. Given this long-term foresight of the implications of the HL-LHC upgrade for energy and carbon footprints, CERN and the WLCG have the unique opportunity to strategically mitigate the expected increase over several years, to grow computing resources in as sustainable a way as possible.

This paper summarizes the current understanding of the likely evolution of the carbon footprint of future compute demand and the ongoing efforts to reduce the overall envi-

* e-mail: atlas.publications@cern.ch

¹ The storage here is network-attached storage and does not include the scratch space generally associated with a compute node, which can be about 20 GB per core or an additional 20 PB worldwide. The compute includes active compute elements and does not include interactive compute, service nodes or compute associated with disk servers, amongst other items. Neither resource includes personal computing (e.g. laptops).

ronmental impact of computing for the ATLAS experiment, through both policy and direct action. Some of these efforts can be directly translated into recommendations for site administrators, developers, and users towards this reduction. These studies are intended to help inform choices in the computing model to be used in the HL-LHC and to guide sites in their purchase decisions. The recommendations and policies are also intended to reduce the environmental impact of computing without having any meaningful impact on the physics program of the experiment. The policies, recommendations, studies, and examples may also prove useful to other experiments considering ways to reduce the environmental impact of their computing.

The sustainability efforts within ATLAS computing can be divided into four groups. First, as described in Sect. 2, efforts have been made to improve awareness among users, developers, site administrators, and managers of issues around sustainability. Second, experiment policies and standard practices can be adjusted to minimize the environmental impact of computing, often without any harm to the physics program of the experiment. These policies are discussed in Sect. 3. Third, there are some actions that site administrators can take to take advantage of the particular patterns and features that ATLAS workloads and site use follow. These actions are described in Sect. 4. Fourth, there are several actions that site administrators can take themselves to lessen the environmental impact of their site. This last category includes things that are not specific to ATLAS, or even to high-energy physics, but are of general applicability. These ideas are being studied and developed by many other groups, like the EE HPC WG [9], the Open Compute Project [10], Green DiSC [11], GreenDIGIT [12], the UKRI Net Zero Digital Research Infrastructure Scoping Project [13], Boavizta [14], or the International Energy Agency [15], and several ATLAS sites have been deeply involved as well. The goal is therefore not to replace or reproduce, but to complement the existing efforts outside of ATLAS. The ongoing investigations associated with ATLAS sites are described in Sect. 5, and possible improvements through knowledge sharing are described in Sect. 6. All of these efforts and approaches must be pursued if the environmental impact of the collaboration is to be minimized.

The most common mechanism for estimating organizational carbon emissions is the Greenhouse Gas (GHG) Protocol Corporate reporting standard [16]. This protocol divides carbon emissions into three “scopes”:

1. Direct emissions, including fuel consumption, refrigerants, chemicals, and gasses;
2. Mainly indirect emissions due to generation of electricity used on site, and also due to the acquisition of steam, cooling or heating from offsite; and
3. Other indirect emissions, a catch-all for any other emissions that occur in an organisation’s value chain, including up-stream and down-stream activities.


Although computing centers have minimal direct (Scope 1) emissions, chip manufacturers’ Scope 1 emissions may be considerably larger. The Scope 1 and 2 emissions of computing hardware manufacturers are generally accounted for as a part of the Scope 3 emissions of a data center. The cost associated with the production of hardware is referred to as “embodied” or “embedded” carbon.

Scope 2 electricity emissions can be reported as “Location-based,” depending on the mean carbon intensity of the grid in the location where it is used, or “Market-based,” reflecting emissions based on intensities of contractual purchases of electricity. Here location-based emissions are used, because they more directly reflect emissions associated with electricity use. Scope 2 emissions present an interesting challenge for ATLAS computing, because the resources are distributed over the world, and the large number and variety of sites each have different power efficiencies, energy practices, and power grid properties. These all vary with time, as well, and the decarbonization of the power grid in various countries is important to account for in extrapolations of the carbon footprint of computing. The lifetime model of typical WLCG sites is also important to consider: sites tend to retain and operate hardware well beyond its period of warranty, producing a significantly different amortization for embodied carbon compared to a commercial data center.

CERN has estimated the Scope 2 emissions from its own computing facilities at around 5% of the total footprint of the laboratory when the LHC is operating, with the majority coming from the cooling of the accelerator complex [2,3]. As CERN operates about 15% of the ATLAS computing resources, and many non-CERN resources rely on power grids with significantly higher carbon intensity than that of France (from which CERN gets its power), the total Scope 2 carbon footprint of the ATLAS computing resources may rival that of CERN itself. The worldwide distribution of ATLAS computing is shown in Fig. 1, based on the amount of CPU provided in HS23, the computational power based on the 2023 HEPscore benchmark (see Sect. 4.1 for the precise definition of HS23).

One critical metric in the discussion of power consumed by computing is the Power Usage Effectiveness (PUE), which is defined as the ratio of power supplied to the building housing the computing resources to that used for computing (i.e. useful energy). Typical PUEs for sites on the WLCG are about 1.5 [4]; a “good” PUE is 1.1, which has been achieved at the new CERN data center, for example [3]. For most of this paper, the units used for estimation are kilograms of CO₂ equivalent, or kgCO₂eq. This is the standard unit used for car-

Fig. 1 The worldwide distribution of ATLAS computing, based on the amount of CPU provided, in HS23 (see Sect. 4.1 for the precise definition of HS23), on average in 2023–2024. Countries in gray did not contribute significant CPU



bon footprints, and requires the translation of all gas outputs into a single unit.

The most commonly discussed and estimated aspect of the environmental impact of computing is due to the power consumption of the hardware and corresponding cooling. Historically, operational (location-based Scope 2) emissions of large data centers have been reported as larger than Scope 3, which originate in particular from hardware manufacturing and infrastructure. More recent studies, however, indicate that the latter may be more significant, especially as grid carbon intensities reduce and more comprehensive and accurate embodied emissions methodologies and data become available (see, for example, Ref. [17]). The implication is that other aspects of the carbon footprint and environmental impact of computing, beyond the power consumption, cannot be overlooked when building a more holistic, forward-looking and comprehensive model. Recommendations based solely on the impact of changes in power consumption, which ignore other aspects like Scope 3 emissions or power grid decarbonization, may prove misleading and inaccurate when a more complete model is considered.² Estimates of embodied carbon often come with significant uncertainties, due to deficiencies or difficulties in manufacturer reporting. As this work progresses, even a comprehensive corporate emissions assessment of Scopes 1–3 might not be sufficient when evaluating the global environmental impact of computing, which also includes water consumption, the use of large volumes of rare or difficult to extract elements, waste management and other socio-environmental impacts [17]. In addition, utilising carbon impact assessment methods more suited to informing policy development and design changes may become more appropriate than the standard corporate reporting approach. For this paper, however, the focus is left on the corporate carbon footprinting approaches, in particular including Scope 2

and Scope 3 emissions associated with embodied emissions of equipment and data center construction.

2 Improving awareness

One of the most important ways to reduce waste in the environmental impact of ATLAS computing is to improve the awareness of environmental issues, and to make members conscious of the positive impacts their decisions and actions can have. Workloads on the grid are classified as either “production” jobs, which include data processing and Monte Carlo simulation dataset production, are run by experienced production teams, and usually follow well-validated patterns, and “user” jobs, which are submitted by any users in the collaboration and may contain almost any code. User jobs typically comprise 10% of ATLAS computing HS23 and 30% of jobs; the jobs are much shorter on average than production jobs. Problematic user workloads are a significant source of waste, and the efforts towards education are meant to serve as a reminder of the consequences of choices made — not to shame users, but to remind them that their decisions are important.

To that end, PanDA [18], the workflow management used by ATLAS, has been equipped to provide estimates of the carbon footprint of all ATLAS jobs. The system draws power grid carbon intensity in $\text{kgCO}_2\text{eq/kWh}$ for each site from Electricity Maps [19], and sites can provide an estimate of the number of Watts per core of compute at their site. This estimate is a stand-in for the entire power draw of a job: the compute itself, as well as the associated share of cooling, storage, networking, and any other relevant hardware. A default value of 10 W per core was assigned, based on Ref. [20].

Users are then provided an estimate of the carbon footprint of their job, along with a link to more information about the calculations [21]. An example of the report that a user receives is shown in Fig. 2. Rather than using the site-specific power grid carbon intensity, the reported carbon footprint estimate uses an average of the global power grid carbon intensity across all ATLAS sites. This is for several reasons:

² In principle, this discussion applies equally to the estimate of the environmental impact of other aspects of ATLAS. However, the focus here is on the relative components of the environmental impact of computing, rather than in the impact of computing relative to the impact from other contributions.



Summary of TaskID: [42140307](#)

Detail	Value
Created	2024-11-21 00:21:50
Ended	2024-11-21 01:13:41.620986
Final Status	done

Total Number of Inputs

Category	Count
Succeeded	1
Failed	0
Cancelled	0
Total	1

Error Dialog: None

Datasets

In	[]
Out	['user.johndoe.prunroottest1_foo.root/']
Log	['user.johndoe.prunroottest1.log/']

Parameters

Command	prun --exec="root -b -q HistTest.C" --nJobs=1 --outputs=foo.root --outDS=user.johndoe.prunroottest1 --rootVer=6.32.06 --cmtConfig=x86_64-el9-gcc13-opt
---------	--

Estimated Carbon Footprint for the Task

Category	gCO2
Succeeded	0.06 gCO2
Failed	0 gCO2
Cancelled	0 gCO2
Total	0.06 gCO2

More details on estimation: https://panda-wms.readthedocs.io/en/latest/advanced/carbon_footprint.html

Fig. 2 An example email of a job report from PanDA including carbon footprint information. The carbon footprint information is boxed at the bottom of the example email, and includes a link to a website providing more information about the calculations [21]

1. Site data is still in its infancy and is relatively unreliable. For example, sites may have solar panels installed or may be paying a premium for renewable energy. The PUE of sites varies significantly, and it is likely that 10 W is an overestimate for some sites and an underestimate for others. Until the site data are complete, the average avoids assigning an incorrect impact to a site. This 10 W estimate also stands in for the power consumption of other IT elements required by the jobs, including network and storage. A more detailed understanding of the real power consumption is required before this can be considered truly reliable.
2. Using the average ensures a more direct connection between the optimization of a user's code and the carbon footprint of that code. Faster running code results in a smaller carbon footprint, ignoring global time variations of power grid carbon intensity. This is important to correctly incentivize and reward positive user actions.
3. Users might respond to site-specific data by redirecting their jobs to sites with lower power grid carbon intensity. In fact, computing resources are not left idle: any cores not taken by user jobs will be consumed by production jobs (e.g. Monte Carlo simulation). Thus, users redirecting their jobs will not affect the number of cores of computing running in any region. Instead, in fact, the impact might be negative: many users redirecting jobs to a single site will cause significant network usage, churning of data on disks, and congestion at the site. The result may be both a longer waiting time for the users and a larger overall environmental impact for the experiment.

Nevertheless, the site-specific numbers are recorded, and are available for internal study and improvement. Although the resources will not sit idle (i.e. if the user job is faster, other workloads will be run), this gives users a useful metric for the carbon footprint of their scientific result.

To ensure that new users are accustomed to seeing these reports, the standard ATLAS analysis software tutorial that most new members attend includes submission of a failing job. The resource waste is minimal, but ensuring that users have some experience seeing what a failing job looks like, analyzing it and understanding the source of the failure is crit-

ical. There were cases identified of users retrying failing jobs an enormous number of times, and the educational program is intended to help avoid such situations in the future. To avoid substantial waste from large numbers of retried jobs, several cutoffs have also been introduced. Configurable limits have been set on the allowed fraction of failed jobs and HS23s, as well as maximum numbers of task retries and attempts for an individual job. Additionally, tasks with very low CPU efficiency are stopped, to prevent jobs from reserving more CPU cores than they require.

The same feature in PanDA can be used to provide reports for large production campaigns run by the experiment. Production managers can see the reported carbon footprint for their preferred selection of jobs. These reports separate the carbon footprint of successful and failing jobs: when a campaign has a significant failure rate, the environmental impact of the failures is visible. Similarly, if a campaign has to be abandoned or redone because of a major software issue, the environmental impact of the lost production is visible. These reports are not meant to prevent the generation of useful data for the experiment, but rather to ensure that appropriate care is taken in validating and setting up new production campaigns.

In a similar manner, the recasting of many experiment activities in an environmental sustainability light has been beneficial towards ensuring that more experiment members consider them thoroughly. Software validation, continuous integration testing, and the validation of the physics output of new software releases can all be understood in sustainability terms: issues missed in validation are costly. Historically, there were cases of software bugs identified far too late into production, only after many millions of events had been processed. The cost of appropriately scaled validation (e.g. frequently-run fast jobs, infrequently-run slower jobs) is insignificant in comparison to the CPU time lost to production with bugs. Furthermore, attempts to optimize software are useful for resource savings, but are also relevant in environmental terms.

Towards raising awareness in site administrators, a survey was done to understand current practices across both large and small sites. One of the issues ATLAS faces is that while most resources are provided by a relatively small number of large sites, the site count is dominated by many small sites, including small computing clusters at institutes. These small sites often lack the capabilities of large sites: more than half of the computing sites surveyed do not monitor their electricity usage, while all the major WLCG sites do. More than 60% of the site administrators indicated a desire for advice and the sharing of best practices towards reducing environmental impact.

3 Experiment computing policies

The ATLAS computing model has many parameters and policies that can be adjusted that affect the environmental impact of the experiment. To date, these policies and the model itself have been mostly examined through the lens of financial savings. Often financial and environmental benefits are connected: simply reducing overall resources is beneficial in both respects. Changes to the policies of the experiment are often the most effective towards sustainability: they can be taken immediately, without requiring site changes or new purchases, for example.

Policies often allow the experiment to consume one resource type and save another, sometimes with an extremely beneficial trade-off. One simple example is the compression of data. Compressing ATLAS data is an obvious thing to do: the enormous savings in storage, often a factor of three or more, hugely outweigh the modest CPU required to compress and decompress the data (1%–10% of the run time, depending on the workload). However, once the initial step has been taken, there is still considerable room for further optimization: the compression algorithm can be changed, with the constraint that it must be sufficiently robust to not risk the long-term integrity of the data, and the algorithm settings can be changed to find the best working point given the expected frequency with which the data will be read [22]. The data layout in storage can also significantly affect the resources required by the experiment. Lossy compression, the dropping of insignificant bits, can save 5%–15% of the storage space required by data analysis formats, but requires careful validation to ensure that physics outputs are not distorted. The adoption of the new ROOT RNTuple format, which has no loss but only improves the on-disk data layout, has the potential to save the experiment 15%–20% of its storage space, and potentially even more for frequently-read analysis data formats [23]. These changes and their rapid adoption after validation can have a significant impact on the experiment's computing resource needs and thereby on its environmental footprint.

One example of a less obvious connection between financial and environmental benefit is the “data carousel” [7] used by the experiment. The data carousel is a means to using tape storage, which is normally purely archival, as a more active storage medium. Infrequently used datasets are archived to tape, deleted from disk, and recalled in the future if they are required. This reduces the overall disk requirement of the experiment at the cost of an increased tape storage requirement. Because tape is a less carbon intensive storage medium per TB in terms of both operational carbon and embodied carbon [24], the policy also reduces the overall environmental impact of the experiment.

Generally speaking, the resources in a single year suffer from the Jevons paradox [25]: the available compute is

intentionally kept busy with scientific work that offers useful improvements of results. This is not only due to the increased demand that can be satisfied by faster or more efficient workloads, but also to honor the cost of investment of the funding agencies. Racks of CPU are expensive, and it is generally thought wise to always keep them busy with lower priority work that provides at least marginal scientific benefit rather than allowing them to sit idle. Idle CPU presents an opportunity to improve the scientific output of the experiment.

Over longer timescales, however, resource requests are made by the experiment based on extrapolations of resource needs. If some software is made significantly more efficient or faster, the resources will still be fully utilized, but a more modest expansion will be requested for subsequent years. It is thereby possible for software optimizations to reduce the total carbon footprint of the experiment over the long term. Continuous investment in software optimization and efficiency improvements is therefore valuable for environmental, as well as financial and scientific, reasons. These efforts continue to pay dividends: the efficiency of 8-core Monte Carlo simulation production jobs, measured by the ratio of CPU time to wall-clock time, has risen by about 5% over the last three years, to over 90%. The time per event required by Monte Carlo simulation has also been reduced by almost 50%, thanks to the validation and deployment of a wide variety of optimizations [7], which is reflected in reduced need for CPU at the WLCG sites. Similarly, significant improvements in event generation software have led to substantial improvements, in some cases 50-fold speed improvements [26,27]. With such substantial speed improvements, the collaboration can decide how much to reinvest in samples with greater physics precision, and by how much to reduce the overall computing resource budget.

3.1 Data reproduction

Another example of an experiment policy that has both a financial and an environmental impact is related to data reproduction, particularly for data formats produced for the final stages of data analysis, called derived analysis object data, or DAODs. These data formats are individually small, but often together comprise more than 25% of the total data that ATLAS retains on disk. New versions are frequently produced, in some cases several times per year due to software or calibration improvements, and old versions are deleted when they are no longer needed. When data analysis teams are close to completing a publication, they will often ask that the particular version used for the preparation of the publication is retained. Often this retention is “just in case”: if during collaboration or journal review the analysis team is asked to check an issue that requires accessing DAODs, they need access to these files. Because these files will be deleted in some relatively short period, it is ineffective to migrate them

to tape: tape repacking is a slow process, and frequent deletion of small datasets from tape creates operational problems at sites.

To reduce the disk footprint of these files that are retained “just in case”, functionality was built into the distributed production system that allows the reproduction of a dataset on demand. This functionality can deliver an identical copy of the deleted data, in a time not much longer than what would be required to recall the data from tape. This 1–2 day production, for modest dataset of a few TB, is sufficiently fast to not induce significant delays in the publication review. The question remains, however, whether the use of dataset reproduction is beneficial from an environmental standpoint.

Because of their different embodied and operational carbon footprints, the relative cost of disk and compute varies by site. Based on data from several sites (see also Ref. [14]), the number of Watt-seconds per CPU core-second is about 100 times the number of Watt-seconds per kB-year of storage. Typical compute resources require about 7.3 W / core ,³ and typical storage requires about 0.06 Ws / kB y (see for example Ref. [28]). That means that, based on the average size of the data formats (30 kB/event) and time for production (20 HS23 s/event including all overheads, or about 1.6 s/event on a typical processor; see Sect. 4.1 for more details on HS23), the creation of an event takes about 11.7 Ws of energy, compared to 1.9 Ws of energy per year to store the event. The break-even point of a replication strategy is the point at which the compute and storage have equal cost. In this case, the break-even point is a replication rate of $1.9/11.7 = 16\%$ for a year of prolonged storage. In other words, saving the data for an additional year costs the same as reproducing 16% of the data. Therefore, any reproduction rate below 16% would mean that reproduction on demand is an environmentally effective solution. If the data are accessed many times, would need only to be saved for a short period, or would have to be reproduced with high probability, reproduction on demand is environmentally harmful.

3.2 Automatic computational waste reduction

The experiment has several policies and automated systems that help to automatically reduce computational waste (unproductive or invalid compute jobs). This can be achieved in several ways, for example by taking sites offline if they are problematic, or by stopping tasks (collections of jobs with the same configuration and either different random number seeds or different input files, for example) if they are problematic. One example of the first is HammerCloud [29]. Among its

functions, HammerCloud is used to continuously test ATLAS computing sites with a variety of well-understood jobs that are expected to succeed. About 0.05% of the total CPU used by ATLAS in 2023 was used for these functionality tests. If HammerCloud detects a problematic number of these test jobs failing, it automatically takes the site offline and notifies the administrators in a process called “auto-exclusion”. Were the site to remain online, large numbers of jobs directed to it would fail. If the compute at the site sits idle, the power consumption of the CPU is reduced by about 50%. If the compute can be used for other purposes (e.g. if a failure only affects ATLAS production jobs, and local user batch system jobs can still be run), then the entire power consumed by the CPU can be considered saved (i.e. not wasted). HammerCloud takes sites offline for about 4.6% of a year on average through auto-exclusion, which can be thought of as an upper bound on the reduction in total energy required thanks to its application. Several issues can reduce its effectiveness, like bad nodes that cause an otherwise good site to be taken offline, or the inability to fill the CPU with other useful work while it is offline for ATLAS. The key to this savings is accurate, precise, and timely auto-exclusion, which is the focus of the ongoing sustainability efforts within HammerCloud.

Within PanDA, there are additional “scout” and “retry” schemes that are designed to prevent waste by reducing failures within a task. For each task, the first 10 jobs, normally a small fraction of the total, are considered “scouts” and are run before the remaining jobs are released into the system. If these scout jobs fail at a high rate, the task can be stopped before significant resources are consumed by failures. The scouts can also be used to estimate the resource needs of the jobs in terms of CPU, memory, and disk space, so that jobs can be more optimally packed into existing computing resources.

When jobs fail, the failures are diagnosed and compared to a set of rules for retrying jobs. Automatic retries help enormously in reducing the load on production managers that would otherwise have to take manual actions. The rules ensure that if a job fails for a reason that appears to be site-related, like a corrupted input file, it can be moved to a different site; if the job fails for a reason that appears random, like an issue that could have been due to memory pressure from other jobs on the node, it might be retried 3–4 times; and if the job fails for a reason that is clearly reproducible, the job is not retried further unless experts intervene. The rules for retrying jobs need regular review and maintenance, but they can help reduce the amount of waste in the system when properly tuned.

3.3 Unused data

Another effort designed to reduce waste is an examination of unused data. These are data that, to the knowledge of the

³ This differs somewhat from the 10 W/core discussed in Sect. 2 because the 10 W/core estimate is a stand-in for additional computing elements like network and job-associated storage, and because this uses site-specific power estimates.

distributed data management system, have never been used as input for a job, downloaded by a user, or transferred by a user to another site. For example, a user might request production of a specific DAOD format for their analysis based on a regular expression pattern. If that expression is too general, DAODs might be produced for datasets that are never examined by the analysis team. Similar issues can arise when DAOD formats are run together in a single job (called a “train”): if a train is run on a dataset that is only needed in some of the DAOD output formats, the other DAOD formats will appear unused. These issues have to be addressed by carefully manually going back through unused data, in each case understanding why it was produced, and tightening any rules or systems through which spurious data were created.

A common source of these unused datasets is connected to software issues in production. If the physics output of a production is inaccurate, a new production might be required. If an analysis team notices an issue after looking at a few datasets, without having looked at the rest, they might request a fresh production. If the original production is not deleted, it will show up later as unused. This can be the case when teams do not report data as invalid to the data management team, and instead simply request a new production and assume that the old data will be deleted in due course. Again, such issues have to be identified, particularly when they result in large quantities of unused data, by systematic follow-up by production managers.

3.4 Alternative resource usage

ATLAS endeavors to make full use of all resources that it has been allocated. For example, the high-level trigger farm, a cluster of CPU deployed at the experiment site for real-time data processing and event selection, comprises almost 2.3 MHS23. This CPU farm is critical for operations, but when ATLAS is not taking data it can be used for processing of Monte Carlo simulation instead [7], saving the experiment the equivalent of an additional 60 000 cores of compute. Similarly, the CERN production site has an allocation dedicated to prompt processing of the data as it streams from the detector. The site, known as the CERN Tier 0, is sized according to the need during peak collider operation. Again, during a significant fraction of the year it is not needed for prompt detector data processing and can therefore be used for other workloads, saving the experiment the equivalent of about 75 000 cores of compute worldwide.

