



Full Detector Simulation with Unprecedented Background Occupancy at a Muon Collider

Nazar Bartosik² · Paolo Andreetto¹ · Laura Buonincontri³ · Massimo Casarsa⁴ · Alessio Gianelle¹ · Simone Pagan Griso⁵ · Sergio Jindariani⁶ · Donatella Lucchesi³ · Federico Meloni⁷ · Nadia Pastrone² · Lorenzo Sestini¹

Received: 6 July 2021 / Accepted: 17 September 2021 / Published online: 5 October 2021
© The Author(s) 2021

Abstract

In recent years, a Muon collider has attracted a lot of interest in the high-energy physics community, thanks to its ability of achieving clean interaction signatures at multi-TeV collision energies in the most cost-effective way. Estimation of the physics potential of such an experiment must take into account the impact of beam-induced background on the detector performance, which has to be carefully evaluated using full detector simulation. Tracing of all the background particles entering the detector region in a single bunch crossing is out of reach for any realistic computing facility due to the unprecedented number of such particles. To make it feasible a number of optimisations have been applied to the detector simulation workflow. This contribution presents an overview of the main characteristics of the beam-induced background at a Muon collider, the detector technologies considered for the experiment and how they are taken into account to strongly reduce the number of irrelevant computations performed during the detector simulation. Special attention is dedicated to the optimisation of track reconstruction with the conformal tracking algorithm in this high-occupancy environment, which is the most computationally demanding part of event reconstruction.

Keywords Muon collider · High energy physics · Detector simulation · Scientific computing · Event reconstruction

Introduction

Nowadays, the high-energy physics community is actively discussing the optimal choice of a successor to the Large Hadron Collider for pushing further the boundaries of the field. Traditionally, two conceptually different options have been considered: linear e^+e^- colliders with clean final states, but a limited energy reach, and circular pp colliders with much higher energy reach, but very complex final states

due to the composite nature of the colliding hadrons. In the recent years, a Muon collider has been attracting increasing level of attention, thanks to its unique ability of providing clean final states by colliding elementary particles at very high energies. Much less synchrotron radiation due to the muon's larger mass allows to accelerate it in a circular collider to much higher energies than electrons, making a Muon collider the most energy-efficient facility at the centre-of-mass energies of 3 TeV and above [1].

An important role in the renewed interest to the Muon collider project was played by the physics performance study based on full detector simulation, which had demonstrated that despite the presence of intensive beam-induced background (BIB) it is possible to determine the Higgs-boson coupling to the b-quark with a precision comparable to projections at the CLIC e^+e^- collider [2]. This work presents an overview of the tools for detector simulation and object reconstruction used for the next iteration of the mentioned studies targeting $\sqrt{s} = 1.5$ TeV and higher, illustrating the major workflow optimisations necessary for maximising the computational efficiency of such simulations in the presence

✉ Nazar Bartosik
nazar.bartosik@to.infn.it

¹ INFN Sezione di Padova, Padua, Italy

² INFN Sezione di Torino, Turin, Italy

³ INFN Sezione di Padova and University of Padova, Padua, Italy

⁴ INFN Sezione di Trieste, Trieste, Italy

⁵ LBNL, Berkeley, USA

⁶ FNAL, Batavia, USA

⁷ DESY, Hamburg, Germany

of beam-induced background and high detector occupancy caused by it.

Beam-induced Background (BIB)

Initial studies of the potential accelerator design and corresponding BIB characterisation have been performed by the MAP collaboration [3] for 750 GeV beams split into bunches of $2 \cdot 10^{12}$ muons/bunch, colliding at the rate of 100 kHz to deliver the design luminosity of $\mathcal{L} = 1.25 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ [4]. A single muon beam with the mentioned parameters would bring $4.28 \cdot 10^5$ decays per meter of the accelerator lattice in a single pass. Interactions of the secondary and tertiary decay products with the machine elements and the machine-detector interface (MDI) create an intense flux of BIB particles reaching the detector volume, as shown in Fig. 1, causing radiation damage in the sensitive detector material,




Fig. 1 Illustration of the machine lattice and MDI used for the MARS15 simulation of the BIB from a 750 GeV muon beam. The cone-shaped shielding nozzles are shown in yellow. Several particle tracks with momenta > 1 GeV are shown by solid lines. The figure is reproduced from Ref. [5]

as well as increasing its occupancy and deteriorating the measurement resolution. Therefore a realistic evaluation of the detector performance requires an explicit simulation of the beam-decays evolution for any given design of the accelerator lattice and the MDI.

