Comput Softw Big Sci (2025) 9:2
https://doi.org/10.1007/s41781-024-00128-x

Research

Total Cost of Ownership and Evaluation of Google Cloud Resources for the ATLAS Experiment at the LHC

The ATLAS Collaboration*

CERN, Geneva, Switzerland

Received: 24 May 2024 / Accepted: 27 September 2024

Abstract The ATLAS Google Project was established as part of an ongoing evaluation of the use of commercial clouds by the ATLAS Collaboration, in anticipation of the potential future adoption of such resources by WLCG grid sites to fulfil or complement their computing pledges. Seamless integration of Google cloud resources into the worldwide ATLAS distributed computing infrastructure was achieved at large scale and for an extended period of time, and hence cloud resources are shown to be an effective mechanism to provide additional, flexible computing capacity to ATLAS. For the first time a total cost of ownership analysis has been performed, to identify the dominant cost drivers and explore effective mechanisms for cost control. Network usage significantly impacts the costs of certain ATLAS workflows, underscoring the importance of implementing such mechanisms. Resource bursting has been successfully demonstrated, whilst exposing the true cost of this type of activity. A follow-up to the project is underway to investigate methods for improving the integration of cloud resources in data-intensive distributed computing environments and reducing costs related to network connectivity, which represents the primary expense when extensively utilising cloud resources.

1 Introduction

The ATLAS experiment [1] at the LHC [2] employs distributed computing resources of up to one million cores of computing, 350 PB of disk and over 500 PB of tape storage. These resources comprise the Tier-0 at CERN, the Tier-1 and Tier-2 Worldwide LHC Computing Grid (WLCG) sites [3,4], opportunistic resources at High Performance Computing (HPC) sites and cloud computing providers, as well as volunteer computing. In recent years such non-WLCG resources, particularly from HPC sites, have made increasingly significant contributions to ATLAS computing, a trend that is expected to continue in the future.

The use of commercial clouds has been investigated by all LHC experiments [5–8], and was recently reviewed by the US ATLAS and CMS communities as part of a blueprint publication [9]. ATLAS has examined the viability of both Google and Amazon cloud resources [10], including the integration of the two core components of the distributed computing environment: the workload management system PanDA [11] and the data management system Rucio [12]. These initial studies showed that such resources can be adopted for both the ATLAS production and user analysis workflows, albeit at a limited capacity, as well as demonstrating the potential for future R&D [10].

As the start of Run 4 and the HL-LHC [13] era approaches, the need for additional CPU and especially disk resources continues to grow, to satisfy the computing requirements of the experiments [14,15]. At the same time, WLCG sites are increasingly exploring methods to utilise new resources, aiming to fulfil their pledges in the most cost-effective and energy-efficient manner while also making emerging technologies accessible to users. With this in mind, the ATLAS Google Project (AGP) was established to continue an evaluation of commercial clouds as a resource for ATLAS.

Whilst basic cost estimates of commercial cloud resources have been done before [7,10], one of the primary goals of the AGP is to develop a Total Cost of Ownership (TCO) model for commercial clouds, which is made possible by the highly granular pricing information available from the cloud provider. By comparison, other types of computing resources employed by ATLAS, namely the grid, HPCs, clouds, and volunteer computing, incur different effective costs to the experiment, which are often difficult to evaluate and compare since funding methodologies vary by resource type, by country, and even between different funding agencies within the same country. Additionally, there are local arrangements at the sites or savings, for example where some sites, typically at universities, do not directly pay electricity or WAN access costs. It is therefore important to note that this TCO

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

Published online: 28 January 2025

evaluation is not attempting to directly compare the cost of running ATLAS jobs in the Google Cloud with the cost of running an ATLAS grid site. A TCO comparison of data centre colocation and commercial cloud solutions has been previously performed in the context of expanding computing resources at CERN [16], using input from the Helix Nebula Science Cloud project [17].

The structure of this paper is the following. "Commercial Cloud Cost Modelling" section delves into the critical elements of commercial cloud cost models relevant to this study, examining the subscription agreement model employed by ATLAS and outlining the methodology and scope of the TCO analysis. "ATLAS Google Site Integration" section briefly outlines the technical aspects of integrating the Google site and "Running the ATLAS Google Site" section offers a detailed account of the various tests conducted at different phases and the resulting conclusions. "Cloud and Network Costs" section more closely examines the role of network traffic in overall costs and underscores the rationale for considering dedicated links. "Feedback from Grid Site Administrators" section presents feedback received from a number of system administrators regarding the adoption of commercial clouds. Finally, "Further Investigations and Future Work" section provides an outline of potential future working directions and a summary is presented in "Summary" section.

2 Commercial Cloud Cost Modelling

In contrast to other computing resources employed by ATLAS, including those deployed via the WLCG grid, commercial cloud resources have well-defined pricing for computing time and services used, for example at Amazon Web Services (AWS) [18] or on the Google Cloud Platform (GCP) [19]. While each cloud may differ in their specific pricing structure, the cost of the individual services used by the client is usually published in an itemised and highly granular way, making it easier to develop future cost models for commercial clouds. In this way, a client can choose from a menu of services, which are charged or billed directly.

While the *list price-based* cost model is suitable for ad hoc or specialised short-term resource needs, it is likely to prove more costly for consistent, long-term usage of resources compared to grid-based, general-purpose offline computing. However, most commercial cloud providers offer discounts and credits for large-scale users, which has been previously explored by ATLAS with varying degrees of success. Such discounts and credits can significantly reduce the cost compared to the list-price. Funding agencies can often negotiate even better deals. The list price-based cost model therefore provides a maximum ceiling for a TCO, but rarely reflects the actual price paid for large-scale cloud services.

The primary obstacle to using commercial clouds is usually egress costs, which are network costs incurred when data is exported from the cloud resource, and may be prohibitively high for the distributed computing workflows used by ATLAS. These workflows often exploit the high interconnectivity of the grid sites and incorporate many data transfers, as data hosted by one site is transferred to be used elsewhere. This high network interconnectivity may incur significant expense, which goes on top of the site budgets, as discussed in "Dedicated Networks" section. Egress costs associated with cloud resources were a concern expressed by several Tier-1 site administrators when interviewed for this report (see "Feedback from Grid Site Administrators" section).

Whilst discounts, credits and subscription plans may increase the complexity in determining future costs, they also provide mechanisms to reduce the total cost of using commercial clouds for ATLAS. It should however still be noted that even with deep discounts and credits, continuous monitoring of expenditure must be done, and as automated as possible, especially for large-scale use.

2.1 Subscription Model Employed by the ATLAS Google Project

For this latest phase of the AGP, a *Google Cloud Service Agreement for Public Sector* contract was negotiated with Google to explore more solid ideas about employing Google as an ATLAS production site. This contract, which uses the latest pricing model, was negotiated via US national laboratories and is based on an initial assumption of an average of 7000 compute cores together with up to 7 PB of storage, and an estimated egress of no more than 0.7 PB per month. As is often the case with commercial cloud resources, there is no charge for data ingress, that is for uploading ATLAS data to the Google Cloud Platform. The contract ran for 15 months, from July 2022 to October 2023, at a flat rate cost of $56,630.54 per month, which resulted in a subscription price of around $1900 per day.

The subscription model does not put a limit on usage once active, and ATLAS can use the available resources at any time at any scale. For example, for 1 month the disk usage could be 10 PB, and the next month 1 PB, or for 1 month ATLAS may use 3000 cores, and the next month use 20,000 cores. This resource elasticity may be very useful for short-duration data-processing campaigns and is explored in dedicated tests.

The TCO for the subscription model is the total price of the contract for 15 months, $849,458, although there is a clear caveat to be considered, namely that the negotiation for any subsequent contract will likely examine the actual average usage during the previous contract. If a significantly higher usage is observed than the initial estimate, it may be the case that the monthly price may be higher in any subsequent contract. A key part of the TCO evaluation is therefore to use the detailed usage breakdown provided by the GCP accounting, without applying the subscription model, to understand which site configurations produce the most significant increase to the total cost according to the list-price, and which particular services drive this increase.

2.2 Total Cost of Ownership Methodology

As previously mentioned, this TCO evaluation is not trying to directly compare the cost of running an ATLAS grid site with the costs associated with running ATLAS jobs in the Google Cloud. This is primarily for two reasons, as described above. Firstly, a significant cost variation exists already among different resources, and among different grid sites, which is at least in part attributable to the variation in service costs, as well as regional variation in labour costs for site administrators. Secondly, relying solely on absolute numbers for Google Cloud costs may not provide meaningful insights, due to the substantial influence of individual agreements and contracts, which include volume discounts that vary case by case and over time.

The focus of this TCO is therefore rather on gaining a comprehensive understanding of how to make the best use of such a resource, including how the configuration may differ from a standard grid site. To accomplish this, the relative contributions of various components within the cloud service to the TCO are analysed, including compute, storage, and network, across different operating models. By identifying the dominant cost drivers and exploring effective cost control mechanisms, it is possible to optimise resource allocation and management to maximise cost efficiency.

In addition to hardware considerations, the invaluable contribution of dedicated personnel at the grid sites should not be overlooked. These individuals not only fulfil the role of system administrators but also play a vital part in maintaining the middleware and the distributed computing infrastructure of the experiment, and in some cases contribute to federated support services, including user-support, and areas such as R&D, outreach and education. Their expertise and involvement are crucial to sustain a high efficiency and effectiveness in the operation of the ATLAS computing infrastructure, more so if sites opt for using cloud computing resources at scale.

3 ATLAS Google Site Integration

The ATLAS data are distributed worldwide across data centres or sites, organised into Tiers with varying capacities and responsibilities under the umbrella of the WLCG [3,4]. The single Tier-0 centre is CERN, and there are ten Tier-1 sites, connected via dedicated National Research and Education Networks (NRENs) typically with between 10 and 100 Gbps, using the LHCOPN and LHCONE overlays [20]. The Tier-1 sites provide both disk and tape storage and function as the perpetual archival of the collision data. Around 50 Tier-2 sites, typically hosted by national universities and laboratories, provide disk storage that is used for data processing and user analysis.

When setting up an ATLAS site using GCP resources the objective is to fully support all ATLAS production workflows. This includes not only processing the RAW collision data into reconstructed data (Analysis Object Data, AOD), but also running all components of the multi-step Monte Carlo (MC) production workflow, such as Event Generation (EVNT), Full or Fast Simulation which produces simulated detector interaction data (HITS), simulated detector output data (Raw Data Object, RDO) and reconstructed simulated data (AOD), as well as creating data and MC derived formats (DAOD) for input to analysis (a process referred to as "Group Production"). Additional formats such as Derived Event Summary Data (DESD), which are tailored to the specific needs of various subdetector and object reconstruction and identification performance groups, may also be produced in data-processing campaigns. The site is also expected to handle user workloads, which can encompass a diverse range of demands. Further details on ATLAS production workflows and data formats can be found elsewhere [21].

In the ATLAS grid site setup, production tasks running ATLAS jobs at various locations consolidate their outputs at a designated site known as the nucleus. The nucleus is chosen during task definition and can be either a Tier-1 site or a substantial and reliable Tier-2 site. Considering its size and expected performance, this responsibility could also be anticipated from the ATLAS Google site and is also examined here.

The integration of compute at cloud-based sites with ATLAS distributed computing relies on Kubernetes [22], where the resource-facing component of PanDA, Harvester [23], utilises the native job controller of the Kubernetes clusters for submitting batch jobs to the PanDA queue associated with the site [24,25]. ATLAS software is provided in the same way as at grid sites via CVMFS [26].

The access to Google storage was configured as a standard *Rucio Storage Element* (RSE), as used by other WLCG storage systems abstracted via standard HTTP/WebDAV using authorisation tokens based on the S3v4 format [27]. At the lowest layer in the stack, the Davix [28] library implements HTTP/WebDAV access to all storage systems used in the WLCG, and already supports chunked transport required by object stores, which are used by ATLAS for cloud-based sites [29]. Further significant development was then required to make access to Google storage work, as described in the following.

The Rucio server was extended to allow only specific accounts to access cloud storage and to generate cloud storage tokens ad hoc when listing replicas, a functionality that is needed both by interactive command line interface users as well as production and analysis jobs. It was also necessary to extend the functionality of the Rucio clients to allow seamless upload/download for object stores, as object stores prohibit several useful functions from the HTTP/WebDAV standard that Rucio uses to ensure safe uploads such as checksum verification. The CERN File Transfer Service (FTS) [30] was also extended to dynamically generate authorisation tokens based on the S3v4 format when enacting a transfer.

Making the Google storage available within the WLCG infrastructure also required a creative authentication solution, as Google is not part of the Interoperable Global Trust Federation (IGTF) [31]. To do this, a dedicated load-balancer was set up on the Google Cloud side with a fake hostname in the CERN DNS server such that a CERN-based X.509 host certificate could be issued and uploaded. This load-balancer was configured based on path-based regular expressions, which allowed two RSEs to be deployed: a DATADISK to store production input and output data and a SCRATCHDISK to temporarily store analysis job outputs. With this setup, the Google storage could be globally integrated into the ATLAS distributed computing infrastructure like any other WLCG site, and could be used by any account with the appropriate permissions.

Due to the proximity to CERN and the low carbon-intensive energy usage in the region, the Google *europe-west1* region in Belgium was chosen to host the primary ATLAS Google site, although during investigations into the network connectivity a second site in the US Google *us-east4* region in Virginia was also employed (see "Cloud and Network Costs" section).

The number of running single-core slots ("job slots") at the ATLAS Google site can be manually configured, and is typically set to either 5000 or 10,000, although additional set-ups such as those used for the evaluation of larger scale "bursts" were also deployed, as described in "Resource Bursting" section. CPU cores are provided as "spot instances" [32], so that allocated resources may be preempted at any time depending on current situation at the cloud, but have the advantage of costing significantly less in real terms. The storage limit at the ATLAS Google site is initially configured to be between 2 and 5 PB.

Thanks to the development of support for cloud-native interfaces, the deployment of an ATLAS site in the Google Cloud capable of scaling to tens of thousands of CPUs and several petabytes of disk could be accomplished within a matter of weeks. Moreover, the operation and maintenance of the site has required only a fraction of a full-time equivalent (FTE) staff with expertise in cloud administration. This achievement paves the way for exploring cost–benefit scenarios in which the agility and scalability of cloud computing can be harnessed to accelerate the ATLAS science programme.

4 Running the ATLAS Google Site

4.1 Initial Phase

The ATLAS Google site was initially configured with a PanDA queue able to run up to 5000 CPU slots, together with a single RSE where data files could be stored. The PanDA queue could be configured to accept certain types of jobs and the initial goal was to try to increase the number of different workloads and eventually test all the job workflows at the site. Figure 1 displays various metrics of the ATLAS Google site during the first 6 months of running.

The number of running jobs is shown in Fig. 1a. Aside from alternating the number of CPU slots between five and ten thousand, the configuration of the site was essentially unchanged during this period. The number of different types of job running in the PanDA queue was progressively increased by changing the brokerage decisions to adjust the job mix of the PanDA queue. The thin spikes that can be seen in the number of running jobs are due to short-term site configuration changes.

In the 3 months from August 24th until November 24th 2022, the overall job mix that was run corresponded to approximately 30% MC Event Generation, 30% MC Full Simulation, 30% MC Reconstruction and 10% Group Production, which is a typical job mix seen on a standard grid site. Analysis workloads were not tested in this period as many changes needed to be made to the ATLAS middleware.

As jobs continued to run at the site, the accumulated data from the various production steps steadily increased at the Google RSE at a rate of approximately 50 TB per day, until it reached 6 PB on November 24th as can be seen in Fig. 1b. At the same time, the availability of these data generated an increasing number of accesses from other ATLAS sites. The egress network traffic from the ATLAS Google site due to production and analysis jobs elsewhere using data at Google as input ramped up from an average of about 20 TB per day in August to about 130 TB per day in October and November. In November, there were periods with egress network traffic over 200 TB per day for several days in a row, as can be seen in Fig. 1c. It is worth noting that this traffic is above the average seen at ATLAS sites, relative to the stored data volume. The ATLAS Google site was generating egress traffic at the level of 4 PB per month in November whilst hosting 5–6 PB data, which is significantly more than the initial estimate of 0.7 PB. By comparison, the MWT2 grid site in the US also generates an average of 4 PB per month of egress traffic, whilst hosting more than 15 PB of data.

Fig. 1d shows the breakdown of the Google list-price cost per service, where the various contributions from compute, storage and egress are presented. It can clearly be seen that the costs of egress traffic and storage can quickly become dominant in cloud resources. To contain these costs, on Novem-

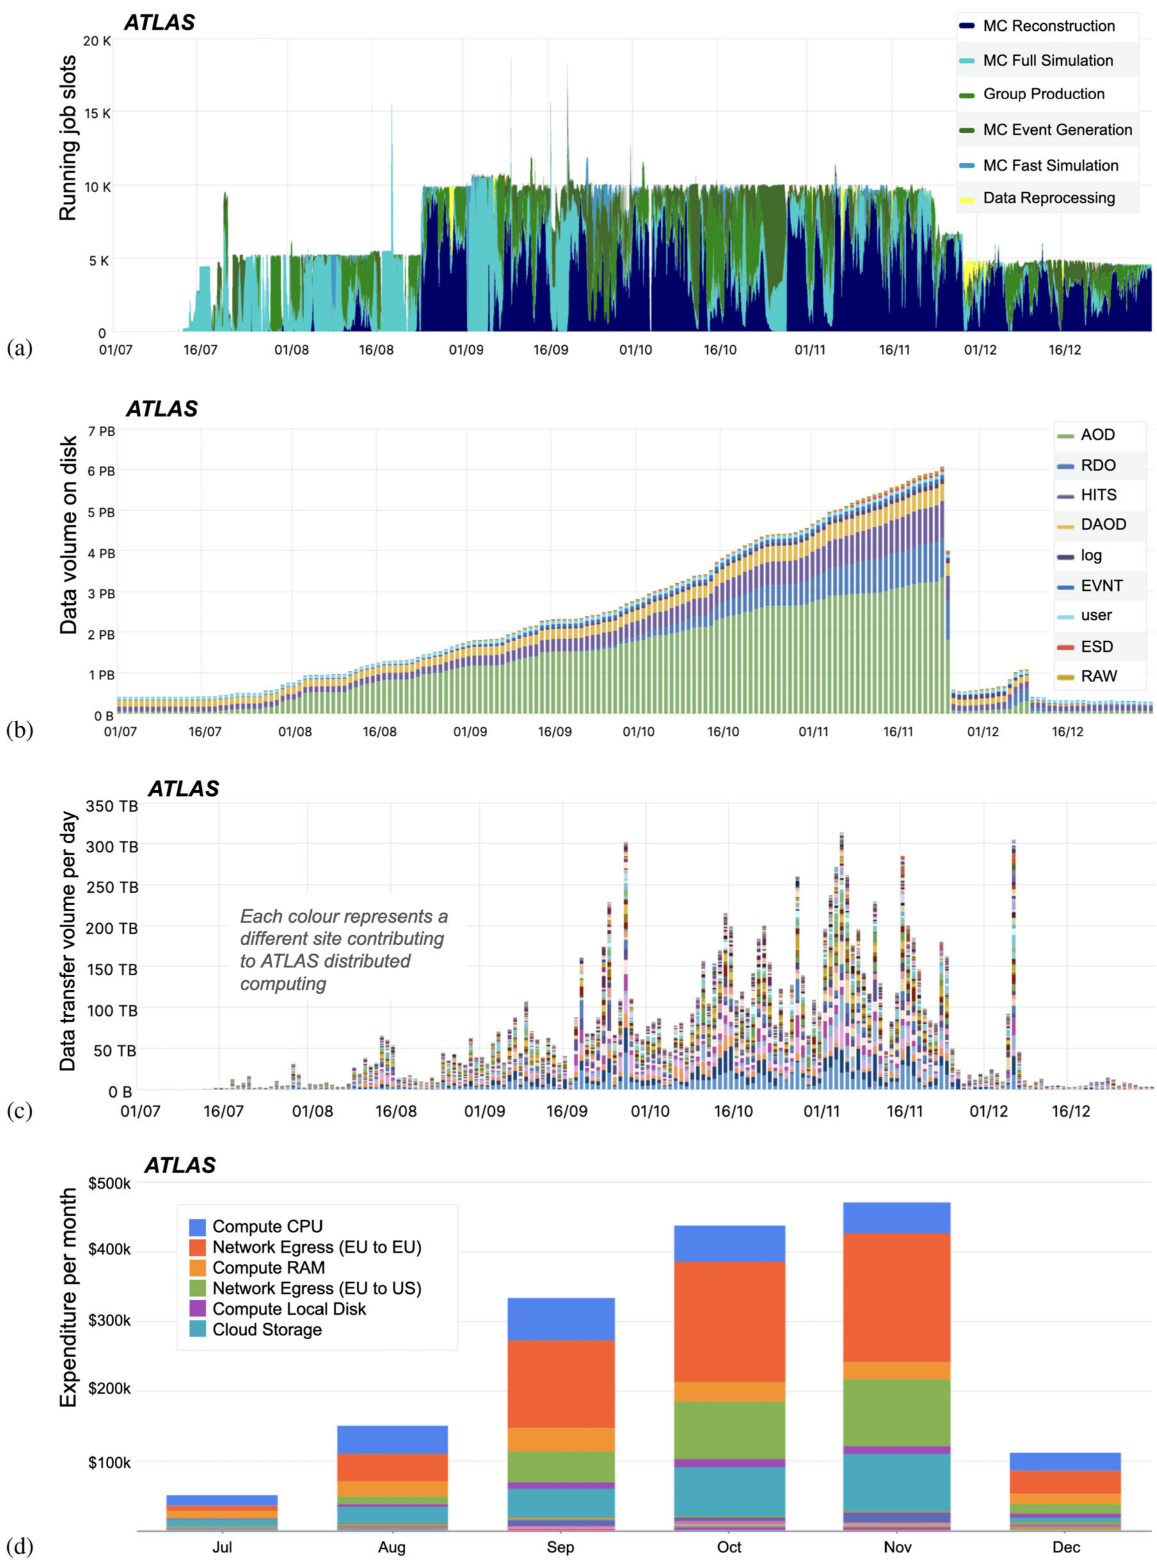

Fig. 1 Monitoring plots for the first 6 months of running at the ATLAS Google site, from July to December 2022. **a** The number of running jobs at the Google site. **b** The accumulated data at the Google RSE split into different formats, the main ones being AOD (green), RDO (blue), HITS (purple) and DAOD (yellow). **c** The daily egress traffic out of the ATLAS Google site, split into the various destination sites. **d** The monthly list-price cost per service from the Google billing console, where the six main components are shown in the legend

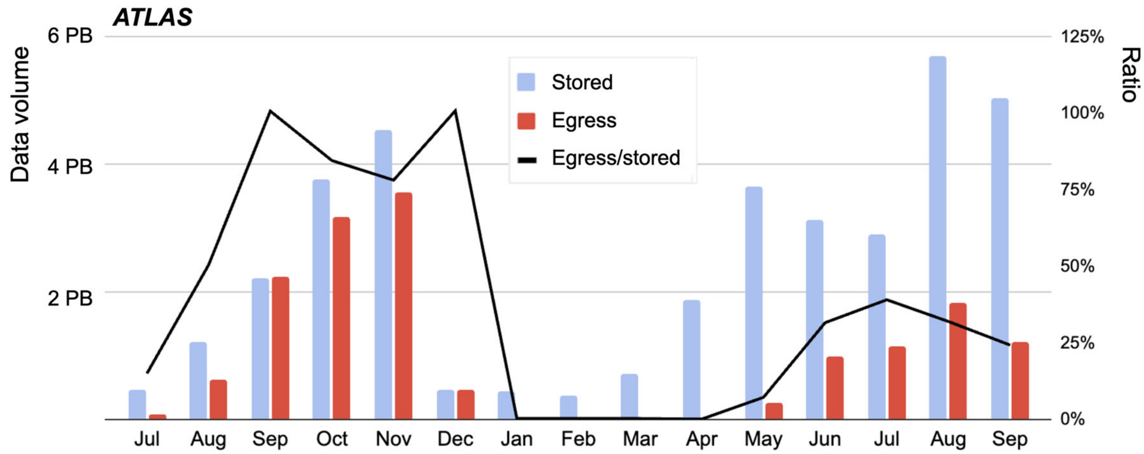

Fig. 2 Data stored (blue) and egressed for job inputs (red) at the ATLAS Google site per month from July 2022 to September 2023. The ratio is also indicated by the black line

ber 25th the configuration of the ATLAS Google site RSE was changed in Rucio by enabling greedy deletion, so that data were deleted as soon as the corresponding replication rules had expired [33]. This had the immediate effect that the 5.5 PB of accumulated cached data were quickly removed, and from then on new temporary data would stay on the RSE for only 1 or 2 weeks. This drastically reduced the egress traffic, since before the change there was significant egress as the cached data were being transferred multiple times to other sites as job inputs.

