Physics Open 25 (2025) 100330

Contents lists available at ScienceDirect

Physics Open

journal homepage: www.elsevier.com/locate/physo

A simulation study to achieve PeV-scale neutrino–proton collisions via staggered beamlines

Tome Anticic

Ruder Boskovic Institute, Bijenicka 54, Zagreb, Croatia

ARTICLE INFO

Handling Editor: Gabriele Coniglio

Keywords:
Neutrinos
Accelerators
FCC
Muon cooling
PeV energy
Staggered beams

ABSTRACT

This paper introduces a novel accelerator-based method to achieve neutrino–proton interactions at energies from 10 to 100 PeV in a controlled laboratory environment. By employing a staggered configuration of a pion decay tunnel and a proton beamline, the method effectively counters neutrino beam divergence, forming repeated overlaps between narrow neutrino beams and the proton beam. Simulations, using parameters from the Future Circular Collider and including muon cooling to enhance event rates, demonstrate the potential to generate hundreds to thousands of events annually at these ultra-high energies.

While the proposed setup presents very significant technological challenges, it could potentially open a new energy range in neutrino physics, enabling precise measurements of neutrino-nucleon cross-sections needed for interpreting high-energy astrophysical neutrino signals and for testing fundamental physics.

Beyond the specific implementation explored here, the staggered beamline technique could be a basis for future novel accelerator designs. Integration with emerging technologies like muon colliders could overcome limitations in beam focusing, emittance, or target durability.

1. Introduction

Precise measurements of neutrino-hadron interaction cross-sections in the PeV energy range and above are crucial for interpreting astrophysical neutrino signals such as those originating from active galactic nuclei, gamma-ray bursts, or evaporating primordial black holes [1–4]. They can also provide insights into small-x QCD dynamics [5], probe the structure of matter at small distance scales and test the Standard Model in previously inaccessible energy domains [6,7].

Neutrino-hadron interactions have been studied over a wide range of energies, but never in a controlled experiment beyond the $\sim$ TeV scale. Conventional accelerator-based neutrino experiments produce neutrinos with energies up to a few hundred GeV via pion decay [8]. Recently, the FASERν experiment at the LHC observed neutrino interactions at energies of a few TeV [9]. Beyond this energy range, neutrino-hadron data come only from rare astrophysical events. The IceCube observatory, for example, has detected neutrinos with energies up to several PeV [10–12], while the KM3NeT neutrino telescope indirectly observed a neutrino of about 200 PeV [13].

Currently, however, there are no laboratory sources of neutrinos above the $\sim$ TeV scale, nor is there an experimental technique capable of generating neutrino-hadron collisions at the PeV energy scale or beyond. Simply increasing the energy of a traditional neutrino beamline is insufficient, as even a multi-TeV proton beam hitting a fixed target produces neutrinos only up to the hundreds-of-GeV scale. These neutrino beamlines pose additional significant challenges. While the charged pion parents can be focused to some extent with magnetic horns [14], the neutrinos themselves cannot be focused. As a result, the neutrino flux rapidly dilutes with distance from the source. This dispersion, combined with the need to compensate for the very low neutrino interaction probability, also necessitates massive, kiloton-scale detector volumes.

This paper explores a new collider-like approach that addresses these challenges, using relativistic kinematics and a novel staggered beamline configuration. One component of the novel approach is to collide a high-energy neutrino beam head-on with a high-energy proton beam, so that in the proton beam's rest frame, the neutrino's lab-frame energy is Lorentz boosted. This boost can, in principle, transform relatively low lab-frame neutrino energies into the multi-PeV energy range. The second key component is arranging the pion decay tunnel and proton beamline in a repeating staggered pattern of short segments, so that each segment of the proton beamline is exposed to a narrow neutrino flux from the short segment of pions just ahead of the protons, as illustrated in Fig. 1.

Simulations using parameters based on the Future Circular Collider (FCC-hh) [15] plans demonstrate the feasibility of this approach to yield detectable event rates and to possibly open a new frontier in high-energy physics.