Extensive studies by the MAP collaboration using the MARS15 software [6] for simulating the ± 200 -m region around the interaction point (IP) resulted in a design of the cone-shaped tungsten shielding nozzles optimised for the 750 GeV muon beams, which strongly reduce the flux of BIB particles by about a factor ~ 500 . A simulated BIB sample is a list of stable particles entering the detector region in a single bunch crossing, which are collected at the surface of a box enclosing the detector and at the outer surface of the shielding nozzles. Further simulation of their interaction with the detector, can therefore, be performed in any other software where the detector geometry is implemented.

A new simulation setup based on FLUKA [7] and FlukaLineBuilder [8] software is currently at the final stage of development, which will be used for further optimisation of the MDI design coherently with detector-performance studies at higher beam energies [9]. All the results presented in this paper were obtained using the existing BIB sample simulated with MARS15 by the MAP collaboration for the $\sqrt{s} = 1.5$ TeV Muon collider [4].

Characteristics of BIB Particles

The BIB particles at the Muon collider have several characteristic properties affecting the detector differently from the BIB or pile-up collisions at e^+e^- or pp colliders. In Fig. 2 the main features of these particles are shown separately for the three dominant particle types: photons, electrons and neutrons.

The first outstanding feature is the extremely large number of particles (about $4 \cdot 10^8$ from the two muon beams)




Fig. 2 Main characteristics of BIB particles at the detector entry point: the momentum (left), the position along the beam line (middle) and the time of arrival with respect to the bunch crossing (right), shown separately for neutrons, photons and electrons

arriving to the detector in a single bunch crossing, making the full simulation of a single event a very challenging and computationally demanding task. The second important feature is that it primarily originates from the outer surface of the nozzles, with a sizeable distribution along the beam, in contrast to particles from typical pile-up collisions that originate from a very small region around the interaction point (IP). Finally, the distribution of arrival time shows that the majority of BIB particles arrive with a substantial delay (~ 1 -10 ns) with respect to the bunch crossing, depending on the particle type, making the timing information a crucial component of BIB suppression at the detector level.

Simulating and reconstructing such a large number of particles in millions of events necessary for a statistically significant physics analysis is nearly impossible with a reasonable amount of computing resources. Yet intelligently using the mentioned kinematic features of the BIB particles it is possible to dramatically reduce the computing resources needed for the full detector simulation of such events, as demonstrated in Section 5.

Simulation Software Stack

A typical simulation of a physics event can be split into the following conceptual stages:

1. generation of stable particles from the collision event and from the BIB;
2. simulation of their interaction with the passive and sensitive material in the detector, producing simulated hits in the sensitive volume: `SimHits`;
3. conversion of simulated hits to reconstructed hits (`RecHits`) according to the detector's response parameters, such as spatial granularity, resolution, electronic noise, dead time, response linearity, etc.;
4. reconstruction of higher-level objects, e.g. charged tracks, hadronic jets, secondary vertices, Particle Flow objects, etc.

Only the first stage, the generation of stable input particles, is performed by external software, whereas the rest of the simulation and reconstruction process is performed within a single `iLCSoft` framework [10] inherited from the CLIC experiment [11]. This framework was chosen because it has all the main components for full simulation of a lepton-collider experiment, is in active use and has a significant overlap with the unified Key4HEP framework [12] being developed for future colliders, to which these studies will migrate in the future.

The interaction of stable particles with the detector material is simulated by the GEANT4 [13] software that

is closely integrated into the `iLCSoft` framework. The three main components of the framework itself are the following:

- `LCIO` - providing a consistent event-data storage model for commonly used objects like `MCParticles`, simulated and reconstructed hits, `Tracks`, `Particle-Flow Objects`, etc. [14];
- `Marlin` - a modular framework with processors for isolated tasks, like hit digitisation, object reconstruction and higher-level analysis that are configured and chained into a flexible data-processing sequence by means of XML configuration files [15];
- `DD4hep` - efficient and highly flexible detector description toolkit that interfaces the unified detector model with GEANT4 and Marlin for detector simulation and object reconstruction respectively [16].