One aspect of the experimental policy has relatively poorly understood environmental implications: the use of HPC systems and commercial clouds in comparison to more traditional high-throughput computing (HTC) at WLCG sites. The use of commercial clouds is still modest, around 1% of the CPU time of the experiment, while the use of HPC sys-

tems is extensive, approaching one third when both pledged and unpledged CPU is included.⁴ The environmental impact of WLCG sites can be controlled to some degree, but the experiment has no control over the standard practices of commercial cloud providers. Some of the standards, like the long-term use of computing hardware in WLCG sites, are not normally disclosed by commercial cloud vendors. HPC systems normally have a relatively short lifetime of order five years, and the equipment is often then sent back to the vendor for reuse, recycling, or disposal. On the other hand, commercial sites tend to procure renewable energy or have on-site sustainable power generation. A detailed comparison of the environmental impacts of these systems would likely prove difficult due to non-disclosure agreements and proprietary information, but as more information becomes publicly available it may be worth undertaking. Environmental considerations are unlikely to dictate more or less extensive use of these resources, as financial considerations are critical, but understanding the implications would be useful nevertheless.

3.5 Open data

Finally, it is important to consider the implications of some aspects of the experiment’s computing policy on the sustainability of the field as a whole. ATLAS has recently begun to release Open Data for research, along with a significant volume of Monte Carlo simulation corresponding to those data. In addition, the experiment is planning to release billions of events, corresponding to hundreds of millions of CPU core-hours, of event generator output to the research community and general public. Oftentimes, substantial effort is wasted and computing resources are consumed within the phenomenology community duplicating the event generation that ATLAS has already run. Small groups tend to work in isolation, sometimes producing similar samples or simply trying to reproduce the samples that the experiment has already created. By releasing these data, ATLAS can help reduce the environmental impact of the high-energy physics community at large [30].

4 ATLAS-specific site and user actions

While some hardware actions are generally beneficial and broadly applicable (see Sect. 5), the particular policies and patterns applied within ATLAS computing offer opportunities for site administrators and developers to work together

⁴ WLCG sites may “pledge” compute resources to ATLAS as an indication of what they expect to be able to provide in the coming year. Any extra resources if a site exceeds that pledge, and all resources from any computing sites that do not pledge resources are considered “unpledged” [7].

to reduce the environmental impact of the experiment. These actions may not be generally applicable, but they offer many opportunities and pathways for other groups to follow to examine their usefulness.

4.1 HEPsScore and power measurement

To evaluate the computing power of resources, ATLAS employs the HEPsScore benchmarking suite [31]. This suite includes seven standard workloads from the LHC experiments, including ATLAS. These workloads are chosen to provide a “typical” mixture of low-level instructions, memory loads, and other features so that the score is a more accurate reflection of a CPU’s capabilities for ATLAS workloads than a standard benchmark like LINPACK [32]. The HEPsScore from 2023 (HS23 score) is important to understand which CPUs and configurations provide the best throughput per Watt, as well as the best throughput for the same cost. It is also one of the key tools the collaboration can use to evaluate the performance of new platforms. The benchmarking suite is intended to test CPU and memory combinations, and therefore to limit disk (input or output) and network usage. For some workloads, particularly user analysis, this makes it less representative, but for the most time-consuming workloads the suite provides a fair characterization of resource use.

Considerable effort had to be put into understanding the power-use profile of the benchmark suite in comparison to standard production, to correctly calculate the score per Watt. For example, the power consumption of a CPU, shown in Fig. 3, shows significant spikes and troughs during some workload initialization periods. If the simple peak power consumption is used for the score, it will not correctly reflect typical production. Similarly, the average is artificially low due to the frequent start/stop periods of the benchmark, which do not occur in normal production jobs that run continuously for 12 hours. Various metrics were examined, and the most stable and simple to calculate metric is “q85”, the 85th-percentile power measurement.

The load on the processor — how much work is running simultaneously — can also affect its performance in ways that depend on the workload. The effectiveness of hyper-threading, for example, in improving the performance per Watt depends on the workload and its use of the CPU. There are also a variety of power regulation mechanisms in modern CPUs that can affect the scaling of power consumption with load. Data from the WLCG sites, shown in Fig. 4, reveal two distinct scaling regimes in the evolution of HS23 score with load. These patterns are tied to specific configuration factors, such as turbo boost settings and thermal shielding, which can vary depending on the CPU brand and generation.

To investigate these aspects more thoroughly, dedicated, isolated tests were conducted on bare-metal servers (i.e. not

running inside of containers) using different CPU frequency governors. These governors set the policy for automatically changing the CPU frequency during operation to improve performance or to save power, depending on the goals of the user. The benchmarking was performed on AMD EPYC 7302 16-Core CPUs and Intel® Xeon® E5-2630 v4 2.20 GHz CPUs. The number of threads was scanned over multiples of four, to approximate the behavior observed in the WLCG site data. The results are shown in Fig. 4, using data from Ref. [33]. Three governors were tested: “performance,” “ondemand,” and “powersave.” For both CPU models, the powersave governor consistently produced about 50% lower performance results, while only reducing power consumption by about 45%. The HS23 per Watt, therefore, is lower when using the powersave governor. Examining the data from the WLCG sites, most servers align with either the performance or ondemand characteristics — two modes that tend to overlap. Notably, the performance governor prioritizes maximum output regardless of power cost, which explains its consistently slightly higher scores than the ondemand governor. The precise behavior varies depending on many factors including turbo boost settings, CPU thermal shielding, and CPU model. This observation underscores how performance can be directly influenced by CPU configuration and governor choice. In environments where energy efficiency is paramount, the powersave or ondemand governors may offer the best trade-off, depending on the specific needs and workloads.

These studies are also key to the evaluation of new hardware platforms. The recent adoption of ARM processors within ATLAS computing was the result of significant work motivated by both financial and power usage considerations [34]. ATLAS computing has also demonstrated the use of commercial cloud resources for development of new platforms [35,36]. Using cloud resources allows the collaboration to carefully evaluate new platforms, run productions at modest scale, and perform detailed validation without having to purchase resources that might be wasted if the validations are unsuccessful or the resources prove more costly than existing options. Eventually, to get detailed power monitoring information, it is currently necessary to procure resources. If in the future commercial cloud providers offer better access to this information, it may be possible to evaluate new platforms entirely in the cloud before purchases have to be made.

Understanding the power usage of CPUs across generations is also critical to understanding the optimal approach to hardware retirement. The issue is still not well understood: whether it is better to run old hardware and amortize the embodied carbon over a longer period, or to retire and donate or recycle old hardware to use only new, more power-efficient CPUs. The optimum depends on the relative importance of operational and embodied carbon, and there-






Fig. 3 The power consumption of a typical processor during the running of the HEPsCore benchmarking suite. Each point represents a power measurement. Seven different sets of three runs of individual workloads chosen to be representative and described in Ref. [31] are distinguished with dashed vertical lines

Fig. 4 The evolution of HS23 score with CPU load for (a) AMD and (b) Intel CPUs. The gray points show data from the WLCG sites, while the colored series connected with lines indicate specific choices of frequency governor on benchmark machines. The data are from Ref. [33]



fore on the carbon intensity of the power grid where the site operates. WLCG sites tend to run hardware well past its warranty lifetime; a modest fraction ($< 10\%$) of CPUs running ATLAS software are more than 10 years old (see for example Refs. [4, 37]). Often site administrators extend the lifetime of hardware by scavenging parts from failing nodes, and in some cases hardware is donated to groups running less resource-hungry workloads. Several efforts are underway to understand the issue of optimal hardware lifetimes in more detail (see e.g. Ref. [38]).

4.2 Power variation and frequency adjustment

One issue that can be addressed through collaboration between ATLAS site administrators and software developers is the variation in power grid availability and carbon intensity. The variation of carbon intensity is visible already in many regions: in certain periods of the day, renewable energy is more broadly available, and the carbon intensity of the power grid is lower. Other periods see reduced production of renewable power, particularly solar and wind. The load on the power grid also varies during the day. As the renewable power sources expand over the coming years, it is expected that during peak production periods the power production from renewables will significantly outstrip the demand. This is shown in Fig. 5, for the case of modeled projections for Germany in 2030, based on data from Refs. [39, 40]. There are several ways in which sites can take advantage of this variation, by increasing production when renewable energy is abundantly available and might even be curtailed if not used, and reducing production and therefore power load when it is not.

Modern CPUs have controls for frequency adjustment. Lower frequencies generally result in reduced work and reduced energy use, but because many ATLAS jobs are restricted by memory access patterns rather than simply CPU cycles, the reduction in throughput is less than linear. Additionally, the power usage of a CPU is proportional to frequency and voltage squared, and when the CPU frequency is reduced the voltage is also reduced. Both of these effects result in the throughput per Watt being significantly better for ATLAS jobs on CPUs with reduced frequency. This improvement can be seen in Fig. 6, showing the HS23 score (as a proxy for throughput) per Watt for various processors as a function of CPU frequency. In all cases, running the CPU below its normal design frequency is beneficial to improving the HS23 score per Watt. The improvement and best set point depends on the CPU model, and therefore must be tested with each procurement.

CPU scaling offers a straightforward way for site administrators to immediately increase the throughput per Watt of their site. Although there is an improvement in HS23 per Watt, the total throughput per second of the site will be

reduced. To provide the same throughput per second, more CPUs will need to be purchased. The optimal working point balances the power costs against the cost of additional hardware, and balances the reduction of environmental impact from reduced power consumption with the increased embodied carbon from the additional hardware. This optimum is site and time dependent and must be more carefully examined for future purchases. The same arguments apply to GPU frequency settings as well [41].

In extreme cases, the variation in power grid availability can result in brownouts (drops in power grid voltage or temporary blackouts) or blackouts of sites, whether scheduled or announced on relatively short notice. In these cases, the site may be taken completely offline. Currently, sites must either “drain” all processing completely or kill all running jobs before such a shut down to avoid leaving many tasks in a corrupted state. ATLAS grid jobs typically run 12–24 hours, with tails on some sites up to 3–4 days. To completely drain a site, no new work can be submitted for at least that period, which can cost hundreds of thousands of CPU-hours for a large site. To avoid the loss of processing, checkpointing (i.e. taking a snapshot of the current state of a running program, storing it on disk, and later restoring the running program from that snapshot) could be implemented within the software. Several ATLAS production workloads are already capable of checkpointing, but the functionality has not been built into the production system to allow sites to send appropriate signals to the jobs, reload their states, and correctly report the job status and resource usage in the monitoring systems. This additional work is non-trivial but could be advanced in preparation for the HL-LHC to improve site operations and efficiency. Checkpointing has other applications as well. For example, if a site is to be taken offline for a regular upgrade, the CPU can continue to run until just before the scheduled upgrade and can be promptly brought back into a running state after an upgrade. If the use of checkpointing produces inconsistencies in the system that causes all checkpointed jobs to fail, the loss of CPU will be no worse than the current system.

Several sites, including commercial cloud providers, are looking into alternative powering mechanisms to resolve the renewable power availability issues. One outstanding question is whether the development of battery systems, currently used for emergency power (e.g. UPS systems) will develop substantial additional capacity in the coming years, such that they could be used to power the entire data center during periods of high-carbon-intensity power production, and then could be recharged during periods of very low-carbon-intensity power availability. Such a battery system could have significant embodied carbon, but might be able to substantially reduce the operational footprint of the site. Similarly, some commercial providers are looking into miniature nuclear reactors to provide local low-carbon energy




Fig. 5 An extrapolation of renewable power sources in Germany and power demand to 2030, based on Refs. [39, 40]. The load (demand), on-shore and off-shore wind, and solar power production are scaled based on data from an example week at the end of May, 2023 and pro-

jections for consumption and renewable energy construction projects. Other sources include hydroelectric, biomass, geothermal, and waste-to-energy production




Fig. 6 The HS23 score (a proxy for throughput) per Watt as a function of the processor frequency for seven CPU models. The 85th percentile power consumption during the benchmark running is shown

sources [42]. In all cases, the commercial development of these options will significantly outpace HEP-specific developments, and ATLAS sites can take a wait-and-see approach to the technological advances.

Two additional advantages of ATLAS computing in comparison to many other sites and operations are consistency and latency. CPUs are kept busy at all times when they are available, including CPUs provided opportunistically to the experiment beyond its standard resources. Excluding resources based on sustainability considerations would impose financial costs that the collaboration could not afford

today. The consistency means that resources can be scoped for typical use, with a small operational safety margin, rather than being scoped for a peak well above typical use. The acceptance of latency means that if demand is high, generally work can be prioritized rather than requiring that resources are expanded temporarily. In the case of networking equipment, for example, the continuous stream of transfers around the world means that their power consumption is steady, and the embodied carbon is only what is required by the typical use pattern and safety margin.

5 General site actions

Many actions that sites can take towards sustainability and reducing environmental impact are not specific to ATLAS, or even to HEP. Groups worldwide are investigating these issues (see for example Refs. [9, 11, 12, 43]). Several of these groups have close connections to the ATLAS experiment: some have used ATLAS sites as examples or case studies; in other cases ATLAS site teams have directly been involved with the studies. Specific examples are included in this section to do with storage optimization, data center construction, and cooling. Many general recommendations are parameterized in terms of common operational parameters like average load and hardware lifetime, and optimal working points may be site- or region-specific.

Several ATLAS sites are also investing effort in broader examinations of power and environmental impact studies. Site power profiles vary substantially depending on the site, its capabilities, and its configuration. Storage often makes up 25%–45% of the total site power consumption. Where tape storage is provided, the system normally consumes another 10% of the site's total power. The remainder is primarily consumed by CPU, with a very small fraction for additional services like networking elements (see, for example, Ref. [44] for a UK study). In Germany there are several ongoing efforts to understand and reduce the carbon footprint and environmental impact of computing centers (see for example Refs. [45, 46]). These efforts are now being harmonized across the WLCG, where dissemination of best practices and standards has been given the highest priority (see for example Ref. [47]).

One example of a generic site action that was motivated by ATLAS studies is the introduction of more flexible memory allocation in site queues. Most sites provide 2 GB of memory per CPU core; historically, jobs requiring 4 GB of memory were assigned two CPU cores, and one of the cores would sit idle. By introducing a greater range of memory requirements in queues, ATLAS can submit low-memory work (e.g. requiring only 1 GB of memory) and maintain the 2 GB per core average while still running a substantial number of higher-memory jobs. This reduces the amount of idle CPU at the site that draws power but does not produce useful work.

The ATLAS WLCG sites are more directly under the supervision of the experiments. They are generally responsive to recommendations and requests from ATLAS, and the administrators are collaborative and deeply connected to the experiment. The diversity of ATLAS resources, however, means that not all sites providing resources are so well integrated. About one third of ATLAS CPU is provided by HPC systems and volunteer resources, many of which were constructed to serve communities and use-cases other than high-energy physics. While the experiment can still offer rec-

ommendations and suggest actions to these sites, a broader community must come together to influence them.

5.1 Storage optimization

The carbon footprint of storage systems remains an understudied area, with relatively few papers focused on reducing Scope 2 (operational) [48–51] and Scope 3 (embodied) [52–54] emissions of storage systems. Recent studies highlight the significant embodied carbon footprint of storage systems, contributing approximately 21%–42% to the data center's embodied carbon footprint [53]. Strategies like extending storage device lifetimes or using denser devices require further investigation to assess their trade-offs with operational carbon emissions. Research on reduction of the operational carbon footprint has primarily concentrated on shifting of foreground and background tasks to use low-carbon power opportunities.

The operational carbon footprint of storage systems can be reduced by temporal shifting of background tasks to low-carbon periods. Background tasks such as deduplication, backup, and scrubbing are essential for maintaining storage reliability and efficiency. These tasks usually handle large volumes of data and are deferrable. In Ref. [48], a workload model for background tasks was built to understand the I/O, compute and power costs of these tasks. In this paper, the background tasks workload model is applied to three HEP site scenarios that follow available site configurations: the Tier-0 (the CERN data center, “CERN DC”), a large Tier-2 site (University of Chicago HEP, “UChicagoHEP”), and all sites worldwide (“WLCG”). These tasks consume 1.3%–2.7% of the overall data center power and 7.3%–17.8% of the storage power, as shown in Fig. 7. The difference between results can be attributed to variations in the amount of storage relative to compute, storage configuration and management policies, which substantially influence the power consumption of background tasks. For instance, the storage capacity of WLCG and the CERN data center is 45–50 times greater than that of the Tier-2 site. Tasks like data scrubbing, which process all stored data, constitute a significant portion of background task power consumption. Many background tasks also do not scale linearly with storage volume, which contributes to the observed differences between sites [48].

The operational carbon footprint of the background tasks in the CERN data center scenario is about 49.4×10^3 kgCO₂eq/year. The CERN data center is powered by the French power grid, where nuclear power is the dominant source of electricity generation and therefore a low average carbon intensity (ACI) is typical. Here, the fluctuations in ACI are lower compared to power grids where renewable energy is the dominant power source. Several other power grids within the United States are considered for modeling the variation in carbon intensity from different power sources:




Fig. 7 Power for background tasks as a percentage of (a) data center power and (b) data center storage power, for three HEP site scenarios

CAISO (California), MISO (Midwest and South), ERCOT (Texas), and SPP (Central South). The carbon emissions of background tasks for a site like the CERN data center would be 3.5 times larger in CAISO and 8.2 times larger in MISO than in France. This is due to different power generation sources. For example, in CAISO, solar energy is the predominant renewable energy source. The ACI is lowest during daylight hours when solar generation peaks. These fluctuations in ACI present an opportunity to reduce the storage system's operational carbon footprint by scheduling background tasks during periods of lower carbon intensity.

Fig. 8 shows the percentage reduction in carbon emissions, calculated using ACI values, relative to the baseline emissions for the CERN data center site scenario based on current scheduling policies, for the aforementioned power grid scenarios. The evaluation is based on the methodology described in Ref. [48]. “Naive” denotes the usual configuration of scheduling background tasks around midnight. “Naive Random” represents scheduling tasks at random times of day. The results of these policies depend on the ACI value at the task running time and increase carbon emissions in some scenarios. “Greedy Constrained” optimizes for carbon while scheduling tasks within a constrained time period (e.g. the tasks must run daily). It achieves a carbon emission reduction of 60.6% in CAISO and 41.3% in SPP. With no task period constraints, “Greedy No Constraint” achieves up to 82.4% reduction of emissions in the CAISO power grid scenario. This is closer to the ideal scenario of “Zero Carbon”, where tasks are scheduled only during zero carbon periods. In 2035, with the rapid decarbonization of power grids [55], the variation in ACI will increase, and these policies achieve greater reductions in emissions relative to the naive policies, up to 96.4% in the CAISO power grid scenario.


In short, the temporal shifting of background tasks to low carbon periods is an effective approach to reducing the operational carbon footprint, as estimated using the ACI, of storage systems. It can achieve a decrease in the carbon emissions of data center storage of up to 7.8% with the “Greedy Constrained” approach and 10.6% with the “Greedy No Constraint” approach.

Storage configurations are highly site-specific, and one of the other ongoing efforts is to understand whether there are opportunities for configuration improvements that might reduce the environmental impact of the sites with minimal operational disruption. Simple differences in the redundancy of storage configurations (RAID levels) and drive layouts can produce significant changes in the power per TB of available storage. Similarly, many sites have reported that the CPU load on disk servers is relatively low. Some therefore run additional production jobs on CPU server head nodes to take advantage of the additional computing power [56]. Although it is unclear whether these studies and actions will generalize to non-HEP workloads, they have already proven useful when applied to LHC or ATLAS sites.

5.2 Data center construction

The carbon footprint of construction is a complex issue, with many variations in estimates of supply chain carbon in particular (see, e.g., Ref. [57]). Particularly in countries with relatively clean energy grids, embodied carbon (construction carbon) may represent a significant fraction of the total footprint of a building, even for office buildings with 50-year life expectancy (much longer than a typical data center) [58]. The primary source of carbon during construction is the sourcing of materials [59], with significant opportunity for the

Fig. 8 The percentage reduction in carbon emissions of background tasks using several approaches to scheduling, for the CERN data center, for (a) the power grid models of 2023 and (b) the power grid models of 2035. The various power grid models and scheduling strategies are defined in the text



improvement of the overall carbon footprint. The machinery used to gather and supply the primary materials, as well as those used to construct the building, also contribute significantly to the total footprint, as these are often difficult to power with green energy [60]. The materials most commonly used in data center construction, steel and concrete, are among the most expensive in terms of carbon footprint.

Many studies were done of the total carbon footprint of building construction, for a wide variety of construction projects [61]. Steel-reinforced concrete buildings, like many data centers, typically have carbon footprints of 200–800 kgCO₂eq/m² [61,62]. Several efforts to characterize and reduce the total carbon footprint of data center construction have been undertaken, including LEED data center certification from the U.S. Green Building Council [63].

When constructing a new data center, the main consideration is the change in PUE. This change leads to a simple relationship for the amortization of a building by comparing the total carbon footprint of the construction of the building to the footprint of the power consumption of the building. The carbon saved per year can be written as:

$$\frac{\Delta \text{PUE}}{\text{PUE}_{\text{old}}} \times \text{Power} \times \text{GCI}, \quad (1)$$

where ΔPUE is the change in the PUE from the old to the new building, PUE_{old} is the PUE of the old data center, Power is the annual power draw of the old building, and GCI is the power grid carbon intensity. This can be compared to the total building construction carbon footprint to estimate the number of years before a new data center “pays off.” This assumes that the computing hardware is simply moved from the old building to the new building rather than buying entirely new hardware, and that the carbon cost of transportation between the two buildings is small (generally true because of the short distances between the two). In some cases, particularly for significant expansions of computing resources, entirely new computing hardware is purchased. This misses the opportunity for carbon savings, but often is justified by other considerations (e.g. the importance of hardware uniformity in the new data center).

Typical large modern data centers include around 10 000 m² of floor space. The GCI varies significantly by country; for French power in 2024 (similar to what CERN uses) generation cost 0.033 kgCO₂eq/kWh; for California in 2024 it was

about 0.227 kgCO₂eq/kWh, typical in the USA [19]. Power draw for a large modern data center is typically around 5 MW. Some examples of annual carbon savings for a new data center are:

- For a 5 MW data center in France, changing the PUE from 1.6 to 1.1:

$$\frac{0.5}{1.6} \times 5000 \text{ kW} \times 8760 \text{ h/year} \times 0.033 \text{ kgCO}_2\text{eq/kWh} = 452\,000 \text{ kgCO}_2\text{eq/year} \quad (2)$$

- For a 5 MW data center in California, changing the PUE from 1.2 to 1.1:

$$\frac{0.1}{1.2} \times 5000 \text{ kW} \times 8760 \text{ h/year} \times 0.227 \text{ kgCO}_2\text{eq/kWh} = 829\,000 \text{ kgCO}_2\text{eq/year} \quad (3)$$

For a 10 000 m² data center, a 500 kgCO₂eq/m² carbon cost, the middle of the range above, gives a total construction carbon footprint of 5×10^6 kgCO₂eq. For the French example, this suggests a total amortization time of about 11 years. For the California example, this suggests a total amortization time about about six years.

As the grid steadily de-carbonizes, the carbon cost from operation will be reduced; although goals are lofty, one can expect that this is unlikely to reach a factor of two in time for the HL-LHC, particularly considering global evolution. However, there are significant opportunities to reduce the carbon cost of new construction [59]. In essence: if a new data center can be brought online with a significant PUE reduction, even of order 0.1 in regions where the GCI is high, it is very likely to be worthwhile in terms of the total carbon footprint of computing. Similarly, any renovations that make significant change to the PUE are likely to have a rapid net benefit on the total carbon footprint of the data center.

5.3 Cooling systems

The PUE of a data center is tightly coupled to the design of the building, as well as the type of cooling systems used both in the building and for the computing. The most common options for computing are currently air cooling, using fans and pressure differentials to force air flow around the computing units, and water cooling, circulating cool water throughout the cooling racks to cool the computing. Various more exotic cooling systems, including liquid immersion cooling, are being investigated worldwide, but these tend to have more significant capital costs and maintenance issues that make them less desirable for standard WLCG sites.