E-mail address: anticic@irb.hr.

https://doi.org/10.1016/j.physo.2025.100330
Received 3 June 2025; Received in revised form 29 September 2025; Accepted 3 October 2025
Available online 8 October 2025

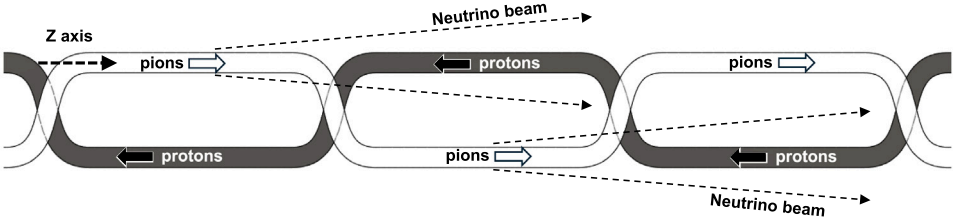

Fig. 1. Proposed staggered arrangement illustrating periodic realignment of pion decay tunnels and proton beamlines.

The proposed method thus uniquely would enable laboratory measurements of ν_μ-p cross-sections in the 10–100 PeV range. Small-x parton distribution functions, which are uncertain by 20–50% [16,17], could be measured with high precision.

Unlike IceCube/KM3NeT, which detect sporadic events/year from uncontrolled astrophysical sources with flavor ambiguities and atmospheric backgrounds, our approach could provide about a hundred controlled charged-current (CC) DIS events/year. This enables flavor-specific studies (primarily $\nu_\mu\bar{\nu}_\mu$) and direct tests of electroweak unification at PeV scales. Accurate lab measurements of PeV neutrino–proton cross-sections through our method would also significantly improve modeling of ultra-high-energy (UHE) cosmogenic neutrino generation and observation. This would be an important step for identifying sources and composition of UHE cosmic rays [18].

While this paper evaluates a specific FCC-inspired configuration, the staggered technique itself is an innovation that could be the basis of future extensions.

Sections 2 and 3 describe the methodology and simulation setup used to evaluate the idea. Section 4 presents the simulation results, demonstrating the gain in event rates due to the staggered arrangement, as well as the achievable neutrino energies. We discuss the implications of the results and the significant engineering challenges in Sections 5 and 6. Finally, we summarize our findings and the outlook for this concept in Section 7.

2. Methodology

2.1. Kinematics

An important component of our proposed method is the use of relativistic kinematics to achieve higher neutrino–proton energies, cross-section enhancement, and larger neutrino fluxes. Specifically, if a neutrino with energy $E_{\nu,\mathrm{lab}}$ strikes a high-energy proton of energy E_p head-on, the neutrino's energy in the proton's rest frame is boosted due to the relativistic Lorentz transformation and is approximately:

$$E_{\nu,\mathrm{rest}} \approx 2\gamma_p E_{\nu,\mathrm{lab}}, \tag{1}$$

where $\gamma_p = E_p/m_p$, and m_p is the proton mass. For example, a 100 GeV neutrino colliding with a 1 TeV proton results in $E_{\nu,\mathrm{rest}} \approx 200\,\mathrm{TeV}$, while for a 50 TeV proton, the same neutrino reaches $E_{\nu,\mathrm{rest}} \approx 10\,\mathrm{PeV}$.

If the neutrino beam itself extends to TeV-scale energies (which is possible from the decay of very high-energy pions), the effective collision energy can reach well into the hundreds of PeV. This increase in neutrino energy also directly translates into a larger neutrino–proton deep-inelastic scattering (DIS) cross-section [16], partially compensating for the low neutrino flux.

Another beneficial effect of high-energy pions is more collimated neutrino beams, since the neutrino angular spread in the lab frame with respect to the parent pion direction is $\approx 1/\gamma_\pi$, where $\gamma_\pi = E_\pi/m_\pi$. As a result, the proton beamline is exposed to larger neutrino fluxes than would be the case for the same number of pions at lower energies.