A number of packages of the framework have been modified or extended for the needs of Muon collider studies, and are maintained in a separate public repository [17]. Revisions of the whole software stack are centrally distributed through Docker and Singularity containers for an easy and coherent use by people performing the simulations across different institutes with independent computing infrastructures, but can also be installed manually.

All the results for computational performance reported in this paper were obtained by running the mentioned software on a machine with Intel Xeon CPU E5-2665 2.40GHz and 32GB of RAM.

Detector Geometry

Taking advantage of the extensive detector-design studies performed by the CLIC collaboration, its latest `CLICdet` geometry [18] was used as a starting point for the Muon collider detector model, `MuColl_v1`, schematically shown in Fig. 3. The main features of this detector include the full-silicon tracker comprising of the Vertex Detector with a double-layer sensor arrangement close to the IP followed by the inner and outer trackers, sampling electromagnetic (ECAL) and hadronic (HCAL) calorimeters, superconducting solenoid ($B = 3.57$ T) and muon detectors based on the RPC technology. Several modifications with respect to `CLICdet` have been implemented in its design for adapting it to the specific BIB conditions at the Muon collider:

- the MDI and the beampipe have been replaced with those designed by the MAP collaboration, including the cone-shaped shielding nozzles;
- inner openings in the endcap region of the tracking detectors, calorimeters and muon stations have been increased to physically fit the larger shielding nozzles inside;




Fig. 3 Layout of the top half of the MuColl_v1 detector geometry featuring the shielding nozzles shown in cyan (top) and a zoomed-in view at the Vertex Detector featuring the double-layer arrangement of Si sensors (bottom)

- the strength of the magnetic field has been reduced to 3.57 T for consistency with the magnetic field assumed in the MARS15 simulation of the BIB particles;
- layout of the vertex detector was optimised to lower the occupancy near the tips of the shielding nozzles, where most of the BIB particles exit into the detector region (see Fig. 4), and to provide more measurement layers close to the interaction point.




Fig. 4 Layout of the vertex detector in R-Z plane (red lines) with sensors arranged in double layers with 2mm spacing, overlaid on the position distribution of the BIB-particle entry point along the beam axis. Positions of the disks of the Vertex Detector are chosen to stay outside of the regions with the highest flux of BIB particles to reduce their occupancy

Optimisation of the Simulation Process

The process of simulating one full event at a Muon collider within the iLCSoft framework is schematically shown in Fig. 5, and comprises of simulating SimHits from the BIB and from the signal process using GEANT4, overlaying the BIB SimHits on top of the signal SimHits and then passing them through the digitisation processors to obtain the RecHits, which are then used for higher-level object reconstruction.

Considering that about $4 \cdot 10^8$ BIB particles have to be added on top of each signal event, performing their GEANT4 simulation in every event is not practical. Instead BIB particles from a single bunch crossing are simulated with GEANT4 only once to obtain the corresponding SimHits stored on a hard disk, which can be efficiently read and merged with SimHits from the signal process.

To avoid statistical biases, a finite number of bunch crossings is simulated, each having a randomised distribution of BIB particles in the azimuthal angle, and the set of BIB SimHits to be merged with a given signal event is picked randomly. A pool of 30 bunch crossings has been simulated for these studies, which is sufficient at the present development stage given the limited amount of computing resources and the focus on technical feasibility rather than on statistical precision of the obtained results.

Hit Digitisation Logics

All SimHits produced by GEANT4 are represented in the LCIO format by two conceptually different classes: SimCalorimeterHit - used for hits in the calorimeters and muon chambers, and SimTrackerHit - used for hits in the tracking detector. Both hit classes keep track of the timing information, because the hit time corrected for the time