After this change, the remaining egress was mainly due to job output being sent to the task nucleus. Another change was then applied to the Google RSE on December 8th that set the distance [34] of the RSE to any other ATLAS site to a very large value. This had the effect of completely eliminating the remaining egress traffic due to job inputs then preferentially being read from other sites. The number of CPU slots at the ATLAS Google site was also reduced back to 5000. The effect of these changes can be seen in the distributions in Fig. 1.

With this new configuration, the stored data at the Google RSE stayed around 300 TB and egress traffic fell below 5–10 TB per day. The solitary spike in transfer volume in December visible in Fig. 1c is from the egress of the outputs of a small data reprocessing running at the ATLAS Google site, visible as the yellow contribution in Fig. 1a.

4.2 First Observations on Network Considerations

Cloud resources have well-defined price structures for each resource type, not only compute and storage, but also network traffic. The first two are the usual capacity metrics that are closely accounted for in a distributed infrastructure like the WLCG. The network capacity is certainly also considered within WLCG, but the planning and provisioning cycle is usually done on a longer time scale. This is in part because the administration domain for networks is typically broader and often spans national or continental institutions beyond the sites themselves. This has probably had an effect in the experiment computing models, whose workflows rely to some extent on a plentiful any-to-any connectivity.

The long-term large-scale test of the ATLAS Google site has provided useful insights into the management of the associated network traffic. Given that the site was newly deployed, it was possible to monitor how the consumption of different types of resources evolved. CPU usage was constant as expected, as a fixed configuration parameter, whereas storage usage increased at an essentially constant rate. The steady increase in egress network traffic was correlated with increased use of storage, albeit with a much larger variability. An interesting outcome of this first period of the ATLAS Google site was that it was possible to reduce the egress network traffic by adjusting a few parameters in Rucio, allowing the cost of cloud resources to be very effectively controlled.

A key question remains about how useful a grid site is with limited egress network traffic. Whilst this question is beyond the scope of this report, it is an important topic that deserves dedicated studies in the future in the context of computing model evolution. The approach taken here is to quantify the tests that were done with the ATLAS Google site. As previously described, after the first 3–4 months of continuous unattended operations, the site had accumulated 6 PB of data on disk. The values for average stored data and total egress of data for job inputs are depicted in Fig. 2. Remarkably, between September and December 2022, between 75% and 100% of all the data at the site was egressed each month for production or analysis job inputs. Comparing this with other large ATLAS sites, typically having storage sizes ranging from 4 to 30 PB, this metric falls to between 15% and 20%. Consequently, the significant amount of new data at the

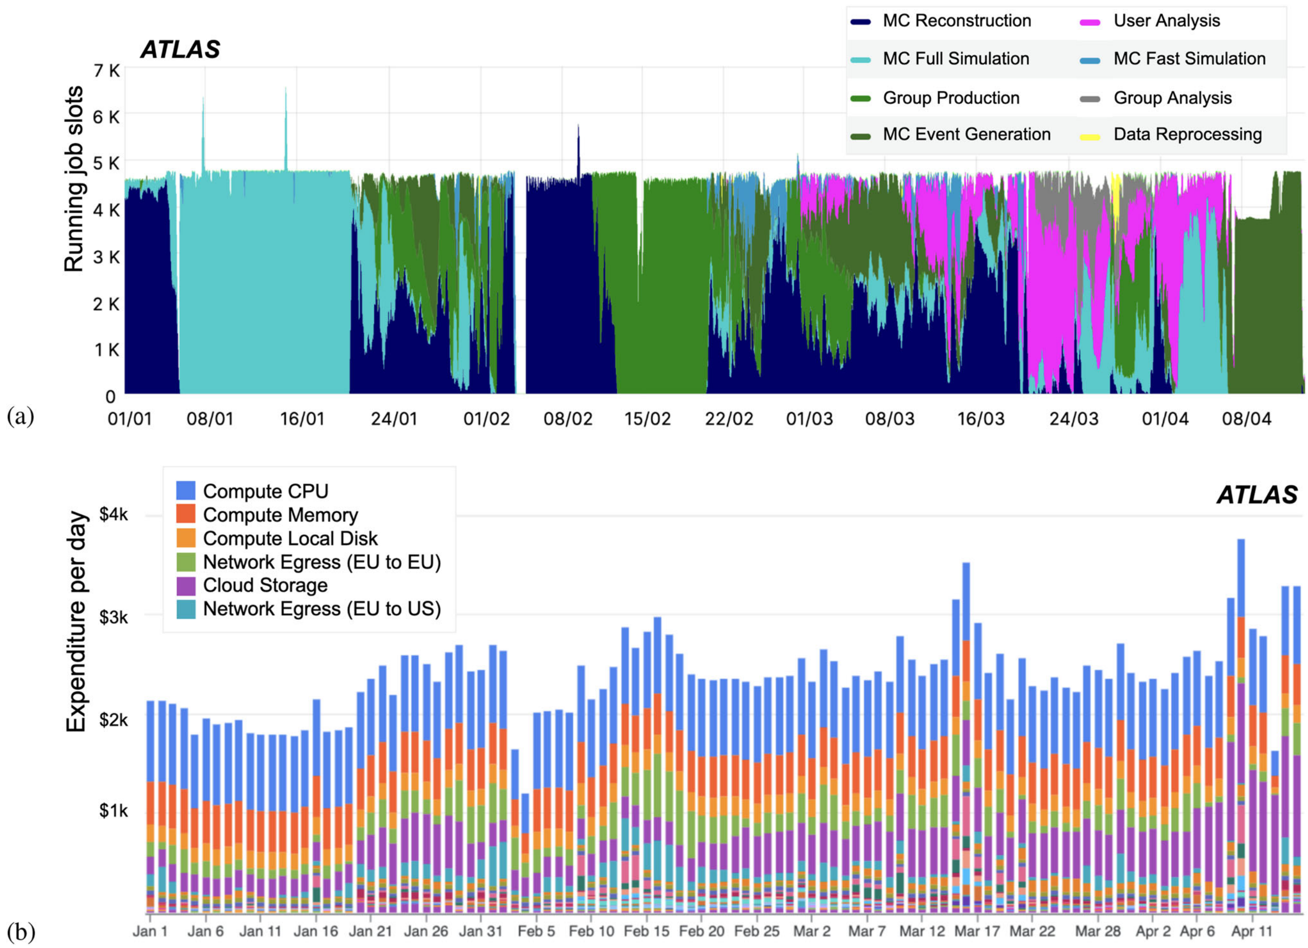

Fig. 3 **a** The variation of workflows running at the ATLAS Google site from January to April 2023 featuring several periods of running with a single workflow. The contribution from user analysis jobs can be seen from March. **b** The daily list-price cost per service from the Google billing console for the period from January to April. The dominant services are compute CPU/RAM (blue/red), local storage on the worker nodes (orange), cloud storage (purple) and network egress (green and turquoise)

ATLAS Google site generated data movement dynamics that significantly deviated from the average. This is relevant when evaluating a cloud resource, since abnormally high egress traffic levels will have an impact on costs.

4.3 Understanding the Cost Impact of Different ATLAS Workflows

Between January and April 2023 the ATLAS Google site ran at an approximately constant CPU capacity of around 5000 job slots. The types of jobs allowed to run were controlled through the fairshare policy parameter associated with the PanDA queue. By adjusting this parameter, four periods of running only one activity were carried out, as shown in Fig. 3a: a 14-day period with only MC Simulation jobs in January, a 5-day period with only MC Reconstruction jobs in early February, a 9-day period with only Group Production jobs in mid-February and a 6-day period of only MC Event Generation jobs at the beginning of April.

The Google billing console provides daily list-price cost information, shown in Fig. 3b, which can be used to infer general trends. It was observed that 80–90% of the cost is consistently dominated by three services: compute, storage and network egress. The remainder comes from a combination of infrastructure overhead and costs associated with orthogonal ATLAS R&D activities at Google. The compute cost (which has three components associated with it: CPU, RAM and local disk) remains pretty much constant throughout this period, as expected from the fact that the number of slots in the ATLAS Google site was set to 5000 job slots and not modified. This cost is consistent with the advertised pricing for the spot *n2-standard-8* instances of around \$0.01 per CPU-h. This situation can be compared to the first months of the project, where egress costs dominated, as shown in Fig. 1d.

Compute is the dominant cost for these four workflows. Storage and network costs for a given day are partially but not fully correlated with the activity that is running at that time. If data are not purged from the storage, the storage costs will reflect the accumulation of any past activity. On the other hand, if there are old data stored at the site that are accessed from outside for any reason, this will generate egress costs that are not correlated with the activity running at that precise moment. Still, the single workflow testing periods were of relatively short duration and the Rucio storage configuration meant that no background egress was present, so any correlation observed should be meaningful.

A final, significant single workflow test was done in July 2023, by running a fraction of the reprocessing of the 2022 data on 10,000 job slots at the ATLAS Google site. About 15% of the proton–proton collision data recorded in 2022 for physics analysis, comprising around 600 M events, were processed into multiple formats for analysis and further, secondary processing between July 11th and 18th. The data carousel mechanism [35] was employed as done for the grid, whereby input files, in this case RAW data, were recalled from ATLAS tape resources and replicated to the Google site, without observing any adverse effects.

Fig. 4a shows the data reprocessing jobs during the period July 11th to July 18th, as well as a small number of associated output data merging jobs. Further data reprocessing jobs can be seen in Fig. 4a after July 18th, which are part of the wider campaign to reprocess the remaining 85% of the 2022 data on all ATLAS distributed computing sites, including the ATLAS Google site. Figure 4b shows the daily data volume transferred out of the ATLAS Google site to the various grid sites, where it can be seen that during the data reprocessing a total of about 100 TB of data were exported daily.

Fig. 4c also shows the daily data volume transferred, but now broken down into the different activities. Production Output (cyan) accounts for approximately half of the egress, which is equivalent to the export to CERN shown as the purple component in Fig. 4b. Most of the remaining egress during this period is attributable to Data Consolidation, which is a rebalancing procedure performed on the ATLAS distributed computing system as a whole. It may be interpreted as other data moved out of the Google storage to make room for the output of the data reprocessing. The egress from the data reprocessing jobs after July 18th is also visible in Fig. 4b and c, albeit at a lower level.

The periods with different types of activities show different trends on the relative cost of the three main services, as shown in Table 1. The average values for the full duration of the project are also displayed. For the first four columns, a few trends are visible in the numbers in the table. In the Group Production period, the egress network activity and its associated costs show a clear increase, averaging 21% over the 9-day testing period. In the 6-day MC Event Generation period, the storage is seen to contribute a significant fraction of the total cost at 30%; however this is from the replication of almost 1 PB of DAOD data sets to the Google site done during the same week. The most striking observation is that the period of Data Reprocessing activity shows a very different pattern compared to the others, with the network egress cost clearly dominating and averaging 63% of the total cost. The relative costs of different workflows are further examined in "Cloud and Network Costs" section when discussing networks.

User analysis workflows have also been running on the ATLAS Google site since March 2023, as can be seen in Fig. 3a. Whilst it would have been desirable to run only user analysis at the ATLAS Google site for some period, this proved difficult, primarily due to the unpredictable nature of such workflows, compared to the rather standard production workflows employed in the single job type periods described in this section. Despite replicating several popular analysis datasets to the site, it was difficult to get enough user jobs in the queue, and ultimately it was not possible to run only analysis workflows on the site, although this was almost achieved in the second half of March.

4.4 Resource Bursting

Cloud computing is intrinsically highly elastic in its nature, and offers the opportunity to acquire a significant number of additional resources, potentially at short notice. In the context of Active Learning [36] this is particularly advantageous, when it is essential to increase the speed of each iteration of MC sample production. Previous studies [25] have shown that it is possible to quickly ramp up many tens of thousands of job slots at Google and process a small number of MC events through all steps in the MC production chain to arrive at the DAOD used for analysis.

Bursting a large amount of additional compute capacity may also be useful if for example a particular MC sample is urgently required, and this scenario was explored at the ATLAS Google site in June 2023. In this case a 50 M event standard top-quark pair production MC sample was chosen to undergo Full Simulation as quickly as possible, by bursting to 100,000 job slots. The task was configured as standard 2000 event 8-core jobs, which each take on average between 6–8 h, so that all 50 M events should be processed within 24 h. The input EVNT data was replicated to the Google storage and the site was drained of all other running jobs before starting the test. Whilst this was not strictly required, it allowed the burst of resources to be isolated, which was useful for monitoring purposes.

The results of the burst test are shown in Fig. 5. The ramp up to 100,000 running job slots was achieved in 1–2 h without issue, as can be seen in Fig. 5a. Wall-clock consumption is shown in Fig. 5b, which revealed a considerable amount of

Table 1 Relative cost of each of the main services during periods of time when only one workflow was running at the ATLAS Google Site: MC Full Simulation, MC Reconstruction, Group Production, MC Event Generation and Data Reprocessing. The final column shows the average values for the full duration of the project, from July 2022 to September 2023

	MC Full Simulation 06/01–19/01	MC Reconstruction 05/02–09/02	Group Production 11/02–19/02	MC Event Generation 07/04–12/04	Data Reprocessing 12/07–16/07	Full Project 07/22–09/23
Compute	73%	65%	52%	47%	21%	28%
Storage	10%	10%	9%	32%	11%	20%
Network egress	7%	11%	26%	3%	63%	46%
Other	10%	14%	14%	19%	4%	6%

lost wall-time in the ramp up phase. This was due to nodes accepting jobs while CVMFS was still being initialised, and hence the burst test was repeated with a slower ramp up profile. In both cases, the overall lost wall-time was 11–13%, somewhat more than observed on the grid, and coming from the ramp up phase and the low and constant level of preemptions throughout both tests. Nevertheless, in both cases all 50 M events were processed within 24 h as expected.

The same MC Full Simulation sample has been processed several times on the various resources currently employed by ATLAS, without draining queues or using a dedicated site or queue. Each of these tasks took between 8 and 10 days to process, even when a significant fraction of the work, between 30% and 75%, was executed on a few powerful sites such as the ATLAS High Level Trigger Farm [37,38] when used in Sim@P1 [39,40] configuration or the Vega [41] and NERSC-Perlmutter [42] HPCs. In this sense, the burst test can be considered a success, whilst at the same time having the advantage of exposing the cost of a well-defined data-processing activity. The list-price cost of each burst run at the ATLAS Google site can be seen in Fig. 5c, where the sum of the compute-based components is around $23,000 each time.

4.5 Scaling Up the Site Again

For the final few weeks of the current project, the size of the ATLAS Google site was scaled up again to between the size of an ATLAS Tier-1 and Tier-2 grid site, with around 5 PB of storage and 10,000 running jobs slots, utilised by all job types. This was done in July 2023, ahead of launching the data reprocessing single workflow period described in "Understanding the Cost Impact of Different ATLAS Workflows" section, and this configuration remained until the end of the project in September. The 15 largest Tier-2 grid sites and about half of the Tier-1 sites typically each provide at least this number of job slots to ATLAS computing. Almost all of the Tier-1 sites and the ten largest Tier-2 sites have at least 5 PB allocated to their DATADISKs, so at the end of July the size of the ATLAS Google Site DATADISK was increased from 2.5 PB to 5 PB. At this time, the site was also reconfigured as a nucleus, so that not only unique data could be stored there but the site could also act as the output destination for a task running jobs at all other ATLAS grid sites. Furthermore, the large distances to other RSEs set in an earlier phase of the project were somewhat reduced, in particular to the German Tier-1 and associated Tier-2 s, which are geographically close to the ATLAS Google site in Belgium. Figure 6 shows the changes to the data stored and the transfers out during the month following these site reconfigurations.

Fig. 6a shows the data stored at the ATLAS Google site, grouped into three different replica types: Persistent, which has a Rucio rule with no lifetime (typically data placed at the site); Temporary, which has a Rucio rule with a lifetime (typically data replicated to the site via PanDA for production jobs, with a lifetime of 2 weeks) and Cached, which is data with no current Rucio rule and therefore may be deleted at any time. The main consequence of the site changes was, as expected, an increase in the Cached data as a result of the increased space available for output data from production tasks running at the site. The volume of Persistent and Temporary data remains roughly constant.

The variety of data types stored at the ATLAS Google site during this period can be seen in Fig. 6b, where a marked increase in RDO files is visible. These data are used as input to MC Reconstruction tasks in combination with HITS, which is the output of MC Simulation. It is likely that the nucleus nature of the site, in combination with the reduced distances, means that more HITS datasets remain on the RSE and are available as favourable input for MC Reconstruction tasks, both at the ATLAS Google site and elsewhere. A steady increase in HITS and AOD (the output of MC Reconstruction) is also observed as well as the presence of some RAW data files used as input to data reprocessing tasks.

Fig. 6c shows the daily transfers out of the ATLAS Google site, which show a steady increase after the end of July. Most of the egress is due to files replicated from the site to use as production input elsewhere, for example AODs to be used as input for the production of analysis-level data (DAOD)

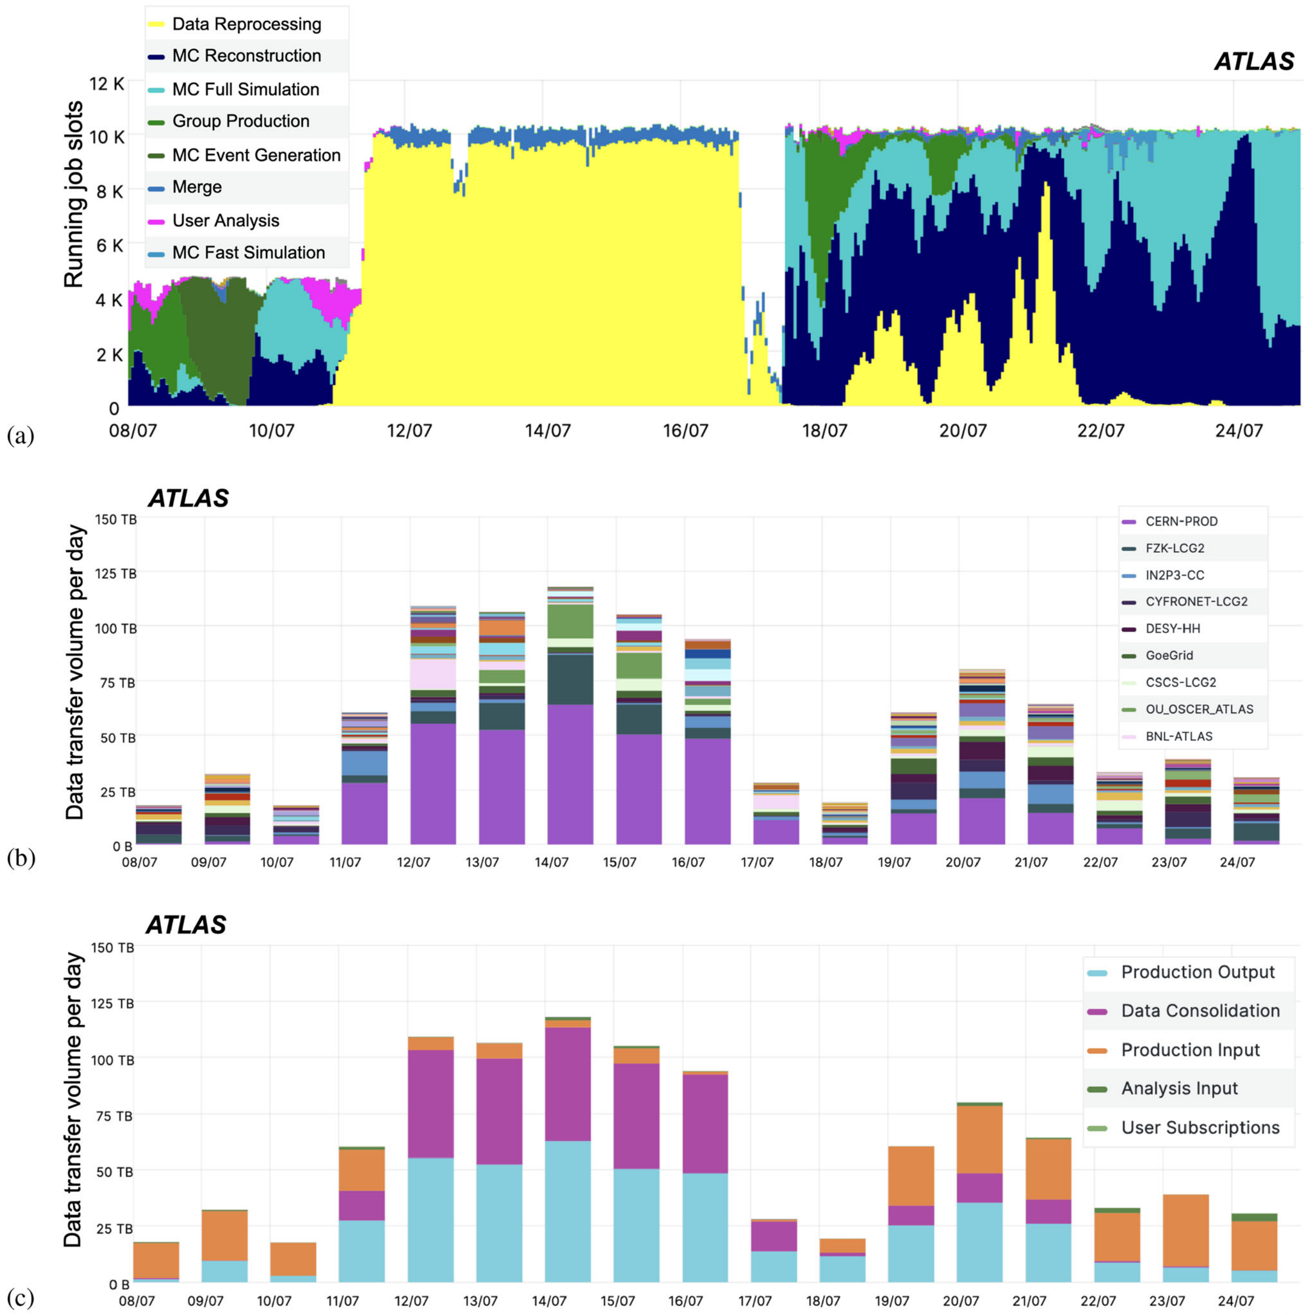

Fig. 4 Distributions covering the data reprocessing campaign performed on the Google ATLAS site. **a** Number of running jobs, where the single job type period from July 11th to July 18th shows the data jobs in yellow, together with a small number of associated merge jobs in blue. **b** The data transfers out of the ATLAS Google site for the same period to different grid sites, where the main contribution in dark purple is the replication of the reprocessing output data to CERN. This is also visible for the data reprocessing jobs after July 18th. **c** The different types of transfers out of the ATLAS Google site for the same period

at another ATLAS grid site. The level of egress is however significantly less than in 2022, which at its peak had a rate of more than 1 PB per week (see Fig. 1c). This reduction is likely due to two main reasons. Firstly, compared to 2022, the short distances set in Rucio between RSEs were limited to only the sites in the German cloud. Secondly, there was a higher fraction of new data at the Google RSE when it was first activated. At the beginning of the project, the RSE was empty and all data were newly written there, whereas in July 2023 there was already around 2 PB stored there before the site was scaled up again.