2.2. Staggered beamline configuration

The kinematic effects of higher energies are insufficient to offset the extremely low density of the proton beam. For example, at the LHC, the density of protons in the beamline is around 10^{14} protons per cubic meter, about 16 orders of magnitude lower than the nucleon density of solid matter (e.g., liquid argon has $\sim 8\times10^{29}$ nucleons/m^3).

This disparity is addressed to a large extent by our concept of a staggered alignment of the pion and proton beams (Fig. 1). In this method the two beams are initially aligned along the z-axis and move in opposite directions. Both beams are at regular intervals deflected transversely with respect to the z-axis by a small distance in opposite directions and then realigned to continue along their respective paths but with a small transverse offset. This way a series of short sections is formed where the neutrino beam from the pion decay tunnel overlaps with the proton beamline.

In effect,as illustrated by the neutrino beam spread represented by the dashed lines in Fig. 1, the proton beam is crossed by a narrow neutrino beam at regular intervals along the full length of the proton beamline, rather than only at its start.

This is in contrast to the standard single-pass setup, where the neutrino flux continuously decreases with distance as the beam diverges. The proposed method thus significantly increases the total neutrino–proton interaction probability over a long distance.

2.3. Neutrino production via pion decay

The pion beam is produced using established methods from neutrino factories [14], with a high-energy proton beam impinging on a target, and a magnetic horn selecting positively charged pions and collimating them. Focusing magnets along the pion decay tunnel further tighten the beam, and the decaying pions then produce a collimated muon neutrino beam in the forward direction.

2.4. Neutrino production via muon decay

In addition to ν_μ, each π^+ also decays to μ^+. These muons then decay via $\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu$, providing additional neutrinos to interact with the proton beam. The $\bar{\nu}_\mu$-proton CC DIS interaction produces a μ^-, just as a ν_μ-proton interaction produces a μ^+, with similar cross-sections in the PeV range [16]. In our simulation, we considered only the events involving muon final states, as the electron showers from electron neutrino–proton interactions are much more challenging to detect. However, future analysis techniques and technology advances might enable the detection of non-muon neutrino channels.

At the TeV energy scale, the 2.2 μs muon lifetime results in decay lengths of thousands of kilometers, potentially enabling ionization cooling. In this process (demonstrated by the MICE collaboration [19]), muons pass through absorbers to lose energy, reducing the transverse and longitudinal emittance, followed by re-acceleration to restore longitudinal momentum. This technique, planned to be used in future muon colliders [20], would result in very low emittance muon beams, and thus in much narrower neutrino beams.

2.5. Integration with a synchrotron accelerator

Our method could be implemented in a future accelerator complex by modifying a proton synchrotron design, such as the FCC-hh. In this

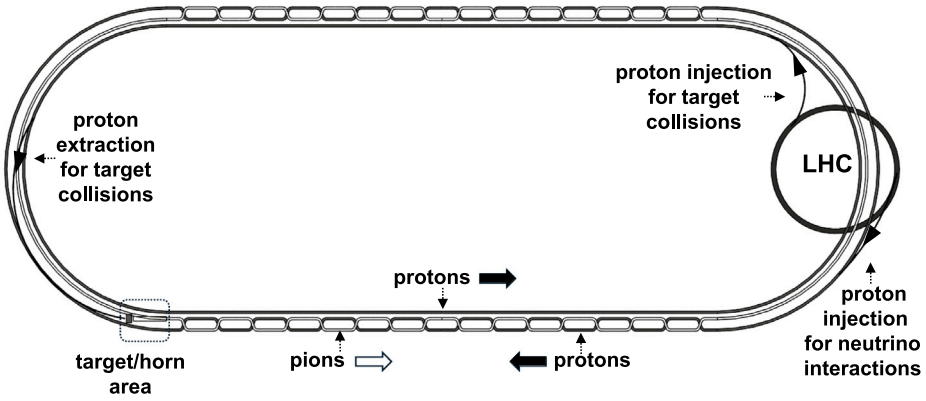

Fig. 2. Schematic of the synchrotron setup with two staggered straight sections between bending arcs. Pions can circulate between sections if their decay length allows, enhancing interaction yield. Protons from one proton beamline are diverted to a fixed target to be the source of pions. The other proton beamline is part of the staggered pattern, together with the pion decay tunnel (in white). The LHC is assumed to be the source of protons fed to the synchrotron.