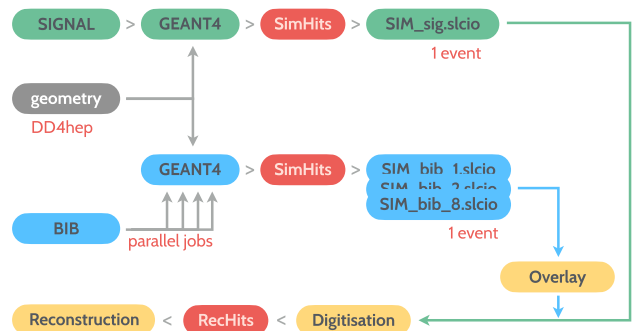


Fig. 5 Schematic diagram of the event simulation process within the iLCSoft framework. Interactions of stable particles with the detector material are simulated by GEANT4. Merging of the BIB and signal hits, followed by their digitisation and reconstruction are performed by individual Marlin processors

of flight of a photon from the IP to the corresponding sensor surface, referenced later in text as time of flight (TOF), is an important discriminating factor between the signal and BIB contributions, as demonstrated in Fig. 6, 7, 8. Yet the two hit classes are treated differently at the digitisation step. `SimCalorimeterHits` reflect the physical granularity of the detector, and a digitised hit is obtained by summing all contributions from `MCParticles` to the corresponding cell during a fixed readout time window. Instead `SimTrackerHits` are treated independently from each other, assuming no physical division of sensor planes into pixels or strips, and the finite spatial and time resolution effects are applied by a Gaussian smearing of their position and time.

A more advanced digitisation processor is being developed for the tracking sensors that takes into account the charge sharing between pixels, realistic hit-time reconstruction and pile-up effects. This more complex approach will unavoidably make the tracker-digitisation process more computationally demanding, and will also need an adjusted selection of input `BIB MCParticles` and `SimTrackerHits` relevant for the digitisation process. Therefore, some of the optimisation strategies described in the following will need to be revised in the future.

The main optimisation steps explored during the course of these studies are summarised in Table 1 together with the approximate effect on the main performance metrics. Detailed description of these and potential future optimisations to be studied is presented in the following subsections.

Simulation of BIB SimHits

Every particle in GEANT4 simulation is processed independently, which allows to easily parallelise this step into an




Fig. 6 Comparison of the TOF-corrected time distributions for hits from signal and BIB particles in the tracking detector, assuming single-hit time resolution of $\sigma_t = 30$ ps (60 ps) in the Vertex Detector (Inner/Outer Tracker). A narrow time window of $\pm 3\sigma_t$ (represented by the dashed lines) allows to reject most of the BIB hits

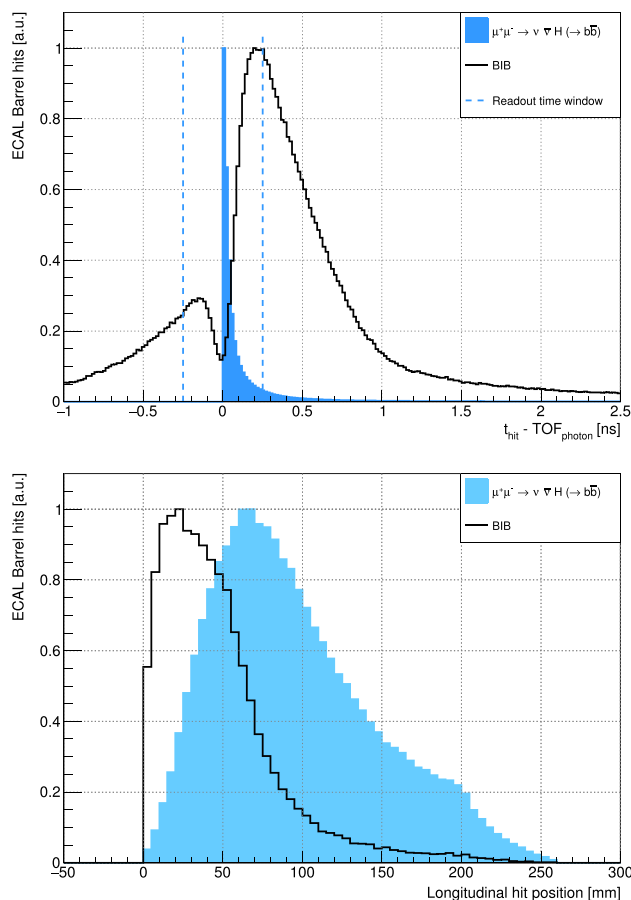


Fig. 7 Comparison of signal and BIB hit properties in the ECAL Barrel: TOF-corrected time (top) and longitudinal position along the barrel radius (bottom). Soft BIB particles have a wider time distribution due to the spread of their origin and time of flight, allowing to suppress their contribution with a narrow readout time window of ± 250 ps. Depth profile of the remaining hits can be used for further subtraction of the average BIB energy deposits