Decommissioning of the ATLAS Google site took place in September 2023, so that all resources employed during the project were effectively switched off from the holistic per-

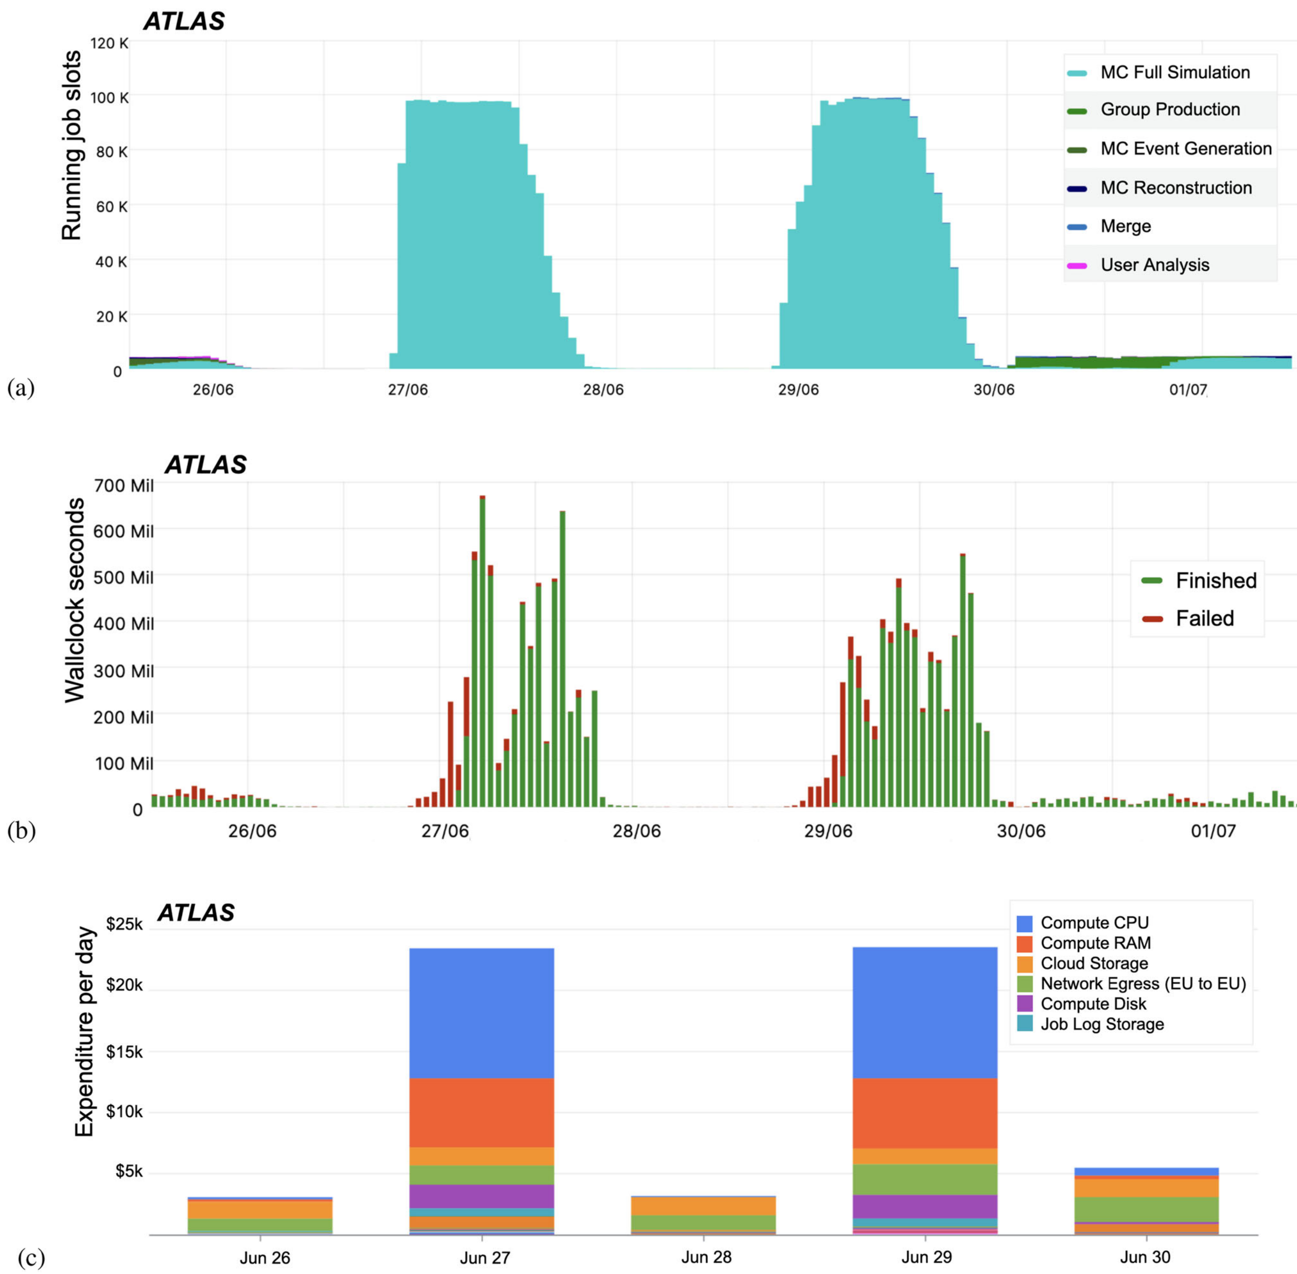

Fig. 5 Distributions covering the resource burst tests done at the ATLAS Google site in June 2023. **a** The running jobs of the two bursts of MC Full Simulation. **b** The wall-clock consumption of the jobs running on the Google site. **c** The daily list-price cost per service from the Google billing console, where the compute contributions are seen to dominate on the burst days

spective of ATLAS distributed computing. The PanDA queue was disabled mid-September and all unique data moved to one of the ATLAS Tier-1 sites. Any remaining user data from R&D projects were removed and the site was fully decommissioned by September 21st. No significant difficulties were encountered during this process, which was similar to the usual decommissioning of a typical ATLAS grid site.

5 Cloud and Network Costs

Deploying grid resources in commercial clouds creates a demand for networking services that can have an impact on performance and potentially also on cost. Egress traffic is a particularly expensive resource in the cloud, due in part to the commercial strategy of providers to incentivise to continue to use their resources and not migrate to other cloud providers. At the time of writing, the list-price for storing data in the

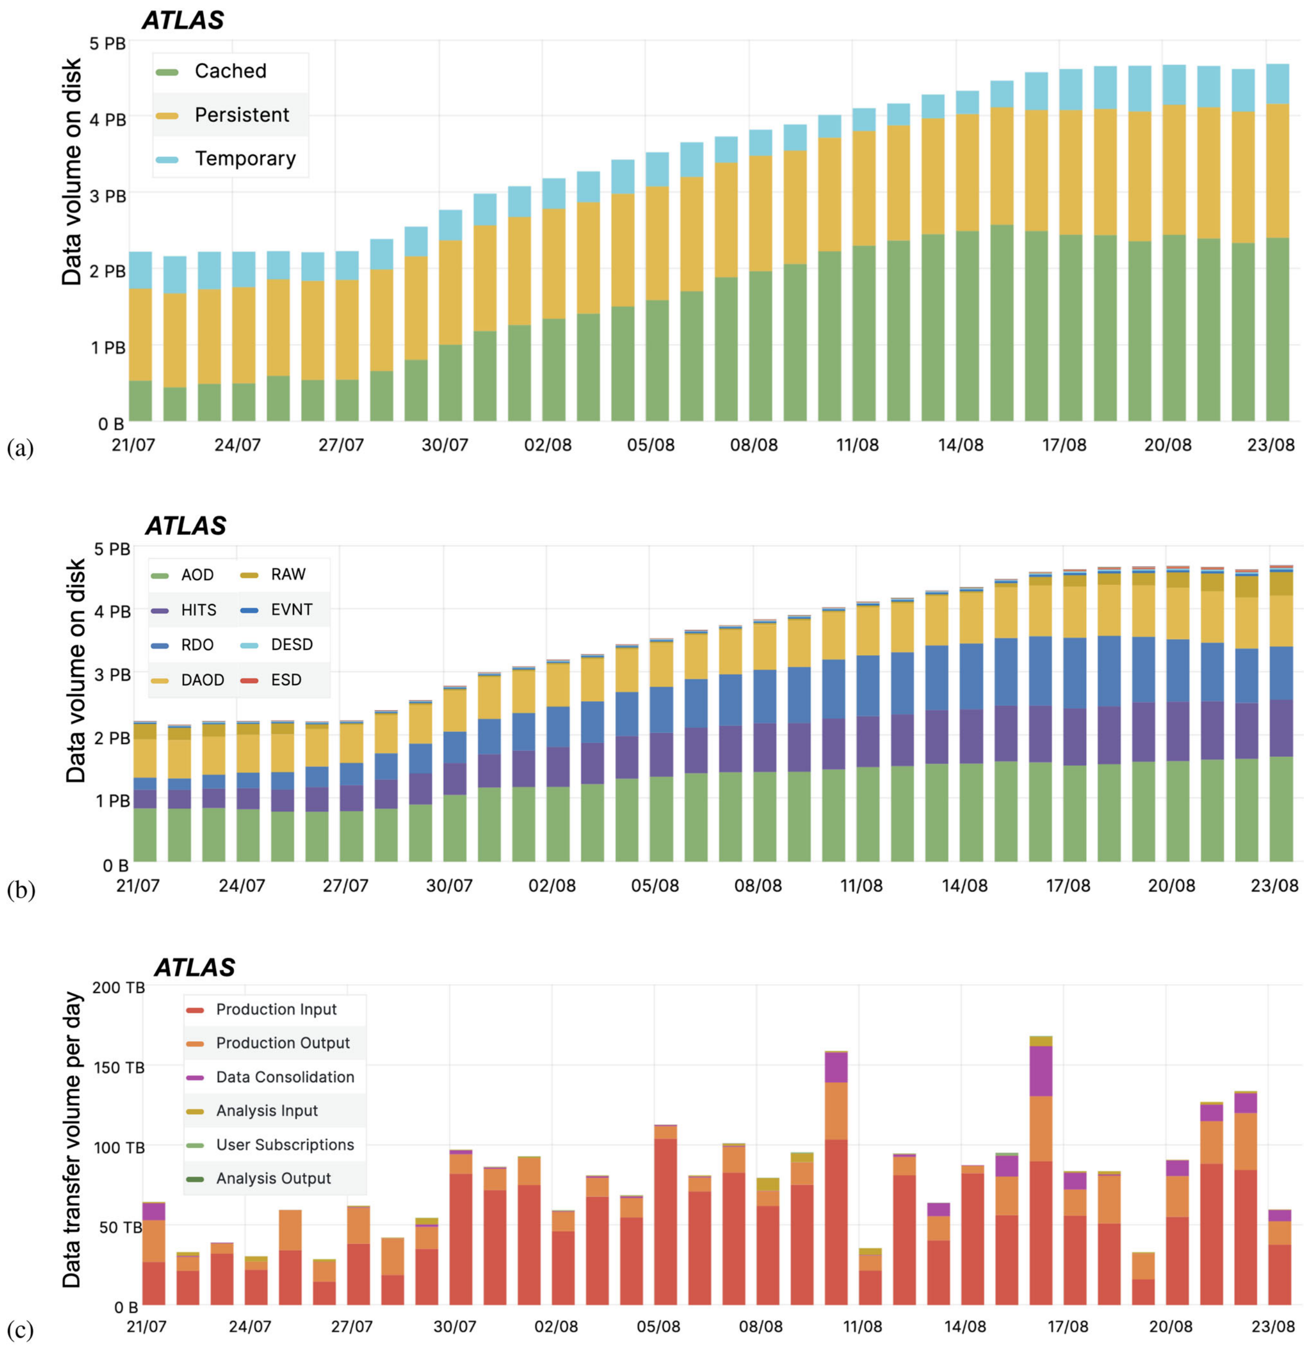

Fig. 6 Data stored at the ATLAS Google site: **a** grouped into different replica types and **b** grouped into different ATLAS data formats. **c** The different types of daily transfers out of the ATLAS Google site

Google *europe-west1* region is $20 per TB per month [43], while the price for egressing data is between $45 and $85 per TB, depending on the volume [44].

As previously described, probably the most important feature of the cloud resources is that the costs involved are heavily dependent on which services are used. This dependence can be seen in the cost breakdown of running the ATLAS Google site, which varied significantly month to month, or day to day, depending on the activity of the site. Figure 7a shows the monthly list-price cost profile for the full 15-month duration of the project. The cost of compute is basically stable, only showing small variations at the times when the number of running job slots at the sites was changed (for example when the site was increased from five to ten thousand job slots from August to November 2022, and the CPU burst test in

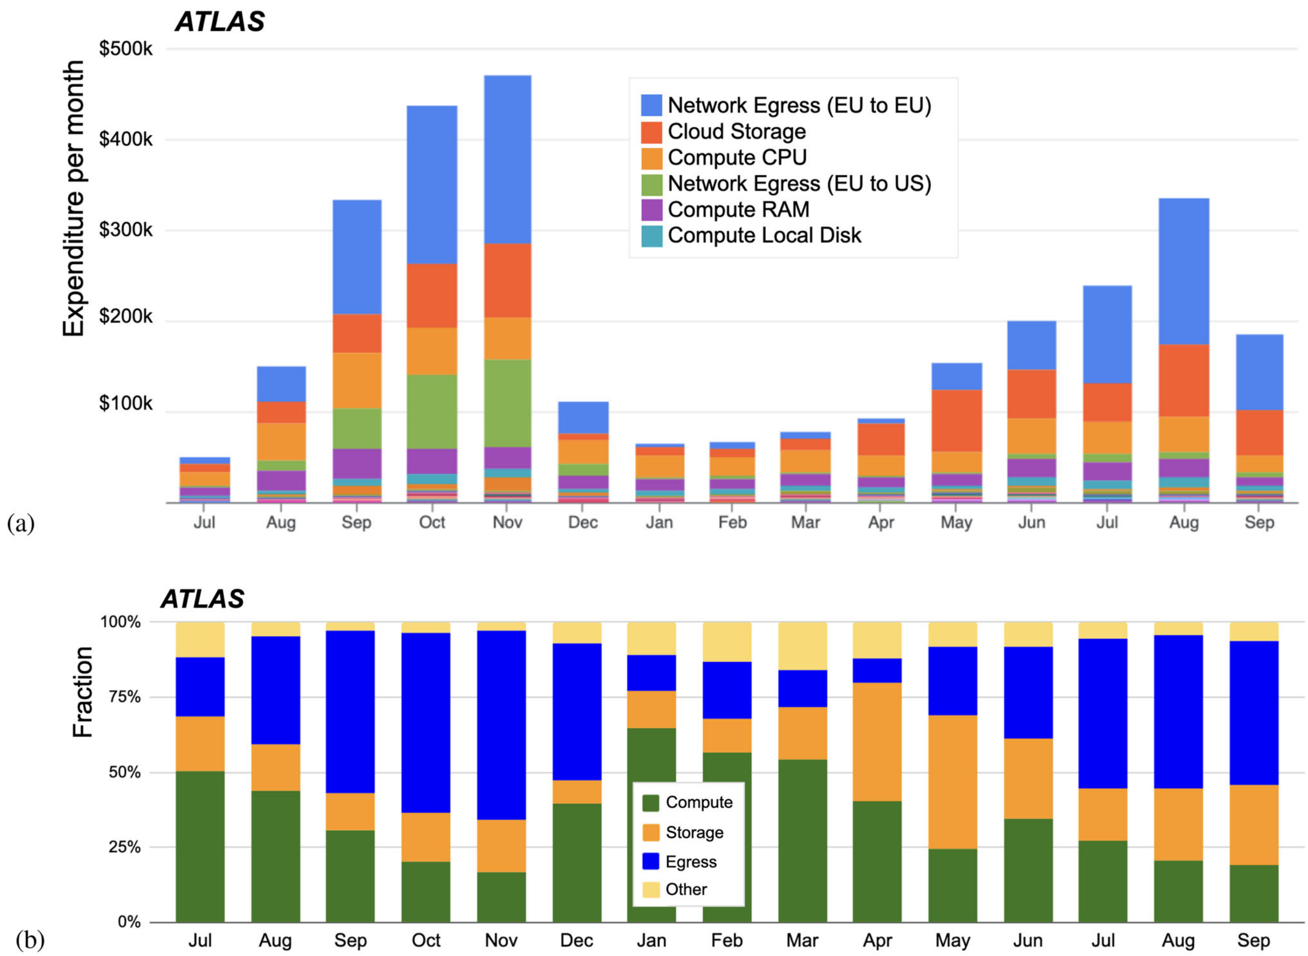

Fig. 7 Cost breakdown for the ATLAS Google site for the period July 2022 to September 2023. **a** The monthly list-price cost per service from the Google billing console. **b** The percentage contribution to the monthly cost from each of the services grouped as indicated in the legend

June 2023). However, the cost of storage and network egress varies considerably, depending on the dominant activity. Fig. 7b shows the relative fraction of each service to the total monthly cost. For the first months of the project, until November 2022, the egress cost increased as new data accumulated at the site and jobs running at other grid sites accessed these data. By November 2022, the costs associated with egress reached 54% of the monthly total. Other patterns can be observed in 2023, such as in April and May when the analysis input data was replicated to the ATLAS Google site, correspondingly increasing the fraction of the total cost spent on storage. The increase in egress due to the data reprocessing performed at Google is also visible in July.

5.1 Dedicated Networks

Egress network traffic is the resource that can have the largest impact on the running costs over short timescales. It is an expensive resource, and its use can increase very rapidly if some data at the site become popular and are suddenly accessed by thousands of jobs at other sites. Moreover, the data going from the ATLAS Google site to other ATLAS grid sites do so through the general purpose internet, as opposed to using the LHCONE/LHCOPN private networks that link most of the ATLAS sites. This has the effect of generating potentially very large traffic over the internet link into the destination sites, which can have cost implications at the destination, since some sites have higher costs or lower available bandwidth associated with their general purpose internet links compared to the dedicated LHCONE links. The utilisation of such links can lead to operational disruptions, potentially impacting a site's availability to its users. This is often due to lower provisioning of general internet bandwidth by the sites, primarily stemming from the associated higher costs, resulting in rapid saturation. It is therefore important to investigate mechanisms that could control and potentially reduce the egress costs. One such mechanism is dedicated network links with the cloud providers.

The price for sending traffic through a dedicated link [45] has two components: a fixed one of the order of $2.4 per hour for a 10 Gbps circuit for instance, plus a variable one that depends on traffic, priced at $20 per TB. Additional costs might also arise, depending on the specific service provider and route of the intermediate connection. According to this, routing the traffic exiting the ATLAS Google site through a dedicated link could potentially reduce the egress costs to less than half if traffic above 3 PB per month is generated. A recent study within the IceCube Collaboration [46] came to a similar conclusion, where by employing dedicated network links the egress costs for data-intensive applications could be reduced by between 50% and 75%. It is however worth emphasising that the scope of the study undertaken by IceCube significantly differs from that of ATLAS. IceCube utilised dedicated links connecting the cloud to a specific site at UW-Madison, whereas ATLAS is currently investigating the feasibility of deploying dedicated links to establish connections between cloud resources and the LHCONE overlay network, facilitating data transfers to numerous sites worldwide. It is therefore important to recognise that the complexity and potentially the cost associated with these two approaches may differ considerably.

To explore this option, an engagement has been started with ESnet [47] to provision a dedicated 10 Gbps link to the Google *us-east4* cloud through their ESnet Cloud Connect service. The aim of this exercise is to test two things: first, to measure and confirm a reduction of the egress costs for large data transfers, and second, to try to route the egress traffic from the ATLAS Google site into LHCONE to avoid downstream costs associated to high-volume traffic through the regular internet links at receiving sites.

The expectation is that dedicated network links will provide a way to lower the networking costs, but at the same time they will also add complexity to the deployment and operations of the cloud resources. Moreover, besides the technical complexity of provisioning the dedicated network links, there is also an organisational complexity that arises from the fact that the implementation of this setup will vary depending on the combination of cloud provider (and even region inside a provider) and NREN that provides the peering. In different countries, NRENs will have different capabilities and conditions for peering with cloud providers. Furthermore, each cloud provider likely offers their own specific tools to deploy and manage the dedicated links, each with very different provisioning procedures. Another consideration would be if multiple experiments, for example ATLAS and CMS, were to provision resources from the same cloud provider whether they could both use the same dedicated network link. There is clearly a significant programme of work in this area, beyond the time frame of the current AGP.

6 Feedback from Grid Site Administrators

As part of this study, feedback was collected from several system administrators representing various ATLAS Tier-1 grid sites. In particular, their experiences with commercial clouds were discussed, to gain some insights into the primary advantages and drawbacks associated with these services.

6.1 Concerns Related to Cloud Computing and Comparisons to this Project

It was a common view among site administrators that cloud is more expensive than on-premises solutions. In particular, concerns were raised about high egress costs, which can significantly impact the overall expenses. Some of these comparisons may have been made against dedicated compute instances rather than spot instances, which have a higher cost. Past experiences [7] with spot instances revealed eviction rates of up to 15%, which administrators considered unacceptably high. This issue becomes particularly critical since many sites support multiple users beyond LHC activities, and these users might be less tolerant to evictions than the LHC experiments, affecting the overall service reliability. During this project preemption was observed to be significantly lower, where only around 20% of all failed jobs were due to evictions. With an overall rate of 5% lost wall-clock from failed jobs over the duration of the project, the resulting eviction rate of between 1–2% compares favourably to the previous result described above. However, it is important to note that eviction rates for spot instances can vary based on several factors, including the time of the year, the cloud provider, the geographical region, the volume, and the types of resources utilised.

Worries were also expressed about unpredictable performance variations over time in cloud environments. Administrators were concerned that cloud providers might change the underlying hardware behind a specific instance type, leading to potential performance fluctuations that could affect the quality of services. The PanDA queue associated with the ATLAS Google site was very stable, experiencing negligible downtime throughout the duration of the project. While some cloud providers do not specify the exact provided CPU models, within CPU families it is possible to define a preference for the newer generations. The masked CPU values (publishing clock frequency and cache size) collected by the workload management system were homogeneous throughout the year, indicating a stable performance. The adoption of the new, more flexible HEPScore [48] benchmarking model should provide further information about this topic in the near future.

Site administrators expressed a general concern about the risks associated with vendor lock-in. They emphasised the importance of maintaining flexibility and the ability to migrate between different cloud providers to avoid being tied to a single vendor and potentially facing challenges with cost escalation, integration or data portability. The solutions implemented by this project to interface Google Cloud resources are essentially cloud-agnostic, not only avoiding vendor lock-in, but also potentially enabling access to other commercial cloud resources, as demonstrated at a lower scale at AWS [18].