setup, one proton beamline serves as the neutrino interaction target, while the counter-rotating proton beamline is the source of protons diverted to a fixed target to create pions, as in conventional neutrino beams [14]. Two long straight sections are placed between the bending halves (see Fig. 2), with the straight sections featuring the staggered arrangement of the pion decay tunnel and the proton beamline. Pions that do not decay in the first straight section are bent around the semicircle to repeat the process, further increasing the neutrino–proton interaction yield. This pattern repeats until most pions decay, utilizing the full pion decay length. The same method would also apply to both cooled and uncooled muons.

3. Simulation

To assess the viability of the proposed staggered pion beam-proton beam collider, we developed a simplified Monte Carlo simulation of the pion decay and neutrino–proton interaction process. The simulation incorporates several idealizations, focusing on the impact of the staggered design on event rates and using parametric models for beam intensities.

The highest planned accelerator energies and rates are those envisaged for the FCC-hh Phase II [15,21], and are thus used in our simulation. Hence, our simulation assumes a 50 TeV energy for both the protons in the proton beamline and the protons hitting the fixed target.

In practice, the high intensities of the proton beams in the considered scenarios would destroy a stationary solid target. An alternative could involve a flowing target, as discussed in Section 6.

The semicircle arc length of the accelerator setup in Fig. 2 is 50 km, while the lengths of the long straight staggered sections are, somewhat arbitrarily, set to 200 km each (larger lengths yield higher rates). The staggered section lengths are set to 100 m. The emittance growth of the pion beam was neglected, with the spot size set to RMS 1 mm throughout its path. This idealized focusing was chosen as an optimistic but potentially achievable target based on current technology projections (see Section 6 for further discussion). Unless otherwise stated, these parameters (see Table 1 are the default in the various scenarios investigated in the simulations.

Pion production: Pions are generated via Pythia8.3 [22] through ROOT [23], with 50 TeV protons colliding with a fixed proton target. This simulates an idealized high-Z target for multiplicity, with positive pions selected from the produced secondaries. The pion transverse p_T and longitudinal p_l momenta were kept unchanged throughout the pion decay tunnel. This p_T, peaking at the Fermi momentum of around ~ 0.3 GeV, sets the angular divergence for the pions once they decay.

Pion beam: The pions are propagated in small dz steps along the tunnel path, with the decay length modeled stochastically.

Staggered design: The pions initially traverse a 1 km section in the $+z$ direction, representing a standard accelerator neutrino beam horn and decay tunnel. The pion beam is then staggered at fixed intervals, with perpendicular offsets moving the pion beam to and from the z-axis. The proton beam, with a momentum in the $-z$ direction, is also staggered, with each segment offset with respect to the pion beam segments (Fig. 1). This staggered pattern is repeated throughout the two straight sections of the accelerator. Distances needed for the bending of the beams are at this point not taken into account, and is addressed in more detail in Section 6.

Proton beam: The proton beam is set to a fixed line density rather than being composed of individual bunches.

Neutrino production: When a pion decays, a neutrino is generated with an energy and angle relative to the pion direction according to isotropic two-body decay kinematics in the pion rest frame. The neutrino's direction with respect to the pion direction in the lab frame thus has an average angle of $1/\gamma_\pi$, with the pion itself already possessing an angle with respect to the horizontal direction of p_T/p_l. The neutrino is tracked in small dz steps along its path and is recorded if, at a distance z from the target, its transverse distance from the proton beamline is within 1 mm. This yields $\Phi_\nu(z, E'_\nu)$, the neutrino flux per proton on target (POT), where E'_ν is the neutrino energy in the proton rest frame.