Fig. 8 Time distribution of simulated tracker and calorimeter hits corrected for the time of flight of a photon from the IP. The maximum hit time relevant for digitisation with realistic readout time windows is marked by the dashed line

Table 1 Summary of optimisation steps and their impact on the computing performance

Optimisation step	Primary metrics	
BIB tracking in GEANT4	CPU time	
All particles	3040 h	
Excluded late particles	480 h	
Excluded slow neutrons	200 h	
Lower-precision physics list	24 h	
BIB overlay	CPU time	RAM
ILCSOFT default	800 s	60 GB
Excluded BIB MCParticles	500 s	25 GB
Trimmed SimHit collections	70 s	8 GB
Track reconstruction	CPU time	
Time-based hit filtering	200 h	
Beamspace-pointing doublets (<i>loose filter</i>)	50 h	
Vertex-pointing doublets (<i>tight filter</i>)	0.03 h	

arbitrary number of threads N , each taking the corresponding $1/N$ fraction of the BIB particles as input. As mentioned in Section 2.1, one of the dominant BIB contributions are neutrons with energies as low as 20 MeV, which require GEANT4 to be run with a high-precision physics list `QGSP_BERT_HP` for accurate results. Simulation of a single bunch crossing in this configuration takes 3040 h (about 127 days) of CPU time, which requires massive parallelisation for obtaining results at a reasonable time scale.

In reality, a detector readout time window of up to 10 ns would be sufficient to detect all the signal particles from the collision event, including tracks from slow particles and hadronic showers, which can take a few nanoseconds to develop. Therefore, all the `SimHits` created after the 10 ns threshold would have no effect on the final result of the event simulation. Figure 8 shows the TOF-corrected time distribution of BIB `SimHits` in the calorimeters and in the tracking detector, which clearly demonstrates that a large fraction of HCAL hits are indeed created later than the mentioned 10 ns threshold and can be safely excluded. Correcting this value for the extra time of flight of signal particles from the IP, particles with initial arrival times greater than 25 ns can be safely skipped from the simulation process without affecting the `SimHits` in the readout time window of interest, effectively reducing the CPU time of simulating a single bunch crossing by a factor 6, down to 480 h on a single thread.

A large fraction of the remaining particles are low-energy neutrons with non-relativistic velocities, which are responsible for the long tail of the HCAL-hit time distribution in Fig. 8, which can reach several microseconds. Considering the direct relation between the momentum and velocity of a given massive particle, a lower threshold on the neutron kinetic energy $E_{\text{kin}} > 150$ MeV can be used to effectively exclude slow neutrons that create calorimeter hits too late, as




Fig. 9 Correlation between the kinetic energy of a neutron and the TOF-corrected time of HCAL hits created by it. Neutrons with energy below 150 MeV arrive to HCAL at $t \geq 10$ ns, which is outside of the readout time window of interest

shown in Fig. 9. This reduces the simulation time by another factor 3, down to 200 h.

Finally the remaining neutrons have sufficiently high momentum to be accurately described by the much faster-performing `QGSP_BERT` physics list, which brings the processing time down to 24 h. Thus, applying thresholds to the minimum momentum of neutrons and maximum arrival time of all particles that are simulated by GEANT4 allows to simulate a full bunch crossing in 1 day on a single thread, making a simulation of multiple bunch crossings for statistically independent event samples perfectly feasible.

Production-Level Optimisations

For a production-level simulation workflow, a more efficient approach would have to be implemented, taking advantage of the extremely high density of particles that have to be simulated. A very promising solution is the so called "Russian roulette" Monte Carlo sampling, which is successfully used by the CMS experiment since 2015 [19]. It reduces the GEANT4 simulation time by tracking only a fraction of randomly sampled individual particles and assigning a correspondingly greater weight to their energy deposits to keep the same total deposited energy. Given the sizeable dimensions of the ECAL and HCAL cells, on average 10–30 individual particles contribute to a single reconstructed hit, which normally are simulated individually by GEANT4 and then summed during the digitisation step. Considering that the main BIB contribution in the ECAL and HCAL are photons and neutrons, respectively, these two types of

particles are suitable candidates for the mentioned sampling technique.