Data ownership and control were significant worries. Another concern was about having critical data solely stored in the cloud, without having direct control over the physical infrastructure where the data is hosted. Ensuring digital sovereignty is seen as crucial when handling data from unique scientific experiments. The Google Cloud data access policy is clearly and strictly defined [49], and user data is fully protected and accessible only to the customer cloud administrators, designated users, and contract managers. Ultimately, ATLAS retains data (and algorithm) ownership when running in the Google Cloud. In addition, an assurance on privacy is made that customer data will not be used for any commercial use or for training purposes. A further, related point of concern is that as cloud resources are essentially leased, there is no opportunity to ensure that the hardware is responsibly and sustainably deployed to its maximum capability and longevity.

According to the feedback gathered from one site administrator, the cost breakdown of operating a grid site generally comprises approximately one-third for personnel, one-third for operational expenses, and one-third for hardware investments. The technical personnel effort required to operate one of the ATLAS Tier-1 centres is estimated to be around 10 FTEs. Some site administrators hold the view that even with a substantial migration of resources to the cloud, the essential operational effort needed to run a site would not experience a significant reduction. This is because only a limited number of hardware-oriented technical positions may no longer be required. However, some issues arise that make this comparison inherently challenging. The experience of the Google site primarily focuses on the operational effort required to provide CPU and storage at scale for a single experiment. In contrast, the roles and responsibilities of technical personnel at on-premises sites may be much broader in scope.

Finally, site administrators emphasised that funding agencies are making substantial investments in building new, energy-efficient data centres, which does not indicate a trend or incentive to increase the utilisation of cloud resources. In this context, Google cloud resources operate with net-zero operational greenhouse gas emissions by neutralising any remainder via investing in carbon offsets. Google data centres have above average Power Usage Effectiveness [50] and provide dynamic information about the carbon intensity of their regions, enabling users to steer their load to minimise emissions. In particular, the site utilised by this project, *europe-west1*, has one of the lowest grid carbon intensities [51] among Google data centres.

6.2 Purchasing Procedures

Public institutions, including those operating grid sites, often encounter the requirement to procure services through public tendering processes. However, when it comes to large-scale purchasing of cloud services, the administrative demands involved can be quite difficult. In response to this challenge, the OCRE [52] project was initiated in 2019 with the aim of streamlining the procurement process for cloud services in Europe.

For organisations contracting cloud services within Europe today, utilising a framework like OCRE becomes a viable option to do their purchasing. OCRE facilitates this purchasing process through NRENs, which in turn maintain lists of country-specific authorised resellers that offer cloud services covered under the OCRE framework agreements. One of the key benefits provided by OCRE is the provisioning of a standardised contract. While this contract serves as a baseline, further negotiations at the country or site level are possible. These negotiations could lead to larger volume discounts, ultimately benefiting the organisations involved.

Whilst OCRE is available for institutions within Europe, it is not applicable in countries other than the 40 members of the framework. The process for contracting cloud services in the USA, for example via HEPCloud [53] or CloudBank [54], or in Asia may be completely different and may result in different pricing conditions compared to those in Europe. It is crucial to recognise that the landscape for cloud service procurement can vary between regions, necessitating tailored approaches based on the specific regulatory and contractual requirements of each region. If multiple sites were to buy into the same cloud provider, some of these differences could be overcome, although remaining administrative hurdles would need to be clarified.

For international organisations like ATLAS seeking to implement a coherent strategy for utilising cloud services, the varying regulations and procurement processes across regions add complexity to the management and planning efforts. Flexibility and adaptability become essential for navigating the diverse cloud landscapes and effectively leveraging cloud resources across different geographical areas.

7 Further Investigations and Future Work

There are several areas that would benefit from an extension of the ATLAS Google Project, to allow further exploration of employing commercial cloud resources for ATLAS.

Firstly, this could involve further investigation of the ATLAS Google Site, evaluating the typical grid site configuration but at a lower level of resources than deployed at the end of this project cycle. As discussed in "Scaling Up the Site Again" section, whilst the egress is more under control it is nevertheless still significant, so a more detailed evaluation of which workflows are suitable for commercial cloud should be done, including any necessary, additional configuration changes. Understanding how the subscription agreement structure works and what a reasonable discount looks like is also an important factor when considering workflow restrictions, site structure adjustments or long-term contract costs. The impact of these studies could also be expanded to potentially evolve the ATLAS workflow and data management systems to take into account the cost of the network.

Secondly, peering with commercial cloud networks would be the main focus for further investigation, as this multi-faceted, time-intensive activity is only just beginning now. The network costs incurred due to transfers via the internet are a critical roadblock to widespread commercial cloud storage adoption for scientific computing overall and in the case of the LHC experiments, transfers need to be routed through the LHCONE overlay to eliminate these costs for the WLCG sites. At the same time, data scheduling should be introduced to reduce the data volumes that incur transfer and storage costs. Widespread peering is necessary to reduce the need to route via complicated paths. There are several NRENs in the US and EU, which should be approached to discuss the technical details of network peering, for example ESnet [55], Internet2 [56], or GÉANT Network [57].

There are various options in the Google Network stack to support this activity, but at the same time the necessary Google documentation does not seem to be publicly available. Detailed discussions would be needed between WLCG and CERN IT network experts together with the corresponding network experts from Google. If necessary, short-term Google premium support could be financed once the peering options with the NRENs have been explored.

There is the potential need to develop new features in the Rucio [12], GFAL [58], Davix [28], and FTS [30] stack to support this R&D and ATLAS would continue to work with the respective teams to discuss the objectives and milestones, and to follow the implementation, deployment, and operations. There are two distinct analyses and respective evolutions of the ATLAS computing model that could be explored. One idea is to improve the necessary workflow and data management policies when using cloud storage. This would involve for example only storing data in the cloud that is currently in use, which is rather different to the current grid storage model. In this way, the complexity of a heterogeneous infrastructure setup reduces the egress volume, but still allows bursting to cloud compute with large data inputs. Alternatively, the necessary capabilities to exploit cloud storage and network features such as bucket-level copy could be developed, to facilitate internal transfers between different cloud regions. This would remove the need to egress the data via FTS, which incurs the usual associated transfer costs.

Thirdly, beyond the ATLAS Google Site configuration and network considerations, many areas of R&D were performed as part of the AGP, using multiple services offered within the Google Cloud [59]. Most of these projects have taken advantage of the elastic availability of special types of resources to ramp up and down ephemeral compute clusters using for example GPUs, large amounts of memory, or ARM CPUs depending on the current need. This was done either through the ATLAS PanDA workflow management system or using interactive compute with Jupyter [60] notebooks and Dask [61] task scheduling.

The usage of such non-standard resources that are not easily available at standard WLCG grid sites has proven to be extremely valuable and effective, helping to develop and expedite new ATLAS data analysis techniques using machine learning and the migration of the ATLAS software to ARM CPUs. Compact data formats with columnar data access have been investigated using Google resources.

ATLAS plans to continue to take advantage of non-standard resources like GPUs and ARM CPUs to work on innovative and novel analysis and software techniques, accelerating the process to discoveries. Another larger focus area could be the development of high-energy physics algorithms using tools not available before, such as Large Language Models and Generative Artificial Intelligence. A continuation of the AGP will facilitate these initiatives, providing proven access to such resources. Like the other activities proposed in this section, they are not expected to be resource or cost intensive.

There are several other topics of interest that might be investigated during a continuation of the AGP. Running analysis jobs at the ATLAS Google Site was only briefly examined and this workflow, which is unique in its highly variable nature, may warrant further scrutiny. In particular, actions to replicate the most highly requested analysis-level data samples were not measurably successful, requiring further understanding of data popularity and data placement. Taking running analysis on the Google Cloud further, the idea of extending the resources of the ATLAS Google site with user-specific Google credits has also been raised. Authentication issues could also be interesting to look into, when considering using the site not only for ATLAS data, but also for hosting Open Data, available to all. On the other hand, it may be worth investing some time in understanding the privacy implications of the data stored in commercial cloud resources, which may contain protected information, at least according to GDPR [62], and whether this must be taken into consideration.

The ATLAS Google Project has provided valuable insights into the use of commercial clouds, and there remain many avenues of investigation that could be pursued. This next phase could feature a significant reduction of resources required, and hence financial expense, as the focus shifts to network connectivity and continued R&D. The ATLAS Google site could nevertheless still continue, albeit at a lower level, although if the network R&D is successful in reducing the egress cost, it could once again be ramped up to verify the savings arising from the implementation of dedicated peering solutions.

8 Summary

While traditional, site-based resources have always formed the backbone of ATLAS computing, commercial clouds may in some cases provide a viable and attractive alternative or addition. Much experience was gained in the integration of the ATLAS Google site into ATLAS distributed computing and no significant technical issue was discovered to prevent the experiment employing such resources in the future. Furthermore, the current workflow and data management tools employed by ATLAS are shown to be adequate for applying changes to the site configuration. The technical solutions implemented are essentially cloud-agnostic, not only avoiding vendor lock-in, but also potentially enabling access to other commercial cloud resources. The subscription pricing model applied in this project has proven to be beneficial to ATLAS, although questions remain how this may change going forward.

The project has shown that commercial cloud sites are an effective mechanism for providing additional, on-demand CPU resources. At the level employed by ATLAS, typically five or ten thousand cores, preemption of the allocated job slots is barely an issue, even when using the spot instance model as is done here. Whilst higher eviction rates were occasionally noticeable, for example during the data reprocessing single workflow period, the overall failure rate was not significantly higher than that observed on the grid. The ATLAS Google site was also shown to be extremely effective as a bursting resource, quickly providing up to one hundred thousand additional job slots, resulting in a significantly faster production turnaround than is possible on the ATLAS grid sites. In addition, the project has enabled parallel R&D efforts to flourish by providing different types of resources, for example GPU or ARM, on an elastic basis, demonstrating rapid integration [59].

Whilst it was also shown that it is possible to integrate, adjust and expand associated storage at the cloud site, this is less trivial than CPU as the intrinsic network costs must be taken into consideration. Storage and in particular network costs are known to dominate the TCO of commercial clouds, so much so this often dissuades sites taking an active interest in employing such resources. The studies performed during the AGP and outlined in this TCO analysis have shown that commercial cloud is a technically viable option for ATLAS distributed computing, albeit with additional costs not necessarily considered when employing traditional grid resources. Within the WLCG model the cost of the network is sometimes hidden, and whilst in reality this is probably rather high and means to reduce it are worth investigating, it can also be considered as irreducible and somewhat independent of ATLAS. Conversely, commercial cloud data transfers over standard internet networks incur significant costs for the data centres.

The TCO evaluation has shown that without the subscription model, the cost of commercial cloud resources is significantly increased. The Google Cloud resources used during this project cost a total of \$3.162M at list-price compared to the \$849,458 paid via the subscription agreement, representing a discount of 73%. Alternatively, ATLAS used 3.72 times more Google Cloud resources than were purchased via the subscription agreement, which means the resources used during this project would have been 272% more expensive at list-price. This is most obvious in the costs associated with the bursting test shown in Fig. 5c, which depicts daily expenditure considerably in excess of the \$1900 per day rate of the subscription agreement.

As shown in Table 1, almost half of the total list-price costs are due to egress. With this in mind, the ATLAS Collaboration is investigating ways to reduce this cost via dedicated network solutions, as outlined in "Dedicated Networks" section. It was also shown that egress costs are workflow dependent, which may be a consideration when employing such resources in the future. In particular, the substantial egress associated with data reprocessing means that for now this workflow is best avoided until further improvements in network connectivity are deployed, and that until then the site cannot be seen as universally suitable for ATLAS. If workflows such as Fast or Full Chain [63] can be employed, where the egress of intermediate MC formats is avoided, this will also help to reduce these costs.

Establishing a viable and cost-effective subscription agreement between experiment and commercial cloud provider is clearly a critical consideration of the TCO, given the large discrepancy between the list-prices and the agreement associated with this project. A collaborative approach may be necessary to obtain the best deal with a large volume discount. One option could be for CERN to make a significant purchase of cloud capacity and give the option to pay to be part of it, which has been done before [64]. This offer could extend not only to multiple sites, but also multiple experiments, and could even be done cooperatively with other international organisations such as EMBL [65]. Another important consideration going forward, especially if commercial clouds are employed by a significant number of sites, is to understand how such resources fit into the WLCG pledge structure.

The initial concerns about a significant loss of on-site personnel when outsourcing computing to the cloud, and the potential wider implications for ATLAS due to additional parallel support roles, appear to be less pronounced. This is because, for the most part, these individuals at the sites would still be needed to contribute to ATLAS distributed computing. As such, at least for larger sites, employing commercial clouds and off-premises resources may not immediately result in significant cost savings.

In summary, commercial cloud computing is an effective technical solution for ATLAS for providing additional CPU resources, and whilst the seamless integration of cloud-based storage was also achieved, network costs may be significant, based on the list-price. Some ATLAS workflows are found to be better than others with respect to egress. Resource bursting was shown to be very effective, albeit at significant cost. Establishing a favourable subscription agreement model that makes sense to both the cloud provider and the client is an advantage. By leveraging the Google Cloud Subscription Agreement pricing model, ATLAS has effectively harnessed between three to four times the resources compared to what the same investment would deliver for the list-price. It is yet to be seen how much this project influences the structure and cost of any potential follow-up deal to be brokered. There is much interest for ATLAS to continue this project, where network connectivity would be the main focus.

Executive Summary

The ATLAS Google Project was established to continue an evaluation of commercial clouds, in anticipation of the potential future adoption of such resources by WLCG grid sites to fulfil or complement their computing pledges to ATLAS. Cost estimates of commercial cloud resources have been done within ATLAS before, but this is the first time a total cost of ownership evaluation was performed for a long-term 15-month period. Whilst commercial cloud pricing structures are usually fine grained, like most clients ATLAS has negotiated a flat rate subscription agreement with Google for the duration of the project. Therefore, the method employed here is to analyse the relative contributions of various components within the cloud service to the total cost of ownership, including compute, storage, and network, for different ATLAS workflows and under different operating conditions, and identify the dominant cost drivers.

A substantial amount of technical development work was dedicated to the seamless integration of cloud resources into the ATLAS distributed computing infrastructure, employing the same software stack, primarily through cloud-native interfaces like Kubernetes and signed HTTP URLs following a cloud signature standard like S3v4. Care was taken to avoid vendor-specific choices and technology to mitigate the risk of potential price volatility in scenarios where one resource or service provider has a monopoly. The ATLAS Google Project has been key to test and validate at scale. Once these interfaces were established, the operation of an ATLAS site in the cloud has proven to be very effective, unlocking capabilities that would be either unfeasible or considerably time-consuming in an on-premises setting. For example, it has enabled rapid scaling of the number of processing jobs within a few hours and the deployment of non-x86 architecture resources on a large scale, tasks that could ordinarily take months to accomplish.

From the technical perspective, the project was a success, demonstrating that an ATLAS site can be deployed in a commercial cloud at a very large scale and in a very effective manner, requiring little additional operational effort. No significant technical issue was discovered to prevent the experiment routinely employing such resources in the future, and the existing workflow and data management services and software packages used by ATLAS are found to be adequate for effectively managing the operation of a significant amount of resources in the Google Cloud Platform.

Over the course of the project, running jobs on the ATLAS Google site may be broken down into several phases. The initial phase, which is described in some detail, was a process of familiarisation, learning how the site behaves, and applying necessary changes to bring the respective costs of the different services under control. This was followed by an extended period of running individual ATLAS workflows. Bursting to a much greater number of cores was then examined, followed by the final phase, where for the last few months of the project both the CPU and storage were expanded to the size of a typical large ATLAS site.

This project has shown that commercial cloud computing is an effective technical mechanism for ATLAS for providing additional CPU resources at the level of a large WLCG site. Resource bursting was successfully demonstrated, where available CPU resources can be increased by 100,000 additional cores within an hour and with no additional operational overhead. The utilisation of cloud-based storage was also demonstrated, and the impact of network costs evaluated. Network egress costs can be very high and currently dominate the overall cost depending on the workloads run at the cloud site. This study has quantified this effect, whilst also confirming that the ATLAS data management software, Rucio, includes features that allow network traffic and derived costs to be effectively controlled. The ATLAS computing model currently relies on a plentiful network connectivity between all sites, but it might evolve to reduce data traffic if the actual network costs were exposed.

By leveraging the Google Cloud Subscription Agreement pricing model, ATLAS has effectively harnessed between three and four times the resources compared to what the same investment would deliver for the list-price. This underscores the vital importance of establishing such agreements with cloud providers, which serve as essential tools for accessing significant volume discounts and ensuring cost predictability, while retaining the flexibility and scalability advantages inherent to cloud services. In essence, these agreements are not merely advantageous but rather a prerequisite for enabling large-scale cloud deployments. As such, the list-price should therefore be seen as the upper limit of the actual price paid for large-scale cloud services.

Whilst many valuable insights into the use of commercial clouds have been provided, it is still possible to further develop the ATLAS workload and data management software stack in order to integrate commercial clouds in the most efficient way. Key areas for future work include evaluating how the private cloud and LHC research network infrastructure can be interconnected and how the orchestration of data and workflows can provide maximal gains in the performance and flexibility of the computing model with minimal additional cost. The full implementation of these developments is important to enable exploring a potential evolution of the ATLAS computing model that tackles the issue of network costs. An extension of the project would also enable the current, wide ranging R&D programme to continue, which exploits the elastic availability of special resources such as GPUs and alternative architectures such as ARM, otherwise not readily available to ATLAS, enabling fast validation and benchmarking of new architectures without the need to make upfront investment in hardware.

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. [66].

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; ISF and Benoziyo Center, Israel; INFN, Italy; MEXT and JSPS, Japan; CNRST, Morocco; NWO, Netherlands; RCN, Norway; MNiSW, Poland; FCT, Portugal; MNE/IFA, Romania; MESTD, Serbia; MSSR, Slovakia; ARRS and MIZŠ, Slovenia; DSI/NRF, South Africa; MICINN, Spain; SRC and Wallenberg Foundation, Sweden; SERI, SNSF and Cantons of Bern and Geneva, Switzerland; MOST, Taipei; TENMAK, Türkiye; STFC, United Kingdom; DOE and NSF, USA.

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 programmes 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 Armenia: Yerevan Physics Institute (FAPERJ); CERN: European Organization for Nuclear Research (CERN PJAS); Chile: Agencia Nacional de Investigación y Desarrollo (FONDECYT 1230812, FONDECYT 1230987, FONDECYT 1240864); China: Chinese Ministry of Science and Technology (MOST-2023YFA1605700), National Natural Science Foundation of China (NSFC - 12175119, NSFC 12275265, NSFC-12075060); Czech Republic: Czech Science Foundation (GACR - 24-11373 S), Ministry of Education Youth and Sports (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), 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-20-CE31-0013, ANR-21-CE31-0013, ANR-21-CE31-0022), Investissements d'Avenir Labex (ANR-11-LABX-0012); Germany: Baden-Württemberg Stiftung (BW Stiftung-Postdoc Eliteprogramme), Deutsche Forschungsgemeinschaft (DFG - 469666862, DFG - CR 312/5-2); Italy: Istituto Nazionale di Fisica Nucleare (ICSC, NextGenerationEU); Japan: Japan Society for the Promotion of Science (JSPS KAKENHI JP22H01227, JSPS KAKENHI JP22H04944, JSPS KAKENHI JP22KK0227, JSPS KAKENHI JP23KK0245); Netherlands: Netherlands Organisation for Scientific Research (NWO Veni 2020 - VI.Veni.202.179); Norway: Research Council of Norway (RCN-314472); Poland: Polish National Agency for Academic Exchange (PPN/PPO/2020/1/00002/U/00001), Polish National Science Centre (NCN 2021/42/E/ST2/00350, NCN OPUS nr 2022/47/B/ST2/03059, NCN UMO-2019/34/E/ST2/00393, UMO-2020/37/B/ST2/01043, UMO-2021/40/C/ST2/00187, UMO-2022/47/O/ST2/00148, UMO-2023/49/B/ST2/04085); Slovenia: Slovenian Research Agency (ARIS grant J1-3010); Spain: Generalitat Valenciana (Artemisa, FEDER, IDIFEDER/2018/048), 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), PROMETEO and GenT Programmes Generalitat Valenciana (CIDEGENT/2019/027); Sweden: Swedish Research Council (Swedish Research Council 2023-04654, VR 2018-00482, VR 2022-03845, VR 2022-04683, VR 2023-03403, VR grant 2021-03651), Knut and Alice Wallenberg Foundation (KAW 2018.0157, KAW 2018.0458, KAW 2019.0447, KAW 2022.0358); Switzerland: Swiss National Science Foundation (SNSF - PCEFP2_194658); United Kingdom: 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), Neubauer Family Foundation.

Author contributions All authors have contributed to the publication, being variously involved in the design and the construction of the detectors, in writing software, calibrating subsystems, operating the detectors and acquiring data, and finally analysing the processed data. The ATLAS Collaboration members discussed and approved the scientific results. The manuscript was prepared by a subgroup of authors appointed by the collaboration and subject to an internal collaboration-wide review process. All authors reviewed and approved the final version of the manuscript.

Funding Open access funding provided by CERN (European Organization for Nuclear Research)

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing interests The authors declare no competing interests.