Muon beam: As with the pions, the muon beam width was fixed to an RMS of 1 mm throughout the decay pipe. After a muon decays, a $\bar{\nu}_\mu$ is generated with an energy and angle relative to the muon direction according to isotropic three-body decay kinematics in the muon rest frame. The rest of the process is the same as for the neutrinos from pion decay. For the cooled muon beam, the beam width and transverse momenta were set to zero, so that the only contribution to the neutrino beam spread is the decay kinematics, with an average angle of $1/\gamma_\mu$, ($\gamma_\mu = E_\mu/m_\mu$) relative to the z-axis.

The total number of events per year is found by numerically integrating over the full path of the proton beamline:

$$N_{\text{events}}/\text{year} = (N_{\text{POT}}/\text{year}) \times \lambda_p \times \int \Phi_\nu(z, E'_\nu) \times \sigma_{\nu p}(E'_\nu)\, dz, \tag{2}$$

where $\sigma_{\nu p}(E'_\nu)$ is the CC neutrino–proton cross-section [16], $N_{\text{POT}}/\text{year}$ is the number of protons on target per year, and λ_p is the proton line density (protons per meter) in the beamline.

The key assumptions in this simulation include:

Table 1
Parameters used in the simulation.

Parameter	Value	Description
Proton energy (beam and on target)	50 TeV (default; varied 10–50 TeV)	FCC-hh baseline
Target material	Proton (PDG 2212)	Fixed target for pion production
Pion beam RMS spot size	1 mm	Focused throughout beamline
Staggered segment length	100 m	Default for overlap reset
Staggered straight sections length	200 km	Two sections (see Fig. 2)
Muon cooling	Yes/No	Zero emittance when cooled

- Negligible emittance growth in the pion beam, with a constant RMS spot size of 1 mm. This represents an optimistic projection of future focusing technology.
- Idealized pion production via Pythia on a proton target. Target details are omitted, and absorption effects or secondary interactions are not simulated.
- No inclusion of bending distances in staggered sections. This underestimates real world space requirements but does not affect the relative gain due to the staggering method.
- Fixed proton line density and no bunch structure, simplifying flux calculations.
- Zero transverse momentum and emittance for cooled muons. These parameters are expected to be very small for muon colliders compared to pion beams [20].
- Accelerator running times are assumed to be during the entire year.

These idealizations allow focus on the staggered configuration's benefits. Real implementations will require addressing them, as discussed in Section 6.

4. Results

The simulation results are presented per one POT/year and per one proton/m in the proton beamline. To find the expected number of events, both realistic (using planned FCC-hh parameters) and more optimistic or alternative scenarios for POT/year and protons/m are then used.

4.1. Viability of the staggered method

Fig. 3 illustrates the effectiveness of the staggered method compared to the standard way of generating accelerator neutrinos. Three scenarios are shown, all of which have an initial 1 km decay tunnel and a proton beamline in front of it:

Staggered configuration: The pion beam first traverses the initial 1 km decay tunnel. For pions generated from a 50 TeV proton beam hitting a fixed target about one third decay by this distance. This ratio drops to about 5% for pions with energies above 100 GeV - energies that are much more likely to result in high energy neutrinos. After the initial 1 km decay tunnel, both the proton beamline and the pion decay tunnel are staggered, in line with our proposed method. The pion beam width is set to a 1 mm RMS size throughout the decay tunnel.

Focused pion beam (non-staggered): Like above, the pion beam starts by remaining focused in the decay tunnel. However, at the end of the tunnel the beam disappears, simulating a beam dump just before the proton beamline.

Conventional (non-staggered, no focusing of pion beam): The pion beam is unfocused in the decay tunnel, and thus diverges due to the pion transverse momenta set at the target. As in the previous case, at the end of the tunnel the beam gets dumped. This represents a typical accelerator neutrino beam configuration at high energy.