Another potential approach could be based on machine learning, e.g. a generative adversarial network (GAN) that would generate already summed *SimHits* from the input list of *MCParticles*, bypassing the particle tracking with GEANT4 in every simulated event. Yet training such a network would require a large sample of bunch crossings to be simulated with GEANT4 in a conventional way, making it justified only if the sample of bunch crossings to be simulated is significantly larger than the sample needed for training such a GAN and when the detector geometry is final.

Efficient BIB Overlay

The first crucial optimisation step is to exclude BIB *MCParticles* and the corresponding *SimHit* relations from the samples to be overlaid, if tracing hits back to the *MCParticles* is not necessary in the simulation studies. This greatly reduces the disk storage and RAM occupied by a single bunch crossing, from about 60 GB to less than 25 GB.

As shown in Fig. 5, to obtain digitised *RecHits* representing the actual signals measured by a real detector, BIB *SimHits* have to be overlaid on top of the signal *SimHits* and processed altogether by the corresponding digitisers, individually for every subdetector type. The 10 ns time window used for selecting relevant particles for the GEANT4 simulation is too conservative for certain subdetectors that use much shorter readout time windows, like the Vertex Detector, in which only hits within the ± 90 ps range are kept, assuming the time resolution of $\sigma_t = 30$ ps. Considering that the majority of BIB *SimHits* are outside of this time window, most of the computation time in the digitiser processor would be spent on hits that will never be kept for analysis, therefore rejecting such hits as early as possible in the data processing chain would make the simulation more efficient. The same is true for calorimeter *SimHits*, a significant fraction of which still have contributions outside of the 10 ns window of interest.

The BIB *SimHits* are added to the event by the overlay processor, as shown in Fig. 5, which merges all the input BIB collections with the corresponding collections of the signal event, optionally filtering them according to *SimHit* acceptance time windows that can be configured individually for each collection. It should be noted that acceptance time windows for calorimeter *SimHits* during the Overlay process are identical to the ones for *RecHits* used during the digitisation step, whereas for tracker hits the windows at the *SimHit* level have to be at least $3\sigma_t$ wider than at the *RecHit* level in order to account for the hit migration due to their time smearing during digitisation.

With *SimHit* collection sizes of the order of 10^5 - 10^8 a simple hit-filtering operation takes several minutes per

event, due to both the CPU processing time and loading data to RAM from disk. Given that the acceptance timing windows are fixed for each collection and don't change from one event to the other, filtering of the *SimHit* collections can be performed only once, storing them in the so called trimmed BIB samples. Overlaying such trimmed samples takes about a factor 70 less time due to minimal CPU processing and much less data read from disk. This also substantially reduces the RAM usage to about 8 GB, allowing to run more simulation processes in parallel.

Track Reconstruction

After all the optimisations mentioned above, the remaining most-CPU-intensive task is the track reconstruction, which is based on the Conformal Tracking algorithm [20]. It is a geometry-agnostic and highly configurable implementation of the cellular automaton track-filtering method [21] that has two conceptual operations during the track-candidate search:

- *track building* - search for compatible hit sequences seeding from every input hit;
- *track extension* - propagate the existing track candidates to successive layers looking for additional hits consistent with the initial track trajectory.

At the Muon collider almost 100% of the track-reconstruction time is spent on track building in the vertex detector due to the very large number of hits and high hit density creating too many possible hit combinations. As the number of possible hit combinations grows exponentially with the number of hits, any further suppression of BIB hits can dramatically speed up the track reconstruction. A distinctive feature of BIB hits in the vertex detector is that they are primarily produced by soft electrons arriving from the tips of the shielding nozzles at shallow angles to the sensor




Fig. 10 Illustration of the doublet hit selection in the vertex detector, exploiting the fact that hits from the BIB particles are not aligned with the interaction point (vertex). Loose doublet selection with wide acceptance angles (magenta) have to be used to account for the spread of the beamspot and displaced vertices. Tight selection (green) can be used for selecting hit pairs aligned with the known vertex position, which allows to reject more BIB hits




Fig. 11 Multiplicity of hits reconstructed at each layer and disk of the Vertex Detector using the three levels of BIB-suppression criteria: only time-based hit selection (gray), time-based + loose doublet selection (magenta) and time-based + tight doublet selection (green)

surface, whereas, signal tracks originate from the IP at the centre of the detector.