References

1. Collaboration ATLAS (2008) The ATLAS experiment at the CERN large Hadron Collider. JINST 3:S08003. https://doi.org/10.1088/1748-0221/3/08/S08003
2. Evans L, Bryant P (2008) LHC machine. JINST 3:S08001. https://doi.org/10.1088/1748-0221/3/08/S08001
3. Bird I et al (2005) LHC computing grid: technical design report, CERN-LHCC-2005-024. https://cds.cern.ch/record/840543
4. Bird I et al (2014) Update of the computing models of the WLCG and the LHC experiments, CERN-LHCC-2014-014. https://cds.cern.ch/record/1695401
5. Úbeda García M et al (2014) Integration of cloud resources in the LHCb distributed computing. J Phys Conf Ser 513:032099. https://doi.org/10.1088/1742-6596/513/3/032099
6. Panitkin S (2015) Look to the clouds and beyond. Nat Phys 11:373. https://doi.org/10.1038/nphys3319
7. Holzman B et al (2017) HEPCloud, a new paradigm for HEP facilities: CMS Amazon web services investigation. Comput Softw Big Sci. https://doi.org/10.1007/s41781-017-0001-9
8. Lončar P (2023) Scalable data processing model of the ALICE experiment in the cloud, PhD thesis: Sveučilište u Splitu. Fakultet elektrotehnike, strojarstva i brodogradnje. Zavod za elektroniku i računarstvo., University of Split. https://cds.cern.ch/record/2874778
9. Megino F Barreiro, Bryant L, Hufnagel D, Anampa K Hurtado (2023) US ATLAS and US CMS HPC and cloud blueprint. arXiv:2304.07376 [physics.comp-ph]
10. Megino F Barreiro et al (2021) Seamless integration of commercial clouds with ATLAS distributed computing. EPJ Web Conf 251:02005. https://doi.org/10.1051/epjconf/202125102005
11. Megino F Barreiro et al (2017) PanDA for ATLAS distributed computing in the next decade. J Phys Conf Ser 898:052002. https://doi.org/10.1088/1742-6596/898/5/052002
12. Barisits M et al (2019) Rucio: scientific data management. Comput Softw Big Sci. https://doi.org/10.1007/s41781-019-0026-3
13. Apollinari G et al (2015) High–luminosity large Hadron Collider (HL–LHC): preliminary design report, CERN-2015-005. https://cds.cern.ch/record/2116337
14. ATLAS Collaboration (2020) ATLAS HL–LHC computing conceptual design report, CERN-LHCC-2020-015, LHCC-G-178. https://cds.cern.ch/record/2729668
15. ATLAS Collaboration (2022) ATLAS software and computing HL–LHC roadmap, CERN-LHCC-2022-005, LHCC-G-182. https://cds.cern.ch/record/2802918 -
16. Devouassoux M (2018) Method to calculate the total cost of ownership of infrastructure as a service, version 2. https://doi.org/10.5281/zenodo.2161088
17. Helix Nebula—the science cloud. https://doi.org/10.3030/687614
18. Amazon Web Services pricing. https://aws.amazon.com/pricing
19. Google Cloud Platform pricing. https://cloud.google.com/products/calculator
20. Martelli E (2024) Evolving the LHCOPN and LHCONE networks to support HL-LHC computing requirements. EPJ Web Conf 295:07016. https://doi.org/10.1051/epjconf/202429507016
21. ATLAS Collaboration (2024) Software and computing for Run 3 of the ATLAS experiment at the LHC. arXiv:2404.06335 [hep-ex]
22. Kubernetes. https://kubernetes.io/docs/home/
23. Maeno T et al (2019) Harvester: an edge service harvesting heterogeneous resources for ATLAS. EPJ Web Conf 214:03030. https://doi.org/10.1051/epjconf/201921403030
24. Megino F Barreiro et al (2020) Using Kubernetes as an ATLAS computing site. EPJ Web Conf 245:07025. https://doi.org/10.1051/epjconf/202024507025
25. Megino F Barreiro et al (2024) The usage of commercial Accelerating science clouds in ATLAS Distributed Computing. EPJ Web Conf 295:07002. https://doi.org/10.1051/epjconf/202429507002
26. Blomer J et al (2020) The CernVM file system, version 2.7.5. https://doi.org/10.5281/zenodo.4114078
27. Authenticating requests (AWS Signature, version 4). https://docs.aws.amazon.com/AmazonS3/latest/API/sig-v4-authenticating-requests.html
28. Davix. https://davix.web.cern.ch/davix/docs/devel/
29. Barisits M et al (2024) Extending Rucio with modern cloud storage support. EPJ Web Conf 295:01030. https://doi.org/10.1051/epjconf/202429501030
30. CERN File Transfer Service. https://fts.web.cern.ch/fts/
31. Interoperable Global Trust Federation. https://www.igtf.net/
32. Google spot VMs. https://cloud.google.com/compute/docs/instances/spot
33. Rucio replica management. https://rucio.cern.ch/documentation/started/concepts/replica_management/
34. Rucio RSE configuration. https://rucio.cern.ch/documentation/started/concepts/rucio_storage_element/
35. Borodin M et al (2021) The ATLAS data carousel project status. EPJ Web Conf 251:02006. https://doi.org/10.1051/epjconf/202125102006
36. ATLAS Collaboration (2022) Active Learning reinterpretation of an ATLAS Dark Matter search constraining a model of a dark Higgs boson decaying to two (b)-quarks, ATL-PHYS-PUB-2022-045, 2022. https://cds.cern.ch/record/2839789
37. ATLAS Collaboration (2017) Performance of the ATLAS trigger system in 2015. Eur Phys J C 77:317. https://doi.org/10.1140/epjc/s10052-017-4852-3. arXiv:1611.09661 [hep-ex]

38. ATLAS Collaboration (2024) The ATLAS trigger system for LHC run 3 and trigger performance in 2022. arXiv:2401.06630 [hep-ex]
39. Berghaus F et al (2020) ATLAS Sim@P1 upgrades during long shutdown two. EPJ Web Conf 245:07044. https://doi.org/10.1051/epjconf/202024507044
40. Glushkov I, Lee C, Di Girolamo A, Walker R, Gottardo CA (2024) Optimization of opportunistic utilization of the ATLAS high-level trigger farm for LHC Run 3. EPJ Web Conf 295:07035. https://doi.org/10.1051/epjconf/202429507035
41. HPC Vega. https://www.izum.si/en/vega-en/
42. HPC Perlmutter. https://docs.nersc.gov/systems/perlmutter/
43. Google storage pricing. https://cloud.google.com/storage/pricing#europe
44. Google network service tiers Pricing. https://cloud.google.com/network-tiers/pricing
45. Google Cloud interconnect pricing. https://cloud.google.com/network-connectivity/docs/interconnect/pricing
46. Sfiligoi I et al (2021) Managing Cloud networking costs for data–intensive applications by provisioning dedicated network links. arXiv:2104.06913 [cs.NI]
47. Energy Sciences Network. https://www.es.net
48. Giordano D et al (2023) HEPScore: a new CPU benchmark for the WLCG; 2023. arXiv:2306.08118 [hep-ex]
49. Google Cloud Platform Terms of Service. https://cloud.google.com/terms
50. Google data centers: efficiency. https://www.google.com/about/datacenters/efficiency/
51. Carbon free energy for Google Cloud regions. https://cloud.google.com/sustainability/region-carbon
52. The OCRE project. https://www.ocre-project.eu/
53. HEPCloud. https://computing.fnal.gov/hep-cloud/
54. CloudBank. https://www.cloudbank.org/
55. Energy Sciences Network peering. https://www.es.net/engineering-services/the-network/peering-connections/
56. Internet2 cloud access. https://internet2.edu/cloud/cloud-access/
57. GÉANT Network. https://network.geant.org/
58. Grid File Access Library, version 2. https://dmc-docs.web.cern.ch/dmc-docs/gfal2/gfal2.html
59. Megino F Barreiro et al (2024) Operational experience and R&D results using the google cloud for high energy physics in the ATLAS experiment. arXiv:2403.15873 [hep-ex]
60. JupyterHub. https://jupyter.org/hub
61. Dask. https://www.dask.org
62. General Data Protection Regulation. https://gdpr-info.eu
63. Javurkova M et al (2021) The fast simulation chain in the ATLAS experiment. EPJ Web Conf 251:03012. https://doi.org/10.1051/epjconf/202125103012
64. CloudBank EU NGI. https://ngiatlantic.eu/funded-experiments/cloudbank-eu-ngi
65. The European Molecular Biology Laboratory. https://www.embl.org
66. ATLAS Collaboration, ATLAS Computing Acknowledgements, ATL-SOFT-PUB-2023-001, 2023. https://cds.cern.ch/record/2869272

The ATLAS Collaboration

G. Aad[104], E. Aakvaag[17], B. Abbott[123], S. Abdelhameed[119a], K. Abeling[56], N. J. Abicht[50], S. H. Abidi[30], M. Aboelela[45], A. Aboulhorma[36e], H. Abramowicz[154], H. Abreu[153], Y. Abulaiti[120], B. S. Acharya[70a,70b,m], A. Ackermann[64a], C. Adam Bourdarios[4], L. Adamczyk[87a], S. V. Addepalli[27], M. J. Addison[103], J. Adelman[118], A. Adiguzel[22c], T. Adye[137], A. A. Affolder[139], Y. Afik[40], M. N. Agaras[13], J. Agarwala[74a,74b], A. Aggarwal[102], C. Agheorghiesei[28c], F. Ahmadov[39,y], W. S. Ahmed[106], S. Ahuja[97], X. Ai[63e], G. Aielli[77a,77b], A. Aikot[166], M. Ait Tamlihat[36e], B. Aitbenchikh[36a], M. Akbiyik[102], T. P. A. Åkesson[100], A. V. Akimov[38], D. Akiyama[171], N. N. Akolkar[25], S. Aktas[22a], K. Al Khoury[42], G. L. Alberghi[24b], J. Albert[168], P. Albicocco[54], G. L. Albouy[61], S. Alderweireldt[53], Z. L. Alegria[124], M. Aleksa[37], I. N. Aleksandrov[39], C. Alexa[28b], T. Alexopoulos[10], F. Alfonsi[24b], M. Algren[57], M. Alhroob[170], B. Ali[135], H. M. J. Ali[93], S. Ali[32], S. W. Alibocus[94], M. Aliev[34c], G. Alimonti[72a], W. Alkakhi[56], C. Allaire[67], B. M. M. Allbrooke[149], J. F. Allen[53], C. A. Allendes Flores[140f], P. P. Allport[21], A. Aloisio[73a,73b], F. Alonso[92], C. Alpigiani[141], Z. M. K. Alsolami[93], M. Alvarez Estevez[101], A. Alvarez Fernandez[102], M. Alves Cardoso[57], M. G. Alviggi[73a,73b], M. Aly[103], Y. Amaral Coutinho[84b], A. Ambler[106], C. Amelung[37], M. Amerl[103], C. G. Ames[111], D. Amidei[108], B. Amini[55], K. J. Amirie[158], S. P. Amor Dos Santos[133a], K. R. Amos[166], D. Amperiadou[155], S. An[85], V. Ananiev[128], C. Anastopoulos[142], T. Andeen[11], J. K. Anders[37], A. C. Anderson[60], S. Y. Andrean[48a,48b], A. Andreazza[72a,72b], S. Angelidakis[9], A. Angerami[42], A. V. Anisenkov[38], A. Annovi[75a], C. Antel[57], E. Antipov[148], M. Antonelli[54], F. Anulli[76a], M. Aoki[85], T. Aoki[156], M. A. Aparo[149], L. Aperio Bella[49], C. Appelt[19], A. Apyan[27], S. J. Arbiol Val[88], C. Arcangeletti[54], A. T. H. Arce[52], J.-F. Arguin[110], S. Argyropoulos[55], J.-H. Arling[49], O. Arnaez[4], H. Arnold[148], G. Artoni[76a,76b], H. Asada[113], K. Asai[121], S. Asai[156], N. A. Asbah[37], R. A. Ashby Pickering[170], K. Assamagan[30], R. Astalos[29a], K. S. V. Astrand[100], S. Atashi[162], R. J. Atkin[34a], M. Atkinson[165], H. Atmani[36f], P. A. Atmasiddha[131], K. Augsten[135], S. Auricchio[73a,73b], A. D. Auriol[21], V. A. Austrup[103], G. Avolio[37], K. Axiotis[57], G. Azuelos[110,ad], D. Babal[29b], H. Bachacou[138], K. Bachas[155,q], A. Bachiu[35], F. Backman[48a,48b], A. Badea[40], T. M. Baer[108], P. Bagnaia[76a,76b], M. Bahmani[19], D. Bahner[55], K. Bai[126], J. T. Baines[137], L. Baines[96], O. K. Baker[175], E. Bakos[16], D. Bakshi Gupta[8], L. E. Balabram Filho[84b], V. Balakrishnan[123], R. Balasubramanian[117], E. M. Baldin[38], P. Balek[87a], E. Ballabene[24b,24a], F. Balli[138], L. M. Baltes[64a], W. K. Balunas[33], J. Balz[102], I. Bamwidhi[119b], E. Banas[88], M. Bandieramonte[132], A. Bandyopadhyay[25], S. Bansal[25], L. Barak[154], M. Barakat[49], E. L. Barberio[107], D. Barberis[58b,58a], M. Barbero[104], M. Z. Barel[117], T. Barillari[112], M.-S. Barisits[37], T. Barklow[146], P. Baron[125], D. A. Baron Moreno[103], A. Baroncelli[63a], A. J. Barr[129], J. D. Barr[98], F. Barreiro[101], J. Barreiro Guimarães da Costa[14], U. Barron[154], M. G. Barros Teixeira[133a], S. Barsov[38], F. Bartels[64a], R. Bartoldus[146], A. E. Barton[93], P. Bartos[29a], A. Basan[102], M. Baselga[50], A. Bassalat[67,b], M. J. Basso[159a], S. Bataju[45], R. Bate[167], R. L. Bates[60], S. Batlamous[101], B. Batool[144], M. Battaglia[139], D. Battulga[19], M. Bauce[76a,76b], M. Bauer[80], P. Bauer[25], L. T. Bazzano Hurrell[31], J. B. Beacham[52], T. Beau[130], J. Y. Beaucamp[92], P. H. Beauchemin[161], P. Bechtle[25], H. P. Beck[20,p], K. Becker[170], A. J. Beddall[83], V. A. Bednyakov[39], C. P. Bee[148], L. J. Beemster[16], T. A. Beermann[37], M. Begalli[84d], M. Begel[30], A. Behera[148], J. K. Behr[49], J. F. Beirer[37], F. Beisiegel[25], M. Belfkir[119b], G. Bella[154], L. Bellagamba[24b], A. Bellerive[35], P. Bellos[21], K. Beloborodov[38], D. Benchekroun[36a], F. Bendebba[36a], Y. Benhammou[154], K. C. Benkendorfer[62], L. Beresford[49], M. Beretta[54], E. Bergeaas Kuutmann[164], N. Berger[4], B. Bergmann[135], J. Beringer[18a], G. Bernardi[5], C. Bernius[146], F. U. Bernlochner[25], F. Bernon[37,104], A. Berrocal Guardia[13], T. Berry[97], P. Berta[136], A. Berthold[51], S. Bethke[112], A. Betti[76a,76b], A. J. Bevan[96], N. K. Bhalla[55], S. Bhatta[148], D. S. Bhattacharya[169], P. Bhattarai[146], K. D. Bhide[55], V. S. Bhopatkar[124], R. M. Bianchi[132], G. Bianco[24b,24a], O. Biebel[111], R. Bielski[126], M. Biglietti[78a], C. S. Billingsley[45], Y. Bimgdi[36f], M. Bindi[56], A. Bingul[22b], C. Bini[76a,76b], G. A. Bird[33], M. Birman[172], M. Biros[136], S. Biryukov[149], T. Bisanz[50], E. Bisceglie[44b,44a], J. P. Biswal[137], D. Biswas[144], I. Bloch[49], A. Blue[60], U. Blumenschein[96], J. Blumenthal[102], V. S. Bobrovnikov[38], M. Boehler[55], B. Boehm[169], D. Bogavac[37], A. G. Bogdanchikov[38], C. Bohm[48a], V. Boisvert[97], P. Bokan[37], T. Bold[87a], M. Bomben[5], M. Bona[96], M. Boonekamp[138], C. D. Booth[97], A. G. Borbély[60], I. S. Bordulev[38], G. Borissov[93], D. Bortoletto[129], D. Boscherini[24b], M. Bosman[13], J. D. Bossio Sola[37], K. Bouaouda[36a], N. Bouchhar[166], L. Boudet[4], J. Boudreau[132], E. V. Bouhova-Thacker[93], D. Boumediene[41], R. Bouquet[58b,58a], A. Boveia[122], J. Boyd[37], D. Boye[30], I. R. Boyko[39], L. Bozianu[57], J. Bracinik[21], N. Brahimi[4], G. Brandt[174], O. Brandt[33],

F. Braren[49], B. Brau[105], J. E. Brau[126], R. Brener[172], L. Brenner[117], R. Brenner[164], S. Bressler[172], G. Brianti[79a,79b], D. Britton[60], D. Britzger[112], I. Brock[25], G. Brooijmans[42], E. M. Brooks[159b], E. Brost[30], L. M. Brown[168], L. E. Bruce[62], T. L. Bruckler[129], P. A. Bruckman de Renstrom[88], B. Brüers[49], A. Bruni[24b], G. Bruni[24b], M. Bruschi[24b], N. Bruscino[76a,76b], T. Buanes[17], Q. Buat[141], D. Buchin[112], A. G. Buckley[60], O. Bulekov[38], B. A. Bullard[146], S. Burdin[94], C. D. Burgard[50], A. M. Burger[37], B. Burghgrave[8], O. Burlayenko[55], J. Burleson[165], J. T. P. Burr[33], J. C. Burzynski[145], E. L. Busch[42], V. Büscher[102], P. J. Bussey[60], J. M. Butler[26], C. M. Buttar[60], J. M. Butterworth[98], W. Buttinger[137], C. J. Buxo Vazquez[109], A. R. Buzykaev[38], S. Cabrera Urbán[166], L. Cadamuro[67], D. Caforio[59], H. Cai[132], Y. Cai[14,114c], Y. Cai[114a], V. M. M. Cairo[37], O. Cakir[3a], N. Calace[37], P. Calafiura[18a], G. Calderini[130], P. Calfayan[69], G. Callea[60], L. P. Caloba[84b], D. Calvet[41], S. Calvet[41], M. Calvetti[75a,75b], R. Camacho Toro[130], S. Camarda[37], D. Camarero Munoz[27], P. Camarri[77a,77b], M. T. Camerlingo[73a,73b], D. Cameron[37], C. Camincher[168], M. Campanelli[98], A. Camplani[43], V. Canale[73a,73b], A. C. Canbay[3a], E. Canonero[97], J. Cantero[166], Y. Cao[165], F. Capocasa[27], M. Capua[44b,44a], A. Carbone[72a,72b], R. Cardarelli[77a], J. C. J. Cardenas[8], G. Carducci[44b,44a], T. Carli[37], G. Carlino[73a], J. I. Carlotto[13], B. T. Carlson[132,r], E. M. Carlson[168,159a], J. Carmignani[94], L. Carminati[72a,72b], A. Carnelli[138], M. Carnesale[76a,76b], S. Caron[116], E. Carquin[140f], S. Carrá[72a], G. Carratta[24b,24a], A. M. Carroll[126], T. M. Carter[53], M. P. Casado[13,j], M. Caspar[49], F. L. Castillo[4], L. Castillo Garcia[13], V. Castillo Gimenez[166], N. F. Castro[133a,133e], A. Catinaccio[37], J. R. Catmore[128], T. Cavaliere[4], V. Cavaliere[30], N. Cavalli[24b,24a], L. J. Caviedes Betancourt[23b], Y. C. Cekmecelioglu[49], E. Celebi[83], S. Cella[37], F. Celli[129], M. S. Centonze[71a,71b], V. Cepaitis[57], K. Cerny[125], A. S. Cerqueira[84a], A. Cerri[149], L. Cerrito[77a,77b], F. Cerutti[18a], B. Cervato[144], A. Cervelli[24b], G. Cesarini[54], S. A. Cetin[83], D. Chakraborty[118], J. Chan[18a], W. Y. Chan[156], J. D. Chapman[33], E. Chapon[138], B. Chargeishvili[152b], D. G. Charlton[21], M. Chatterjee[20], C. Chauhan[136], Y. Che[114a], S. Chekanov[6], S. V. Chekulaev[159a], G. A. Chelkov[39,a], A. Chen[108], B. Chen[154], B. Chen[168], H. Chen[114a], H. Chen[30], J. Chen[63c], J. Chen[145], M. Chen[129], S. Chen[156], S. J. Chen[114a], X. Chen[63c], X. Chen[15,ac], Y. Chen[63a], C. L. Cheng[173], H. C. Cheng[65a], S. Cheong[146], A. Cheplakov[39], E. Cheremushkina[49], E. Cherepanova[117], R. Cherkaoui El Moursli[36e], E. Cheu[7], K. Cheung[66], L. Chevalier[138], V. Chiarella[54], G. Chiarelli[75a], N. Chiedde[104], G. Chiodini[71a], A. S. Chisholm[21], A. Chitan[28b], M. Chitishvili[166], M. V. Chizhov[39], K. Choi[11], Y. Chou[141], E. Y. S. Chow[116], K. L. Chu[172], M. C. Chu[65a], X. Chu[14,114c], Z. Chubinidze[54], J. Chudoba[134], J. J. Chwastowski[88], D. Cieri[112], K. M. Ciesla[87a], V. Cindro[95], A. Ciocio[18a], F. Cirotto[73a,73b], Z. H. Citron[172], M. Citterio[72a], D. A. Ciubotaru[28b], A. Clark[57], P. J. Clark[53], N. Clarke Hall[98], C. Clarry[158], J. M. Clavijo Columbie[49], S. E. Clawson[49], C. Clement[48a,48b], Y. Coadou[104], M. Cobal[70a,70c], A. Coccaro[58b], R. F. Coelho Barrue[133a], R. Coelho Lopes De Sa[105], S. Coelli[72a], B. Cole[42], J. Collot[61], P. Conde Muiño[133a,133g], M. P. Connell[34c], S. H. Connell[34c], E. I. Conroy[129], F. Conventi[73a,ae], H. G. Cooke[21], A. M. Cooper-Sarkar[129], F. A. Corchia[24b,24a], A. Cordeiro Oudot Choi[130], L. D. Corpe[41], M. Corradi[76a,76b], F. Corriveau[106,x], A. Cortes-Gonzalez[19], M. J. Costa[166], F. Costanza[4], D. Costanzo[142], B. M. Cote[122], J. Couthures[4], G. Cowan[97], K. Cranmer[173], D. Cremonini[24b,24a], S. Crépé-Renaudin[61], F. Crescioli[130], M. Cristinziani[144], M. Cristoforetti[79a,79b], V. Croft[117], J. E. Crosby[124], G. Crosetti[44b,44a], A. Cueto[101], H. Cui[98], Z. Cui[7], W. R. Cunningham[60], F. Curcio[166], J. R. Curran[53], P. Czodrowski[37], M. J. Da Cunha Sargedas De Sousa[58b,58a], J. V. Da Fonseca Pinto[84b], C. Da Via[103], W. Dabrowski[87a], T. Dado[37], S. Dahbi[151], T. Dai[108], D. Dal Santo[20], C. Dallapiccola[105], M. Dam[43], G. D'amen[30], V. D'Amico[111], J. Damp[102], J. R. Dandoy[35], D. Dannheim[37], M. Danninger[145], V. Dao[148], G. Darbo[58b], S. J. Das[30,af], F. Dattola[49], S. D'Auria[72a,72b], A. D'avanzo[73a,73b], C. David[34a], T. Davidek[136], I. Dawson[96], H. A. Day-hall[135], K. De[8], R. De Asmundis[73a], N. De Biase[49], S. De Castro[24b,24a], N. De Groot[116], P. de Jong[117], H. De la Torre[118], A. De Maria[114a], A. De Salvo[76a], U. De Sanctis[77a,77b], F. De Santis[71a,71b], A. De Santo[149], J. B. De Vivie De Regie[61], D. V. Dedovich[39], J. Degens[94], A. M. Deiana[45], F. Del Corso[24b,24a], J. Del Peso[101], F. Del Rio[64a], L. Delagrange[130], F. Deliot[138], C. M. Delitzsch[50], M. Della Pietra[73a,73b], D. Della Volpe[57], A. Dell'Acqua[37], L. Dell'Asta[72a,72b], M. Delmastro[4], P. A. Delsart[61], S. Demers[175], M. Demichev[39], S. P. Denisov[38], L. D'Eramo[41], D. Derendarz[88], F. Derue[130], P. Dervan[94], K. Desch[25], C. Deutsch[25], F. A. Di Bello[58b,58a], A. Di Ciaccio[77a,77b], L. Di Ciaccio[4], A. Di Domenico[76a,76b], C. Di Donato[73a,73b], A. Di Girolamo[37], G. Di Gregorio[37], A. Di Luca[79a,79b], B. Di Micco[78a,78b], R. Di Nardo[78a,78b], K. F. Di Petrillo[40], M. Diamantopoulou[35], F. A. Dias[117], T. Dias Do Vale[145], M. A. Diaz[140a,140b], F. G. Diaz Capriles[25], A. R. Didenko[39], M. Didenko[166], E. B. Diehl[108], S. Díez Cornell[49], C. Diez Pardos[144], C. Dimitriadi[164], A. Dimitrievska[21], J. Dingfelder[25],