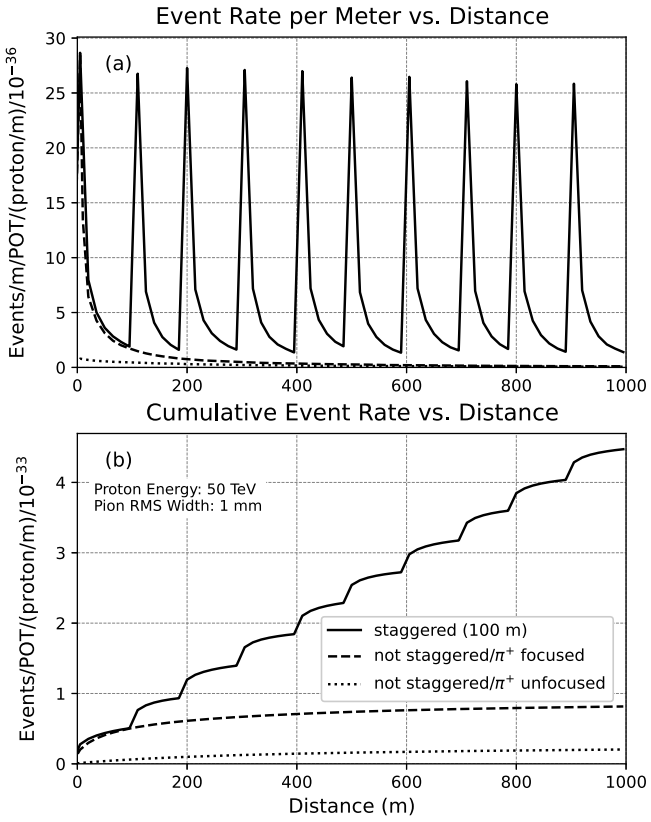

Fig. 3. (a) Neutrino–proton events per meter vs. distance along the proton beam for three cases: focused pions with 1 mm RMS width and staggered beams; focused pions with beam dump; unfocused pions with beam dump. (b) Cumulative events for the same cases.

We plot the results for only the first km of the proton beamline to more clearly visualize the effects of the staggered method. In the conventional case (non-staggered, non-focused), the event rate falls off roughly exponentially with distance. The focused but non-staggered case has a similar exponential drop, but with a several times larger cumulative number of events due to the narrow pion beam forming a more concentrated neutrino beam. For both cases the event rates drop to effectively zero after only a kilometer along the proton beamline.

In contrast to the conventional and focused non-staggered cases, the staggered configuration maintains a high interaction rate by periodically resetting the neutrino beam's transverse width the proton beam is exposed to. In fact, by extending the simulation up to 2000 km (at which point most of the pions have decayed), we find the total neutrino–proton interaction count to be $\approx 2 \times 10^3$ times higher than the focused non-staggered case, and $\approx 10^4$ times larger than the conventional case.

4.2. Sensitivity to beam parameters

Fig. 4 presents the dependence of the interaction rate on the staggered segment length on a log–log scale, demonstrating an inverse

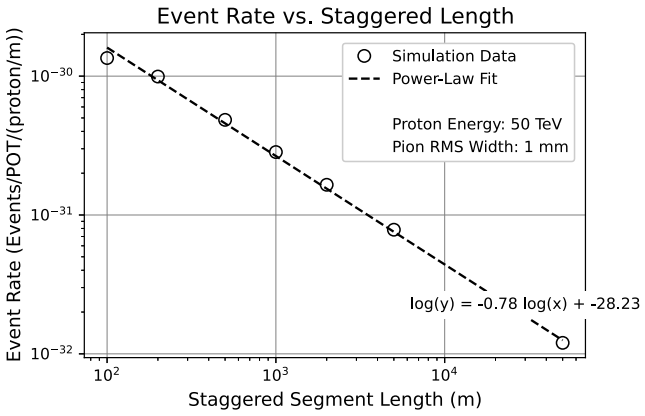

Fig. 4. Log–log plot of neutrino–proton event rates versus staggered section lengths. Open circles represent simulation data, and the dashed line shows a linear fit in log–log space with the equation displayed.

power relation. This correlation is expected, as shorter segment lengths result in more segments, and thus more places where protons interact with a neutrino beam reset to a narrow transverse width. Furthermore, the neutrino beam spread is smaller for shorter distances. However, the mechanical and engineering complexity of frequent beam bending and alignment limits how short the staggered segment lengths can be.