The double-layer arrangement of Si sensors of the Vertex Detector allows to select hit doublets aligned in azimuthal angle with the IP, as illustrated in Fig. 10. The filtering of hits using this angular requirement was implemented in a separate Marlin processor with the angular acceptance individually configured for each double layer based on the spatial resolution of the sensors, the layer's distance from the IP and the assumed precision of the vertex position along the beam. To estimate the effectiveness of this BIB-suppression method the hit filtering was performed with two different doublet-selection criteria:

- *loose* - assuming no prior information about the IP position, which was smeared by a Gaussian distribution with $\sigma_z = 10$ mm along the beam, corresponding to the expected beamspot size at the $\sqrt{s} = 1.5$ TeV Muon collider;
- *tight* - assuming the exact knowledge of the IP position, which was fixed at the geometrical centre of the detector.

The effectiveness of the two doublet-selection criteria is reflected in the multiplicity of reconstructed hits shown in Fig. 11, which clearly demonstrates the dramatic suppression of BIB hits in the vertex detector when the position of the IP is known precisely. This reduction of the number of accepted hits brings the CPU time of reconstructing a single event from 2 days in the case of loose doublet selection to 3 minutes in the case of tight selection. Thus, it is evident that a standard approach of first reconstructing all tracks in the event, which are then used for the reconstruction of primary and secondary vertices is not the optimal solution in this case. A much more computationally efficient approach would be to first estimate approximate

positions of vertices with some faster algorithm, then filter hits using tighter doublet selection consistent with each vertex, and finally perform the full-scale track reconstruction using the nominal algorithm.

The baseline configuration of the conformal tracking algorithm, which starts by building tracks in the whole vertex detector and then extends the found candidates towards the inner and outer trackers, is extremely slow at the Muon collider due to the very large number of hits and corresponding large combinatorics in the vertex detector. Looking at the hit multiplicity in Fig. 11 it is clear that the two innermost barrel layers and most of the endcap disks are bad candidates for track seeding. Instead, seeding from the barrel layers 2–7 and the outermost disks allows to reconstruct tracks in the central and very forward region in less than a minute, which can be sufficient for reconstructing the vertex positions needed for tight doublet selection. Detailed performance studies of this approach are still in progress.

Conclusions

The unique properties of a multi-TeV Muon collider come at a cost of a high-intensity diffused beam-induced background. It has been shown that the unprecedentedly high flux of particles arriving to the detector in a single bunch crossing of a $\sqrt{s} = 1.5$ TeV Muon collider makes it practically impossible to perform a full detector simulation and event reconstruction with the simulation tools of the CLIC experiment if no dedicated workflow optimisations are applied. Yet the characteristic kinematic properties of these BIB particles allow to dramatically reduce the CPU time and memory usage of the simulation and reconstruction sequence by rearranging it in such a way that particles and hits that are irrelevant for the end results are excluded from the process at the earliest stage possible. The presented results also demonstrate how strongly the computation performance of these studies depends on the accelerator and MDI design, the choice of detector technologies and on the used reconstruction algorithms. Therefore, these and future simulation studies at high-energy Muon colliders require the different stages of the experiment design to be carried out in a highly organised and coherent manner.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are