T. Dingley[129], I.-M. Dinu[28b], S. J. Dittmeier[64b], F. Dittus[37], M. Divisek[136], F. Djama[104], T. Djobava[152b], C. Doglioni[103,100], A. Dohnalova[29a], J. Dolejsi[136], Z. Dolezal[136], K. Domijan[87a], K. M. Dona[40], M. Donadelli[84d], B. Dong[109], J. Donini[41], A. D'Onofrio[73a,73b], M. D'Onofrio[94], J. Dopke[137], A. Doria[73a], N. Dos Santos Fernandes[133a], P. Dougan[103], M. T. Dova[92], A. T. Doyle[60], M. A. Draguet[129], E. Dreyer[172], I. Drivas-koulouris[10], M. Drnevich[120], M. Drozdova[57], D. Du[63a], T. A. du Pree[117], F. Dubinin[38], M. Dubovsky[29a], E. Duchovni[172], G. Duckeck[111], O. A. Ducu[28b], D. Duda[53], A. Dudarev[37], E. R. Duden[27], M. D'uffizi[103], L. Duflot[67], M. Dührssen[37], I. Duminica[28g], A. E. Dumitriu[28b], M. Dunford[64a], S. Dungs[50], K. Dunne[48a,48b], A. Duperrin[104], H. Duran Yildiz[3a], M. Düren[59], A. Durglishvili[152b], B. L. Dwyer[118], G. I. Dyckes[18a], M. Dyndal[87a], B. S. Dziedzic[37], Z. O. Earnshaw[149], G. H. Eberwein[129], B. Eckerova[29a], S. Eggebrecht[56], E. Egidio Purcino De Souza[84e], L. F. Ehrke[57], G. Eigen[17], K. Einsweiler[18a], T. Ekelof[164], P. A. Ekman[100], S. El Farkh[36b], Y. El Ghazali[63a], H. El Jarrari[37], A. El Moussaouy[36a], V. Ellajosyula[164], M. Ellert[164], F. Ellinghaus[174], N. Ellis[37], J. Elmsheuser[30], M. Elsawy[119a], M. Elsing[37], D. Emeliyanov[137], Y. Enari[85], I. Ene[18a], S. Epari[13], P. A. Erland[88], D. Ernani Martins Neto[88], M. Errenst[174], M. Escalier[67], C. Escobar[166], E. Etzion[154], G. Evans[133a], H. Evans[69], L. S. Evans[97], A. Ezhilov[38], S. Ezzarqtouni[36a], F. Fabbri[24b,24a], L. Fabbri[24b,24a], G. Facini[98], V. Fadeyev[139], R. M. Fakhrutdinov[38], D. Fakoudis[102], S. Falciano[76a], L. F. Falda Ulhoa Coelho[37], F. Fallavollita[112], G. Falsetti[44b,44a], J. Faltova[136], C. Fan[165], Y. Fan[14], Y. Fang[14,114c], M. Fanti[72a,72b], M. Faraj[70a,70b], Z. Farazpay[99], A. Farbin[8], A. Farilla[78a], T. Farooque[109], S. M. Farrington[53], F. Fassi[36e], D. Fassouliotis[9], M. Faucci Giannelli[77a,77b], W. J. Fawcett[33], L. Fayard[67], P. Federic[136], P. Federicova[134], O. L. Fedin[38,a], M. Feickert[173], L. Feligioni[104], D. E. Fellers[126], C. Feng[63b], Z. Feng[117], M. J. Fenton[162], L. Ferencz[49], R. A. M. Ferguson[93], S. I. Fernandez Luengo[140f], P. Fernandez Martinez[13], M. J. V. Fernoux[104], J. Ferrando[93], A. Ferrari[164], P. Ferrari[117,116], R. Ferrari[74a], D. Ferrere[57], C. Ferretti[108], D. Fiacco[76a,76b], F. Fiedler[102], P. Fiedler[135], A. Filipčič[95], E. K. Filmer[1], F. Filthaut[116], M. C. N. Fiolhais[133a,133c,c], L. Fiorini[166], W. C. Fisher[109], T. Fitschen[103], P. M. Fitzhugh[138], I. Fleck[144], P. Fleischmann[108], T. Flick[174], M. Flores[34d,aa], L. R. Flores Castillo[65a], L. Flores Sanz De Acedo[37], F. M. Follega[79a,79b], N. Fomin[33], J. H. Foo[158], A. Formica[138], A. C. Forti[103], E. Fortin[37], A. W. Fortman[18a], M. G. Foti[18a], L. Fountas[9,k], D. Fournier[67], H. Fox[93], P. Francavilla[75a,75b], S. Francescato[62], S. Franchellucci[57], M. Franchini[24b,24a], S. Franchino[64a], D. Francis[37], L. Franco[116], V. Franco Lima[37], L. Franconi[49], M. Franklin[62], G. Frattari[27], Y. Y. Frid[154], J. Friend[60], N. Fritzsche[37], A. Froch[55], D. Froidevaux[37], J. A. Frost[129], Y. Fu[63a], S. Fuenzalida Garrido[140f], M. Fujimoto[104], K. Y. Fung[65a], E. Furtado De Simas Filho[84e], M. Furukawa[156], J. Fuster[166], A. Gaa[56], A. Gabrielli[24b,24a], A. Gabrielli[158], P. Gadow[37], G. Gagliardi[58b,58a], L. G. Gagnon[18a], S. Gaid[163], S. Galantzan[154], E. J. Gallas[129], B. J. Gallop[137], K. K. Gan[122], S. Ganguly[156], Y. Gao[53], F. M. Garay Walls[140a,140b], B. Garcia[30], C. García[166], A. Garcia Alonso[117], A. G. Garcia Caffaro[175], J. E. García Navarro[166], M. Garcia-Sciveres[18a], G. L. Gardner[131], R. W. Gardner[40], N. Garelli[161], D. Garg[81], R. B. Garg[146], J. M. Gargan[53], C. A. Garner[158], C. M. Garvey[34a], V. K. Gassmann[161], G. Gaudio[74a], V. Gautam[13], P. Gauzzi[76a,76b], J. Gavranovic[95], I. L. Gavrilenko[38], A. Gavrilyuk[38], C. Gay[167], G. Gaycken[126], E. N. Gazis[10], A. A. Geanta[28b], C. M. Gee[139], A. Gekow[122], C. Gemme[58b], M. H. Genest[61], A. D. Gentry[115], S. George[97], W. F. George[21], T. Geralis[47], P. Gessinger-Befurt[37], M. E. Geyik[174], M. Ghani[170], K. Ghorbanian[96], A. Ghosal[144], A. Ghosh[162], A. Ghosh[7], B. Giacobbe[24b], S. Giagu[76a,76b], T. Giani[117], A. Giannini[63a], S. M. Gibson[97], M. Gignac[139], D. T. Gil[87b], A. K. Gilbert[87a], B. J. Gilbert[42], D. Gillberg[35], G. Gilles[117], L. Ginabat[130], D. M. Gingrich[2,ad], M. P. Giordani[70a,70c], P. F. Giraud[138], G. Giugliarelli[70a,70c], D. Giugni[72a], F. Giuli[37], I. Gkialas[9,k], L. K. Gladilin[38], C. Glasman[101], G. R. Gledhill[126], G. Glemža[49], M. Glisic[126], I. Gnesi[44b,f], Y. Go[30], M. Goblirsch-Kolb[37], B. Gocke[50], D. Godin[110], B. Gokturk[22a], S. Goldfarb[107], T. Golling[57], M. G. D. Gololo[34g], D. Golubkov[38], J. P. Gombas[109], A. Gomes[133a,133b], G. Gomes Da Silva[144], A. J. Gomez Delegido[166], R. Gonçalo[133a], L. Gonella[21], A. Gongadze[152c], F. Gonnella[21], J. L. Gonski[146], R. Y. González Andana[53], S. González de la Hoz[166], R. Gonzalez Lopez[94], C. Gonzalez Renteria[18a], M. V. Gonzalez Rodrigues[49], R. Gonzalez Suarez[164], S. Gonzalez-Sevilla[57], L. Goossens[37], B. Gorini[37], E. Gorini[71a,71b], A. Gorišek[95], T. C. Gosart[131], A. T. Goshaw[52], M. I. Gostkin[39], S. Goswami[124], C. A. Gottardo[37], S. A. Gotz[111], M. Gouighri[36b], V. Goumarre[49], A. G. Goussiou[141], N. Govender[34c], R. P. Grabarczyk[129], I. Grabowska-Bold[87a], K. Graham[35], E. Gramstad[128], S. Grancagnolo[71a,71b], C. M. Grant[1,138], P. M. Gravila[28f], F. G. Gravili[71a,71b], H. M. Gray[18a], M. Greco[71a,71b], M. J. Green[1], C. Grefe[25], A. S. Grefsrud[17], I. M. Gregor[49], K. T. Greif[162], P. Grenier[146], S. G. Grewe[112], A. A. Grillo[139], K. Grimm[32], S. Grinstein[13,t], J.-F. Grivaz[67], E. Gross[172], J. Grosse-Knetter[56],

J. C. Grundy[129], L. Guan[108], J. G. R. Guerrero Rojas[166], G. Guerrieri[37], R. Gugel[102], J. A. M. Guhit[108], A. Guida[19], E. Guilloton[170], S. Guindon[37], F. Guo[14,114c], J. Guo[63c], L. Guo[49], Y. Guo[108], R. Gupta[132], S. Gurbuz[25], S. S. Gurdasani[55], G. Gustavino[76a,76b], P. Gutierrez[123], L. F. Gutierrez Zagazeta[131], M. Gutsche[51], C. Gutschow[98], C. Gwenlan[129], C. B. Gwilliam[94], E. S. Haaland[128], A. Haas[120], M. Habedank[49], C. Haber[18a], H. K. Hadavand[8], A. Hadef[51], S. Hadzic[112], A. I. Hagan[93], J. J. Hahn[144], E. H. Haines[98], M. Haleem[169], J. Haley[124], J. J. Hall[142], G. D. Hallewell[104], L. Halser[20], K. Hamano[168], M. Hamer[25], G. N. Hamity[53], E. J. Hampshire[97], J. Han[63b], K. Han[63a], L. Han[114a], L. Han[63a], S. Han[18a], Y. F. Han[158], K. Hanagaki[85], M. Hance[139], D. A. Hangal[42], H. Hanif[145], M. D. Hank[131], J. B. Hansen[43], P. H. Hansen[43], D. Harada[57], T. Harenberg[174], S. Harkusha[38], M. L. Harris[105], Y. T. Harris[129], J. Harrison[13], N. M. Harrison[122], P. F. Harrison[170], N. M. Hartman[112], N. M. Hartmann[111], R. Z. Hasan[97,137], Y. Hasegawa[143], F. Haslbeck[129], S. Hassan[17], R. Hauser[109], C. M. Hawkes[21], R. J. Hawkings[37], Y. Hayashi[156], D. Hayden[109], C. Hayes[108], R. L. Hayes[117], C. P. Hays[129], J. M. Hays[96], H. S. Hayward[94], F. He[63a], M. He[14,114c], Y. He[49], Y. He[98], N. B. Heatley[96], V. Hedberg[100], A. L. Heggelund[128], N. D. Hehir[96,*], C. Heidegger[55], K. K. Heidegger[55], J. Heilman[35], S. Heim[49], T. Heim[18a], J. G. Heinlein[131], J. J. Heinrich[126], L. Heinrich[112,ab], J. Hejbal[134], A. Held[173], S. Hellesund[17], C. M. Helling[167], S. Hellman[48a,48b], R. C. W. Henderson[93], L. Henkelmann[33], A. M. Henriques Correia[37], H. Herde[100], Y. Hernández Jiménez[148], L. M. Herrmann[25], T. Herrmann[51], G. Herten[55], R. Hertenberger[111], L. Hervas[37], M. E. Hesping[102], N. P. Hessey[159a], M. Hidaoui[36b], N. Hidic[136], E. Hill[158], S. J. Hillier[21], J. R. Hinds[109], F. Hinterkeuser[25], M. Hirose[127], S. Hirose[160], D. Hirschbuehl[174], T. G. Hitchings[103], B. Hiti[95], J. Hobbs[148], R. Hobincu[28e], N. Hod[172], M. C. Hodgkinson[142], B. H. Hodkinson[129], A. Hoecker[37], D. D. Hofer[108], J. Hofer[49], T. Holm[25], M. Holzbock[37], L. B. A. H. Hommels[33], B. P. Honan[103], J. J. Hong[69], J. Hong[63c], T. M. Hong[132], B. H. Hooberman[165], W. H. Hopkins[6], M. C. Hoppesch[165], Y. Horii[113], S. Hou[151], A. S. Howard[95], J. Howarth[60], J. Hoya[6], M. Hrabovsky[125], A. Hrynevich[49], T. Hryn'ova[4], P. J. Hsu[66], S.-C. Hsu[141], T. Hsu[67], M. Hu[18a], Q. Hu[63a], S. Huang[65b], X. Huang[14,114c], Y. Huang[142], Y. Huang[102], Y. Huang[14], Z. Huang[103], Z. Hubacek[135], M. Huebner[25], F. Huegging[25], T. B. Huffman[129], C. A. Hugli[49], M. Huhtinen[37], S. K. Huiberts[17], R. Hulsken[106], N. Huseynov[12,h], J. Huston[109], J. Huth[62], R. Hyneman[146], G. Iacobucci[57], G. Iakovidis[30], L. Iconomidou-Fayard[67], J. P. Iddon[37], P. Iengo[73a,73b], R. Iguchi[156], Y. Iiyama[156], T. Iizawa[129], Y. Ikegami[85], N. Ilic[158], H. Imam[84c], M. Ince Lezki[57], T. Ingebretsen Carlson[48a,48b], J. M. Inglis[96], G. Introzzi[74a,74b], M. Iodice[78a], V. Ippolito[76a,76b], R. K. Irwin[94], M. Ishino[156], W. Islam[173], C. Issever[19,49], S. Istin[22a,ah], H. Ito[171], R. Iuppa[79a,79b], A. Ivina[172], J. M. Izen[46], V. Izzo[73a], P. Jacka[134], P. Jackson[1], C. S. Jagfeld[111], G. Jain[159a], P. Jain[49], K. Jakobs[55], T. Jakoubek[172], J. Jamieson[60], W. Jang[156], M. Javurkova[105], P. Jawahar[103], L. Jeanty[126], J. Jejelava[152a,z], P. Jenni[55,g], C. E. Jessiman[35], C. Jia[63b], J. Jia[148], X. Jia[14,114c], Z. Jia[114a], C. Jiang[53], S. Jiggins[49], J. Jimenez Pena[13], S. Jin[114a], A. Jinaru[28b], O. Jinnouchi[157], P. Johansson[142], K. A. Johns[7], J. W. Johnson[139], F. A. Jolly[49], D. M. Jones[149], E. Jones[49], K. S. Jones[8], P. Jones[33], R. W. L. Jones[93], T. J. Jones[94], H. L. Joos[56,37], R. Joshi[122], J. Jovicevic[16], X. Ju[18a], J. J. Junggeburth[105], T. Junkermann[64a], A. Juste Rozas[13,t], M. K. Juzek[88], S. Kabana[140e], A. Kaczmarska[88], M. Kado[112], H. Kagan[122], M. Kagan[146], A. Kahn[131], C. Kahra[102], T. Kaji[156], E. Kajomovitz[153], N. Kakati[172], I. Kalaitzidou[55], C. W. Kalderon[30], N. J. Kang[139], D. Kar[34g], K. Karava[129], M. J. Kareem[159b], E. Karentzos[55], O. Karkout[117], S. N. Karpov[39], Z. M. Karpova[39], V. Kartvelishvili[93], A. N. Karyukhin[38], E. Kasimi[155], J. Katzy[49], S. Kaur[35], K. Kawade[143], M. P. Kawale[123], C. Kawamoto[89], T. Kawamoto[63a], E. F. Kay[37], F. I. Kaya[161], S. Kazakos[109], V. F. Kazanin[38], Y. Ke[148], J. M. Keaveney[34a], R. Keeler[168], G. V. Kehris[62], J. S. Keller[35], A. S. Kelly[98], J. J. Kempster[149], P. D. Kennedy[102], O. Kepka[134], B. P. Kerridge[137], S. Kersten[174], B. P. Kerševan[95], L. Keszeghova[29a], S. Ketabchi Haghighat[158], R. A. Khan[132], A. Khanov[124], A. G. Kharlamov[38], T. Kharlamova[38], E. E. Khoda[141], M. Kholodenko[133a], T. J. Khoo[19], G. Khoriauli[169], J. Khubua[152b,*], Y. A. R. Khwaira[130], B. Kibirige[34g], D. Kim[6], D. W. Kim[48a,48b], Y. K. Kim[40], N. Kimura[98], M. K. Kingston[56], A. Kirchhoff[56], C. Kirfel[25], F. Kirfel[25], J. Kirk[137], A. E. Kiryunin[112], C. Kitsaki[10], O. Kivernyk[25], M. Klassen[161], C. Klein[35], L. Klein[169], M. H. Klein[45], S. B. Klein[57], U. Klein[94], P. Klimek[37], A. Klimentov[30], T. Klioutchnikova[37], P. Kluit[117], S. Kluth[112], E. Kneringer[80], T. M. Knight[158], A. Knue[50], D. Kobylianskii[172], S. F. Koch[129], M. Kocian[146], P. Kodyš[136], D. M. Koeck[126], P. T. Koenig[25], T. Koffas[35], O. Kolay[51], I. Koletsou[4], T. Komarek[88], K. Köneke[55], A. X. Y. Kong[1], T. Kono[121], N. Konstantinidis[98], P. Kontaxakis[57], B. Konya[100], R. Kopeliansky[42], S. Koperny[87a], K. Korcyl[88], K. Kordas[155,d], A. Korn[98], S. Korn[56], I. Korolkov[13], N. Korotkova[38], B. Kortman[117], O. Kortner[112],

S. Kortner[112], W. H. Kostecka[118], V. V. Kostyukhin[144], A. Kotsokechagia[37], A. Kotwal[52], A. Koulouris[37], A. Kourkoumeli-Charalampidi[74a,74b], C. Kourkoumelis[9], E. Kourlitis[112,ab], O. Kovanda[126], R. Kowalewski[168], W. Kozanecki[138], A. S. Kozhin[38], V. A. Kramarenko[38], G. Kramberger[95], P. Kramer[102], M. W. Krasny[130], A. Krasznahorkay[37], A. C. Kraus[118], J. W. Kraus[174], J. A. Kremer[49], T. Kresse[51], L. Kretschmann[174], J. Kretzschmar[94], K. Kreul[19], P. Krieger[158], M. Krivos[136], K. Krizka[21], K. Kroeninger[50], H. Kroha[112], J. Kroll[134], J. Kroll[131], K. S. Krowpman[109], U. Kruchonak[39], H. Krüger[25], N. Krumnack[82], M. C. Kruse[52], O. Kuchinskaia[38], S. Kuday[3a], S. Kuehn[37], R. Kuesters[55], T. Kuhl[49], V. Kukhtin[39], Y. Kulchitsky[38,a], S. Kuleshov[140d,140b], M. Kumar[34g], N. Kumari[49], P. Kumari[159b], A. Kupco[134], T. Kupfer[50], A. Kupich[38], O. Kuprash[55], H. Kurashige[86], L. L. Kurchaninov[159a], O. Kurdysh[67], Y. A. Kurochkin[38], A. Kurova[38], M. Kuze[157], A. K. Kvam[105], J. Kvita[125], T. Kwan[106], N. G. Kyriacou[108], L. A. O. Laatu[104], C. Lacasta[166], F. Lacava[76a,76b], H. Lacker[19], D. Lacour[130], N. N. Lad[98], E. Ladygin[39], A. Lafarge[41], B. Laforge[130], T. Lagouri[175], F. Z. Lahbabi[36a], S. Lai[56], J. E. Lambert[168], S. Lammers[69], W. Lampl[7], C. Lampoudis[155,d], G. Lamprinoudis[102], A. N. Lancaster[118], E. Lançon[30], U. Landgraf[55], M. P. J. Landon[96], V. S. Lang[55], O. K. B. Langrekken[128], A. J. Lankford[162], F. Lanni[37], K. Lantzsch[25], A. Lanza[74a], J. F. Laporte[138], T. Lari[72a], F. Lasagni Manghi[24b], M. Lassnig[37], V. Latonova[134], A. Laurier[153], S. D. Lawlor[142], Z. Lawrence[103], R. Lazaridou[170], M. Lazzaroni[72a,72b], B. Le[103], E. M. Le Boulicaut[52], L. T. Le Pottier[18a], B. Leban[24b,24a], A. Lebedev[82], M. LeBlanc[103], F. Ledroit-Guillon[61], S. C. Lee[151], S. Lee[48a,48b], T. F. Lee[94], L. L. Leeuw[34c], H. P. Lefebvre[97], M. Lefebvre[168], C. Leggett[18a], G. Lehmann Miotto[37], M. Leigh[57], W. A. Leight[105], W. Leinonen[116], A. Leisos[155,s], M. A. L. Leite[84c], C. E. Leitgeb[19], R. Leitner[136], K. J. C. Leney[45], T. Lenz[25], S. Leone[75a], C. Leonidopoulos[53], A. Leopold[147], R. Les[109], C. G. Lester[33], M. Levchenko[38], J. Levêque[4], L. J. Levinson[172], G. Levrini[24b,24a], M. P. Lewicki[88], C. Lewis[141], D. J. Lewis[4], A. Li[5], B. Li[63b], C. Li[63a], C.-Q. Li[112], H. Li[63a], H. Li[63b], H. Li[114a], H. Li[15], H. Li[63b], J. Li[63c], K. Li[141], L. Li[63c], M. Li[14,114c], S. Li[14,114c], S. Li[63d,63c], T. Li[5], X. Li[106], Z. Li[129], Z. Li[156], Z. Li[14,114c], Z. Li[63a], S. Liang[14,114c], Z. Liang[14], M. Liberatore[138], B. Liberti[77a], K. Lie[65c], J. Lieber Marin[84e], H. Lien[69], H. Lin[108], K. Lin[109], R. E. Lindley[7], J. H. Lindon[2], J. Ling[62], E. Lipeles[131], A. Lipniacka[17], A. Lister[167], J. D. Little[69], B. Liu[14], B. X. Liu[114b], D. Liu[63d,63c], E. H. L. Liu[21], J. B. Liu[63a], J. K. K. Liu[33], K. Liu[63d], K. Liu[63d,63c], M. Liu[63a], M. Y. Liu[63a], P. Liu[14], Q. Liu[63d,141,63c], X. Liu[63a], X. Liu[63b], Y. Liu[114b,114c], Y. L. Liu[63b], Y. W. Liu[63a], S. L. Lloyd[96], E. M. Lobodzinska[49], P. Loch[7], T. Lohse[19], K. Lohwasser[142], E. Loiacono[49], M. Lokajicek[134,*], J. D. Lomas[21], J. D. Long[165], I. Longarini[162], R. Longo[165], I. Lopez Paz[68], A. Lopez Solis[49], N. A. Lopez-canelas[7], N. Lorenzo Martinez[4], A. M. Lory[111], M. Losada[119a], G. Löschcke Centeno[149], O. Loseva[38], X. Lou[48a,48b], X. Lou[14,114c], A. Lounis[67], P. A. Love[93], G. Lu[14,114c], M. Lu[67], S. Lu[131], Y. J. Lu[66], H. J. Lubatti[141], C. Luci[76a,76b], F. L. Lucio Alves[114a], F. Luehring[69], I. Luise[148], O. Lukianchuk[67], O. Lundberg[147], B. Lund-Jensen[147,*], N. A. Luongo[6], M. S. Lutz[37], A. B. Lux[26], D. Lynn[30], R. Lysak[134], E. Lytken[100], V. Lyubushkin[39], T. Lyubushkina[39], M. M. Lyukova[148], M.Firdaus M. Soberi[53], H. Ma[30], K. Ma[63a], L. L. Ma[63b], W. Ma[63a], Y. Ma[124], J. C. MacDonald[102], P. C. Machado De Abreu Farias[84e], R. Madar[41], T. Madula[98], J. Maeda[86], T. Maeno[30], H. Maguire[142], V. Maiboroda[138], A. Maio[133a,133b,133d], K. Maj[87a], O. Majersky[49], S. Majewski[126], N. Makovec[67], V. Maksimovic[16], B. Malaescu[130], Pa. Malecki[88], V. P. Maleev[38], F. Malek[61,o], M. Mali[95], D. Malito[97], U. Mallik[81], S. Maltezos[10], S. Malyukov[39], J. Mamuzic[13], G. Mancini[54], M. N. Mancini[27], G. Manco[74a,74b], J. P. Mandalia[96], S. S. Mandarry[149], I. Mandić[95], L. Manhaes de Andrade Filho[84a], I. M. Maniatis[172], J. Manjarres Ramos[91], D. C. Mankad[172], A. Mann[111], S. Manzoni[37], L. Mao[63c], X. Mapekula[34c], A. Marantis[155,s], G. Marchiori[5], M. Marcisovsky[134], C. Marcon[72a], M. Marinescu[21], S. Marium[49], M. Marjanovic[123], A. Markhoos[55], M. Markovitch[67], E. J. Marshall[93], Z. Marshall[18a], S. Marti-Garcia[166], J. Martin[98], T. A. Martin[137], V. J. Martin[53], B. Martin dit Latour[17], L. Martinelli[76a,76b], M. Martinez[13,t], P. Martinez Agullo[166], V. I. Martinez Outschoorn[105], P. Martinez Suarez[13], S. Martin-Haugh[137], G. Martinovicova[136], V. S. Martoiu[28b], A. C. Martyniuk[98], A. Marzin[37], D. Mascione[79a,79b], L. Masetti[102], J. Masik[103], A. L. Maslennikov[38], P. Massarotti[73a,73b], P. Mastrandrea[75a,75b], A. Mastroberardino[44b,44a], T. Masubuchi[127], T. Mathisen[164], J. Matousek[136], J. Maurer[28b], A. J. Maury[67], B. Maček[95], D. A. Maximov[38], A. E. May[103], R. Mazini[151], I. Maznas[118], M. Mazza[109], S. M. Mazza[139], E. Mazzeo[72a,72b], C. Mc Ginn[30], J. P. Mc Gowan[168], S. P. Mc Kee[108], C. C. McCracken[167], E. F. McDonald[107], A. E. McDougall[117], J. A. Mcfayden[149], R. P. McGovern[131], R. P. Mckenzie[34g], T. C. Mclachlan[49], D. J. Mclaughlin[98], S. J. McMahon[137], C. M. Mcpartland[94], R. A. McPherson[168,x], S. Mehlhase[111], A. Mehta[94], D. Melini[166],