Liquid (water) cooling systems are generally much more efficient than air cooling systems, and they can cool racks with significantly higher total thermal design power. There are significant capital costs associated with the transition to liquid cooling, but these can be overcome by cost savings from reduced power needs and careful planning [64]. Many individual examples are available, most indicating some significant savings in power from the transition to liquid cooling [65,66]. For modern buildings that are designed with energy efficiency in mind, where the climate is temperate, and therefore the data center PUE is quite low, the transition from air to water cooling may not result in a significant savings. However, for older data centers with higher PUE, or in places with warmer climates even for part of the year, the transition to water cooling may result in quite significant carbon and cost savings. Staged transitions to liquid cooling solutions can also be beneficial. For example, direct liquid cooling with perimeter computer room air handlers can save 15%–20% of the total data center energy consumption compared to air cooling alone. Moving from perimeter computer room air handlers to close-coupled air-cooling can save an additional 15%–20% of the data center energy consumption [67].

The transition from air to water cooling has some ancillary benefits as well. For example, water-cooled systems tend to be much quieter, and the lack of a “hot aisle” makes them much more pleasant to access and maintain. Air-cooled systems often run around 30° C, making the reuse of heat from these systems considerably more difficult and expensive (see also Sect. 5.4). The transition to water cooling seems to have no significant down-sides, and should be undertaken wherever practical.

A transition from air to liquid cooling is being explored at one of the largest ATLAS Tier-1 sites, hosted at Brookhaven National Laboratory. The pilot project also includes the reuse of data center heat for buildings and domestic hot water. The eventual deployment of liquid cooling, after optimization and heat re-use, is expected to save 700 MWh of energy and 150 000 kgCO₂eq/year [68]. In these cooling systems, there is significant scope for optimization by adjusting temperature set-points.

Generally speaking, running a warmer system requires less energy for cooling. However, modern CPUs include complex temperature regulation functionality that automatically lowers the frequency to avoid overheating. One concern, therefore, is that a higher temperature set-point would result in lower data throughput. To investigate this issue, the HEPsScore benchmarking suite was run on a test rack with 26 Dell x86 compute nodes, where rear-door heat-exchanger fan speeds were varied from 20% to 100% of normal. With the lowest fan speed the ambient air remained just below the maximum threshold defined by the CPU manufacturer of 35° C. As shown in Fig. 9, in all configurations the HS23

score was identical, implying the higher temperatures had no significant impact on data throughput. By raising the data center cooling set point by 10%, which allows temperatures of all compute nodes to increase, a 15% savings in energy can be achieved without compromising computing throughput.

5.4 Waste heat usage

One of the few opportunities for computing to truly save carbon is in the reuse of waste heat. The computing systems generate significant heat when running. Historically, this heat was expelled into the atmosphere or, in colder climates, used to warm the data center itself, particularly offices associated with the data center. More recently, systems have been designed to capture waste heat and put it to use for warming buildings or for warming water to be distributed more broadly, particularly using liquid cooling [69]. This is a practical solution only for data centers that are in reasonably well-populated areas; the transmission of waste heat over long distances may require significant time to be of value, particularly after accounting for the embodied carbon of the additional infrastructure. Still, data centers can provide a stable, year-round heat source, and the heat exchange systems can be quite efficient, recapturing 80% of the heat generated by the computing systems [70]. In some regions of the world, heat reuse, or at least the exploration thereof, is required by law for new data centers; grants are also often available to assist with the infrastructure costs [69].

The carbon savings from waste heat usage depends on several factors, for example:

- The local environment, and how much heat is needed domestically or for local office buildings;
- The typical local sources of heat, and how carbon-intense those sources are (e.g. an area heated by coal-fired plants would see much larger savings than an area heated by geothermal activity);
- The efficiency of the heat exchange system, which depends on the cooling system used for the data center.

One way to quantify the impact of waste heat usage is through a modification of the PUE, sometimes also referred to as the Energy Reuse Efficiency (ERE), defined as [65]:

$$\begin{aligned} \text{ERE} &= \frac{\text{Building Power} - \text{Extracted Energy}}{\text{IT Power Consumption}} \\ &= \text{PUE} - \frac{\text{Extracted Energy}}{\text{IT Power Consumption}}, \end{aligned} \quad (4)$$

where the Extracted Energy represents a correction to the PUE and assumes that the carbon cost for heat generation is the same as the carbon cost for energy generation more

broadly. Using this metric, it is possible to achieve an effective ERE below 0.6 — a significant energy savings overall.

This discussion generally assumes that the infrastructure costs (embodied carbon) for a heat exchange system is comparable to other waste heat disposal infrastructure, or that it is a relatively small component of the total embodied carbon cost of building infrastructure when a new data center is constructed or significant renovations are undertaken. Based on the components involved, this is likely a reasonable assumption in many areas of the world relevant to the WLCG.

6 Sharing knowledge


One final source of computational resource waste that can be reduced to some degree is the rate of failing jobs. Simply examining the failure rates across sites is insufficient to understand where issues lie, however. It is critical in examining these failures to distinguish between “site failures” that are the responsibility of the site administrator or local configuration and “software failures” that are the responsibility of the production teams or users that have submitted the jobs.

Some sites offer unusual resources like very high memory nodes or GPUs, and therefore see different failure rates owing to different mixtures of workflows. Similarly, some sites are configured to accept jobs at different rates: the Tier-0 computing site at CERN accepts mostly very well understood job configurations and therefore has a failure rate below 2%, while general user analysis jobs lose about 15% of the total wall clock time to failures. Across the distributed production system, non-user jobs lose about 6% of wall clock time to failures, an improvement from 10% three to five years ago.

Worldwide, sites have many different hardware and software configurations that can affect operations. The use of network-mounted storage or the size of a cache for software on the worker nodes, for example, can affect the running of jobs. Comparing the rate of failures on different sites can therefore provide useful information to site administrators about what configuration changes might be made to the site to improve its efficiency. To make this comparison fair, however, the variability of workloads across sites has to be accounted for.

Fig. 10 shows a comparison of the wall clock time in HS23s at ATLAS sites lost to failing production jobs over the last year, both in absolute terms and as a fraction of the total compute time provided to ATLAS by the sites. In addition to the observed lost wall clock time, an expected loss is shown, which is derived based on an average failure rate. For each site, this expectation is calculated using the number of hours of each workload+core-count combination run at the site (e.g. 8-core data reconstruction or 64-core detector simulation) and applying a world-average failure rate for that workload+core-count combination. If a site runs only well-

Fig. 9 Ratio of the modified HS23 score to the nominal for a selection of data center nodes, for several fan heat exchanger operating speeds relative to the nominal



understood workloads, the expected loss to failures should be very low. If a site runs complex workloads known to be prone to failure, the expected loss to failures will be high. If a site runs unique workload+core-count configurations, as might be the case for HPC systems that are capable of running jobs with very high core counts, then its expected failures will match its observed failures. Because of their variability, user jobs are excluded from this comparison.

It is clear from this comparison that some sites have failures significantly below, and some significantly above, what would be expected based on their job composition. This first study is sufficient to identify sites that are far from expectation. The next step is to go in some detail through the site configurations for these sites to try to identify issues that might be leading to significantly higher- or lower-than-expected failures, and to produce a set of recommendations for the sites based on these comparisons.

7 Summary and ongoing studies


The studies described in this paper represent the current state of understanding of issues around sustainability and the environmental impact of computing in ATLAS. Many more studies are underway. This work is not in isolation: collaborations with a wide variety of groups are ongoing, including work within the wider WLCG, with the goal of broader recommendations in line with those from Sect. 5, and with an eye towards a general life-cycle assessment for LHC computing. As companies recognize the importance of these issues, embodied carbon estimates should become more reliable, allowing a more nuanced and complex model of carbon emissions and environmental impact. Recognizing that many ongoing efforts and strategies in the experiment have an environmental impact is one of the key aspects of these examinations. In preparation for the expansion of resources that will come over the next few years, the policies and prac-

tices of the experiment and the plans of the computing sites can be examined through the lens of environmental sustainability to provide a more complete picture of the impact of changes and new development. The efforts in computing are only one aspect of the broader sustainability efforts within the ATLAS Collaboration, which are already proving profitable in other areas [71].


The computing landscape is also changing, and the environmental “best path” for the experiment is expected to evolve with some of those changes. For example, as power grids worldwide decarbonize, the importance of embodied carbon (Scope 3) in comparison to operational carbon (Scope 2) will continue to rise. Changes to hardware production may also change the calculation: chip makers worldwide are improving their sustainability practices as well as the efficiency of their hardware. More significant changes may come as well, like changes to chip packaging to integrate more components inside a single package, which can improve power efficiency at the cost of modularity and hardware lifetimes. Within the WLCG, further effort is needed to develop concrete scenarios for the evolution of these external factors, and to monitor their evolution in the coming years.

As the experiment diversifies its computing resources, it is important to consider carbon as one of the metrics by which new technologies are judged. GPUs, for example, offer significant promise for the acceleration of workloads and for increasing the density of compute. However, if they are not used with a sufficient load over their lifetime, the total carbon cost may be higher than the cost of simply buying additional CPUs. The rise of machine learning and AI is driving significant increases in data center environmental impact worldwide, and several efforts are ongoing to understand and mitigate these impacts (see for example Refs. [15, 55, 72]). If commercial cloud providers are willing to offer more detailed power monitoring information, it might be possible to thoroughly evaluate new hardware technologies in power and carbon footprint terms without having to purchase and deploy

Fig. 10 (a) The computing time in HS23s at a given site lost to failing production (non-user) jobs in the last year. The observed loss is shown along with an expected loss based on the weighted world-average failure rates. (b) The same loss as a fraction of the HS23s provided by the site. In both cases, sites are ordered in the number of HS23 they contribute to ATLAS, from (left) most to (right) least



(a)



(b)

them in a WLCG site. In terms of site diversification, ATLAS already runs software on “traditional” high-throughput compute farms, HPC centers, commercial clouds, and personal computing resources through volunteer programs. The carbon intensity of these various types of site should be evaluated where possible, in as complete a way as possible. For example, HPC centers offer dense, high-performance computing, but on machines that often have short lifetimes.

The environmental impact of computing in ATLAS is expected to feature more prominently as the upgrades for the HL-LHC approach. The large-scale expansion of resources offers an opportunity for major changes in purchasing patterns, usage, and policies. These practices also serve as a testbed for future experiments, where computing challenges may be even more complex.

Acknowledgements We thank CERN for the very successful operation of the LHC and its injectors, as well as the support staff at CERN and at our institutions worldwide without whom ATLAS could not be operated efficiently. The crucial computing support from all WLCG partners is acknowledged gratefully, in particular from CERN, the ATLAS Tier-1 facilities at TRIUMF/SFU (Canada), NDGF (Denmark, Norway,

Sweden), CC-IN2P3 (France), KIT/GridKA (Germany), INFN-CNAF (Italy), NL-T1 (Netherlands), PIC (Spain), RAL (UK) and BNL (USA), the Tier-2 facilities worldwide and large non-WLCG resource providers. Major contributors of computing resources are listed in Ref. [73]. We gratefully acknowledge the support of ANPCyT, Argentina; YerPhI, Armenia; ARC, Australia; BMWFW and FWF, Austria; ANAS, Azerbaijan; CNPq and FAPESP, Brazil; NSERC, NRC and CFI, Canada; CERN; ANID, Chile; CAS, MOST and NSFC, China; Minciencias, Colombia; MEYS CR, Czech Republic; DNRf and DNSRC, Denmark; IN2P3-CNRS and CEA-DRF/IRFU, France; SRNSFG, Georgia; BMBF, HGF and MPG, Germany; GSRI, Greece; RGC and Hong Kong SAR, China; ICHEP and Academy of Sciences and Humanities, Israel; INFN, Italy; MEXT and JSPS, Japan; CNRST, Morocco; NWO, Netherlands; RCN, Norway; MNiSW, Poland; FCT, Portugal; MNE/IFA, Romania; MSTDI, Serbia; MSSR, Slovakia; ARIS and MVZI, Slovenia; DSI/NRF, South Africa; MICIU/AEI, Spain; SRC and Wallenberg Foundation, Sweden; SERI, SNSF and Cantons of Bern and Geneva, Switzerland; NSTC, Taipei; TENMAK, Türkiye; STFC/UKRI, United Kingdom; DOE and NSF, United States of America. Individual groups and members have received support from BCKDF, CANARIE, CRC and DRAC, Canada; CERN-CZ, FORTE and PRIMUS, Czech Republic; COST, ERC, ERDF, Horizon 2020, ICSC-NextGenerationEU and Marie Skłodowska-Curie Actions, European Union; Investissements d’Avenir Labex, Investissements d’Avenir Idex and ANR, France; DFG and AvH Foundation, Germany; Herakleitos, Thales and Aristeia pro-

grammes co-financed by EU-ESF and the Greek NSRF, Greece; BSF-NSF and MINERVA, Israel; NCN and NAWA, Poland; La Caixa Banking Foundation, CERCA Programme Generalitat de Catalunya and PROMETEO and GenT Programmes Generalitat Valenciana, Spain; Göran Gustafssons Stiftelse, Sweden; The Royal Society and Leverhulme Trust, United Kingdom. In addition, individual members wish to acknowledge support from CERN: European Organization for Nuclear Research (CERN DOCT); Chile: Agencia Nacional de Investigación y Desarrollo (FONDECYT 1230812, FONDECYT 1230987, FONDECYT 1240864); China: Chinese Ministry of Science and Technology (MOST-2023YFA1605700, MOST-2023YFA1609300), National Natural Science Foundation of China (NSFC - 12175119, NSFC 12275265); Czech Republic: Czech Science Foundation (GACR - 24-11373S), Ministry of Education Youth and Sports (ERC-CZ-LL2327, FORTE CZ.02.01.01/00/22_008/0004632), PRIMUS Research Programme (PRIMUS/21/SCI/017); EU: H2020 European Research Council (ERC - 101002463); European Union: European Research Council (ERC - 948254, ERC 101089007, ERC, BARD, 101116429), European Regional Development Fund (SMASH COFUND 101081355, SLO ERDF), Horizon 2020 Framework Programme (MUCCA - CHIST-ERA-19-XAI-00), European Union, Future Artificial Intelligence Research (FAIR-NextGenerationEU PE00000013), Italian Center for High Performance Computing, Big Data and Quantum Computing (ICSC, NextGenerationEU); France: Agence Nationale de la Recherche (ANR-21-CE31-0022, ANR-22-EDIR-0002); Germany: Baden-Württemberg Stiftung (BW Stiftung-Postdoc Eliteprogramme), Deutsche Forschungsgemeinschaft (DFG - 469666862, DFG - CR 312/5-2); China: Research Grants Council (GRF); Italy: Istituto Nazionale di Fisica Nucleare (ICSC, NextGenerationEU), Ministero dell'Università e della Ricerca (NextGenEU PRIN20223N7F8K M4C2.1.1); Japan: Japan Society for the Promotion of Science (JSPS KAKENHI JP22H01227, JSPS KAKENHI JP22H04944, JSPS KAKENHI JP22KK0227, JSPS KAKENHI JP25H01291, JSPS KAKENHI JP25K01023); Norway: Research Council of Norway (RCN-314472); Poland: Ministry of Science and Higher Education (IDUB AGH, POB8, D4 no 9722), Polish National Science Centre (NCN 2021/42/E/ST2/00350, NCN OPUS 2023/51/B/ST2/02507, NCN OPUS nr 2022/47/B/ST2/03059, NCN UMO-2019/34/E/ST2/00393, UMO-2022/47/O/ST2/00148, UMO-2023/49/B/ST2/04085, UMO-2023/51/B/ST2/00920, UMO-2024/53/N/ST2/00869); Portugal: Foundation for Science and Technology (FCT); Spain: Ministry of Science and Innovation (MCIN & NextGenEU PCI2022-135018-2, MICIN & FEDER PID2021-125273NB, RYC2019-028510-I, RYC2020-030254-I, RYC2021-031273-I, RYC2022-038164-I); Sweden: Carl Trygger Foundation (Carl Trygger Foundation CTS 22:2312), Swedish Research Council (Swedish Research Council 2023-04654, VR 2021-03651, VR 2022-03845, VR 2022-04683, VR 2023-03403, VR 2024-05451), Knut and Alice Wallenberg Foundation (KAW 2018.0458, KAW 2022.0358, KAW 2023.0366); Switzerland: Swiss National Science Foundation (SNSF - PCEFP2_194658); United Kingdom: EPSRC (CANDISE, EP/X038858/1), Leverhulme Trust (Leverhulme Trust RPG-2020-004), Royal Society (NIF-R1-231091); United States of America: U.S. Department of Energy (ECA DE-AC02-76SF00515), NSF (CNS-1901466), Neubauer Family Foundation; VMware University Research Fund.

Data Availability Statement My manuscript has no associated data. [Author's comment: The ATLAS Collaboration's Athena software, including the configuration of the event generators, is open source (<http://gitlab.cern.ch/atlas/athena>)].

Code Availability Statement My manuscript has no associated code/software. [Author's comment: The public release of data supporting the findings of this article will follow the CERN Open Data Policy [74]. Inquiries about plots and tables associated with this article can be addressed to atlas.publications@cern.ch.]

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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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 licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP³.

References

1. P.A. Arias et al., 'Technical Summary', Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, vol. 35 (Cambridge University Press, Cambridge, 2023). <https://doi.org/10.1017/9781009157896.002>
2. CERN Occupational Health and Safety and Environmental Protection unit and the CERN Education, Communications and Outreach group, CERN ENVIRONMENT REPORT 2019-2020 (2021). <https://hse.cern/environment-report-2019-2020>
3. CERN Occupational Health and Safety and Environmental Protection unit and the CERN Education, Communications and Outreach group, CERN ENVIRONMENT REPORT 2021-2022 (2023). <https://hse.web.cern.ch/environment-report-2021-2022/energy>
4. D. Britton, S. Campana, B. Panzer-Stradel, A holistic study of the WLCG energy needs for the LHC scientific program. EPJ Web Conf. **295**, 04001 (2024). <https://doi.org/10.1051/epjconf/202429504001>
5. ATLAS Collaboration, The ATLAS experiment at the CERN large hadron collider. JINST **3**, S08003 (2008). <https://doi.org/10.1088/1748-0221/3/08/S08003>
6. ATLAS Collaboration, The ATLAS experiment at the CERN large hadron collider: a description of the detector configuration for Run 3. JINST **19**, P05063 (2024). <https://doi.org/10.1088/1748-0221/19/05/P05063>. arXiv: 2305.16623 [physics.ins-det]
7. ATLAS Collaboration, Software and computing for Run 3 of the ATLAS experiment at the LHC. Eur. Phys. J. C **85**, 234 (2025). <https://doi.org/10.1140/epjc/s10052-024-13701-w>. arXiv: 2404.06335 [hep-ex]
8. ATLAS Collaboration, ATLAS Software and Computing HL-LHC Roadmap, CERN-LHCC-2022-005 (2022). <http://cds.cern.ch/record/2802918>
9. EE HPC WG and Center of Expertise for Energy Efficiency in Data Centers, Center of Expertise for Energy Efficiency in Data Centers: Data Center Master List of Energy Efficiency Actions. <https://datacenters.lbl.gov/resources/data-center-master-list-energy>
10. Open Compute Project (2024). <https://www.opencompute.org>
11. Software Sustainability Institute, Green DiSC: a Digital Sustainability Certification (2024). <https://www.software.ac.uk/GreenDiSC>
12. GreenDIGIT (2024). <https://greendigit-project.eu>
13. M. Juckes et al., UKRI Net Zero Digital Research Infrastructure Scoping Project (2021). <https://net-zero-dri.ceda.ac.uk>
14. T. Simon et al., BoaviztAPI: a bottom-up model to assess the environmental impacts of cloud services (2024). <https://hal.science/hal-04621947v3>
15. IEA, Energy and AI (2025). <https://www.iea.org/reports/energy-and-ai>