included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Long KR, Lucchesi D, Palmer MA, Pastrone N, Schulte D, Shiltsev V (2021) Muon colliders to expand frontiers of particle physics. *Nat Phys*. <https://doi.org/10.1038/s41567-020-01130-x>
- Bartosik N, Bertolin A, Buonincontri L, Casarsa M, Collamati F, Ferrari A, Striganov S (2020) Detector and physics performance at a muon collider. *J Instrum* 15(05):P05001. <https://doi.org/10.1088/1748-0221/15/05/p05001>
- The Muon Accelerator Program (2021). <https://map.fnal.gov/>. Accessed 20 Feb 2021
- Mokhov NV, Striganov SI (2012) Detector background at muon colliders. arXiv preprint [arXiv:1204.6721](https://arxiv.org/abs/1204.6721)
- Di Benedetto V, Gatto C, Mazzacane A, Mokhov NV, Striganov SI, Terentiev NK (2018) A study of muon collider background rejection criteria in silicon vertex and tracker detectors. *JINST* 13:P09004. <https://doi.org/10.1088/1748-0221/13/09/p09004>
- Mokhov NV, James CC (2017) The MARS code system users guide version 15(2016). FNAL. <https://doi.org/10.2172/1462233>
- Battistoni G, Bauer J, Boehlen TT, Cerutti F, Chin MPW, Dos Santos Augusto R, Ferrari A, Ortega PG, Kozłowska W, Magro G, Mairani A, Parodi K, Sala PR, Schoofs P, Tessonnier T, Vlachoudis V (2016) The FLUKA code: an accurate simulation tool for particle therapy. *Front Oncol* 6:116
- Mereghetti A, Boccone V, Cerutti F, Versaci R, Vlachoudis V (2012) FLUKA simulations of the operational injection losses in TI8/IR8. *Conf Proc C1205201:2687*
- Collamati F, Curatolo C, Lucchesi D, Mereghetti A, Mokhov N, Palmer M, Sala P (2021) Advanced assessment of Beam Induced Background at a Muon Collider. arXiv preprint [arXiv:2105.09116](https://arxiv.org/abs/2105.09116)
- Ilc Soft (2021). <https://ilcsoft.desy.de/portal>. Accessed 20 Feb 2021
- Linssen L, Miyamoto A, Stanitzki M, Weerts H (2012) Physics and detectors at clic: clic conceptual design report. CERN-2012-003. arXiv preprint [arXiv:1202.5940](https://arxiv.org/abs/1202.5940)
- Key4hep Documentation (2021). <https://key4hep.github.io/key4hep-doc/>. Accessed 20 Feb 2021
- Allison J et al (2016) Recent developments in Geant4. *NIMA* (835):186–225. <https://doi.org/10.1016/j.nima.2016.06.125>
- Gaede F, Behnke T, Graf N, Johnson T (2003) LCIO: a persistency framework for linear collider simulation studies. *Tech. Rep. physics/0306114*, DESY, Hamburg. <https://cds.cern.ch/record/622146>. Accessed 3 July 2021
- Marlin Framework (2021). <https://github.com/iLCSoft/Marlin>. Accessed 20 Feb 2021
- Frank M, Gaede F, Petric M, Sailer A (2018) Aidasoft/dd4hep. <https://doi.org/10.5281/zenodo.592244>
- Muon Collider Repository (2021). <https://github.com/MuonColliderSoft>. Accessed 20 Feb 2021
- Alipour TN, Blaising JJ, Cure B, Dannheim D, Duarte RF, Elsener K, Gaddi A, Gerwig H, Green S, Greife C, Hynds D, Klempt W, Linssen L, Nikiforou N, Nurnberg AM, Marshall JS, Petric M, Redford S, Roloff PG, Sailer A, Sefkow F, Sicking E, Siegrist N, Simon FR, Simoniello R, Spannagel S, Sroka SK, Strom LR, Weber MA (2017). <https://cds.cern.ch/record/2254048>
- Lange DJ, Hildreth M, Ivantchenko VN (2015) Upgrades for the CMS simulation. *J Phys* 608:012056. <https://doi.org/10.1088/1742-6596/608/1/012056>
- Brondolin E, Leogrande E, Hynds D, Gaede F, Petrič M, Sailer A, Simoniello R (2020) Conformal tracking for all-silicon trackers at future electron–positron colliders. *NIMA* (956):163304. <https://doi.org/10.1016/j.nima.2019.163304>
- Glazov A, Kisel I, Konotopskaya E, Ososkov G (1993) Filtering tracks in discrete detectors using a cellular automaton. *NIMA* (329):1–2, 262–268. [https://doi.org/10.1016/0168-9002\(93\)90945-E](https://doi.org/10.1016/0168-9002(93)90945-E)

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