B. R. Mellado Garcia[34g], A. H. Melo[56], F. Meloni[49], A. M. Mendes Jacques Da Costa[103], H. Y. Meng[158], L. Meng[93], S. Menke[112], M. Mentink[37], E. Meoni[44b,44a], G. Mercado[118], S. Merianos[155], G. Merino Arevalo[e], C. Merlassino[70a,70c], L. Merola[73a,73b], C. Meroni[72a,72b], J. Metcalfe[6], A. S. Mete[6], E. Meuser[102], C. Meyer[69], J.-P. Meyer[138], R. P. Middleton[137], L. Mijović[53], G. Mikenberg[172], M. Mikestikova[134], M. Mikuž[95], H. Mildner[102], A. Milic[37], D. W. Miller[40], E. H. Miller[146], L. S. Miller[35], A. Milov[172], D. A. Milstead[48a,48b], T. Min[114a], A. A. Minaenko[38], I. A. Minashvili[152b], L. Mince[60], A. I. Mincer[120], B. Mindur[87a], M. Mineev[39], Y. Mino[89], L. M. Mir[13], M. Miralles Lopez[60], M. Mironova[18a], M. C. Missio[116], A. Mitra[170], V. A. Mitsou[166], Y. Mitsumori[113], O. Miu[158], P. S. Miyagawa[96], T. Mkrtchyan[64a], M. Mlinarevic[98], T. Mlinarevic[98], M. Mlynarikova[37], S. Mobius[20], P. Mogg[111], M. H. Mohamed Farook[115], A. F. Mohammed[14,114c], S. Mohapatra[42], G. Mokgatitswane[34g], L. Moleri[172], B. Mondal[144], S. Mondal[135], K. Mönig[49], E. Monnier[104], L. Monsonis Romero[166], J. Montejo Berlingen[13], A. Montella[48a,48b], M. Montella[122], F. Montereali[78a,78b], F. Monticelli[92], S. Monzani[70a,70c], A. Morancho Tarda[43], N. Morange[67], A. L. Moreira De Carvalho[49], M. Moreno Llácer[166], C. Moreno Martinez[57], P. Morettini[58b], S. Morgenstern[37], M. Morii[62], M. Morinaga[156], F. Morodei[76a,76b], L. Morvaj[37], P. Moschovakos[37], B. Moser[129], M. Mosidze[152b], T. Moskalets[45], P. Moskvitina[116], J. Moss[32,l], P. Moszkowicz[87a], A. Moussa[36d], E. J. W. Moyse[105], O. Mtintsilana[34g], S. Muanza[104], J. Mueller[132], D. Muenstermann[93], R. Müller[37], G. A. Mullier[164], A. J. Mullin[33], J. J. Mullin[131], D. P. Mungo[158], D. Munoz Perez[166], F. J. Munoz Sanchez[103], M. Murin[103], W. J. Murray[170,137], M. Muškinja[95], C. Mwewa[30], A. G. Myagkov[38,a], A. J. Myers[8], G. Myers[108], M. Myska[135], B. P. Nachman[18a], O. Nackenhorst[50], K. Nagai[129], K. Nagano[85], J. L. Nagle[30,af], E. Nagy[104], A. M. Nairz[37], Y. Nakahama[85], K. Nakamura[85], K. Nakkalil[5], H. Nanjo[127], E. A. Narayanan[115], I. Naryshkin[38], L. Nasella[72a,72b], M. Naseri[35], S. Nasri[119b], C. Nass[25], G. Navarro[23a], J. Navarro-Gonzalez[166], R. Nayak[154], A. Nayaz[19], P. Y. Nechaeva[38], S. Nechaeva[24b,24a], F. Nechansky[49], L. Nedic[129], T. J. Neep[21], A. Negri[74a,74b], M. Negrini[24b], C. Nellist[117], C. Nelson[106], K. Nelson[108], S. Nemecek[134], M. Nessi[37,i], M. S. Neubauer[165], F. Neuhaus[102], J. Neundorf[49], P. R. Newman[21], C. W. Ng[132], Y. W. Y. Ng[49], B. Ngair[119a], H. D. N. Nguyen[110], R. B. Nickerson[129], R. Nicolaidou[138], J. Nielsen[139], M. Niemeyer[56], J. Niermann[56], N. Nikiforou[37], V. Nikolaenko[38,a], I. Nikolic-Audit[130], K. Nikolopoulos[21], P. Nilsson[30], I. Ninca[49], G. Ninio[154], A. Nisati[76a], N. Nishu[2], R. Nisius[112], J.-E. Nitschke[51], E. K. Nkadimeng[34g], T. Nobe[156], T. Nommensen[150], M. B. Norfolk[142], B. J. Norman[35], M. Noury[36a], J. Novak[95], T. Novak[95], L. Novotny[135], R. Novotny[115], L. Nozka[125], K. Ntekas[162], N. M. J. Nunes De Moura Junior[84b], J. Ocariz[130], A. Ochi[86], I. Ochoa[133a], S. Oerdek[49,u], J. T. Offermann[40], A. Ogrodnik[136], A. Oh[103], C. C. Ohm[147], H. Oide[85], R. Oishi[156], M. L. Ojeda[49], Y. Okumura[156], L. F. Oleiro Seabra[133a], I. Oleksiyuk[57], S. A. Olivares Pino[140d], G. Oliveira Correa[13], D. Oliveira Damazio[30], J. L. Oliver[162], Ö. O. Öncel[55], A. P. O'Neill[20], A. Onofre[133a,133e], P. U. E. Onyisi[11], M. J. Oreglia[40], G. E. Orellana[92], D. Orestano[78a,78b], N. Orlando[13], R. S. Orr[158], L. M. Osojnak[131], R. Ospanov[63a], G. Otero y Garzon[31], H. Otono[90], P. S. Ott[64a], G. J. Ottino[18a], M. Ouchrif[36d], F. Ould-Saada[128], T. Ovsiannikova[141], M. Owen[60], R. E. Owen[137], V. E. Ozcan[22a], F. Ozturk[88], N. Ozturk[8], S. Ozturk[83], H. A. Pacey[129], A. Pacheco Pages[13], C. Padilla Aranda[13], G. Padovano[76a,76b], S. Pagan Griso[18a], G. Palacino[69], A. Palazzo[71a,71b], J. Pampel[25], J. Pan[175], T. Pan[65a], D. K. Panchal[11], C. E. Pandini[117], J. G. Panduro Vazquez[137], H. D. Pandya[1], H. Pang[15], P. Pani[49], G. Panizzo[70a,70c], L. Panwar[130], L. Paolozzi[57], S. Parajuli[165], A. Paramonov[6], C. Paraskevopoulos[54], D. Paredes Hernandez[65b], A. Pareti[74a,74b], K. R. Park[42], T. H. Park[158], M. A. Parker[33], F. Parodi[58b,58a], E. W. Parrish[118], V. A. Parrish[53], J. A. Parsons[42], U. Parzefall[55], B. Pascual Dias[110], L. Pascual Dominguez[101], E. Pasqualucci[76a], S. Passaggio[58b], F. Pastore[97], P. Patel[88], U. M. Patel[52], J. R. Pater[103], T. Pauly[37], C. I. Pazos[161], J. Pearkes[146], M. Pedersen[128], R. Pedro[133a], S. V. Peleganchuk[38], O. Penc[37], E. A. Pender[53], S. Peng[15], G. D. Penn[175], K. E. Penski[111], M. Penzin[38], B. S. Peralva[84d], A. P. Pereira Peixoto[141], L. Pereira Sanchez[146], D. V. Perepelitsa[30,af], G. Perera[105], E. Perez Codina[159a], M. Perganti[10], H. Pernegger[37], S. Perrella[76a,76b], O. Perrin[41], K. Peters[49], R. F. Y. Peters[103], B. A. Petersen[37], T. C. Petersen[43], E. Petit[104], V. Petousis[135], C. Petridou[155,d], T. Petru[136], A. Petrukhin[144], M. Pettee[18a], A. Petukhov[38], K. Petukhova[37], R. Pezoa[140f], L. Pezzotti[37], G. Pezzullo[175], T. M. Pham[173], T. Pham[107], P. W. Phillips[137], G. Piacquadio[148], E. Pianori[18a], F. Piazza[126], R. Piegaia[31], D. Pietreanu[28b], A. D. Pilkington[103], M. Pinamonti[70a,70c], J. L. Pinfold[2], B. C. Pinheiro Pereira[133a], A. E. Pinto Pinoargote[138,138], L. Pintucci[70a,70c], K. M. Piper[149], A. Pirttikoski[57], D. A. Pizzi[35], L. Pizzimento[65b], A. Pizzini[117], M.-A. Pleier[30], V. Plesko[136], E. Plotnikova[39], G. Poddar[96], R. Poettgen[100], L. Poggioli[130], I. Pokharel[56], S. Polacek[136], G. Polesello[74a], A. Poley[145,159a],

A. Polini[24b], C. S. Pollard[170], Z. B. Pollock[122], E. Pompa Pacchi[76a,76b], N. I. Pond[98], D. Ponomarenko[116], L. Pontecorvo[37], S. Popa[28a], G. A. Popeneciu[28d], A. Poreba[37], D. M. Portillo Quintero[159a], S. Pospisil[135], M. A. Postill[142], P. Postolache[28c], K. Potamianos[170], P. A. Potepa[87a], I. N. Potrap[39], C. J. Potter[33], H. Potti[150], J. Poveda[166], M. E. Pozo Astigarraga[37], A. Prades Ibanez[77a,77b], J. Pretel[168], D. Price[103], M. Primavera[71a], L. Primomo[70a,70c], M. A. Principe Martin[101], R. Privara[125], T. Procter[60], M. L. Proffitt[141], N. Proklova[131], K. Prokofiev[65c], G. Proto[112], J. Proudfoot[6], M. Przybycien[87a], W. W. Przygoda[87b], A. Psallidas[47], J. E. Puddefoot[142], D. Pudzha[55], D. Pyatiizbyantseva[38], J. Qian[108], D. Qichen[103], Y. Qin[13], T. Qiu[53], A. Quadt[56], M. Queitsch-Maitland[103], G. Quetant[57], R. P. Quinn[167], G. Rabanal Bolanos[62], D. Rafanoharana[55], F. Raffaeli[77a,77b], F. Ragusa[72a,72b], J. L. Rainbolt[40], J. A. Raine[57], S. Rajagopalan[30], E. Ramakoti[38], L. Rambelli[58b,58a], I. A. Ramirez-Berend[35], K. Ran[49,114c], D. S. Rankin[131], N. P. Rapheeha[34g], H. Rasheed[28b], V. Raskina[130], D. F. Rassloff[64a], A. Rastogi[18a], S. Rave[102], S. Ravera[58b,58a], B. Ravina[56], I. Ravinovich[172], M. Raymond[37], A. L. Read[128], N. P. Readioff[142], D. M. Rebuzzi[74a,74b], G. Redlinger[30], A. S. Reed[112], K. Reeves[27], J. A. Reidelsturz[174], D. Reikher[126], A. Rej[50], C. Rembser[37], M. Renda[28b], F. Renner[49], A. G. Rennie[162], A. L. Rescia[49], S. Resconi[72a], M. Ressegotti[58b,58a], S. Rettie[37], J. G. Reyes Rivera[109], E. Reynolds[18a], O. L. Rezanova[38], P. Reznicek[136], H. Riani[36d], N. Ribaric[93], E. Ricci[79a,79b], R. Richter[112], S. Richter[48a,48b], E. Richter-Was[87b], M. Ridel[130], S. Ridouani[36d], P. Rieck[120], P. Riedler[37], E. M. Riefel[48a,48b], J. O. Rieger[117], M. Rijssenbeek[148], M. Rimoldi[37], L. Rinaldi[24b,24a], P. Rincke[56,164], T. T. Rinn[30], M. P. Rinnagel[111], G. Ripellino[164], I. Riu[13], J. C. Rivera Vergara[168], F. Rizatdinova[124], E. Rizvi[96], B. R. Roberts[18a], S. S. Roberts[139], S. H. Robertson[106,x], D. Robinson[33], M. Robles Manzano[102], A. Robson[60], A. Rocchi[77a,77b], C. Roda[75a,75b], S. Rodriguez Bosca[37], Y. Rodriguez Garcia[23a], A. Rodriguez Rodriguez[55], A. M. Rodríguez Vera[118], S. Roe[37], J. T. Roemer[37], A. R. Roepe-Gier[139], O. Røhne[128], R. A. Rojas[105], C. P. A. Roland[130], J. Roloff[30], A. Romaniouk[38], E. Romano[74a,74b], M. Romano[24b], A. C. Romero Hernandez[165], N. Rompotis[94], L. Roos[130], S. Rosati[76a], B. J. Rosser[40], E. Rossi[129], E. Rossi[73a,73b], L. P. Rossi[62], L. Rossini[55], R. Rosten[122], M. Rotaru[28b], B. Rottler[55], C. Rougier[91], D. Rousseau[67], D. Rousso[49], A. Roy[165], S. Roy-Garand[158], A. Rozanov[104], Z. M. A. Rozario[60], Y. Rozen[153], A. Rubio Jimenez[166], A. J. Ruby[94], V. H. Ruelas Rivera[19], T. A. Ruggeri[1], A. Ruggiero[129], A. Ruiz-Martinez[166], A. Rummler[37], Z. Rurikova[55], N. A. Rusakovich[39], H. L. Russell[168], G. Russo[76a,76b], J. P. Rutherfoord[7], S. Rutherford Colmenares[33], M. Rybar[136], E. B. Rye[128], A. Ryzhov[45], J. A. Sabater Iglesias[57], H.F.-W. Sadrozinski[139], F. Safai Tehrani[76a], B. Safarzadeh Samani[137], S. Saha[1], M. Sahinsoy[83], A. Saibel[166], M. Saimpert[138], M. Saito[156], T. Saito[156], A. Sala[72a,72b], D. Salamani[37], A. Salnikov[146], J. Salt[166], A. Salvador Salas[154], D. Salvatore[44b,44a], F. Salvatore[149], A. Salzburger[37], D. Sammel[55], E. Sampson[93], D. Sampsonidis[155,d], D. Sampsonidou[126], J. Sánchez[166], V. Sanchez Sebastian[166], H. Sandaker[128], C. O. Sander[49], J. A. Sandesara[105], M. Sandhoff[174], C. Sandoval[23b], L. Sanfilippo[64a], D. P. C. Sankey[137], T. Sano[89], A. Sansoni[54], L. Santi[37,76b], C. Santoni[41], H. Santos[133a,133b], A. Santra[172], E. Sanzani[24b,24a], K. A. Saoucha[163], J. G. Saraiva[133a,133d], J. Sardain[7], O. Sasaki[85], K. Sato[160], C. Sauer[64b], E. Sauvan[4], P. Savard[158,ad], R. Sawada[156], C. Sawyer[137], L. Sawyer[99], C. Sbarra[24b], A. Sbrizzi[24b,24a], T. Scanlon[98], J. Schaarschmidt[141], U. Schäfer[102], A. C. Schaffer[67,45], D. Schaile[111], R. D. Schamberger[148], C. Scharf[19], M. M. Schefer[20], V. A. Schegelsky[38], D. Scheirich[136], M. Schernau[162], C. Scheulen[56], C. Schiavi[58b,58a], M. Schioppa[44b,44a], B. Schlag[146,n], K. E. Schleicher[55], S. Schlenker[37], J. Schmeing[174], M. A. Schmidt[174], K. Schmieden[102], C. Schmitt[102], N. Schmitt[102], S. Schmitt[49], L. Schoeffel[138], A. Schoening[64b], P. G. Scholer[35], E. Schopf[129], M. Schott[25], J. Schovancova[37], S. Schramm[57], T. Schroer[57], H.-C. Schultz-Coulon[64a], M. Schumacher[55], B. A. Schumm[139], Ph. Schune[138], A. J. Schuy[141], H. R. Schwartz[139], A. Schwartzman[146], T. A. Schwarz[108], Ph. Schwemling[138], R. Schwienhorst[109], F. G. Sciacca[20], A. Sciandra[30], G. Sciolla[27], F. Scuri[75a], C. D. Sebastiani[94], K. Sedlaczek[118], S. C. Seidel[115], A. Seiden[139], B. D. Seidlitz[42], C. Seitz[49], J. M. Seixas[84b], G. Sekhniaidze[73a], L. Selem[61], N. Semprini-Cesari[24b,24a], D. Sengupta[57], V. Senthilkumar[166], L. Serin[67], M. Sessa[77a,77b], H. Severini[123], F. Sforza[58b,58a], A. Sfyrla[57], Q. Sha[14], E. Shabalina[56], A. H. Shah[33], R. Shaheen[147], J. D. Shahinian[131], D. Shaked Renous[172], L. Y. Shan[14], M. Shapiro[18a], A. Sharma[37], A. S. Sharma[167], P. Sharma[81], P. B. Shatalov[38], K. Shaw[149], S. M. Shaw[103], Q. Shen[63c], D. J. Sheppard[145], P. Sherwood[98], L. Shi[98], X. Shi[14], S. Shimizu[85], C. O. Shimmin[175], J. D. Shinner[97], I. P. J. Shipsey[129], S. Shirabe[90], M. Shiyakova[39,v], M. J. Shochet[40], D. R. Shope[128], B. Shrestha[123], S. Shrestha[122,ag], M. J. Shroff[168], P. Sicho[134], A. M. Sickles[165], E. Sideras Haddad[34g], A. C. Sidley[117], A. Sidoti[24b], F. Siegert[51], Dj. Sijacki[16], F. Sili[92], J. M. Silva[53], I. Silva Ferreira[84b], M. V. Silva Oliveira[30],