16. World Business Council for Sustainable Development and World Resources Institute, Greenhouse Gas (GHG) Protocol Corporate reporting standard (2024). <https://ghgprotocol.org/corporate-standard>
17. U. Gupta et al., Chasing carbon: the elusive environmental footprint of computing. *IEEE Micro* **42**, 37 (2022). <https://doi.org/10.1109/MM.2022.3163226>
18. T. Maeno et al., PanDA: production and distributed analysis system. *Comput. Softw. Big Sci.* **8**, 4 (2024). <https://doi.org/10.1007/s41781-024-00114-3>
19. Electricity Maps, Electricity Maps API documentation (2024). https://api-portal.electricitymaps.com/catalog/prd_8e3ood60rcv8f51h
20. L. Lannelongue, J. Grealey, M. Inouye, Green algorithms: quantifying the carbon footprint of computation. *Adv. Sci.* **8**, 2100707 (2021). <https://doi.org/10.1002/adv.202100707>
21. PanDA Collaboration, Estimation of PanDA job carbon footprint (2024). https://panda-wms.readthedocs.io/en/latest/advanced/carbon_footprint.html
22. C. Marcon, A.S. Mete, P. Van Gemmeren, L. Carminati, Optimizing ATLAS data storage: the impact of compression algorithms on ATLAS physics analysis data formats. *EPJ Web Conf.* **295**, 03027 (2024). <https://doi.org/10.1051/epjconf/202429503027>
23. F. de Geus, J. López-Gómez, J. Blomer, M. Nowak, P. van Gemmeren, Integration of RNTuple in ATLAS Athena. *EPJ Web Conf.* **295**, 06013 (2024). <https://doi.org/10.1051/epjconf/202429506013>
24. C. Ledo, V. Choinsky, Sustainability Benefits of Deploying LTO Technology in Modern IT Ecosystems, ESG First Look (2022). <https://www.lto.org/wp-content/uploads/2022/08/ESG-First-Look-Sustainable-IT-with-LTO-Jul-2022.pdf>
25. D. Bauer, J. Polimeni, K. Mayumi, M. Giampietro, B. Alcott, The Jevons paradox and the myth of resource efficiency improvements. *Sustain. Sci. Pract. Policy* **5**, 48 (2009). <https://doi.org/10.1080/15487733.2009.11908028>
26. E. Bothmann et al., Accelerating LHC event generation with simplified pilot runs and fast PDFs. *Eur. Phys. J. C* **82**, 1128 (2022). [arXiv: 2209.00843](https://arxiv.org/abs/2209.00843) [hep-ph]. DOI: 10.1140/epjc/s10052-022-11087-1
27. E. Bothmann et al., Efficient precision simulation of processes with many-jet final states at the LHC. *Phys. Rev. D* **109**, 014013 (2024). <https://doi.org/10.1103/PhysRevD.109.014013>. [arXiv: 2309.13154](https://arxiv.org/abs/2309.13154) [hep-ph]
28. D. Spiteri et al., Sustainable computing with RF2.0 at DESY, WLCG Environmental Sustainability Workshop (2024). <https://indico.cern.ch/event/1450885/contributions/6251300>
29. A. Lory, M. Boehler, G. Duckeck, B. Rottler, Enhancements and resource optimisations for ATLAS use of HammerCloud, ATL-SOFT-PROC-2025-023 (2025). <https://cds.cern.ch/record/2923294>
30. R. Mahbubani, Z. Marshall, Open Event Generation, Contribution to the Reinterpretation Forum Workshop (2025). <https://indico.cern.ch/event/1466101/contributions/6292062>
31. D. Giordano et al., HEPscore: a new CPU benchmark for the WLCG. *EPJ Web Conf.* **295**, 07024 (2024). <https://doi.org/10.1051/epjconf/202429507024>. [arXiv: 2306.08118](https://arxiv.org/abs/2306.08118) [hep-ex]
32. J.J. Dongarra, P. Luszczyk, A. Petitet, The LINPACK benchmark: past, present and future. *Concurr. Comput. Pract. Exp.* **15**, 803–820 (2003). <https://doi.org/10.1002/cpe.728>
33. N. Szczepanek et al., HEP Benchmark Suite: Enhancing Efficiency and Sustainability in Worldwide LHC Computing Infrastructures (2024). <https://doi.org/10.48550/arXiv.2408.12445>. [arXiv: 2408.12445](https://arxiv.org/abs/2408.12445) [hep-ex]
34. J. Elmsheuser et al., The ATLAS experiment software on ARM. *EPJ Web Conf.* **295**, 05019 (2024). <https://doi.org/10.1051/epjconf/202429505019>
35. F.B. Megino et al., Operational experience and R&D results using the Google Cloud for High-Energy Physics in the ATLAS experiment. *Int. J. Mod. Phys. A* **39**, 2450054 (2024). <https://doi.org/10.1142/S0217751X24500544>. [arXiv: 2403.15873](https://arxiv.org/abs/2403.15873) [hep-ex]
36. ATLAS Collaboration, Total Cost of Ownership and Evaluation of Google Cloud Resources for the ATLAS Experiment at the LHC. *Comput. Softw. Big Sci.* **9**, 2 (2025). <https://doi.org/10.1007/s41781-024-00128-x>. [arXiv: 2405.13695](https://arxiv.org/abs/2405.13695) [cs.DC]
37. M. Boehler on behalf of the ATLAS Collaboration, A comparison of HEPSPC benchmark performance on ATLAS Grid-Sites versus ideal conditions, ATL-SOFT-PROC-2023-001 (2023). <https://cds.cern.ch/record/2850802/>
38. W. Vanderbauwhede, M. Wadenstein, LCA model for servers in a data centre (2025). <https://codeberg.org/wimvanderbauwhede/low-carbon-computing/src/branch/master/LCA-model-equations>
39. Federal Ministry for Economic Affairs and Climate Action, Renewable Energy. <https://www.bmwk.de/Redaktion/DE/Dossier/erneuerbare-energien.html>. Last Accessed 18 Aug 2024
40. F.I. for Solar Energy Systems ISE, Energy-Charts (2024). <https://www.energy-charts.info/charts/power/chart.htm?l=en&c=DE>
41. A. Portelli, Optimisation of lattice simulations energy efficiency. Zenodo (2022). <https://doi.org/10.5281/zenodo.7057318>
42. Autorité de Sûreté Nucléaire, The Small Modular Reactors (SMR) (2025). <https://www.french-nuclear-safety.fr/oversight/small-modular-reactors>
43. T. Ilsche, S. Schrader, R. Schöne, Optimizing Idle power of HPC systems: practical insights and methods. *IEEE Int. Conf. Clust. Comput.* (2024). <https://doi.org/10.1109/CLUSTERWorkshops61563.2024.00014>
44. D. Schien, N. Zilberman, D. Greenwood, P. Shabajee, A. Dewhurst, Preliminary Release: CarbonQuandRI: Decarbonising UKRI Electricity Use—A Case Study of GridPP Tier-1 (2023). <https://doi.org/10.5281/zenodo.7692132>
45. B. Bruers et al., Resource-aware research on universe and matter: call-to-action in digital transformation (2023). [arXiv: 2311.01169](https://arxiv.org/abs/2311.01169) [physics.comp-ph]
46. D.P. Adan et al., Humboldt Highway II-computer cluster on renewable energies, DESY-24-052 (2024). <https://doi.org/10.48550/arXiv.2404.07665>. [arXiv: 2404.07665](https://arxiv.org/abs/2404.07665) [physics.soc-ph]
47. WLCG Environmental Sustainability Workshop (2024). <https://indico.cern.ch/event/1450885/>
48. V. Rao, A.A. Chien, Understanding the Operational Carbon Footprint of Storage Reliability and Management, HotCarbon Workshop on Sustainable Computer Systems (2024). <https://hotcarbon.org/assets/2024/pdf/hotcarbon24-final61.pdf>
49. D. Li et al., Workload scheduling for massive storage systems with arbitrary renewable supply. *IEEE Trans. Parallel Distrib. Syst.* **29**, 2373 (2018). <https://doi.org/10.1109/TPDS.2018.2820070>
50. Z. Xu, N. Deng, C. Stewart, X. Wang, Cadre: carbon-aware data replication for geo-diverse services. *IEEE Int. Conf. Auton. Comput.* (2015). <https://doi.org/10.1109/ICAC.2015.15>
51. W. Katsak, Í. Goiri, R. Bianchini, T.D. Nguyen, Greencassandra: using renewable energy in distributed structured storage systems. *IGSC. 1* (2015). <https://doi.org/10.1109/IGCC.2015.7393711>
52. J. Wang et al., Designing cloud servers for lower carbon. *ISCA* (2024). <https://doi.org/10.1109/ISCA59077.2024.00041>
53. B. Li et al., Toward sustainable HPC: carbon footprint estimation and environmental implications of HPC systems, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '23* (2023). <https://doi.org/10.1145/3581784.3607035>
54. S. Tannu, P.J. Nair, The dirty secret of SSDs: embodied carbon. *SIGENERGY Energy Inform. Rev.* **3**, 4 (2023). <https://doi.org/10.1145/3630614.3630616>

55. A. Chien et al., Reducing the carbon impact of generative AI inference (today and in 2035), in *Proceedings of the 2nd Workshop on Sustainable Computer Systems*, vol. 11, 1 (2023). <https://doi.org/10.1145/3604930.3605705>
56. H. Tollefsen, Evaluating performance impact of performing computations on storage nodes—Batch on Eos Extra Resources (BEER), Master's thesis (Norwegian U. Sci. Tech., 2018). <https://cds.cern.ch/record/2841037>
57. Y.H. Labaran, V.S. Mathur, S.U. Muhammad, A.A. Musa, Carbon footprint management: a review of construction industry. *Clean. Eng. Technol.* **9**, 100531 (2022). <https://doi.org/10.1016/j.clet.2022.100531>
58. M. Airaksinen, P. Matilainen, A carbon footprint of an office building. *Energies* **4**, 1197 (2011). <https://doi.org/10.3390/en4081197>
59. B. Sizerici, Y. Fseha, C.-S. Cho, I. Yildiz, Y.-J. Byon, A review of carbon footprint reduction in construction industry, from design to operation. *Materials* **14**, 6094 (2021). <https://doi.org/10.3390/ma14206094>
60. W. Wu, P. Sun, H. Zhou, I.O.P. Conf, Ser. Earth Environ. Sci. **371**, 022011 (2019). <https://doi.org/10.1088/1755-1315/371/2/022011>
61. Y.H. Labaran, V.S. Mathur, M.M. Farouq, The carbon footprint of construction industry: a review of direct and indirect emission. *J. Sustain. Constr. Mater. Technol.* **6**, 101 (2021). <https://doi.org/10.29187/jscmt.2021.66>
62. Circular Ecology: Net Zero Carbon Buildings. <https://circularecology.com/net-zero-carbon-buildings.html>
63. U.S. Green Building Council, LEED v4.1 (2025). <https://www.usgbc.org/leed/v41>
64. D. Martinez, D. Sickinger, “How-to” Guide for Transitioning from Air to Liquid-cooled High Performance Computing Systems (2024). <https://sites.google.com/lbl.gov/ee-hpc-wg-liquid-cooling/home>
65. V. Ljungdahl, M. Jradi, C. Veje, A decision support model for waste heat recovery systems design in Data Center and High-Performance Computing clusters utilizing liquid cooling and Phase Change Materials. *Appl. Therm. Eng.* **201**, 117671 (2022). <https://doi.org/10.1016/j.applthermaleng.2021.117671>
66. Bytes2Heat, Best-Practice-Overview: Learning from inspiring projects for waste heat recovery from data centers (2024). https://www.bytes2heat.de/static/files/20240403_BP%C3%9C_V2404_eng.pdf
67. R. Curtis, T. Shedd, E. Clark, Performance comparison of five data center server thermal management technologies, in *2023 39th Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*, 1 (2023). <https://doi.org/10.23919/SEMI-THERM59981.2023.10267908>
68. I. Latif et al., Advancing sustainability in data centers: evaluation of hybrid air/liquid cooling schemes for IT payload using sea water. *IEEE Trans. Cloud Comput.* **13**, 184 (2025). <https://doi.org/10.1109/TCC.2024.3521666>
69. C. Pecchioli, J. Comella, D. Sickinger, O. VanGeet, DATA CENTERS HEAT REUSE 101 (2023). <https://www.opencompute.org/documents/20230623-data-centers-heatreuse-101-3-2-docx-pdf>
70. S. Zimmermann et al., Aquasar: a hot water cooled data center with direct energy reuse, *Energy* **43**, 237 (2012). 2nd International Meeting on Cleaner Combustion (CM0901-Detailed Chemical Models for Cleaner Combustion). <https://doi.org/10.1016/j.energy.2012.04.037>
71. S. Simsek on behalf of the ATLAS Collaboration, Mitigation of the ATLAS RPC environmental impact. *PoS* **476**, 1238 (2024). <https://doi.org/10.22323/1.476.1238>
72. I. Latif et al., Single-Node Power Demand During AI Training: Measurements on an 8-GPU NVIDIA H100 System, *IEEE Access* 61740 (2025). <https://doi.org/10.1109/ACCESS.2025.3554728>
73. ATLAS Collaboration, ATLAS Computing Acknowledgements, ATL-SOFT-PUB-2025-001 (2025). <https://cds.cern.ch/record/2922210>
74. CERN-OPEN-2020-013: <https://cds.cern.ch/record/2745133/export/hx?ln=en>

ATLAS Collaboration*

G. Aad¹⁰⁴, E. Aakvaag¹⁷, B. Abbott¹²³, S. Abdelhameed^{119a}, K. Abeling⁵⁵, N. J. Abicht⁴⁹, S. H. Abidi³⁰, M. Aboelela⁴⁵, A. Aboulhorma^{36c}, H. Abramowicz¹⁵⁷, Y. Abulaiti¹²⁰, B. S. Acharya^{69a,69b,m}, A. Ackermann^{63a}, C. Adam Bourdarios⁴, L. Adamczyk^{86a}, S. V. Addepalli¹⁴⁹, M. J. Addison¹⁰³, J. Adelman¹¹⁸, A. Adiguzel^{22c}, T. Adye¹³⁷, A. A. Affolder¹³⁹, Y. Afik⁴⁰, M. N. Agaras¹³, A. Aggarwal¹⁰², C. Agheorghiesei^{28c}, F. Ahmadov^{39,ad}, S. Ahuja⁹⁷, X. Ai^{143b}, G. Aielli^{76a,76b}, A. Aikot¹⁶⁹, M. Ait Tamliah^{36c}, B. Aitbenkikh^{36a}, M. Akbiyik¹⁰², T. P. A. Åkesson¹⁰⁰, A. V. Akimov¹⁵¹, D. Akiyama¹⁷⁴, N. N. Akolkar²⁵, S. Aktas^{22a}, G. L. Alberghi^{24b}, J. Albert¹⁷¹, U. Alberti²⁰, P. Albicocco⁵³, G. L. Albouy⁶⁰, S. Alderweireldt⁵², Z. L. Alegria¹²⁴, M. Aleksa^{37a}, I. N. Aleksandrov³⁹, C. Alexa^{28b}, T. Alexopoulos¹⁰, F. Alfonsi^{24b}, M. Algren⁵⁶, M. Alhroob¹⁷³, B. Ali¹³⁵, H. M. J. Ali^{93,w}, S. Ali³², S. W. Alibocus⁹⁴, M. Aliev^{34c}, G. Alimonti^{71a}, W. Alkakh⁵⁵, C. Allaire⁶⁶, B. M. M. Allbrooke¹⁵², J. S. Allen¹⁰³, J. F. Allen⁵², P. P. Allport²¹, A. Aloisio^{72a,72b}, F. Alonso⁹², C. Alpigiani¹⁴², Z. M. K. Alsolami⁹³, A. Alvarez Fernandez¹⁰², M. Alves Cardoso⁵⁶, M. G. Alviggi^{72a,72b}, M. Aly¹⁰³, Y. Amaral Coutinho^{83b}, A. Ambler¹⁰⁶, C. Amelung^{37a}, M. Ameri¹⁰³, C. G. Ames¹¹¹, T. Amezza¹³⁰, D. Amidei¹⁰⁸, B. Amini⁵⁴, K. Amirie¹⁶¹, A. Amirkhanov³⁹, S. P. Amor Dos Santos^{133a}, K. R. Amos¹⁶⁹, D. Amperiadou¹⁵⁸, S. An⁸⁴, C. Anastopoulos¹⁴⁵, T. Andeen¹¹, J. K. Anders⁹⁴, A. C. Anderson⁵⁹, A. Andreazza^{71a,71b}, S. Angelidakis⁹, A. Angerami⁴², A. V. Anisenkov³⁹, A. Annovi^{74a}, C. Antel^{37a}, E. Antipov¹⁵¹, M. Antonelli⁵³, F. Anulli^{75a}, M. Aoki⁸⁴, T. Aoki¹⁵⁹, M. A. Aparo¹⁵², L. Aperio Bella⁴⁸, M. Apicella³¹, C. Appelt¹⁵⁷, A. Apyan²⁷, M. Arampatzi¹⁰, S. J. Arbiol Val⁸⁷, C. Arcangeletti⁵³, A. T. H. Arce⁵¹, J-F. Arguin¹¹⁰, S. Argyropoulos¹⁵⁸, J.-H. Arling⁴⁸, O. Arnaez⁴, H. Arnold¹⁵¹, G. Artoni^{75a,75b}, H. Asada¹¹³, K. Asai¹²¹, S. Asatryan¹⁷⁹, N. A. Asbah^{37a}, R. A. Ashby Pickering¹⁷³, A. M. Aslam⁹⁷, K. Assamagan³⁰, R. Astalos^{29a}, K. S. V. Astrand¹⁰⁰, S. Atashi¹⁶⁵, R. J. Atkin^{34a}, H. Atmani^{36f}, P. A. Atmasiddha¹³¹, K. Augsten¹³⁵, A. D. Aurio⁴¹, V. A. Austrup¹⁰³, G. Avolio^{37a}, K. Axiotis⁵⁶, A. Azzam¹³, D. Babal^{29b}, H. Bachacou¹³⁸, K. Bachas^{158,g}, A. Bachi³⁵, E. Bachmann⁵⁰, M. J. Backes^{63a}, A. Badea⁴⁰, T. M. Baer¹⁰⁸, P. Bagnaia^{75a,75b}, M. Bahmani¹⁹, D. Bahner⁵⁴, K. Bai¹²⁶, J. T. Baines¹³⁷, L. Baines⁹⁶, O. K. Baker¹⁷⁸, E. Bakos¹⁶, D. Bakshi Gupta⁸, L. E. Balabram Filho^{83b}, V. Balakrishnan¹²³, R. Balasubramanian⁴, E. M. Baldin³⁸, P. Balek^{86a}, E. Ballabene^{24a,24b}, F. Balli¹³⁸, L. M. Baltes^{63a}, W. K. Balunas³³, J. Balz¹⁰², I. Bamwidhi^{119b}, E. Banas⁸⁷, M. Bandieramonte¹³², A. Bandyopadhyay²⁵, S. Bansal²⁵, L. Barak¹⁵⁷, M. Barakat⁴⁸, E. L. Barberio¹⁰⁷, D. Barberis^{18b}, M. Barbero¹⁰⁴, M. Z. Barel¹¹⁷, T. Barillari¹¹², M-S. Barisits^{37a}, T. Barklow¹⁴⁹, P. Baron¹³⁶, D. A. Baron Moreno¹⁰³, A. Baroncelli⁶², A. J. Barr¹²⁹, J. D. Barr⁹⁸, F. Barreiro¹⁰¹, J. Barreiro Guimarães da Costa¹⁴, F. H. Barreiro Megino⁸, M. G. Barros Teixeira^{133a}, S. Barsov³⁸, F. Bartels^{63a}, R. Bartoldus¹⁴⁹, A. E. Barton⁹³, P. Bartos^{29a}, A. Basan¹⁰², M. Baselga⁴⁹, S. Bashiri⁸⁷, A. Bassalat^{66,b}, M. J. Basso^{162a}, S. Bataju⁴⁵, R. Bate¹⁷⁰, R. L. Bates⁵⁹, S. Batlamous¹⁰¹, M. Battaglia¹³⁹, D. Battulga¹⁹, M. Bauce^{75a,75b}, M. Bauer⁷⁹, P. Bauer²⁵, L. T. Bayer⁴⁸, L. T. Bazzano Hurrell³¹, J. B. Beacham¹¹², T. Beau¹³⁰, J. Y. Beaucamp⁹², P. H. Beauchemin¹⁶⁴, P. Bechtel²⁵, H. P. Beck^{20,p}, K. Becker¹⁷³, A. J. Beddall⁸², V. A. Bednyakov³⁹, C. P. Bee¹⁵¹, L. J. Beemster¹⁶, M. Begalli^{83d}, M. Begel³⁰, J. K. Behr⁴⁸, J. F. Beirer^{37a}, F. Beisiegel²⁵, M. Belfkir^{119b}, G. Bella¹⁵⁷, L. Bellagamba^{24b}, A. Bellerive³⁵, C. D. Bellgraph⁶⁸, P. Bellos²¹, K. Beloborodov³⁸, D. Benchechroun^{36a}, F. Bendebba^{36a}, Y. Benhammou¹⁵⁷, K. C. Benkendorfer⁶¹, L. Beresford⁴⁸, M. Beretta⁵³, E. Bergeas Kuutmann¹⁶⁷, N. Berger⁴, B. Bergmann¹³⁵, J. Beringer^{18a}, G. Bernardi⁵, C. Bernius¹⁴⁹, F. U. Bernlochner²⁵, F. Bernon^{37a}, A. Berrocal Guardia¹³, T. Berry⁹⁷, P. Berta¹³⁶, A. Berthold⁵⁰, A. Berti^{133a}, R. Bertrand¹⁰⁴, S. Bethke¹¹², A. Betti^{75a,75b}, A. J. Bevan⁹⁶, L. Bezio⁵⁶, N. K. Bhalla⁵⁴, S. Bharthuar¹¹², S. Bhatta¹⁵¹, P. Bhattacharya¹⁴⁹, Z. M. Bhatti¹²⁰, K. D. Bhide⁵⁴, V. S. Bhopatkar¹²⁴, R. M. Bianchi¹³², G. Bianco^{24a,24b}, O. Biebel¹¹¹, M. Biglietti^{77a}, C. S. Billingsley⁴⁵, Y. Bimgdi^{36f}, M. Bindi⁵⁵, A. Bingham¹⁷⁷, A. Bingul^{22b}, C. Bini^{75a,75b}, G. A. Bird³³, M. Birman¹⁷⁵, M. Biros¹³⁶, S. Biryukov¹⁵², T. Bisanz⁴⁹, E. Bisceglie^{24a,24b}, J. P. Biswal¹³⁷, D. Biswas¹⁴⁷, I. Bloch⁴⁸, A. Blue⁵⁹, U. Blumenschein⁹⁶, J. Blumenthal¹⁰², V. S. Bobrovnikov³⁹, L. Boccardo^{57a,57b}, M. Boehler⁵⁴, B. Boehm¹⁷², D. Bogavac¹³, A. G. Bogdanchikov³⁸, L. S. Boggia¹³⁰, V. Boisvert⁹⁷, P. Bokan^{37a}, T. Bold^{86a}, M. Bomben⁵, M. Bona⁹⁶, M. Boonekamp¹³⁸, A. G. Borbely⁵⁹, I. S. Bordulev³⁸, G. Borissov⁹³, D. Bortoletto¹²⁹, D. Boscherini^{24b}, M. Bosman¹³, K. Bouaouda^{36a}, N. Bouchhar¹⁶⁹, L. Boudet⁴, J. Boudreau¹³², E. V. Bouhova-Thacker⁹³, D. Boumediene⁴¹, R. Bouquet^{57a,57b}, A. Boveia¹²², J. Boyd^{37a}, D. Boye³⁰, I. R. Boyko³⁹, L. Bozianu⁵⁶, J. Bracinik²¹, N. Brahimi⁴, G. Brandt¹⁷⁷, O. Brandt³³, B. Brau¹⁰⁵, J. E. Brau¹²⁶, R. Brenner¹⁷⁵, L. Brenner¹¹⁷, R. Brenner¹⁶⁷, S. Bressler¹⁷⁵, G. Brianti^{78a,78b}