S. B. Silverstein[48a], S. Simion[67], R. Simoniello[37], E. L. Simpson[103], H. Simpson[149], L. R. Simpson[108], N. D. Simpson[100], S. Simsek[83], S. Sindhu[56], P. Sinervo[158], S. Singh[158], S. Sinha[49], S. Sinha[103], M. Sioli[24b,24a], I. Siral[37], E. Sitnikova[49], J. Sjölin[48a,48b], A. Skaf[56], E. Skorda[21], P. Skubic[123], M. Slawinska[88], V. Smakhtin[172], B. H. Smart[137], S.Yu. Smirnov[38], Y. Smirnov[38], L. N. Smirnova[38,a], O. Smirnova[100], A. C. Smith[42], D. R. Smith[162], E. A. Smith[40], H. A. Smith[129], J. L. Smith[103], R. Smith[146], M. Smizanska[93], K. Smolek[135], A. A. Snesarev[38], S. R. Snider[158], H. L. Snoek[117], S. Snyder[30], R. Sobie[168,x], A. Soffer[154], C. A. Solans Sanchez[37], E.Yu. Soldatov[38], U. Soldevila[166], A. A. Solodkov[38], S. Solomon[27], A. Soloshenko[39], K. Solovieva[55], O. V. Solovyanov[41], P. Sommer[51], A. Sonay[13], W. Y. Song[159b], A. Sopczak[135], A. L. Sopio[98], F. Sopkova[29b], J. D. Sorenson[115], I. R. Sotarriva Alvarez[157], V. Sothilingam[64a], O. J. Soto Sandoval[140c,140b], S. Sottocornola[69], R. Soualah[163], Z. Soumaimi[36e], D. South[49], N. Soybelman[172], S. Spagnolo[71a,71b], M. Spalla[112], D. Sperlich[55], G. Spigo[37], B. Spisso[73a,73b], D. P. Spiteri[60], M. Spousta[136], E. J. Staats[35], R. Stamen[64a], A. Stampekis[21], M. Standke[25], E. Stanecka[88], W. Stanek-Maslouska[49], M. V. Stange[51], B. Stanislaus[18a], M. M. Stanitzki[49], B. Stapf[49], E. A. Starchenko[38], G. H. Stark[139], J. Stark[91], P. Staroba[134], P. Starovoitov[64a], S. Stärz[106], R. Staszewski[88], G. Stavropoulos[47], P. Steinberg[30], B. Stelzer[145,159a], H. J. Stelzer[132], O. Stelzer-Chilton[159a], H. Stenzel[59], T. J. Stevenson[149], G. A. Stewart[37], J. R. Stewart[124], M. C. Stockton[37], G. Stoicea[28b], M. Stolarski[133a], S. Stonjek[112], A. Straessner[51], J. Strandberg[147], S. Strandberg[48a,48b], M. Stratmann[174], M. Strauss[123], T. Strebler[104], P. Strizenec[29b], R. Ströhmer[169], D. M. Strom[126], R. Stroynowski[45], A. Strubig[48a,48b], S. A. Stucci[30], B. Stugu[17], J. Stupak[123], N. A. Styles[49], D. Su[146], S. Su[63a], W. Su[63d], X. Su[63a], D. Suchy[29a], K. Sugizaki[156], V. V. Sulin[38], M. J. Sullivan[94], D. M. S. Sultan[129], L. Sultanaliyeva[38], S. Sultansoy[3b], T. Sumida[89], S. Sun[173], O. Sunneborn Gudnadottir[164], N. Sur[104], M. R. Sutton[149], H. Suzuki[160], M. Svatos[134], M. Swiatlowski[159a], T. Swirski[169], I. Sykora[29a], M. Sykora[136], T. Sykora[136], D. Ta[102], K. Tackmann[49,u], A. Taffard[162], R. Tafirout[159a], J. S. Tafoya Vargas[67], Y. Takubo[85], M. Talby[104], A. A. Talyshev[38], K. C. Tam[65b], N. M. Tamir[154], A. Tanaka[156], J. Tanaka[156], R. Tanaka[67], M. Tanasini[148], Z. Tao[167], S. Tapia Araya[140f], S. Tapprogge[102], A. Tarek Abouelfadl Mohamed[109], S. Tarem[153], K. Tariq[14], G. Tarna[28b], G. F. Tartarelli[72a], M. J. Tartarin[91], P. Tas[136], M. Tasevsky[134], E. Tassi[44b,44a], A. C. Tate[165], G. Tateno[156], Y. Tayalati[36e,w], G. N. Taylor[107], W. Taylor[159b], R. Teixeira De Lima[146], P. Teixeira-Dias[97], J. J. Teoh[158], K. Terashi[156], J. Terron[101], S. Terzo[13], M. Testa[54], R. J. Teuscher[158,x], A. Thaler[80], O. Theiner[57], N. Themistokleous[53], T. Theveneaux-Pelzer[104], O. Thielmann[174], D. W. Thomas[97], J. P. Thomas[21], E. A. Thompson[18a], P. D. Thompson[21], E. Thomson[131], R. E. Thornberry[45], C. Tian[63a], Y. Tian[56], V. Tikhomirov[38,a], Yu.A. Tikhonov[38], S. Timoshenko[38], D. Timoshyn[136], E. X. L. Ting[1], P. Tipton[175], A. Tishelman-Charny[30], S. H. Tlou[34g], K. Todome[157], S. Todorova-Nova[136], S. Todt[51], L. Toffolin[70a,70c], M. Togawa[85], J. Tojo[90], S. Tokár[29a], K. Tokushuku[85], O. Toldaiev[69], M. Tomoto[85,113], L. Tompkins[146,n], K. W. Topolnicki[87b], E. Torrence[126], H. Torres[91], E. Torró Pastor[166], M. Toscani[31], C. Tosciri[40], M. Tost[11], D. R. Tovey[142], I. S. Trandafir[28b], T. Trefzger[169], A. Tricoli[30], I. M. Trigger[159a], S. Trincaz-Duvoid[130], D. A. Trischuk[27], B. Trocmé[61], A. Tropina[39], L. Truong[34c], M. Trzebinski[88], A. Trzupek[88], F. Tsai[148], M. Tsai[108], A. Tsiamis[155,d], P. V. Tsiareshka[38], S. Tsigaridas[159a], A. Tsirigotis[155,s], V. Tsiskaridze[158], E. G. Tskhadadze[152a], M. Tsopoulou[155], Y. Tsujikawa[89], I. I. Tsukerman[38], V. Tsulaia[18a], S. Tsuno[85], K. Tsuri[121], D. Tsybychev[148], Y. Tu[65b], A. Tudorache[28b], V. Tudorache[28b], A. N. Tuna[62], S. Turchikhin[58b,58a], I. Turk Cakir[3a], R. Turra[72a], T. Turtuvshin[39], P. M. Tuts[42], S. Tzamarias[155,d], E. Tzovara[102], F. Ukegawa[160], P. A. Ulloa Poblete[140c,140b], E. N. Umaka[30], G. Unal[37], A. Undrus[30], G. Unel[162], J. Urban[29b], P. Urrejola[140a], G. Usai[8], R. Ushioda[157], M. Usman[110], Z. Uysal[83], V. Vacek[135], B. Vachon[106], T. Vafeiadis[37], A. Vaitkus[98], C. Valderanis[111], E. Valdes Santurio[48a,48b], M. Valente[159a], S. Valentinetti[24b,24a], A. Valero[166], E. Valiente Moreno[166], A. Vallier[91], J. A. Valls Ferrer[166], D. R. Van Arneman[117], T. R. Van Daalen[141], A. Van Der Graaf[50], P. Van Gemmeren[6], M. Van Rijnbach[37], S. Van Stroud[98], I. Van Vulpen[117], P. Vana[136], M. Vanadia[77a,77b], W. Vandelli[37], E. R. Vandewall[124], D. Vannicola[154], L. Vannoli[54], R. Vari[76a], E. W. Varnes[7], C. Varni[18b], T. Varol[151], D. Varouchas[67], L. Variale[166], K. E. Varvell[150], M. E. Vasile[28b], L. Vaslin[85], G. A. Vasquez[168], A. Vasyukov[39], L. M. Vaughan[124], R. Vavricka[102], T. Vazquez Schroeder[37], J. Veatch[32], V. Vecchio[103], M. J. Veen[105], I. Veliscek[30], L. M. Veloce[158], F. Veloso[133a,133c], S. Veneziano[76a], A. Ventura[71a,71b], S. Ventura Gonzalez[138], A. Verbytskyi[112], M. Verducci[75a,75b], C. Vergis[96], M. Verissimo De Araujo[84b], W. Verkerke[117], J. C. Vermeulen[117], C. Vernieri[146], M. Vessella[105], M. C. Vetterli[145,ad], A. Vgenopoulos[102], N. Viaux Maira[140f], T. Vickey[142], O. E. Vickey Boeriu[142], G. H. A. Viehhauser[129], L. Vigani[64b],

M. Vigl[112], M. Villa[24b,24a], M. Villaplana Perez[166], E. M. Villhauer[53], E. Vilucchi[54], M. G. Vincter[35], A. Visibile[117], C. Vittori[37], I. Vivarelli[24b,24a], E. Voevodina[112], F. Vogel[111], J. C. Voigt[51], P. Vokac[135], Yu. Volkotrub[87b], J. Von Ahnen[49], E. Von Toerne[25], B. Vormwald[37], V. Vorobel[136], K. Vorobev[38], M. Vos[166], K. Voss[144], M. Vozak[117], L. Vozdecky[123], N. Vranjes[16], M. Vranjes Milosavljevic[16], M. Vreeswijk[117], N. K. Vu[63d,63c], R. Vuillermet[37], O. Vujinovic[102], I. Vukotic[40], S. Wada[160], C. Wagner[105], J. M. Wagner[18a], W. Wagner[174], S. Wahdan[174], H. Wahlberg[92], J. Walder[137], R. Walker[111], W. Walkowiak[144], A. Wall[131], E. J. Wallin[100], T. Wamorkar[6], A. Z. Wang[139], C. Wang[102], C. Wang[11], H. Wang[18a], J. Wang[65c], P. Wang[98], R. Wang[62], R. Wang[6], S. M. Wang[151], S. Wang[63b], S. Wang[14], T. Wang[63a], W. T. Wang[81], W. Wang[14], X. Wang[114a], X. Wang[165], X. Wang[63c], Y. Wang[63d], Y. Wang[114a], Y. Wang[63a], Z. Wang[108], Z. Wang[63d,52,63c], Z. Wang[108], A. Warburton[106], R. J. Ward[21], N. Warrack[60], S. Waterhouse[97], A. T. Watson[21], H. Watson[60], M. F. Watson[21], E. Watton[60,137], G. Watts[141], B. M. Waugh[98], J. M. Webb[55], C. Weber[30], H. A. Weber[19], M. S. Weber[20], S. M. Weber[64a], C. Wei[63a], Y. Wei[55], A. R. Weidberg[129], E. J. Weik[120], J. Weingarten[50], C. Weiser[55], C. J. Wells[49], T. Wenaus[30], B. Wendland[50], T. Wengler[37], N. S. Wenke[112], N. Wermes[25], M. Wessels[64a], A. M. Wharton[93], A. S. White[62], A. White[8], M. J. White[1], D. Whiteson[162], L. Wickremasinghe[127], W. Wiedenmann[173], M. Wielers[137], C. Wiglesworth[43], D. J. Wilbern[123], H. G. Wilkens[37], J. J. H. Wilkinson[33], D. M. Williams[42], H. H. Williams[131], S. Williams[33], S. Willocq[105], B. J. Wilson[103], P. J. Windischhofer[40], F. I. Winkel[31], F. Winklmeier[126], B. T. Winter[55], J. K. Winter[103], M. Wittgen[146], M. Wobisch[99], T. Wojtkowski[61], Z. Wolffs[117], J. Wollrath[162], M. W. Wolter[88], H. Wolters[133a,133c], M. C. Wong[139], E. L. Woodward[42], S. D. Worm[49], B. K. Wosiek[88], K. W. Woźniak[88], S. Wozniewski[56], K. Wraight[60], C. Wu[21], M. Wu[114b], M. Wu[116], S. L. Wu[173], X. Wu[57], Y. Wu[63a], Z. Wu[4], J. Wuerzinger[112,ab], T. R. Wyatt[103], B. M. Wynne[53], S. Xella[43], L. Xia[114a], M. Xia[15], M. Xie[63a], S. Xin[14,114c], A. Xiong[126], J. Xiong[18a], D. Xu[14], H. Xu[63a], L. Xu[63a], R. Xu[131], T. Xu[108], Y. Xu[15], Z. Xu[53], Z. Xu[114a], B. Yabsley[150], S. Yacoob[34a], Y. Yamaguchi[85], E. Yamashita[156], H. Yamauchi[160], T. Yamazaki[18a], Y. Yamazaki[86], J. Yan[63c], S. Yan[60], Z. Yan[105], H. J. Yang[63c,63d], H. T. Yang[63a], S. Yang[63a], T. Yang[65c], X. Yang[37], X. Yang[14], Y. Yang[45], Y. Yang[63a], Z. Yang[63a], W.-M. Yao[18a], H. Ye[114a], H. Ye[56], J. Ye[14], S. Ye[30], X. Ye[63a], Y. Yeh[98], I. Yeletskikh[39], B. K. Yeo[18b], M. R. Yexley[98], T. P. Yildirim[129], P. Yin[42], K. Yorita[171], S. Younas[28b], C. J. S. Young[37], C. Young[146], C. Yu[14,114c], Y. Yu[63a], J. Yuan[14,114c], M. Yuan[108], R. Yuan[63d,63c], L. Yue[98], M. Zaazoua[63a], B. Zabinski[88], E. Zaid[53], Z. K. Zak[88], T. Zakareishvili[166], S. Zambito[57], J. A. Zamora Saa[140d,140b], J. Zang[156], D. Zanzi[55], O. Zaplatilek[135], C. Zeitnitz[174], H. Zeng[14], J. C. Zeng[165], D. T. Zenger Jr[27], O. Zenin[38], T. Ženiš[29a], S. Zenz[96], S. Zerradi[36a], D. Zerwas[67], M. Zhai[14,114c], D. F. Zhang[142], J. Zhang[63b], J. Zhang[6], K. Zhang[14,114c], L. Zhang[63a], L. Zhang[114a], P. Zhang[14,114c], R. Zhang[173], S. Zhang[108], S. Zhang[91], T. Zhang[156], X. Zhang[63c], X. Zhang[63b], Y. Zhang[63c], Y. Zhang[98], Y. Zhang[114a], Z. Zhang[18a], Z. Zhang[63b], Z. Zhang[67], H. Zhao[141], T. Zhao[63b], Y. Zhao[139], Z. Zhao[63a], Z. Zhao[63a], A. Zhemchugov[39], J. Zheng[114a], K. Zheng[165], X. Zheng[63a], Z. Zheng[146], D. Zhong[165], B. Zhou[108], H. Zhou[7], N. Zhou[63c], Y. Zhou[15], Y. Zhou[114a], Y. Zhou[7], C. G. Zhu[63b], J. Zhu[108], X. Zhu[63d], Y. Zhu[63c], Y. Zhu[63a], X. Zhuang[14], K. Zhukov[69], N. I. Zimine[39], J. Zinsser[64b], M. Ziolkowski[144], L. Živković[16], A. Zoccoli[24b,24a], K. Zoch[62], T. G. Zorbas[142], O. Zormpa[47], W. Zou[42], L. Zwalinski[37]

[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
[9] Physics Department, National and Kapodistrian University of Athens, Athens, Greece
[10] Physics Department, National Technical University of Athens, Zografou, Greece
[11] Department of Physics, University of Texas at Austin, Austin, 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 Investigaciónes, 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)]National Institute for Research and Development of Isotopic and Molecular Technologies, Physics Department, Cluj-Napoca, Romania; [(e)]National University of Science and Technology Politechnica, 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, Los Angeles, 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), Quezon City, Philippines; [(e)]University of South Africa, Department of Physics, Pretoria, South Africa; [(f)]University of Zululand, KwaDlangezwa, South Africa; [(g)]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, Réseau Universitaire de Physique des Hautes Energies-Université Hassan II, 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] CERN, Geneva, Switzerland
[38] Affiliated with an institute 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] Physics Department, University of Texas at Dallas, Richardson, TX, USA
[47] National Centre for Scientific Research “Demokritos”, Agia Paraskevi, Greece
[48] [(a)]Department of Physics, Stockholm University, Stockholm, Sweden; [(b)]Oskar Klein Centre, Stockholm, Sweden
[49] Deutsches Elektronen-Synchrotron DESY, Hamburg and Zeuthen, Zeuthen, Germany
[50] Fakultät Physik , Technische Universität Dortmund, Dortmund, Germany
[51] Institut für Kern- und Teilchenphysik, Technische Universität Dresden, Dresden, Germany
[52] Department of Physics, Duke University, Durham, NC, USA
[53] SUPA-School of Physics and Astronomy, University of Edinburgh, Edinburgh, UK
[54] INFN e Laboratori Nazionali di Frascati, Frascati, Italy
[55] Physikalisches Institut, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany
[56] II. Physikalisches Institut, Georg-August-Universität Göttingen, Göttingen, Germany
[57] Département de Physique Nucléaire et Corpusculaire, Université de Genève, Genève, Switzerland
[58] [(a)]Dipartimento di Fisica, Università di Genova, Genova, Italy; [(b)]INFN Sezione di Genova, Genova, Italy
[59] II. Physikalisches Institut, Justus-Liebig-Universität Giessen, Giessen, Germany
[60] SUPA-School of Physics and Astronomy, University of Glasgow, Glasgow, UK
[61] LPSC, Université Grenoble Alpes, CNRS/IN2P3, Grenoble INP, Grenoble, France
[62] Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge, MA, USA
[63] [(a)]Department of Modern Physics and State Key Laboratory of Particle Detection and Electronics, University of Science and Technology of China, Hefei, China; [(b)]Institute of Frontier and Interdisciplinary Science and Key Laboratory of Particle Physics and Particle Irradiation (MOE), Shandong University, Qingdao, China; [(c)]School of Physics and Astronomy, Shanghai Jiao Tong University, Key Laboratory for Particle Astrophysics and Cosmology (MOE), SKLPPC, Shanghai, China; [(d)]Tsung-Dao Lee Institute, Shanghai, China; [(e)]School of Physics and Microelectronics, Zhengzhou University, China
[64] [(a)]Kirchhoff-Institut für Physik, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany; [(b)]Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
[65] [(a)]Department of Physics, Chinese University of Hong Kong, Shatin, N.T., Hong Kong; [(b)]Department of Physics, University of Hong Kong, Pok Fu Lam, 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
[66] Department of Physics, National Tsing Hua University, Hsinchu, Taiwan
[67] IJCLab, Université Paris-Saclay, CNRS/IN2P3, 91405 Orsay, France
[68] Centro Nacional de Microelectrónica (IMB-CNM-CSIC), Barcelona, Spain
[69] Department of Physics, Indiana University, Bloomington, IN, USA
[70] [(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
[71] [(a)]INFN Sezione di Lecce, Lecce, Italy; [(b)]Dipartimento di Matematica e Fisica, Università del Salento, Lecce, Italy
[72] [(a)]INFN Sezione di Milano, Segrate, Italy; [(b)]Dipartimento di Fisica, Università di Milano, Milano, Italy
[73] [(a)]INFN Sezione di Napoli, Napoli, Italy; [(b)]Dipartimento di Fisica, Università di Napoli, Napoli, Italy
[74] [(a)]INFN Sezione di Pavia, Pavia, Italy; [(b)]Dipartimento di Fisica, Università di Pavia, Pavia, Italy
[75] [(a)]INFN Sezione di Pisa, Pisa, Italy; [(b)]Dipartimento di Fisica E. Fermi, Università di Pisa, Pisa, Italy
[76] [(a)]INFN Sezione di Roma, Roma, Italy; [(b)]Dipartimento di Fisica, Sapienza Università di Roma, Roma, Italy
[77] [(a)]INFN Sezione di Roma Tor Vergata, Roma, Italy; [(b)]Dipartimento di Fisica, Università di Roma Tor Vergata, Roma, Italy
[78] [(a)]INFN Sezione di Roma Tre, Roma, Italy; [(b)]Dipartimento di Matematica e Fisica, Università Roma Tre, Roma, Italy
[79] [(a)]INFN-TIFPA, Povo, Italy; [(b)]Università degli Studi di Trento, Trento, Italy
[80] Department of Astro and Particle Physics, Universität Innsbruck, Innsbruck, Austria
[81] University of Iowa, Iowa City, IA, USA
[82] Department of Physics and Astronomy, Iowa State University, Ames, IA, USA
[83] Istinye University, Sariyer, Istanbul, Türkiye
[84] [(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
[85] KEK, High Energy Accelerator Research Organization, Tsukuba, Japan

[86] Graduate School of Science, Kobe University, Kobe, Japan
[87] [(a)] Faculty of Physics and Applied Computer Science, AGH University of Krakow, Krakow, Poland; [(b)] Marian Smoluchowski Institute of Physics, Jagiellonian University, Krakow, Poland
[88] Institute of Nuclear Physics Polish Academy of Sciences, Krakow, Poland
[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] School 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 Teorica 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, München, Germany
[112] Max-Planck-Institut für Physik (Werner-Heisenberg-Institut), München, 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, Shenzhen, 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, Netherlands
[117] Nikhef National Institute for Subatomic Physics and University of Amsterdam, Amsterdam, 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, Osaka University, 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, 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)]Universidad Andres Bello, Department of Physics, Santiago, Chile; [(e)]Instituto de Alta Investigación, Universidad de Tarapacá, Arica, Chile; [(f)]Departamento de Física, Universidad Técnica Federico Santa María, Valparaíso, Chile
[141] Department of Physics, University of Washington, Seattle, WA, USA
[142] Department of Physics and Astronomy, University of Sheffield, Sheffield, UK
[143] Department of Physics, Shinshu University, Nagano, Japan
[144] Department Physik, Universität Siegen, Siegen, Germany
[145] Department of Physics, Simon Fraser University, Burnaby, BC, Canada
[146] SLAC National Accelerator Laboratory, Stanford, CA, USA
[147] Department of Physics, Royal Institute of Technology, Stockholm, Sweden
[148] Departments of Physics and Astronomy, Stony Brook University, Stony Brook, NY, USA
[149] Department of Physics and Astronomy, University of Sussex, Brighton, UK
[150] School of Physics, University of Sydney, Sydney, Australia
[151] Institute of Physics, Academia Sinica, Taipei, Taiwan
[152] [(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
[153] Department of Physics, Technion, Israel Institute of Technology, Haifa, Israel
[154] Raymond and Beverly Sackler School of Physics and Astronomy, Tel Aviv University, Tel Aviv, Israel
[155] Department of Physics, Aristotle University of Thessaloniki, Thessaloniki, Greece
[156] International Center for Elementary Particle Physics and Department of Physics, University of Tokyo, Tokyo, Japan
[157] Department of Physics, Tokyo Institute of Technology, Tokyo, Japan
[158] Department of Physics, University of Toronto, Toronto, ON, Canada
[159] [(a)]TRIUMF, Vancouver, BC, Canada; [(b)]Department of Physics and Astronomy, York University, Toronto, ON, Canada
[160] Division of Physics and Tomonaga Center for the History of the Universe, Faculty of Pure and Applied Sciences, University of Tsukuba, Tsukuba, Japan
[161] Department of Physics and Astronomy, Tufts University, Medford, MA, USA
[162] Department of Physics and Astronomy, University of California Irvine, Irvine, CA, USA
[163] University of Sharjah, Sharjah, United Arab Emirates
[164] Department of Physics and Astronomy, University of Uppsala, Uppsala, Sweden
[165] Department of Physics, University of Illinois, Urbana, IL, USA
[166] Instituto de Física Corpuscular (IFIC), Centro Mixto Universidad de Valencia-CSIC, Valencia, Spain
[167] Department of Physics, University of British Columbia, Vancouver, BC, Canada
[168] Department of Physics and Astronomy, University of Victoria, Victoria, BC, Canada
[169] Fakultät für Physik und Astronomie, Julius-Maximilians-Universität Würzburg, Würzburg, Germany
[170] Department of Physics, University of Warwick, Coventry, UK
[171] Waseda University, Tokyo, Japan
[172] Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot, Israel
[173] Department of Physics, University of Wisconsin, Madison, WI, USA
[174] Fakultät für Mathematik und Naturwissenschaften, Fachgruppe Physik, Bergische Universität Wuppertal, Wuppertal, Germany
[175] Department of Physics, Yale University, New Haven, CT, USA

[a] Also Affiliated with an institute 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] Associated at Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas, Spain
[f] Also at Centro Studi e Ricerche Enrico Fermi, Roma, Italy
[g] Also at CERN, Geneva, Switzerland
[h] Also at CMD-AC UNEC Research Center, Azerbaijan State University of Economics (UNEC), Baku, Azerbaijan
[i] Also at Département de Physique Nucléaire et Corpusculaire, Université de Genève, Genève, Switzerland
[j] Also at Departament de Fisica de la Universitat Autonoma de Barcelona, Barcelona, Spain
[k] Also at Department of Financial and Management Engineering, University of the Aegean, Chios, Greece
[l] Also at Department of Physics, California State University, Sacramento, USA
[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 Hellenic Open University, Patras, Greece
[t] Also at Institucio Catalana de Recerca i Estudis Avancats, ICREA, Barcelona, Spain
[u] Also at Institut für Experimentalphysik, Universität Hamburg, Hamburg, Germany
[v] Also at Institute for Nuclear Research and Nuclear Energy (INRNE) of the Bulgarian Academy of Sciences, Sofia, Bulgaria
[w] Also at Institute of Applied Physics, Mohammed VI Polytechnic University, Ben Guerir, Morocco
[x] Also at Institute of Particle Physics (IPP), Toronto, Canada
[y] Also at Institute of Physics, Azerbaijan Academy of Sciences, Baku, Azerbaijan
[z] Also at Institute of Theoretical Physics, Ilia State University, Tbilisi, Georgia
[aa] Also at National Institute of Physics, University of the Philippines Diliman (Philippines), Quezon City, Philippines
[ab] Also at Technical University of Munich, Munich, Germany
[ac] Also at The Collaborative Innovation Center of Quantum Matter (CICQM), Beijing, China
[ad] Also at TRIUMF, Vancouver, BC, Canada
[ae] Also at Università di Napoli Parthenope, Napoli, Italy
[af] Also at University of Colorado Boulder, Department of Physics, Colorado, USA
[ag] Also at Washington College, Chestertown, MD, USA
[ah] Also at Yeditepe University, Physics Department, Istanbul, Türkiye
* Deceased