D. Britton⁵⁹, D. Britzger¹¹², I. Brock²⁵, R. Brock¹⁰⁹, G. Brooijmans⁴², A. J. Brooks⁶⁸, E. M. Brooks^{162b}, E. Brost³⁰, L. M. Brown^{162a,171}, L. E. Bruce⁶¹, T. L. Bruckler¹²⁹, P. A. Bruckman de Renstrom⁸⁷, B. Brüers⁴⁸, A. Bruni^{24b}, G. Bruni^{24b}, D. Brunner^{47a,47b}, M. Bruschi^{24b}, N. Bruscino^{75a,75b}, T. Buanes¹⁷, Q. Buat¹⁴², D. Buchin¹¹², A. G. Buckley⁵⁹, O. Bulekov⁸², B. A. Bullard¹⁴⁹, S. Burdin⁹⁴, C. D. Burgard⁴⁹, A. M. Burger⁹¹, B. Burghgrave⁸, O. Burlayenko⁵⁴, J. Bureson¹⁶⁸, J. C. Burzynski¹⁴⁸, E. L. Busch⁴², V. Büscher¹⁰², P. J. Bussey⁵⁹, J. M. Butler²⁶, C. M. Buttar⁵⁹, J. M. Butterworth⁹⁸, W. Buttinger¹³⁷, C. J. Buxo Vazquez¹⁰⁹, A. R. Buzykaev³⁹, S. Cabrera Urbán¹⁶⁹, L. Cadamuro⁶⁶, H. Cai¹³², Y. Cai^{24a,24b,114c}, Y. Cai^{114a}, V. M. M. Cairo^{37a}, O. Cakir^{3a}, N. Calace^{37a}, P. Calafiura^{18a}, G. Calderini¹³⁰, P. Calfayan³⁵, L. Calic¹⁰⁰, G. Callea⁵⁹, L. P. Caloba^{83b}, D. Calvet⁴¹, S. Calvet⁴¹, R. Camacho Toro¹³⁰, S. Camarda^{37a}, D. Camarero Munoz²⁷, P. Camarri^{76a,76b}, C. Camincher¹⁷¹, M. Campanelli⁹⁸, A. Camplani⁴³, V. Canale^{72a,72b}, A. C. Canbay^{3a}, E. Canonero⁹⁷, J. Cantero¹⁶⁹, Y. Cao¹⁶⁸, F. Capocasa²⁷, M. Capua^{44a,44b}, A. Carbone^{71a,71b}, R. Cardarelli^{76a}, J. C. J. Cardenas⁸, M. P. Cardiff²⁷, G. Carducci^{44a,44b}, T. Carli^{37a}, G. Carlino^{72a}, J. I. Carlotto¹³, B. T. Carlson^{132,r}, E. M. Carlson¹⁷¹, J. Carmignani⁹⁴, L. Carminati^{71a,71b}, A. Carnelli⁴, M. Carnesale^{37a}, S. Caron¹¹⁶, E. Carquin^{140g}, I. B. Carr¹⁰⁷, S. Carrá^{73a,73b}, G. Carratta^{24a,24b}, C. Carrion Martinez¹⁶⁹, A. M. Carroll¹²⁶, M. P. Casado^{13,h}, P. Casolaro^{72a,72b}, M. Caspar⁴⁸, F. L. Castillo⁴, L. Castillo Garcia¹³, V. Castillo Gimenez¹⁶⁹, N. F. Castro^{133a,133e}, A. Catinaccio^{37a}, J. R. Catmore¹²⁸, T. Cavaliere⁴, V. Cavaliere³⁰, L. J. Caviedes Betancourt^{23b}, E. Celebi⁸², S. Cella^{37a}, V. Cepaitis⁵⁶, K. Cerny¹²⁵, A. S. Cerqueira^{83a}, A. Cerri^{74a,74b,ak}, L. Cerrito^{76a,76b}, F. Cerutti^{18a}, B. Cervato^{71a,71b}, A. Cervelli^{24b}, G. Cesarini⁵³, S. A. Cetin⁸², P. M. Chabrilat¹³⁰, R. Chakkappai⁶⁶, S. Chakraborty¹⁷³, J. Chan^{18a}, W. Y. Chan¹⁵⁹, J. D. Chapman³³, E. Chapon¹³⁸, B. Chargeishvili^{155b}, D. G. Charlton²¹, C. Chauhan¹³⁶, Y. Che^{114a}, S. Chekanov⁶, S. V. Chekulaev^{162a}, G. A. Chelkov^{39,a}, B. Chen¹⁵⁷, B. Chen¹⁷¹, H. Chen^{114a}, H. Chen³⁰, J. Chen^{144a}, J. Chen¹⁴⁸, M. Chen¹²⁹, S. Chen⁸⁹, S. J. Chen^{114a}, X. Chen^{144a}, X. Chen^{15,ag}, Z. Chen⁶², C. L. Cheng¹⁷⁶, H. C. Cheng^{64a}, S. Cheong¹⁴⁹, A. Cheplakov³⁹, E. Cherepanova¹¹⁷, R. Cherkaoui El Moursli^{36e}, E. Cheu⁷, K. Cheung⁶⁵, L. Chevalier¹³⁸, V. Chiarella⁵³, G. Chiarelli^{74a}, A. Chien^{an,ao}, G. Chiodini^{70a}, A. S. Chisholm²¹, A. Chitan^{28b}, M. Chitishvili¹⁶⁹, M. V. Chizhov^{39,s}, K. Choi¹¹, Y. Chou¹⁴², E. Y. S. Chow¹¹⁶, K. L. Chu¹⁷⁵, M. C. Chu^{64a}, X. Chu^{14,114c}, Z. Chubinizde⁵³, J. Chudoba¹³⁴, J. J. Chwastowski⁸⁷, D. Cieri¹¹², K. M. Ciesla^{86a}, V. Cindro⁹⁵, A. Ciochio^{18a}, F. Ciotto^{72a,72b}, Z. H. Citron¹⁷⁵, M. Citterio^{71a}, D. A. Ciubotaru^{28b}, A. Clark⁵⁶, P. J. Clark⁵², N. Clarke Hall⁹⁸, C. Clarry¹⁶¹, S. E. Clawson⁴⁸, C. Clement^{47a,47b}, Y. Coadou¹⁰⁴, M. Cobal^{69a,69c}, A. Coccaro^{57b}, R. F. Coelho Barrue^{133a}, R. Coelho Lopes De Sa¹⁰⁵, S. Coelli^{71a}, L. S. Colangeli¹⁶¹, B. Cole⁴², P. Collado Soto¹⁰¹, J. Collot⁶⁰, R. Coluccia^{70a,70b}, P. Conde Muño^{133a,133g}, M. P. Connell^{34c}, S. H. Connell^{34c}, E. I. Conroy¹²⁹, F. Conventi^{72a,ai}, A. M. Cooper-Sarkar¹²⁹, L. Corazzina^{75a,75b}, F. A. Corchia^{24a,24b}, A. Cordeiro Oudot Choi¹⁴², L. D. Corpe⁴¹, M. Corradi^{75a,75b}, F. Corriveau^{106,ab}, A. Cortes-Gonzalez¹⁵⁹, M. J. Costa¹⁶⁹, F. Costanza⁴, D. Costanzo¹⁴⁵, B. M. Cote¹²², J. Couthures⁴, G. Cowan⁹⁷, K. Cranmer¹⁷⁶, L. Cremer⁴⁹, D. Cremonini^{24a,24b}, S. Crépe-Renaudin⁶⁰, F. Crescioli¹³⁰, T. Cresta^{73a,73b}, M. Cristinziani¹⁴⁷, M. Cristoforetti^{78a,78b}, V. Croft¹¹⁷, J. E. Crosby¹²⁴, G. Crossetti^{44a,44b}, A. Cueto¹⁰¹, H. Cui⁹⁸, Z. Cui⁷, B. M. Cunnett¹⁵², W. R. Cunningham⁵⁹, F. Curcio¹⁶⁹, J. R. Curran⁵², M. J. Da Cunha Sargedas De Sousa^{57a,57b}, J. V. Da Fonseca Pinto^{83b}, C. Da Via¹⁰³, W. Dabrowski^{86a}, T. Dado^{37a}, S. Dahbi¹⁵⁴, T. Dai¹⁰⁸, D. Dal Santo²⁰, C. Dallapiccola¹⁰⁵, M. Dam⁴³, G. D'amen³⁰, V. D'Amico¹¹¹, J. Damp¹⁰², J. R. Dandoy³⁵, M. D'andrea^{57a,57b}, D. Dannheim^{37a}, G. D'anniballe^{74a,74b}, M. Danninger¹⁴⁸, V. Dao¹⁵¹, G. Darbo^{57b}, S. J. Das³⁰, F. Dattola⁴⁸, S. D'Auria^{71a,71b}, A. D'Avanzo^{72a,72b}, T. Davidek¹³⁶, J. Davidson¹⁷³, I. Dawson⁹⁶, K. De⁸, C. De Almeida Rossi¹⁶¹, R. De Asmundis^{72a}, N. De Biase⁴⁸, S. De Castro^{24a,24b}, N. De Groot¹¹⁶, P. de Jong¹¹⁷, H. De la Torre¹¹⁸, A. De Maria^{114a}, A. De Salvo^{75a}, U. De Sanctis^{76a,76b}, F. De Santis^{70a,70b}, A. De Santo¹⁵², J. B. De Vivie De Regie⁶⁰, J. Debevc⁹⁵, D. V. Dedovich³⁹, J. Degens⁹⁴, A. M. Deiana⁴⁵, J. Del Peso¹⁰¹, L. Delagrange¹³⁰, F. Deliot¹³⁸, C. M. Delitzsch⁴⁹, M. Della Pietra^{72a,72b}, D. Della Volpe⁵⁶, A. Dell'Acqua^{37a}, L. Dell'Asta^{71a,71b}, M. Delmastro⁴, C. C. Delogu¹⁰², P. A. Delsart⁶⁰, S. Demers¹⁷⁸, M. Demichev³⁹, S. P. Denisov³⁸, H. Denizli^{22a,1}, L. D'Eramo⁴¹, D. Derendarz⁸⁷, F. Derue¹³⁰, P. Dervan⁹⁴, A. M. Desai¹, K. Desch²⁵, A. Dewhurst¹³⁷, F. A. Di Bello^{57a,57b}, A. Di Ciaccio^{76a,76b}, L. Di Ciaccio⁴, A. Di Domenico^{75a,75b}, C. Di Donato^{72a,72b}, A. Di Girolamo^{37a}, G. Di Gregorio^{37a}, A. Di Luca^{78a,78b}, B. Di Micco^{77a,77b}, R. Di Nardo^{77a,77b}, K. F. Di Petrillo⁴⁰, M. Diamantopoulou³⁵, F. A. Dias¹¹⁷, M. A. Diaz^{140a,140b}, A. R. Didenko³⁹, M. Didenko¹⁶⁹, S. D. Diefenbacher^{18a}, E. B. Diehl¹⁰⁸, S. Díez Cornell⁴⁸, C. Díez Pardos¹⁴⁷, C. Dimitriadis¹⁵⁰, A. Dimitrievska²¹, A. Dimri¹⁵¹, Y. Ding⁶², J. Dingfelder²⁵, T. Dingley¹²⁹, I-M. Dinu^{28b}

A. Guida¹⁹, E. Guillon¹⁷³, S. Guindon^{37a}, F. Guo^{14,114c}, J. Guo^{144a}, L. Guo⁴⁸, L. Guo^{114b,v}, Y. Guo¹⁰⁸, A. Gupta⁴⁹, R. Gupta¹³², S. Gupta²⁷, S. Gurbuz²⁵, S. S. Gurdasani⁴⁸, G. Gustavino^{75a,75b}, P. Gutierrez¹²³, L. F. Gutierrez Zagazeta¹³¹, M. Gutsche⁵⁰, C. Gutschow⁹⁸, C. Gwenlan¹²⁹, C. B. Gwilliam⁹⁴, E. S. Haaland¹²⁸, A. Haas¹²⁰, M. Habedank⁵⁹, C. Haber^{18a}, H. K. Hadavand⁸, A. Haddad⁴¹, A. Hadeef⁵⁰, A. I. Hagan⁹³, J. J. Hahn¹⁴⁷, E. H. Haines⁹⁸, M. Haleem¹⁷², J. Haley¹²⁴, G. D. Halleswell¹⁰⁴, K. Hamano¹⁷¹, H. Hamdaoui¹⁶⁷, M. Hamer²⁵, S. E. D. Hammoud⁶⁶, E. J. Hampshire⁹⁷, J. Han^{143a}, L. Han^{114a}, L. Han⁶², S. Han¹⁴, K. Hanagaki⁸⁴, M. Hance¹³⁹, D. A. Hangal⁴², H. Hanif¹⁴⁸, M. D. Hank¹³¹, J. B. Hansen⁴³, P. H. Hansen⁴³, D. Harada⁵⁶, T. Harenberg¹⁷⁷, S. Harkusha¹⁷⁹, M. L. Harris¹⁰⁵, Y. T.
Harris²⁵, J. Harrison¹³, N. M. Harrison¹²², P. F. Harrison¹⁷³, M. L. E. Hart⁹⁸, N. M. Hartman¹¹², N. M. Hartmann¹¹¹, R. Z. Hasan^{97,137}, Y. Hasegawa¹⁴⁶, F. Haslbeck¹²⁹, S. Hassan¹⁷, R. Hauser¹⁰⁹, M. Haviernik¹³⁶, C. M. Hawkes²¹, R. J. Hawkings^{37a}, Y. Hayashi¹⁵⁹, D. Hayden¹⁰⁹, C. Hayes¹⁰⁸, R. L. Hayes¹¹⁷, C. P. Hays¹²⁹, J. M. Hays⁹⁶, H. S. Hayward⁹⁴, M. He^{14,114c}, Y. He⁴⁸, Y. He⁹⁸, N. B. Heatley⁹⁶, V. Hedberg¹⁰⁰, C. Heidegger⁵⁴, K. K. Heidegger⁵⁴, J. Heilman³⁵, S. Heim⁴⁸, T. Heim^{18a}, J. G. Heinlein¹³¹, J. J. Heinrich¹²⁶, L. Heinrich¹¹², J. Hejbal¹³⁴, M. Helbig⁵⁰, A. Held¹⁷⁶, S. Hellesund¹⁷, C. M. Helling¹⁷⁰, S. Hellman^{47a,47b}, A. M. Henriques Correia^{37a}, H. Herde¹⁰⁰, Y. Hernández Jiménez¹⁵¹, L. M. Herrmann²⁵, T. Herrmann⁵⁰, G. Herten⁵⁴, R. Hertenberger¹¹¹, L. Hervas^{37a}, M. E. Hespington¹⁰², N. P. Hessey^{162a}, J. Hessler¹¹², M. Hidaoui^{36b}, N.
Hidic¹³⁶, E. Hill¹⁶¹, T. S. Hillersoy¹⁷, S. J. Hillier²¹, J. R. Hinds¹⁰⁹, F. Hinterkeuser²⁵, M. Hirose¹²⁷, S. Hirose¹⁶³, D. Hirschbuehl¹⁷⁷, T. G. Hitchings¹⁰³, B. Hiti⁹⁵, J. Hobbs¹⁵¹, R. Hobincu^{28c}, N. Hod¹⁷⁵, A. M. Hodges¹⁶⁸, M. C. Hodgkinson¹⁴⁵, B. H. Hodgkinson¹²⁹, A. Hoecker^{37a}, D. D. Hofer¹⁰⁸, J. Hofer¹⁶⁹, M. Holzbock^{37a}, L. B. A. H. Hommels³³, V. Homsak¹²⁹, B. P. Honan¹⁰³, J. J. Hong⁶⁸, T. M. Hong¹³², B. H. Hooberman¹⁶⁸, W. H. Hopkins⁶, M. C. Hoppesch¹⁶⁸, Y. Horii¹¹³, M. E. Horstmann¹¹², S. Hou¹⁵⁴, M. R. Housenga¹⁶⁸, J. Howarth⁵⁹, J. Hoya⁶, M. Hrabovsky¹²⁵, T. Hryn'ova⁴, P. J. Hsu⁶⁵, S.-C. Hsu¹⁴², T. Hsu⁶⁶, M. Hu^{18a}, Q. Hu⁶², S. Huang³³, X. Huang^{14,114c}, Y. Huang¹³⁶, Y. Huang^{114b}, Y. Huang¹⁰², Y. Huang¹⁴, Z. Huang⁶⁶, Z. Hubacek¹³⁵, M. Huebner²⁵, F. Huegging²⁵, T. B. Huffman¹²⁹, M. Hufnagel Maranhã De
Faria^{83a}, C. A. Hugli⁴⁸, M. Huhtinen^{37a}, S. K. Huiberts¹⁷, R. Hulsken¹⁰⁶, C. E. Hultquist^{18a}, D. L. Humphreys¹⁰⁵, N. Huseynov¹², J. Huston¹⁰⁹, J. Huth⁶¹, R. Hyneman⁷, G. Iacobucci⁵⁶, G. Iakovidis³⁰, L. Iconomidou-Fayard⁶⁶, J. P. Iddon^{37a}, P. Iengo^{72a,72b}, R. Iguchi¹⁵⁹, Y. Iiyama¹⁵⁹, T. Iizawa¹⁵⁹, Y. Ikegami⁸⁴, D. Iliadis¹⁵⁸, N. Ilic¹⁶¹, H. Imam^{36a}, G. Inacio Goncalves^{83d}, S. A. Infante Cabanas^{140c}, T. Ingebrechtsen Carlson^{47a,47b}, J. M. Inglis⁹⁶, G. Introzzi^{73a,73b}, M. Iodice^{77a}, V. Ippolito^{75a,75b}, R. K. Irwin⁹⁴, M. Ishino¹⁵⁹, W. Islam¹⁷⁶, C. Issever¹⁹, S. Istin^{22a,am}, K. Itabashi⁸⁴, H. Ito¹⁷⁴, R. Iuppa^{78a,78b}, A. Ivina¹⁷⁵, V. Izzo^{72a}, P. Jacka¹³⁵, P. Jackson¹, P. Jain⁴⁸, K. Jakobs⁵⁴, T. Jakoubek¹⁷⁵, J. Jamieson⁵⁹, W. Jang¹⁵⁹, S. Jankovych¹³⁶, M. Javurkova¹⁰⁵, P. Jawahar¹⁰³, L. Jeanty¹²⁶, J. Jejelava^{155a,ae}, P.
Jenni^{54,f}, C. E. Jessiman³⁵, C. Jia^{143a}, H. Jia¹⁷⁰, J. Jia¹⁵¹, X. Jia^{112,114c}, Z. Jia^{114a}, C. Jiang⁵², Q. Jiang^{64b}, S. Jiggins⁴⁸, M. Jimenez Ortega¹⁶⁹, J. Jimenez Pena¹³, S. Jin^{114a}, A. Jinaru^{28b}, O. Jinnouchi¹⁴¹, P. Johansson¹⁴⁵, K. A. Johns⁷, J. W. Johnson¹³⁹, F. A. Jolly⁴⁸, D. M. Jones¹⁵², E. Jones⁴⁸, K. S. Jones⁸, P. Jones³³, R. W. L. Jones⁹³, T. J. Jones⁹⁴, H. L. Joos⁵⁵, R. Joshi¹²², J. Jovicevic¹⁶, X. Ju^{18a}, J. J. Junggeburth^{37a}, T. Junkermann^{63a}, A. Juste Rozas^{13,x}, M. K. Juzek⁸⁷, S. Kabana^{140f}, A. Kaczmarzka⁸⁷, S. A. Kadir¹⁴⁹, M. Kado¹¹², H. Kagan¹²², M. Kagan¹⁴⁹, A. Kahn¹³¹, C. Kahra¹⁰², T. Kaji¹⁵⁹, E. Kajomovitz¹⁵⁶, N. Kakati¹⁷⁵, N. Kakoty¹³, I. Kalaitzidou⁵⁴, S. Kandel⁸, N. J. Kang¹³⁹, D. Kar^{34h}, E. Karentzos²⁵, O. Karkout¹¹⁷, S. N. Karpov³⁹, Z. M. Karpova³⁹, V.
Kartvelishvili⁹³, A. N. Karyukhin³⁸, E. Kasimi¹⁵⁸, J. Katzy⁴⁸, S. Kaur³⁵, K. Kawade¹⁴⁶, M. P. Kawale¹²³, C. Kawamoto⁸⁹, T. Kawamoto⁶², E. F. Kay^{37a}, F. I. Kaya¹⁶⁴, S. Kazakos¹⁰⁹, V. F. Kazanin³⁸, J. M. Keaveney^{34a}, R. Keeler¹⁷¹, G. V. Kehris⁶¹, J. S. Keller³⁵, J. M. Kelly¹⁷¹, J. J. Kempster¹⁵², O. Kepka¹³⁴, J. Kerr^{162b}, B. P. Kerridge¹³⁷, B. P. Kerševan⁹⁵, L. Keszeghova^{29a}, R. A. Khan¹³², A. Khanov¹²⁴, A. G. Kharlamov³⁸, T. Kharlamova³⁸, E. E. Khoda¹⁴², M. Kholodenko^{133a}

O. Kortner¹¹² , S. Kortner¹¹² , W. H. Kostecka¹¹⁸ , M. Kostov^{29a} , V. V. Kostyukhin¹⁴⁷ , A. Kotsokechagia^{37a} , A. Kotwal⁵¹ , A. Koulouris^{37a} , A. Kourkouveli-Charalampidi^{73a,73b} , C. Kourkouvelis⁹ , E. Kourlitis¹¹² , O. Kovanda¹²⁶ , R. Kowalewski¹⁷¹ , W. Kozanecki¹²⁶ , A. S. Kozhin³⁸ , V. A. Kramarenko³⁸ , G. Kramberger⁹⁵ , P. Kramer²⁵ , M. W. Krasny¹³⁰ , A. Krasznahorkay¹⁰⁵ , A. C. Kraus¹¹⁸ , J. W. Kraus¹⁷⁷ , J. A. Kremer⁴⁸ , N. B. Kregel¹⁴⁷ , T. Kresse⁵⁰ , L. Kretschmann¹⁷⁷ , J. Kretschmar⁹⁴ , P. Krieger¹⁶¹ , K. Krizka²¹ , K. Kroeninger⁴⁹ , H. Kroha¹¹² , J. Kroll¹³⁴ , J. Kroll¹³¹ , K. S. Krowpman¹⁰⁹ , U. Kruchonak³⁹ , H. Krüger²⁵ , N. Krumnack⁸¹ , M. C. Kruse⁵¹ , O. Kuchinskaia³⁹ , S. Kuday^{3a} , S. Kuehn^{37a} , R. Kuesters⁵⁴ , T. Kuhl⁴⁸ , V. Kukhtin³⁹ , Y. Kulchitsky³⁹ , S. Kuleshov^{140b,140d} , J. Kull¹ , E. V. Kumar¹¹¹ , M. Kumar^{34h} , N. Kumari⁴⁸ , P. Kumari^{162b} , A. Kupco¹³⁴ , T. Kupfer⁴⁹

, A. Kupich³⁸ , O. Kuprash⁵⁴ , H. Kurashige⁸⁵ , L. L. Kurchaninov^{162a} , O. Kurdysh⁴ , Y. A. Kurochkin³⁸ , A. Kurova³⁸ , E. S. Kuwertz^{37b} , M. Kuze¹⁴¹ , A. K. Kvam¹⁰⁵ , J. Kvita¹²⁵ , N. G. Kyriacou¹⁰⁸ , C. Lacasta¹⁶⁹ , F. Lacava^{75a,75b} , H. Lacker¹⁹ , D. Lacour¹³⁰ , N. N. Lad⁹⁸ , E. Ladygin³⁹ , A. Lafarge⁴¹ , B. Laforge¹³⁰ , T. Lagouri¹⁷⁸ , F. Z. Lahbabi^{36a} , S. Lai⁵⁵ , W. S. Lai⁹⁸ , J. E. Lambert¹⁷¹ , S. Lammers⁶⁸ , W. Lampl⁷ , C. Lampoudis^{158,d} , G. Lamprinoudis¹⁰² , A. N. Lancaster¹¹⁸ , E. Lançon³⁰ , U. Landgraf⁵⁴ , M. P. J. Landon⁹⁶ , V. S. Lang⁵⁴ , O. K. B. Langrekken¹²⁸ , A. J. Lankford¹⁶⁵ , F. Lanni^{37a} , K. Lantzscher²⁵ , A. Lanza^{73a} , M. Lanzac Berrocal¹⁶⁹ , J. F. Laporte¹³⁸ , T. Lari^{71a} , D. Larsen¹⁷ , L. Larson¹¹ , F. Lasagni Manghi^{24b} , M. Lassnig^{37a} , I. Latif³⁰ , S. D. Lawlor¹⁴⁵ , R. Lazaridou¹⁶⁵ , M. Lazzaroni^{71a,71b} , H. D. M. Le¹⁰⁹ , E. M. Le Boulicaut¹⁷⁸ , L. T. Le Pottier^{18a}

, B. Leban^{24a,24b} , F. Ledroit-Guillon⁶⁰ , T. F. Lee^{162b} , L. L. Leeuw^{34c} , M. Lefebvre¹⁷¹ , C. Leggett^{18a} , G. Lehmann Miotto^{37a} , M. Leigh⁵⁶ , W. A. Leight¹⁰⁵ , W. Leinonen¹¹⁶ , A. Leisos^{158,u} , M. A. L. Leite^{83c} , C. E. Leitgeb¹⁹ , R. Leitner¹³⁶ , K. J. C. Leney⁴⁵ , T. Lenz²⁵ , S. Leone^{74a} , C. Leonidopoulos⁵² , A. Leopold¹⁵⁰ , J. H. Lepage Bourbonnais³⁵ , R. Les¹⁰⁹ , C. G. Lester³³ , M. Levchenko³⁸ , J. Levêque⁴ , L. J. Levinson¹⁷⁵ , G. Levrini^{24a,24b} , M. P. Lewicki⁸⁷ , C. Lewis¹⁴² , D. J. Lewis⁴ , L. Lewitt¹⁴⁵ , A. Li³⁰ , B. Li^{143a} , C. Li¹⁰⁸ , C-Q. Li¹¹² , H. Li^{143a} , H. Li¹⁰³ , H. Li¹⁵ , H. Li⁶² , H. Li^{143a} , J. Li^{144a} , K. Li¹⁴ , L. Li^{144a} , R. Li¹⁷⁸ , S. Li^{14,114c} , S. Li^{144a,144b} , T. Li⁵ , X. Li¹⁰⁶ , Z. Li¹⁵⁹ , Z. Li^{14,114c} , Z. Li⁶² , S. Liang^{14,114c} , Z. Liang¹⁴ , M. Liberatore¹³⁸ , B. Liberti^{76a}

, K. Lie^{64c} , J. Lieber Marin^{83e} , H. Lien⁶⁸ , H. Lin¹⁰⁸ , S. F. Lin¹⁵¹ , L. Linden¹¹¹ , R. E. Lindley⁷ , J. H. Lindon^{37a} , J. Ling⁶¹ , E. Lipeles¹³¹ , A. Lipniacka¹⁷ , A. Lister¹⁷⁰ , J. D. Little⁶⁸ , B. Liu¹⁴ , B. X. Liu^{114b} , D. Liu^{144a,144b} , D. Liu¹³⁹ , E. H. L. Liu²¹ , J. K. K. Liu¹²⁰ , K. Liu^{144b} , K. Liu^{144a,144b} , M. Liu⁶² , M. Y. Liu⁶² , P. Liu¹⁴ , Q. Liu^{142,144a,144b} , X. Liu⁶² , X. Liu^{143a} , Y. Liu^{114b,114c} , Y. L. Liu^{143a} , Y. W. Liu⁶² , Z. Liu^{66,k} , S. L. Lloyd⁹⁶ , E. M. Lobodzinska⁴⁸ , P. Loch⁷ , E. Lodhi¹⁶¹ , K. Lohwasser¹⁴⁵ , E. Loiacono⁴⁸ , J. D. Lomas²¹ , J. D. Long⁴² , I. Longarini¹⁶⁵ , R. Longo¹⁶⁸ , A. Lopez Solis¹³ , N. A. Lopez-canelas⁷ , N. Lorenzo Martinez⁴ , A. M. Lory¹¹¹ , M. Losada^{119a} , G. Löschke Centeno¹⁵² , X. Lou^{47a,47b} , X. Lou^{14,114c} , A. Lounis⁶⁶ , P. A. Love⁹³ , M. Lu⁶⁶ , S. Lu¹³¹ , Y. J.
Lu¹⁵⁴ , H. J. Lubatti¹⁴² , C. Luci^{75a,75b} , F. L. Lucio Alves^{114a} , F. Luehring⁶⁸ , B. S. Lunday¹³¹ , O. Lundberg¹⁵⁰ , J. Lunde^{37a} , N. A. Luongo⁶ , M. S. Lutz^{37a} , A. B. Lux²⁶ , D. Lynn³⁰ , R. Lysak¹³⁴ , V. Lysenko¹³⁵ , E. Lytken¹⁰⁰ , V. Lyubushkin³⁹ , T. Lyubushkina³⁹ , M. M. Lyukova¹⁵¹ , M. Firdaus M. Soberi⁵² , H. Ma³⁰ , K. Ma⁶² , L. L. Ma^{143a} , W. Ma⁶² , Y. Ma¹²⁴ , J. C. MacDonald¹⁰² , P. C. Machado De Abreu Farias^{83e} , R. Madar⁴¹ , T. Madula⁹⁸ , J. Maeda⁸⁵ , T. Maeno³⁰ , P. T. Mafa^{34c,j} , H. Maguire¹⁴⁵ , M. Maheshwari³³ , V. Maiboroda⁶⁶ , A. Maio^{133a,133b,133d} , K. Maj^{86a} , O. Majersky⁴⁸ , S. Majewski¹²⁶ , R. Makhmanazarov³⁸ , N. Makovec⁶⁶ , V. Maksimovic¹⁶ , B. Malaescu¹³⁰ , J. Malamant¹²⁸ , Pa. Malecki⁸⁷ , V. P. Maleev³⁸ , F. Malek^{60,o} , M. Mali⁹⁵ , D. Malito⁹⁷ , U. Mallik^{80,*} , A. Maloizel⁵ , S. Maltezos¹⁰ , A. Malvezzi Lopes^{83d} , S. Malyukov³⁹

, J. Mamuzic⁹⁵ , G. Mancini⁵³ , M. N. Mancini²⁷ , G. Manco^{73a,73b} , J. P. Mandalia⁹⁶ , S. S. Mandary¹⁵² , I. Mandić⁹⁵ , L. Manhaes de Andrade Filho^{83a} , I. M. Maniatis¹⁷⁵ , J. Manjarres Ramos⁹¹ , D. C. Mankad¹⁷⁵ , A. Mann¹¹¹ , T. Manoussos^{37a} , M. N. Mantinan⁴⁰ , S. Manzoni^{37a} , L. Mao^{144a} , X. Mapekula^{34c} , A. Marantis¹⁵⁸ , R. R. Marcelo Gregorio⁹⁶ , G. Marchiori⁵

E. F. McDonald¹⁰⁷, A. E. McDougall¹¹⁷, L. F. Mcelhinney⁹³, J. A. Mcfayden¹⁵², R. P. McGovern¹³¹, R. P. McKenzie^{34h}, T. C. McLachlan⁴⁸, D. J. McLaughlin⁹⁸, S. J. McMahon¹³⁷, C. M. Mcpartland⁹⁴, R. A. McPherson^{171,ab}, S. Mehlhase¹¹¹, A. Mehta⁹⁴, D. Melini¹⁶⁹, B. R. Mellado Garcia^{34h}, A. H. Melo⁵⁵, F. Meloni⁴⁸, A. M. Mendes Jacques Da Costa¹⁰³, L. Meng⁹³, S. Menke¹¹², M. Mentink^{37a}, E. Meoni^{44a,44b}, G. Mercado¹¹⁸, S. Merianos¹⁵⁸, C. Merlassino^{69a,69c}, C. Meroni^{71a,71b}, J. Metcalfe⁶, A. S. Mete⁶, E. Meuser¹⁰², C. Meyer⁶⁸, J.-P. Meyer¹³⁸, Y. Miao^{114a}, R. P. Middleton¹³⁷, M. Mihovilovic⁶⁶, L. Mijović⁵², G. Mikenberg¹⁷⁵, M. Mikestikova¹³⁴, M. Mikuž⁹⁵, H. Mildner¹⁰², A. Milic^{37a}, D. W. Miller⁴⁰, E. H. Miller¹⁴⁹, A. Milov¹⁷⁵, D. A. Milstead^{47a,47b}, T. Min^{114a}, A. A. Minaenko³⁸, I. A. Minashvili^{155b}, A. I. Mincer¹²⁰, B. Mindur^{86a}, M. Mineev³⁹, Y. Mino⁸⁹, L. M. Mir¹³, M. Miralles Lopez⁵⁹, M. Mironova^{18a}, M. C. Missio¹¹⁶, A. Mitra¹⁷³, V. A. Mitsou¹⁶⁹, Y. Mitsumori¹¹³, O. Miu¹⁶¹, P. S. Miyagawa⁹⁶, T. Mkrtchyan^{63a}, M. Mlinarevic⁹⁸, T. Mlinarevic⁹⁸, M. Mlynarikova¹³⁶, L. Mlynarska^{86a}, S. Mobius²⁰, M. H. Mohamed Farook¹¹⁵, S. Mohapatra⁴², S. Mohiuddin¹²⁴, G. Mokgatitswane^{34h}, L. Moleri¹⁷⁵, U. Molinatti¹²⁹, L. G. Mollier²⁰, B. Mondal¹⁴⁷, S. Mondal¹³⁵, K. Mönig⁴⁸, E. Monnier¹⁰⁴, L. Monsonis Romero¹⁶⁹, J. Montejo Berlingen¹³, A. Montella^{47a,47b}, M. Montella¹²², F. Montereali^{77a,77b}, F. Monticelli⁹², S. Monzani^{69a,69c}, A. Morancho Tarda⁴³, N. Morange⁶⁶, A. L. Moreira De Carvalho⁴⁸, M. Moreno Llácer¹⁶⁹, C. Moreno Martinez⁵⁶, J. M. Moreno Perez^{23b}, P. Moretini^{57b}, S. Morgenstern^{37a}, M. Morii⁶¹, M. Morinaga¹⁵⁹, M. Moritsu⁹⁰, F. Morodei^{75a,75b}, P. Moschovakos^{37a}, B. Moser⁵⁴, M. Mosidze^{155b}, T. Moskalets⁴⁵, P. Moskvitina¹¹⁶, J. Moss³², P. Moszkowicz^{86a}, A. Moussa^{36d}, Y. Moyal¹⁷⁵, H. Moyano Gomez¹³, E. J. W. Moyse¹⁰⁵, T. G. Mroz⁸⁷, O. Mtintsilana^{34h}, S. Muanza¹⁰⁴, M. Mucha²⁵, J. Mueller¹³², R. Müller^{37a}, G. A. Mullier¹⁶⁷, A. J. Mullin³³, J. J. Mullin⁵¹, A. C. Mullins⁴⁵, A. E. Mulski⁹⁵, D. P. Mungo¹⁶¹, D. Munoz Perez¹⁶⁹, F. J. Munoz Sanchez¹⁰³, W. J. Murray^{137,173}, M. Muškinja⁹⁵, C. Mwewa⁴⁸, A. G. Myagkov^{38a}, A. J. Myers⁸, G. Myers¹⁰⁸, M. Myska¹³⁵, B. P. Nachman^{18a}, K. Nagai¹²⁹, K. Nagano⁸⁴, R. Nagasaka¹⁵⁹, J. L. Nagle^{30,aj}, E. Nagy¹⁰⁴, A. M. Nairz^{37a}, Y. Nakahama⁸⁴, K. Nakamura⁸⁴, K. Nakkalil⁵, A. Nandi^{63b}, H. Nanjo¹²⁷, E. A. Narayanan⁴⁵, Y. Narukawa¹⁵⁹, I. Naryshkin³⁸, L. Nasella^{71a,71b}, S. Nasri^{119b}, C. Nass²⁵, G. Navarro^{23a}, J. Navarro-Gonzalez¹⁶⁹, A. Nayaz¹⁹, P. Y. Nechaeva³⁸, S. Nechaeva^{24a,24b}, F. Nechansky¹³⁴, L. Nedic¹²⁹, T. J. Neep²¹, A. Negri^{73a,73b}, M. Negrini^{24b}, C. Nellist¹¹⁷, C. Nelson¹⁰⁶, K. Nelson¹⁰⁸, S. Nemecek¹³⁴, M. Nessi^{37a,g}, M. S. Neubauer¹⁶⁸, J. Newell⁹⁴, P. R. Newman²¹, Y. W. Y. Ng¹⁶⁸, B. Ngair^{119a}, H. D. N. Nguyen¹¹⁰, J. D. Nichols¹²³, R. B. Nickerson¹²⁹, R. Nicolaidou¹³⁸, J. Nielsen¹³⁹, M. Niemeyer⁵⁵, J. Niermann^{37a}, N. Nikiforou^{37a}, V. Nikolaenko^{38a}, I. Nikolic-Audit¹³⁰, P. Nilsson³⁰, I. Ninca⁴⁸, G. Ninio¹⁵⁷, A. Nisati^{75a}, R. Nisius¹¹², N. Nitika^{69a,69c}, J.-E. Nitschke⁵⁰, E. K. Nkadimeng^{34b}, T. Nobe¹⁵⁹, D. Noll^{18a}, T. Nommensen¹⁵³, M. B. Norfolk¹⁴⁵, B. J. Norman³⁵, M. Noury^{36a}, J. Novak⁹⁵, T. Novak⁹⁵, R. Novotny¹³⁵, L. Nozka¹²⁵, K. Ntekas¹⁶⁵, N. M. J. Nunes De Moura Junior^{83b}, J. Ocariz¹³⁰, A. Ochi⁸⁵, I. Ochoa^{133a}, S. Oerdek^{48,y}, J. T. Offermann⁴⁰, A. Ogrodnik¹³⁶, A. Oh¹⁰³, C. C. Ohm¹⁵⁰, H. Oide⁸⁴, M. L. Ojeda^{37a}, Y. Okumura¹⁵⁹, L. F. Oleiro Seabra^{133a}, I. Oleksiyuk⁵⁶, G. Oliveira Correa¹³, D. Oliveira Damazio³⁰, J. L. Oliver¹⁶⁵, R. Omar⁶⁸, Ö. O. Öncel⁵⁴, A. P. O'Neill²⁰, A. Onofre^{133a,133e,e}, P. U. E. Onyisi¹¹, M. J. Oreglia⁴⁰, D. Orestano^{77a,77b}, R. Orlandini^{77a,77b}, R. S. Orr¹⁶¹, L. M. Osojnak¹³¹, Y. Osumi¹¹³, G. Otero y Garzon³¹, H. Otono⁹⁰, M. Ouchrif^{36d}, F. Ould-Saada¹²⁸, T. Ovsianikova¹⁴², M. Owen⁵⁹, R. E. Owen¹³⁷, V. E. Ozcan^{22a}, F. Ozturk⁸⁷, N. Ozturk⁸, S. Ozturk⁸², H. A. Pacey¹²⁹, K. Pachal^{162a}, A. Pacheco Pages¹³, C. Padilla Aranda¹³, G. Padovano^{75a,75b}, S. Pagan Griso^{18a}, G. Palacino⁶⁸, A. Palazzo^{70a,70b}, J. Pampel²⁵, J. Pan¹⁷⁸, T. Pan^{64a}, D. K. Panchal¹¹, C. E. Pandini⁶⁰, J. G. Panduro Vazquez¹³⁷, H. D. Pandya¹, H. Pang¹³⁸, P. Pani⁴⁸, G. Panizzo^{69a,69c}, L. Panwar¹³⁰, L. Paolozzi⁵⁶, S. Parajuli¹⁶⁸, A. Paramonov⁶, C. Paraskevopoulos⁵³, D. Paredes Hernandez^{64b}, A. Pareti^{73a,73b}, K. R. Park⁴², T. H. Park¹¹², F. Parodi^{57a,57b}, J. A. Parsons⁴², U. Parzefall⁵⁴, B. Pascual Dias⁴¹, L. Pascual Dominguez¹⁰¹, E. Pasqualucci^{75a}, S. Passaggio^{57b}, F. Pastore⁹⁷, P. Patel⁸⁷, U. M. Patel⁵¹, J. R. Pater¹⁰³, T. Pauly^{37a}, F. Pauwels¹³⁶, C. I. Pazos¹⁶⁴, M. Pedersen¹²⁸, R. Pedro^{133a}, S. V. Peleganchuk³⁸, O. Penc¹³⁴, E. A. Pender⁵², S. Peng¹⁵, G. D. Penn¹⁷⁸, K. E. Penski¹¹¹, M. Penzin³⁸, B. S. Peralva^{83d}, A. P. Pereira Peixoto¹⁴², L. Pereira Sanchez¹⁴⁹, D. V. Perepelitsa^{30,aj}, G. Perera¹⁰⁵, E. Perez Codina^{37a}, M. Perganti¹⁰, H. Pernegger^{37a}, S. Perrella^{75a,75b}, K. Peters⁴⁸, R. F. Y. Peters¹⁰³, B. A. Petersen^{37a}, T. C. Petersen⁴³, E. Petit¹⁰⁴, V. Petousis¹³⁵, A. R. Petri^{71a,71b}, C. Petridou^{158,d}, T. Petru¹³⁶, A. Petrukhin¹⁴⁷, M. Pettee^{18a}, A. Petukhov⁸², K. Petukhova^{37a}, R. Pezoa^{140g}, L. Pezzotti^{24a,24b}, G. Pezzullo¹⁷⁸, L. Pfaffenbichler^{37a}, A. J. Pfleger⁷⁹, T. M. Pham¹⁷⁶, T. Pham¹⁰⁷, P. W. Phillips¹³⁷, G. Piacquadio¹⁵¹, E. Pianori^{18a}, F. Piazza¹²⁶, R. Piegaia³¹, D. Pietreanu^{28b}, A. D. Pilkington¹⁰³, M. Pinamonti^{69a,69c}, J. L. Pinfold²,

B. C. Pinheiro Pereira^{133a}, J. Pinol Bel¹³, A. E. Pinto Pinoargote¹³⁰, L. Pintucci^{69a,69c}, K. M. Piper¹⁵², A. Pirttikoski⁵⁶, D. A. Pizzi³⁵, L. Pizzimento^{64b}, A. Plebani³³, M.-A. Pleier³⁰, V. Pleskot¹³⁶, E. Plotnikova³⁹, G. Poddar⁹⁶, R. Poettgen¹⁰⁰, L. Poggioli¹³⁰, S. Polacek¹³⁶, G. Polesello^{73a}, A. Poley¹⁴⁸, A. Polini^{24b}, C. S. Pollard¹⁷³, Z. B. Pollock¹²², E. Pompa Pacchi¹²³, N. I. Pond⁹⁸, D. Ponomarenko⁶⁸, L. Pontecorvo^{37a}, S. Popa^{28a}, G. A. Popeneciu^{28d}, A. Poreba^{37a}, D. M. Portillo Quintero^{162a}, S. Pospisil¹³⁵, M. A. Postill¹⁴⁵, P. Postolache^{28c}, K. Potamianos¹⁷³, P. A. Potepa^{86a}, I. N. Potrap³⁹, C. J. Potter³³, H. Potti¹⁵³, J. Poveda¹⁶⁹, M. E. Pozo Astigarraga^{37a}, R. Pozzi^{37a}, A. Prades Ibanez^{76a,76b}, S. R. Pradhan¹⁴⁵, J. Pretel¹⁷¹, D. Price¹⁰³, M. Primavera^{70a}, L. Primomo^{69a,69c}, M. A. Principe Martin¹⁰¹, R. Privara¹²⁵, T. Procter^{86b}, M. L. Proffitt¹⁴², N. Proklova¹³¹, K. Prokofiev^{64c}, G. Proto¹¹², J. Proudfoot⁶, M. Przybycien^{86a}, W. W. Przygoda^{86b}, A. Psallidas⁴⁶, J. E. Puddefoot¹⁴⁵, D. Pudzha⁵³, H. I. Purnell¹, D. Pyatiybyantseva¹¹⁶, J. Qian¹⁰⁸, R. Qian¹⁰⁹, D. Qichen¹²⁹, Y. Qin¹³, T. Qiu⁵², A. Quadt⁵⁵, M. Queitsch-Maitland¹⁰³, G. Quetant⁵⁶, R. P. Quinn¹⁷⁰, G. Rabanal Bolanos⁶¹, D. Rafanoharana¹¹², F. Raffaeli^{76a,76b}, F. Ragusa^{71a,71b}, J. L. Rainbolt⁴⁰, S. Rajagopalan³⁰, E. Ramakoti³⁹, L. Rambelli^{57a,57b}, I. A. Ramirez-Berend³⁵, K. Ran^{48,114c}, D. S. Rankin¹³¹, V. Rao^{an}, N. P. Rapheeha^{34b}, H. Rasheed^{28b}, D. F. Rassloff^{63a}, A. Rastogi^{18a}, S. Rave¹⁰², S. Ravera^{57a,57b}, B. Ravina^{37a}, I. Ravinovitch¹⁷⁵, M. Raymond^{37a}, A. L. Read¹²⁸, N. P. Readioff¹⁴⁵, D. M. Rebuffi^{73a,73b}, A. S. Reed⁵⁹, K. Reeves²⁷, J. A. Reidelsturz¹⁷⁷, D. Reikher¹²⁶, A. Rej⁴⁹, C. Rembser^{37a}, H. Ren⁶², M. Renda^{28b}, F. Renner⁴⁸, A. G. Rennie⁵⁹, A. L. Rescia^{57a,57b}, S. Resconi^{71a}, M. Ressegotti^{57a,57b}, S. Rettie^{37a}, W. F. Rettie³⁵, M. M. Revering³³, E. Reynolds^{18a}, O. L. Rezanova³⁹, P. Reznicek¹³⁶, H. Riani^{36d}, N. Ribaric⁵¹, B. Ricci^{69a,69c}, E. Ricci^{78a,78b}, R. Richter¹¹², S. Richter^{47a,47b}, E. Richter-Was^{86b}, M. Ridel¹³⁰, S. Ridouani^{36d}, P. Rieck¹²⁰, P. Riedler^{37a}, E. M. Riefel^{47a,47b}, J. O. Rieger¹¹⁷, M. Rijssenbeek¹⁵¹, M. Rimoldi^{37a}, L. Rinaldi^{24a,24b}, P. Rincke^{55,167}, G. Ripellino¹⁶⁷, I. Riu¹³, J. C. Rivera Vergara¹⁷¹, F. Rizatdinova¹²⁴, E. Rizvi⁹⁶, B. R. Roberts^{18a}, S. S. Roberts¹³⁹, D. Robinson³³, M. Robles Manzano¹⁰², A. Robson⁵⁹, A. Rocchi^{76a,76b}, C. Roda^{74a,74b}, S. Rodriguez Bosca^{37a}, Y. Rodriguez Garcia^{23a}, A. M. Rodríguez Vera¹¹⁸, S. Roe^{37a}, J. T. Roemer^{37a}, O. Røhne¹²⁸, R. A. Rojas^{37a}, C. P. A. Roland¹³⁰, A. Romaniouk⁷⁹, E. Romano^{73a,73b}, M. Romano^{24b}, A. C. Romero Hernandez¹⁶⁸, N. Rompotis⁹⁴, L. Roos¹³⁰, S. Rosati^{75a}, B. J. Rosser⁴⁰, E. Rossi¹²⁹, E. Rossi^{72a,72b}, L. P. Rossi⁶¹, L. Rossini⁵⁴, R. Rosten¹²², M. Rotaru^{28b}, B. Rottler⁵⁴, D. Rousseau⁶⁶, D. Rouso⁴⁸, S. Roy-Garand¹⁶¹, A. Rozanov¹⁰⁴, Z. M. A. Rozario⁵⁹, Y. Rozen¹⁵⁶, A. Rubio Jimenez¹⁶⁹, V. H. Ruelas Rivera¹⁹, T. A. Ruggeri¹, A. Ruggiero¹²⁹, A. Ruiz-Martinez¹⁶⁹, A. Rummler^{37a}, Z. Rurikova⁵⁴, N. A. Rusakovich³⁹, S. Ruscelli⁴⁹, H. L. Russell¹⁷¹, G. Russo^{75a,75b}, J. P. Rutherford⁷, S. Rutherford Colmenares³³, M. Rybar¹³⁶, P. Rybczynski^{86a}, A. Ryzhov⁴⁵, J. A. Sabater Iglesias⁵⁶, H. F.-W. Sadrozinski¹³⁹, F. Safai Tehrani^{75a}, S. Saha¹, M. Sahinsoy⁸², B. Sahoo¹⁷⁵, A. Saibel¹⁶⁹, B. T. Saifuddin¹²³, M. Saimpert¹³⁸, G. T. Saito^{83c}, M. Saito¹⁵⁹, T. Saito¹⁵⁹, A. Sala^{71a,71b}, A. Salnikov¹⁴⁹, J. Salt¹⁶⁹, A. Salvador Salas¹⁵⁷, F. Salvatore¹⁵², A. Salzburger^{37a}, D. Sammel⁵⁴, E. Sampson⁹³, D. Sampsonidis^{158,d}, D. Sampsonidou¹²⁶, J. Sánchez¹⁶⁹, V. Sanchez Sebastian¹⁶⁹, H. Sandaker¹²⁸, C. O. Sander⁴⁸, J. A. Sandesara¹⁷⁶, M. Sandhoff¹⁷⁷, C. Sandoval^{23b}, L. Sanfilippo^{63a}, D. P. C. Sankey¹³⁷, T. Sano⁸⁹, A. Sansoni⁵³, M. Santana Queiroz^{18b}, L. Santi^{37a}, C. Santoni⁴¹, H. Santos^{133a,133b}, A. Santra¹⁷⁵, E. Sanzani^{24a,24b}, K. A. Saoucha^{88b}, J. G. Saraiva^{133a,133d}, J. Sardain⁷, O. Sasaki⁸⁴, K. Sato¹⁶³, C. Sauer^{37a}, E. Sauvan⁴, P. Savard^{161,ah}, R. Sawada¹⁵⁹, C. Sawyer¹³⁷, L. Sawyer⁹⁹, C. Sbarra^{24b}, A. Sbrizzi^{24a,24b}, T. Scanlon⁹⁸, J. Schaarschmidt¹⁴², U. Schäfer¹⁰², A. C. Schaffer^{45,66}, D. Schaile¹¹¹, R. D. Schamberger¹⁵¹, C. Scharf¹⁹, M. M. Schefer²⁰, V. A. Schegelsky³⁸, D. Scheirich¹³⁶, M. Schernau^{140f}, C. Scheulen⁵⁶, C. Schiavi^{57a,57b}, D. Schien^{ap}, M. Schioppa^{44a,44b}, B. Schlag¹⁴⁹, S. Schlenker^{37a}, J. Schmeing¹⁷⁷, E. Schmidt¹¹², M. A. Schmidt¹⁷⁷, K. Schmieden¹⁰², C. Schmitt¹⁰², N. Schmitt¹⁰², S. Schmitt⁴⁸, N. A. Schneider¹¹¹, L. Schoeffel¹³⁸, A. Schoening^{63b}, P. G. Scholer³⁵, E. Schopf¹⁴⁷, M. Schott²⁵, S. Schramm⁵⁶, T. Schroer⁵⁶, H.-C. Schultz-Coulon^{63a}, M. Schumacher⁵⁴, B. A. Schumm¹³⁹, Ph. Schune¹³⁸, H. R. Schwartz⁷, A. Schwartzman¹⁴⁹, T. A. Schwarz¹⁰⁸, Ph. Schwemling¹³⁸, R. Schwienhorst¹⁰⁹, F. G. Sciaccia²⁰, A. Sciandra³⁰, G. Sciolla²⁷, F. Scuri^{74a}, C. D. Sebastiani^{37a}, K. Sedlaczek¹¹⁸, S. C. Seidel¹¹⁵, A. Seiden¹³⁹, B. D. Seidlitz⁴², C. Seitz⁴⁸, J. M. Seixas^{83b}, G. Sekhniaidze^{72a}, L. Selem⁶⁰, N. Semprini-Cesari^{24a,24b}, A. Semushin¹⁷⁹, D. Sengupta⁵⁶, V. Senthilkumar¹⁶⁹, L. Serin⁶⁶, M. Sessa^{72a,72b}, H. Severini¹²³, F. Sforza^{57a,57b}, A. Sfyrila⁵⁶, Q. Sha¹⁴, P. Shabajee^{ap}, E. Shabalina⁵⁵, H. Shaddix¹¹⁸, A. H. Shah³³, R. Shaheen¹⁵⁰, J. D. Shahinian¹³¹, M. Shamim^{37a}, L. Y. Shan¹⁴, M. Shapiro^{18a}, A. Sharma^{37a}, A. S. Sharma¹⁷⁰, P. Sharma³⁰, P. B. Shatalov³⁸, K. Shaw¹⁵², S. M. Shaw¹⁰³, Q. Shen^{144a}, D. J. Sheppard¹⁴⁸, P. Sherwood⁹⁸, L. Shi⁹⁸, X. Shi¹⁴, S. Shimizu⁸⁴, C. O. Shimmin¹⁷⁸, I. P. J. Shipsey^{129,*}, S. Shirabe⁹⁰, M. Shiyakova^{39,z}, M. J. Shochet⁴⁰,

D. R. Shope¹²⁸, B. Shrestha¹²³, S. Shrestha^{122,al}, I. Shreyber³⁹, M. J. Shroff¹⁷¹, P. Sicho¹³⁴, A. M. Sickles¹⁶⁸, E. Sideras Haddad^{34h,166}, A. C. Sidley¹¹⁷, A. Sidoti^{24b}, F. Siegert⁵⁰, Dj. Sijacki¹⁶, F. Sili⁹², J. M. Silva⁵², I. Silva Ferreira^{83b}, M. V. Silva Oliveira³⁰, S. B. Silverstein^{47a}, E. Simili⁵⁹, S. Simion⁶⁶, R. Simoniello^{37a}, E. L. Simpson¹⁰³, H. Simpson¹⁵², L. R. Simpson⁶, S. Simsek⁸², S. Sindhu⁵⁵, P. Sinervo¹⁶¹, S. N. Singh²⁷, S. Singh³⁰, S. Sinha⁴⁸, S. Sinha¹⁰³, M. Sioli^{24a,24b}, K. Sioulas⁹, I. Siral^{37a}, E. Sitnikova⁴⁸, J. Sjölin^{47a,47b}, A. Skaf⁵⁵, E. Skorda²¹, P. Skubic¹²³, M. Slawinska⁸⁷, I. Slazyk¹⁷, I. Sliusar¹²⁸, V. Smakhtin¹⁷⁵, B. H. Smart¹³⁷, S. Yu. Smirnov^{140b}, Y. Smirnov⁸², L. N. Smirnova^{38,a}, O. Smirnova¹⁰⁰, A. C. Smith⁴², D. R. Smith¹⁶⁵, J. L. Smith¹⁰³, M. B. Smith³⁵, R. Smith¹⁴⁹, H. Smitmanns¹⁰², M. Smizanska⁹³, K. Smolek¹³⁵, P. Smolyanskiy¹³⁵, A. A. Snesarev³⁹, H. L. Snoek¹¹⁷, S. Snyder³⁰, R. Sobie^{171,ab}, A. Soffer¹⁵⁷, C. A. Solans Sanchez^{37a}, E. Yu. Soldatov³⁹, U. Soldevila¹⁶⁹, A. A. Solodkov^{34h}, S. Solomon²⁷, A. Soloshenko³⁹, K. Solovieva⁵⁴, O. V. Solovyanov⁴¹, P. Sommer⁵⁰, A. Sonay¹³, A. Sopczak¹³⁵, A. L. Sopio⁵², F. Sopkova^{29b}, J. D. Sorenson¹¹⁵, I. R. Sotarriva Alvarez¹⁴¹, V. Sothilingam^{63a}, O. J. Soto Sandoval^{140b,140c}, S. Sottocornola⁶⁸, R. Soualah^{88a}, Z. Soumami^{36e}, D. South⁴⁸, N. Soybelman¹⁷⁵, S. Spagnolo^{70a,70b}, M. Spalla¹¹², D. Sperlich⁵⁴, B. Spisso^{72a,72b}, D. P. Spiteri⁵⁹, L. Splendori¹⁰⁴, M. Spousta¹³⁶, E. J. Staats³⁵, R. Stamen^{63a}, E. Stanecka⁸⁷, W. Stanek-Maslouska⁴⁸, M. V. Stange⁵⁰, B. Stanislaus^{18a}, M. M. Stanitzki⁴⁸, B. Stapf⁴⁸, E. A. Starchenko³⁸, G. H. Stark¹³⁹, J. Stark⁹¹, P. Staroba¹³⁴, P. Starovoitov^{88b}, R. Staszewski⁸⁷, C. Stauch¹¹¹, G. Stavropoulos⁴⁶, A. Steff^{37a}, A. Stein¹⁰², P. Steinberg³⁰, B. Stelzer^{148,162a}, H. J. Stelzer¹³², O. Stelzer-Chilton^{162a}, H. Stenzel⁵⁸, T. J. Stevenson¹⁵², G. Stewart⁵⁹, G. A. Stewart^{37a}, J. R. Stewart¹²⁴, G. Stoicescu^{28b}, M. Stolarski^{133a}, S. Stonjek¹¹², A. Straessner⁵⁰, J. Strandberg¹⁵⁰, S. Strandberg^{47a,47b}, M. Stratmann¹⁷⁷, M. Strauss¹²³, T. Streblner¹⁰⁴, P. Strizenc^{29b}, R. Ströhmer¹⁷², D. M. Strom¹²⁶, R. Stroyanowski⁴⁵, A. Strubig^{47a,47b}, S. A. Stucci³⁰, B. Stugu¹⁷, J. Stupak¹²³, N. A. Styles⁴⁸, D. Su¹⁴⁹, S. Su⁶², X. Su⁶², D. Suchy^{29a}, K. Sugizaki¹³¹, V. V. Sulin³⁸, D. M. S. Sultan¹²⁹, L. Sultanaliyeva³⁸, S. Sultansoy^{3b}, S. Sun¹⁷⁶, W. Sun¹⁴, O. Sunneborn Gudnadottir¹⁶⁷, N. Sur¹⁰⁰, M. R. Sutton¹⁵², M. Svatos¹³⁴, P. N. Swallow³³, M. Swiatkowski^{162a}, T. Swirski¹⁷², A. Swoboda^{37a}, I. Sykora^{29a}, M. Sykora¹³⁶, T. Sykora¹³⁶, N. D. Szczepanek^{37a}, D. Ta¹⁰², K. Tackmann^{48,y}, A. Taffard¹⁶⁵, R. Tahirout^{162a}, Y. Takubo⁸⁴, M. Talby¹⁰⁴, A. A. Talyshev³⁸, K. C. Tam^{64b}, N. M. Tamir¹⁵⁷, A. Tanaka¹⁵⁹, J. Tanaka¹⁵⁹, R. Tanaka⁶⁶, M. Tanasini¹⁵¹, Z. Tao¹⁷⁰, S. Tapia Araya^{140g}, S. Tapprogge¹⁰², A. Tarek Abouelfadl Mohamed¹⁰⁹, S. Tarem¹⁵⁶, K. Tariq¹⁴, G. Tarna^{37a}, G. F. Tartarelli^{71a}, M. J. Tartarin⁹¹, P. Tas¹³⁶, M. Tasevsky¹³⁴, E. Tassi^{44a,44b}, A. C. Tate¹⁶⁸, Y. Tayalati^{36e,aa}, G. N. Taylor¹⁰⁷, W. Taylor^{162b}, A. S. Tegetmeier⁹¹, P. Teixeira-Dias⁹⁷, J. J. Teoh¹⁶¹, K. Terashi¹⁵⁹, J. Terron¹⁰¹, S. Terzo¹³, M. Testa⁵³, R. J. Teuscher^{161,ab}, A. Thaler⁷⁹, O. Theiner⁵⁶, T. Theveniaux-Pelzer¹⁰⁴, D. W. Thomas⁹⁷, J. P. Thomas²¹, E. A. Thompson^{18a}, P. D. Thompson²¹, E. Thomson¹³¹, R. E. Thornberry⁴⁵, C. Tian⁶², Y. Tian⁵⁶, V. Tikhomirov⁸², Yu. A. Tikhonov³⁹, S. Timoshenko³⁸, D. Timoshyn¹³⁶, E. X. L. Ting¹, P. Tipton¹⁷⁸, A. Tishelman-Charny³⁰, K. Todome¹⁴¹, S. Todorova-Nova¹³⁶, L. Toffolin^{69a,69c}, M. Togawa⁸⁴, J. Tojo⁹⁰, S. Tokár^{29a}, O. Toldaiev⁶⁸, G. Tolkachev¹⁰⁴, M. Tomoto⁸⁴, L. Tompkins^{149,n}, E. Torrence¹²⁶, H. Torres⁹¹, E. Torró Pastor¹⁶⁹, M. Toscani³¹, C. Toscirri⁴⁰, M. Tost¹¹, D. R. Tovey¹⁴⁵, T. Trefzger¹⁷², P. M. Tricarico¹³, A. Tricoli³⁰, I. M. Trigger^{162a}, S. Trincas-Duvold¹³⁰, D. A. Trischuk²⁷, A. Tropina³⁹, L. Truong^{34c}, M. Trzebinski⁸⁷, A. Trzupek⁸⁷, F. Tsai¹⁵¹, M. Tsai¹⁰⁸, A. Tsiamis¹⁵⁸, P. V. Tsiareshka³⁹, S. Tsigaridas^{162a}, A. Tsigotis^{158,u}, V. Tsiskaridze^{155a}, E. G. Tskhadadze^{155a}, M. Tsopoulou¹⁵⁸, Y. Tsujikawa⁸⁹, I. I. Tsukerman³⁸, V. Tsulaia^{18a}, S. Tsuno⁸⁴, K. Tsuru¹²¹, D. Tsybychev¹⁵¹, Y. Tu^{64b}, A. Tudorache^{28b}, V. Tudorache^{28b}, S. B. Tuncay¹²⁹, S. Turchikhin^{57a,57b}, I. Turk Cakir^{3a}, R. Turra^{71a}, T. Turtuvshin^{39,ac}, P. M. Tuts⁴², S. Tzamarias^{158,d}, E. Tzovara¹⁰², Y. Uematsu⁸⁴, F. Ukegawa¹⁶³, P. A. Ulloa Poblete^{140b,140c}, E. N. Umaka³⁰, G. Unal^{37a}, A. Undrus³⁰, G. Unel¹⁶⁵, J. Urban^{29b}, P. Urrejola^{140a}, G. Usai⁸, R. Ushioda¹⁶⁰, M. Usman¹¹⁰, F. Ustuner⁵², Z. Uysal⁸², V. Vacek¹³⁵, B. Vachon¹⁰⁶, T. Vafeiadis^{37a}, A. Vaitkus⁹⁸, C. Valderanis¹¹¹, E. Valdes Santurio^{47a,47b}, M. Valente^{37a}, S. Valentinetti^{24a,24b}, A. Valero¹⁶⁹, E. Valiente Moreno¹⁶⁹, A. Vallier⁹¹, J. A. Valls Ferrer¹⁶⁹, D. R. Van Arneman¹¹⁷, A. Van Der Graaf⁴⁹, H. Z. Van Der Schyf^{34h}, P. Van Gemmeren⁶, M. Van Rijnbach^{37a}, S. Van Stroud⁹⁸, I. Van Vulpen¹¹⁷, P. Vana¹³⁶, M. Vanadia^{76a,76b}, U. M. Vande Voorde¹⁵⁰, W. Vandelli^{37a}, E. R. Vandewall¹²⁴, D. Vannicola¹⁵⁷, L. Vannoli⁵³, R. Vari^{75a}, M. Varma¹⁷⁸, E. W. Varnes⁷, C. Varni¹¹⁸, D. Varouchas⁶⁶, L. Varriale¹⁶⁹, K. E. Varvell¹⁵³, M. E. Vasile^{28b}, L. Vaslin⁸⁴, M. D. Vassilev¹⁴⁹, A. Vasyukov³⁹, L. M. Vaughan¹²⁴, R. Vavricka¹³⁶, T. Vazquez Schroeder¹³, J. Veatch³², V. Vecchio¹⁰³, M. J. Veen¹⁰⁵, I. Veliscek³⁰, I. Velkovska⁹⁵, L. M. Veloce¹⁶¹, F. Veloso^{133a,133c}, S. Veneziano^{75a}, A. Ventura^{70a,70b}, A. Verbytskyi¹¹², M. Verducci^{74a,74b}, C. Vergis⁹⁶, M. Verissimo De Araujo^{83b}

W. Verkerke¹¹⁷ , J. C. Vermeulen¹¹⁷ , C. Vernieri¹⁴⁹ , M. Vessella¹⁶⁵ , M. C. Vetterli^{148,ah} , A. Vgenopoulos¹⁰² , N. Viaux Maira^{140g} , T. Vickey¹⁴⁵ , O. E. Vickey Boeriu¹⁴⁵ , G. H. A. Viehhauser¹²⁹ , L. Viganì^{63b} , M. Vigi¹¹² , M. Villa^{24a,24b} , M. Villaplana Perez¹⁶⁹ , E. M. Villhauer⁴⁰ , E. Vilucchi⁵³ , M. Vincent¹⁶⁹ , M. G. Vinciter³⁵ , A. Visibile¹¹⁷ , C. Vittori^{37a} , I. Vivarelli^{24a,24b} , E. Voevodina¹¹² , F. Vogel¹¹¹ , J. C. Voigt⁵⁰ , P. Vokac¹³⁵ , Yu. Volkotrub^{86b} , L. Vomberg²⁵ , E. Von Toerne²⁵ , B. Vormwald^{37a} , K. Vorobev⁵¹ , M. Vos¹⁶⁹ , K. Voss¹⁴⁷ , M. Vozak^{37a} , L. Vozdecky¹²³ , N. Vranjes¹⁶ , M. Vranjes Milosavljevic¹⁶ , M. Vreeswijk¹¹⁷ , N. K. Vu^{144a,144b} , R. Vuillermet^{37a} , O. Vujanovic¹⁰² , I. Vukotic⁴⁰ , I. K. Vyas³⁵ , J. F. Wack³³ , S. Wada¹⁶³ , C. Wagner¹⁴⁹ , J. M. Wagner^{18a} , W. Wagner¹⁷⁷ , S. Wahdan¹⁷⁷ , H. Wahlberg⁹² , C. H. Waits¹²³ , J. Walder¹³⁷ , R. Walker¹¹¹ , K. Walkingshaw Pass⁵⁹ , W. Walkowiak¹⁴⁷ , A. Wall¹³¹ , E. J. Wallin¹⁰⁰ , T. Wamorkar^{18a} , K. Wandall-Christensen¹⁶⁹ , A. Wang⁶² , A. Z. Wang¹³⁹ , C. Wang¹⁰² , C. Wang¹¹ , H. Wang^{18a} , J. Wang^{64c} , P. Wang¹⁰³ , P. Wang⁹⁸ , R. Wang⁶¹ , R. Wang⁶ , S. M. Wang¹⁵⁴ , S. Wang¹⁴ , T. Wang¹¹⁶ , T. Wang⁶² , W. T. Wang⁸⁰ , W. Wang¹⁴ , X. Wang¹⁶⁸ , X. Wang^{144a} , X. Wang⁴⁸ , Y. Wang^{114a} , Y. Wang⁶² , Z. Wang¹⁰⁸ , Z. Wang^{144b} , C. Wanotayaroj⁸⁴ , A. Warburton¹⁰⁶ , A. L. Warnerbring¹⁴⁷ , S. Waterhouse⁹⁷ , A. T. Watson²¹ , H. Watson⁵² , M. F. Watson²¹ , E. Watton⁵⁹ , G. Watts¹⁴² , B. M. Waugh⁹⁸ , J. M. Webb⁵⁴ , C. Weber³⁰ , H. A. Weber¹⁹ , M. S. Weber²⁰ , S. M. Weber^{63a} , C. Wei⁶² , Y. Wei⁵⁴ , A. R. Weidberg¹²⁹ , E. J. Weik¹²⁰ , J. Weingarten⁴⁹ , C. Weiser⁵⁴ , C. J. Wells⁴⁸ , T. Wenaus³⁰ , T. Wengler^{37a} , N. S. Wenke¹¹² , N. Wermes²⁵ , M. Wessels^{63a} , A. M. Wharton⁹³ , A. S. White⁶¹ , A. White⁸ , M. J. White¹ , D. Whiteson¹⁶⁵ , L. Wickremasinghe¹²⁷ , W. Wiedenmann¹⁷⁶ , M. Wielers¹³⁷ , R. Wierda¹⁵⁰ , C. Wigglesworth⁴³ , H. G. Wilkens^{37a} , J. J. H. Wilkinson³³ , D. M. Williams⁴² , H. H. Williams¹³¹ , S. Williams³³ , S. Willocq¹⁰⁵ , B. J. Wilson¹⁰³ , D. J. Wilson¹⁰³ , P. J. Windischhofer⁴⁰ , F. I. Winkel³¹ , F. Winklmeier¹²⁶ , B. T. Winter⁵⁴ , M. Wittgen¹⁴⁹ , M. Wobisch⁹⁹ , T. Wojtkowski⁶⁰ , Z. Wolffs¹¹⁷ , J. Wollrath^{37a} , M. W. Wolter⁸⁷ , H. Wolters^{133a,133c} , M. C. Wong¹³⁹ , E. L. Woodward⁴² , S. D. Worm⁴⁸ , B. K. Wosiek⁸⁷ , K. W. Woźniak⁸⁷ , S. Wozniowski⁵⁵ , K. Wraight⁵⁹ , C. Wu¹⁶¹ , C. Wu²¹ , J. Wu¹⁵⁹ , M. Wu^{114b} , M. Wu¹¹⁶ , S. L. Wu¹⁷⁶ , S. Wu¹⁴ , X. Wu⁶² , Y. Wu⁶² , Z. Wu⁴ , Z. Wu^{114a} , J. Wuerzinger¹¹² , T. R. Wyatt¹⁰³ , B. M. Wynne⁵² , S. Xella⁴³ , L. Xia^{114a} , M. Xia¹⁵ , M. Xie⁶² , A. Xiong¹²⁶ , J. Xiong^{18a} , D. Xu¹⁴ , H. Xu⁶² , L. Xu⁶² , R. Xu¹³¹ , T. Xu¹⁰⁸ , Y. Xu¹⁴² , Z. Xu⁵² , R. Xue¹³² , B. Yabsley¹⁵³ , S. Yacoub^{34a} , Y. Yamaguchi⁸⁴ , E. Yamashita¹⁵⁹ , H. Yamauchi¹⁶³ , T. Yamazaki^{18a} , Y. Yamazaki⁸⁵ , S. Yan⁵⁹ , Z. Yan¹⁰⁵ , H. J. Yang^{144a,144b} , H. T. Yang⁶² , S. Yang⁶² , T. Yang^{64c} , X. Yang^{37a} , X. Yang¹⁴ , Y. Yang¹⁵⁹ , Y. Yang⁶² , W-M. Yao^{18a} , C. L. Yardley¹⁵² , J. Ye¹⁴ , S. Ye³⁰ , X. Ye⁶² , Y. Yeh⁹⁸ , I. Yeletsikh³⁹ , B. Yeo^{18b} , M. R. Yexley⁹⁸ , T. P. Yildirim¹²⁹ , K. Yorita¹⁷⁴ , C. J. S. Young^{37a} , C. Young¹⁴⁹ , N. D. Young¹²⁶ , Y. Yu⁶² , J. Yuan^{14,114c} , M. Yuan¹⁰⁸ , R. Yuan^{144a,144b} , L. Yue⁹⁸ , M. Zaazoua⁶² , B. Zabinski⁸⁷ , I. Zahir^{36a} , A. Zai^{57a,57b} , Z. K. Zak⁸⁷ , T. Zakareishvili¹⁶⁹ , S. Zambito⁵⁶ , J. A. Zamora Saa^{140d} , J. Zang¹⁵⁹ , R. Zanzottera^{71a,71b} , O. Zaplatilek¹³⁵ , C. Zeitnitz¹⁷⁷ , H. Zeng¹⁴ , J. C. Zeng¹⁶⁸ , D. T. Zenger Jr²⁷ , O. Zenin³⁸ , T. Ženiš^{29a} , S. Zenz⁹⁶ , D. Zerwas⁶⁶ , M. Zhai^{14,114c} , D. F. Zhang¹⁴⁵ , G. Zhang¹⁴ , J. Zhang^{143a} , J. Zhang⁶ , K. Zhang^{14,114c} , L. Zhang⁶² , L. Zhang^{114a} , P. Zhang^{14,114c} , R. Zhang^{114a} , S. Zhang⁹¹ , T. Zhang¹⁵⁹ , Y. Zhang¹⁴² , Y. Zhang⁹⁸ , Y. Zhang⁶² , Y. Zhang^{114a} , Z. Zhang^{143a} , Z. Zhang⁶⁶ , H. Zhao¹⁴² , T. Zhao^{143a} , Y. Zhao³⁵ , Z. Zhao⁶² , Z. Zhao⁶² , A. Zhemchugov³⁹ , J. Zheng^{114a} , K. Zheng¹⁶⁸ , X. Zheng⁶² , Z. Zheng¹⁴⁹ , D. Zhong¹⁶⁸ , B. Zhou¹⁰⁸ , H. Zhou⁷ , N. Zhou^{144a} , Y. Zhou¹⁵ , Y. Zhou^{114a} , Y. Zhou⁷ , C. G. Zhu^{143a} , J. Zhu¹⁰⁸ , X. Zhu^{144b} , Y. Zhu^{144a} , Y. Zhu⁶² , X. Zhuang¹⁴ , K. Zhukov⁶⁸ , N. I. Zimine³⁹ , J. Zinsser^{63b} , M. Ziolkowski¹⁴⁷ , L. Živković¹⁶ , A. Zoccoli^{24a,24b} , K. Zoch⁶¹ , A. Zografos^{37a} , T. G. Zorbas¹⁴⁵ , O. Zormpa⁴⁶ , L. Zwalinski^{37a}

¹ Department of Physics, University of Adelaide, Adelaide, Australia

² Department of Physics, University of Alberta, Edmonton, AB, Canada

³ (a) Department of Physics, Ankara University, Ankara, Türkiye; (b) Division of Physics, TOBB University of Economics and Technology, Ankara, Türkiye

⁴ LAPP, Université Savoie Mont Blanc, CNRS/IN2P3, Annecy, France

⁵ APC, Université Paris Cité, CNRS/IN2P3, Paris, France

⁶ High Energy Physics Division, Argonne National Laboratory, Argonne, IL, USA

⁷ Department of Physics, University of Arizona, Tucson, AZ, USA

⁸ Department of Physics, University of Texas at Arlington, Arlington, TX, USA

- ¹² Institute of Physics, Azerbaijan Academy of Sciences, Baku, Azerbaijan
- ¹³ Institut de Física d'Altes Energies (IFAE), Barcelona Institute of Science and Technology, Barcelona, Spain
- ¹⁴ Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China
- ¹⁵ Physics Department, Tsinghua University, Beijing, China
- ¹⁶ Institute of Physics, University of Belgrade, Belgrade, Serbia
- ¹⁷ Department for Physics and Technology, University of Bergen, Bergen, Norway
- ¹⁸ (a) Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA; (b) University of California, Berkeley, CA, USA
- ¹⁹ Institut für Physik, Humboldt Universität zu Berlin, Berlin, Germany
- ²⁰ Albert Einstein Center for Fundamental Physics and Laboratory for High Energy Physics, University of Bern, Bern, Switzerland
- ²¹ School of Physics and Astronomy, University of Birmingham, Birmingham, UK
- ²² (a) Department of Physics, Bogazici University, Istanbul, Türkiye; (b) Department of Physics Engineering, Gaziantep University, Gaziantep, Türkiye; (c) Department of Physics, Istanbul University, Istanbul, Türkiye
- ²³ (a) Facultad de Ciencias y Centro de Investigaciones, Universidad Antonio Nariño, Bogotá, Colombia; (b) Departamento de Física, Universidad Nacional de Colombia, Bogotá, Colombia
- ²⁴ (a) Dipartimento di Fisica e Astronomia A. Righi, Università di Bologna, Bologna, Italy; (b) INFN Sezione di Bologna, Bologna, Italy
- ²⁵ Physikalisches Institut, Universität Bonn, Bonn, Germany
- ²⁶ Department of Physics, Boston University, Boston, MA, USA
- ²⁷ Department of Physics, Brandeis University, Waltham, MA, USA
- ²⁸ (a) Transilvania University of Brasov, Brasov, Romania; (b) Horia Hulubei National Institute of Physics and Nuclear Engineering, Bucharest, Romania; (c) Department of Physics, Alexandru Ioan Cuza University of Iasi, Iasi, Romania; (d) Physics Department, National Institute for Research and Development of Isotopic and Molecular Technologies, Cluj-Napoca, Romania; (e) National University of Science and Technology Politehnica, Bucharest, Romania; (f) West University in Timisoara, Timisoara, Romania; (g) Faculty of Physics, University of Bucharest, Bucharest, Romania
- ²⁹ (a) Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava, Slovakia; (b) Department of Subnuclear Physics, Institute of Experimental Physics of the Slovak Academy of Sciences, Kosice, Slovak Republic
- ³⁰ Physics Department, Brookhaven National Laboratory, Upton, NY, USA
- ³¹ Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física, y CONICET, Instituto de Física de Buenos Aires (IFIBA), Buenos Aires, Argentina
- ³² California State University, Long Beach, CA, USA
- ³³ Cavendish Laboratory, University of Cambridge, Cambridge, UK
- ³⁴ (a) Department of Physics, University of Cape Town, Cape Town, South Africa; (b) iThemba Labs, Western Cape, South Africa; (c) Department of Mechanical Engineering Science, University of Johannesburg, Johannesburg, South Africa; (d) National Institute of Physics, University of the Philippines, Diliman, Philippines; (e) Department of Physics, Stellenbosch University, Matieland, South Africa; (f) Department of Physics, University of South Africa, Pretoria, South Africa; (g) University of Zululand, KwaDlangezwa, South Africa; (h) School of Physics, University of the Witwatersrand, Johannesburg, South Africa
- ³⁵ Department of Physics, Carleton University, Ottawa, ON, Canada
- ³⁶ (a) Faculté des Sciences Ain Chock, Université Hassan II de Casablanca, Casablanca, Morocco; (b) Faculté des Sciences, Université Ibn-Tofail, Kénitra, Morocco; (c) Faculté des Sciences Semlalia, Université Cadi Ayyad, LPHEA-Marrakech, Marrakech, Morocco; (d) LPMR, Faculté des Sciences, Université Mohamed Premier, Oujda, Morocco; (e) Faculté des sciences, Université Mohammed V, Rabat, Morocco; (f) Institute of Applied Physics, Mohammed VI Polytechnic University, Ben Guerir, Morocco
- ³⁷ (a) CERN, Geneva, Switzerland; (b) CERN Tier-0, Geneva, Switzerland
- ³⁸ Affiliated with an Institute Formerly Covered by a Cooperation Agreement with CERN, Geneva, Switzerland
- ³⁹ Affiliated with an International Laboratory Covered by a Cooperation Agreement with CERN, Geneva, Switzerland
- ⁴⁰ Enrico Fermi Institute, University of Chicago, Chicago, IL, USA
- ⁴¹ LPC, Université Clermont Auvergne, CNRS/IN2P3, Clermont-Ferrand, France
- ⁴² Nevis Laboratory, Columbia University, Irvington, NY, USA
- ⁴³ Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark

- 44 (a)Dipartimento di Fisica, Università della Calabria, Rende, Italy; (b)INFN Gruppo Collegato di Cosenza, Laboratori Nazionali di Frascati, Frascati, Italy
- 45 Physics Department, Southern Methodist University, Dallas, TX, USA
- 46 National Centre for Scientific Research “Demokritos”, Agia Paraskevi, Greece
- 47 (a)Department of Physics, Stockholm University, Stockholm, Sweden; (b)Oskar Klein Centre, Stockholm, Sweden
- 48 Deutsches Elektronen-Synchrotron DESY, Hamburg and Zeuthen, Germany
- 49 Fakultät Physik, Technische Universität Dortmund, Dortmund, Germany
- 50 Institut für Kern- und Teilchenphysik, Technische Universität Dresden, Dresden, Germany
- 51 Department of Physics, Duke University, Durham, NC, USA
- 52 SUPA-School of Physics and Astronomy, University of Edinburgh, Edinburgh, UK
- 53 INFN e Laboratori Nazionali di Frascati, Frascati, Italy
- 54 Physikalisches Institut, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany
- 55 II. Physikalisches Institut, Georg-August-Universität Göttingen, Göttingen, Germany
- 56 Département de Physique Nucléaire et Corpusculaire, Université de Genève, Geneva, Switzerland
- 57 (a)Dipartimento di Fisica, Università di Genova, Genoa, Italy; (b)INFN Sezione di Genova, Genoa, Italy
- 58 II. Physikalisches Institut, Justus-Liebig-Universität Giessen, Giessen, Germany
- 59 SUPA-School of Physics and Astronomy, University of Glasgow, Glasgow, UK
- 60 LPSC, Université Grenoble Alpes, CNRS/IN2P3, Grenoble INP, Grenoble, France
- 61 Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge, MA, USA
- 62 Department of Modern Physics and State Key Laboratory of Particle Detection and Electronics, University of Science and Technology of China, Hefei, China
- 63 (a)Kirchhoff-Institut für Physik, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany; (b)Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
- 64 (a)Department of Physics, Chinese University of Hong Kong, Shatin, N.T., Hong Kong, China; (b)Department of Physics, University of Hong Kong, Hong Kong, China; (c)Department of Physics and Institute for Advanced Study, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China
- 65 Department of Physics, National Tsing Hua University, Hsinchu, Taiwan
- 66 IJCLab, Université Paris-Saclay, CNRS/IN2P3, 91405 Orsay, France
- 67 Centro Nacional de Microelectrónica (IMB-CNM-CSIC), Barcelona, Spain
- 68 Department of Physics, Indiana University, Bloomington, IN, USA
- 69 (a)INFN Gruppo Collegato di Udine, Sezione di Trieste, Udine, Italy; (b)ICTP, Trieste, Italy; (c)Dipartimento Politecnico di Ingegneria e Architettura, Università di Udine, Udine, Italy
- 70 (a)INFN Sezione di Lecce, Lecce, Italy; (b)Dipartimento di Matematica e Fisica, Università del Salento, Lecce, Italy
- 71 (a)INFN Sezione di Milano, Milan, Italy; (b)Dipartimento di Fisica, Università di Milano, Milan, Italy
- 72 (a)INFN Sezione di Napoli, Naples, Italy; (b)Dipartimento di Fisica, Università di Napoli, Naples, Italy
- 73 (a)INFN Sezione di Pavia, Pavia, Italy; (b)Dipartimento di Fisica, Università di Pavia, Pavia, Italy
- 74 (a)INFN Sezione di Pisa, Pisa, Italy; (b)Dipartimento di Fisica E. Fermi, Università di Pisa, Pisa, Italy
- 75 (a)INFN Sezione di Roma, Rome, Italy; (b)Dipartimento di Fisica, Sapienza Università di Roma, Rome, Italy
- 76 (a)INFN Sezione di Roma Tor Vergata, Rome, Italy; (b)Dipartimento di Fisica, Università di Roma Tor Vergata, Rome, Italy
- 77 (a)INFN Sezione di Roma Tre, Rome, Italy; (b)Dipartimento di Matematica e Fisica, Università Roma Tre, Rome, Italy
- 78 (a)INFN-TIFPA, Povo, Italy; (b)Università degli Studi di Trento, Trento, Italy
- 79 Department of Astro and Particle Physics, Universität Innsbruck, Innsbruck, Austria
- 80 University of Iowa, Iowa City, IA, USA
- 81 Department of Physics and Astronomy, Iowa State University, Ames, IA, USA
- 82 Istinye University, Sariyer, Istanbul, Türkiye
- 83 (a)Departamento de Engenharia Elétrica, Universidade Federal de Juiz de Fora (UFJF), Juiz de Fora, Brazil; (b)Universidade Federal do Rio de Janeiro COPPE/EE/IF, Rio de Janeiro, Brazil; (c)Instituto de Física, Universidade de São Paulo, São Paulo, Brazil; (d)Rio de Janeiro State University, Rio de Janeiro, Brazil; (e)Federal University of Bahia, Bahia, Brazil
- 84 KEK, High Energy Accelerator Research Organization, Tsukuba, Japan
- 85 Graduate School of Science, Kobe University, Kobe, Japan

- 86 (a) AGH University of Krakow, Faculty of Physics and Applied Computer Science, Krakow, Poland; (b) Marian Smoluchowski Institute of Physics, Jagiellonian University, Krakow, Poland
- 87 Institute of Nuclear Physics Polish Academy of Sciences, Krakow, Poland
- 88 (a) Khalifa University of Science and Technology, Abu Dhabi, United Arab Emirates; (b) University of Sharjah, Sharjah, United Arab Emirates
- 89 Faculty of Science, Kyoto University, Kyoto, Japan
- 90 Research Center for Advanced Particle Physics and Department of Physics, Kyushu University, Fukuoka, Japan
- 91 L2IT, Université de Toulouse, CNRS/IN2P3, UPS, Toulouse, France
- 92 Instituto de Física La Plata, Universidad Nacional de La Plata and CONICET, La Plata, Argentina
- 93 Physics Department, Lancaster University, Lancaster, UK
- 94 Oliver Lodge Laboratory, University of Liverpool, Liverpool, UK
- 95 Department of Experimental Particle Physics, Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana, Slovenia
- 96 Department of Physics and Astronomy, Queen Mary University of London, London, UK
- 97 Department of Physics, Royal Holloway University of London, Egham, UK
- 98 Department of Physics and Astronomy, University College London, London, UK
- 99 Louisiana Tech University, Ruston, LA, USA
- 100 Fysiska institutionen, Lunds universitet, Lund, Sweden
- 101 Departamento de Física Teórica C-15 and CIAFF, Universidad Autónoma de Madrid, Madrid, Spain
- 102 Institut für Physik, Universität Mainz, Mainz, Germany
- 103 School of Physics and Astronomy, University of Manchester, Manchester, UK
- 104 CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France
- 105 Department of Physics, University of Massachusetts, Amherst, MA, USA
- 106 Department of Physics, McGill University, Montreal, QC, Canada
- 107 School of Physics, University of Melbourne, Victoria, Australia
- 108 Department of Physics, University of Michigan, Ann Arbor, MI, USA
- 109 Department of Physics and Astronomy, Michigan State University, East Lansing, MI, USA
- 110 Group of Particle Physics, University of Montreal, Montreal, QC, Canada
- 111 Fakultät für Physik, Ludwig-Maximilians-Universität München, Munich, Germany
- 112 Max-Planck-Institut für Physik (Werner-Heisenberg-Institut), Munich, Germany
- 113 Graduate School of Science and Kobayashi-Maskawa Institute, Nagoya University, Nagoya, Japan
- 114 (a) Department of Physics, Nanjing University, Nanjing, China; (b) School of Science, Shenzhen Campus of Sun Yat-sen University, Guangzhou, China; (c) University of Chinese Academy of Science (UCAS), Beijing, China
- 115 Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM, USA
- 116 Institute for Mathematics, Astrophysics and Particle Physics, Radboud University/Nikhef, Nijmegen, The Netherlands
- 117 Nikhef National Institute for Subatomic Physics and University of Amsterdam, Amsterdam, The Netherlands
- 118 Department of Physics, Northern Illinois University, DeKalb, IL, USA
- 119 (a) New York University Abu Dhabi, Abu Dhabi, United Arab Emirates; (b) United Arab Emirates University, Al Ain, United Arab Emirates
- 120 Department of Physics, New York University, New York, NY, USA
- 121 Ochanomizu University, Otsuka, Bunkyo-ku, Tokyo, Japan
- 122 Ohio State University, Columbus, OH, USA
- 123 Homer L. Dodge Department of Physics and Astronomy, University of Oklahoma, Norman, OK, USA
- 124 Department of Physics, Oklahoma State University, Stillwater, OK, USA
- 125 Palacký University, Joint Laboratory of Optics, Olomouc, Czech Republic
- 126 Institute for Fundamental Science, University of Oregon, Eugene, OR, USA
- 127 Graduate School of Science, University of Osaka, Osaka, Japan
- 128 Department of Physics, University of Oslo, Oslo, Norway
- 129 Department of Physics, Oxford University, Oxford, UK
- 130 LPNHE, Sorbonne Université, Université Paris Cité, CNRS/IN2P3, Paris, France
- 131 Department of Physics, University of Pennsylvania, Philadelphia, PA, USA
- 132 Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, PA, USA

- 133 (a) Laboratório de Instrumentação e Física Experimental de Partículas-LIP, Lisbon, Portugal; (b) Departamento de Física, Faculdade de Ciências, Universidade de Lisboa, Lisbon, Portugal; (c) Departamento de Física, Universidade de Coimbra, Coimbra, Portugal; (d) Centro de Física Nuclear da Universidade de Lisboa, Lisbon, Portugal; (e) Departamento de Física, Escola de Ciências, Universidade do Minho, Braga, Portugal; (f) Departamento de Física Teórica y del Cosmos, Universidad de Granada, Granada, Spain; (g) Departamento de Física, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal
- 134 Institute of Physics of the Czech Academy of Sciences, Prague, Czech Republic
- 135 Czech Technical University in Prague, Prague, Czech Republic
- 136 Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic
- 137 Particle Physics Department, Rutherford Appleton Laboratory, Didcot, UK
- 138 IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette, France
- 139 Santa Cruz Institute for Particle Physics, University of California Santa Cruz, Santa Cruz, CA, USA
- 140 (a) Departamento de Física, Pontificia Universidad Católica de Chile, Santiago, Chile; (b) Millennium Institute for Subatomic physics at high energy frontier (SAPHIR), Santiago, Chile; (c) Instituto de Investigación Multidisciplinario en Ciencia y Tecnología y Departamento de Física, Universidad de La Serena, La Serena, Chile; (d) Department of Physics, Universidad Andres Bello, Santiago, Chile; (e) Universidad San Sebastian, Recoleta, Chile; (f) Instituto de Alta Investigación, Universidad de Tarapacá, Arica, Chile; (g) Departamento de Física, Universidad Técnica Federico Santa María, Valparaíso, Chile
- 141 Department of Physics, Institute of Science, Tokyo, Japan
- 142 Department of Physics, University of Washington, Seattle, WA, USA
- 143 (a) Institute of Frontier and Interdisciplinary Science and Key Laboratory of Particle Physics and Particle Irradiation (MOE), Shandong University, Qingdao, China; (b) School of Physics, Zhengzhou University, Zhengzhou, China
- 144 (a) State Key Laboratory of Dark Matter Physics, School of Physics and Astronomy, Shanghai Jiao Tong University, Key Laboratory for Particle Astrophysics and Cosmology (MOE), SKLPPC, Shanghai, China; (b) State Key Laboratory of Dark Matter Physics, Tsung-Dao Lee Institute, Shanghai Jiao Tong University, Shanghai, China
- 145 Department of Physics and Astronomy, University of Sheffield, Sheffield, UK
- 146 Department of Physics, Shinshu University, Nagano, Japan
- 147 Department Physik, Universität Siegen, Siegen, Germany
- 148 Department of Physics, Simon Fraser University, Burnaby, BC, Canada
- 149 SLAC National Accelerator Laboratory, Stanford, CA, USA
- 150 Department of Physics, Royal Institute of Technology, Stockholm, Sweden
- 151 Departments of Physics and Astronomy, Stony Brook University, Stony Brook, NY, USA
- 152 Department of Physics and Astronomy, University of Sussex, Brighton, UK
- 153 School of Physics, University of Sydney, Sydney, Australia
- 154 Institute of Physics, Academia Sinica, Taipei, Taiwan
- 155 (a) E. Andronikashvili Institute of Physics, Iv. Javakhishvili Tbilisi State University, Tbilisi, Georgia; (b) High Energy Physics Institute, Tbilisi State University, Tbilisi, Georgia; (c) University of Georgia, Tbilisi, Georgia
- 156 Department of Physics, Technion, Israel Institute of Technology, Haifa, Israel
- 157 Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv, Israel
- 158 Department of Physics, Aristotle University of Thessaloniki, Thessaloniki, Greece
- 159 International Center for Elementary Particle Physics and Department of Physics, University of Tokyo, Tokyo, Japan
- 160 Graduate School of Science and Technology, Tokyo Metropolitan University, Tokyo, Japan
- 161 Department of Physics, University of Toronto, Toronto, ON, Canada
- 162 (a) TRIUMF, Vancouver, BC, Canada; (b) Department of Physics and Astronomy, York University, Toronto, ON, Canada
- 163 Division of Physics and Tomonaga Center for the History of the Universe, Faculty of Pure and Applied Sciences, University of Tsukuba, Tsukuba, Japan
- 164 Department of Physics and Astronomy, Tufts University, Medford, MA, USA
- 165 Department of Physics and Astronomy, University of California Irvine, Irvine, CA, USA
- 166 University of West Attica, Athens, Greece
- 167 Department of Physics and Astronomy, University of Uppsala, Uppsala, Sweden
- 168 Department of Physics, University of Illinois, Urbana, IL, USA
- 169 Instituto de Física Corpuscular (IFIC), Centro Mixto Universidad de Valencia - CSIC, Valencia, Spain
- 170 Department of Physics, University of British Columbia, Vancouver, BC, Canada

- ¹⁷¹ Department of Physics and Astronomy, University of Victoria, Victoria, BC, Canada
- ¹⁷² Fakultät für Physik und Astronomie, Julius-Maximilians-Universität Würzburg, Würzburg, Germany
- ¹⁷³ Department of Physics, University of Warwick, Coventry, UK
- ¹⁷⁴ Waseda University, Tokyo, Japan
- ¹⁷⁵ Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot, Israel
- ¹⁷⁶ Department of Physics, University of Wisconsin, Madison, WI, USA
- ¹⁷⁷ Fakultät für Mathematik und Naturwissenschaften, Fachgruppe Physik, Bergische Universität Wuppertal, Wuppertal, Germany
- ¹⁷⁸ Department of Physics, Yale University, New Haven, CT, USA
- ¹⁷⁹ Yerevan Physics Institute, Yerevan, Armenia
- ^a Also at Affiliated with an Institute Formerly Covered by a Cooperation Agreement with CERN, Geneva, Switzerland
- ^b Also at An-Najah National University, Nablus, Palestine
- ^c Also at Borough of Manhattan Community College, City University of New York, New York, NY, USA
- ^d Also at Center for Interdisciplinary Research and Innovation (CIRI-AUTH), Thessaloniki, Greece
- ^e Also at Centre of Physics of the Universities of Minho and Porto (CF-UM-UP), Gualtar, Portugal
- ^f Also at CERN, Geneva, Switzerland
- ^g Also at Département de Physique Nucléaire et Corpusculaire, Université de Genève, Geneva, Switzerland
- ^h Also at Departament de Física de la Universitat Autònoma de Barcelona, Barcelona, Spain
- ⁱ Also at Department of Financial and Management Engineering, University of the Aegean, Chios, Greece
- ^j Also at Department of Mathematical Sciences, University of South Africa, Johannesburg, South Africa
- ^k Also at Department of Modern Physics and State Key Laboratory of Particle Detection and Electronics, University of Science and Technology of China, Hefei, China
- ^l Also at Department of Physics, Bolu Abant İzzet Baysal University, Bolu, Türkiye
- ^m Also at Department of Physics, King's College London, London, UK
- ⁿ Also at Department of Physics, Stanford University, Stanford, CA, USA
- ^o Also at Department of Physics, Stellenbosch University, Stellenbosch, South Africa
- ^p Also at Department of Physics, University of Fribourg, Fribourg, Switzerland
- ^q Also at Department of Physics, University of Thessaly, Volos, Greece
- ^r Also at Department of Physics, Westmont College, Santa Barbara, USA
- ^s Also at Faculty of Physics, Sofia University, 'St. Kliment Ohridski', Sofia, Bulgaria
- ^t Also at Faculty of Physics, University of Bucharest, Bucharest, Romania
- ^u Also at Hellenic Open University, Patras, Greece
- ^v Also at Henan University, Kaifeng, China
- ^w Also at Imam Mohammad Ibn Saud Islamic University, Riyadh, Saudi Arabia
- ^x Also at Institutio Catalana de Recerca i Estudis Avancats, ICREA, Barcelona, Spain
- ^y Also at Institut für Experimentalphysik, Universität Hamburg, Hamburg, Germany
- ^z Also at Institute for Nuclear Research and Nuclear Energy (INRNE) of the Bulgarian Academy of Sciences, Sofia, Bulgaria
- ^{aa} Also at Institute of Applied Physics, Mohammed VI Polytechnic University, Ben Guerir, Morocco
- ^{ab} Also at Institute of Particle Physics (IPP), Toronto, Canada
- ^{ac} Also at Institute of Physics and Technology, Mongolian Academy of Sciences, Ulaanbaatar, Mongolia
- ^{ad} Also at Institute of Physics, Azerbaijan Academy of Sciences, Baku, Azerbaijan
- ^{ae} Also at Institute of Theoretical Physics, Ilia State University, Tbilisi, Georgia
- ^{af} Also at National Institute of Physics, University of the Philippines, Diliman, Philippines
- ^{ag} Also at The Collaborative Innovation Center of Quantum Matter (CICQM), Beijing, China
- ^{ah} Also at TRIUMF, Vancouver, BC, Canada
- ^{ai} Also at Università di Napoli Parthenope, Naples, Italy
- ^{aj} Also at Department of Physics, University of Colorado, Boulder, CO, USA
- ^{ak} Also at University of Siena, Siena, Italy
- ^{al} Also at Washington College, Chestertown, MD, USA
- ^{am} Also at Physics Department, Yeditepe University, Istanbul, Türkiye
- ^{an} Associated to Department of Computer Science, University of Chicago, Chicago, IL, USA

^{ao} Associated to Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, USA

^{ap} Associated to School of Computer Science, University of Bristol, Bristol, UK

*Deceased