Fig. 5 shows the dependence of the interaction rate on the pion beam RMS spot size on a log–log plot. An inverse power relation is observed for beam spot sizes above 1 mm. Beam spot sizes below 1 mm result in only minimal additional gain due to the inherent neutrino decay kinematic angle of $1/\gamma_\pi$, and the pion spread at decay caused by its transverse momentum.

4.3. Neutrino energies

Fig. 6 shows the neutrino energy spectrum in the proton rest frame for neutrinos that cross the proton beamline. We find neutrino energies up to several hundred PeV, with the bulk of interactions in the tens of PeV range. The evolution of the average neutrino energy with distance along the proton beamline is illustrated in Fig. 7, rising from about 150 PeV at the start to more than 300 PeV at the end of the beamline. This increase occurs because higher-energy pions, due to time dilation, have longer decay lengths, and thus are able to interact at greater distances.

Higher proton energies are expected to lead to increased event rates, as well as to larger average neutrino energies in the proton rest frame. Indeed, Figs. 8 and 9 indicate that both approximately scale as the square of the proton energy.

4.4. Contributions from muon decay

Including $\bar{\nu}_\mu$ from muons doubles the event rate compared to events with ν_μ from pions, with cooled muons leading to a fourfold increase.

4.5. Final event rates

Table 2 provides the number of neutrino–proton events per year and the average neutrino energy in the proton beam rest frame, based on FCC-hh Phase II parameters [21]. Scenarios for three proton energies and proton beam currents are considered. Final rates are found using Eq. (2) and calculated for neutrinos decaying from pions and from cooled muons, with:

$$\lambda_p = I_{\text{beamline}}/(\text{e} \times \text{c}) \tag{3}$$

$$N_{\text{POT}}/\text{year} = I_{\text{beamline}}/(\text{e} \times \text{frequency}) \times (365{*}24/t_{\text{turnaround}}), \tag{4}$$

where I_{beamline} is the proton beam current, frequency is the revolution frequency at FCC (3000Hz), and $t_{\text{turnaround}}$ is the turnaround time for one spill in hours.

The first scenario uses a 50 TeV proton energy. The current of the proton beam is set to $I = 0.5\,\text{A}$, resulting in a proton line density $\lambda_p \approx 1 \times 10^{10}$ protons/m. The total number of protons in a beam is then $\approx 1 \times 10^{15}$. The minimum envisaged turnaround time at FCC-hh of $t_{\text{cycle}} \approx 1.8$ hours (including ramp-up, ramp-down, and fill time) is used to obtain $N_{\text{POT}}/\text{year} \approx 5 \times 10^{19}$. These calculations assume no downtime.

The FCC-hh has a current ceiling of 0.5 A due to large proton beam synchrotron radiation [21]. This limit prevents an increase in event rates through raising the current. However, synchrotron radiation power scales with energy as E^4, so, for example, a 20 TeV proton in an FCC-hh setup could instead have a 20 A current.

The turnaround time may decrease to approximately 1 h per cycle at 20 TeV (from 1.8 h at 50 TeV). This estimate is based on the FCC-hh fill time of 32 min and an assumed linear scaling with energy. This is the second scenario given in Table 2.

The third scenario is for a 10 TeV proton beam energy and again assumes that synchrotron radiation is the only limitation for the current and that ramping times scale linearly with energy.

5. Discussion

The baseline scenario, using 50 TeV protons, yields an average neutrino energy in the proton rest frame of approximately 250 PeV, with event rates of about 0.1 neutrino–proton interactions per year, increasing to around 0.5 events with muon cooling. Although these rates are too small for detection, they are within a factor of ten to a hundred of being observable, indicating the potential viability of the staggered method for high-energy neutrino–proton studies.

In the 20 TeV proton scenario, the event rate improves significantly to approximately 30 events per year without cooling and over 100 with muon cooling, albeit at a reduced average neutrino energy of about 45 PeV. This trade-off between energy and event yield suggests that a 20 TeV setup could result in detectable rates at energies not achievable in any other foreseeable accelerator experiment.

Further reducing the proton energy to 10 TeV results in thousands of events per year, at energies of 5 to 10 PeV. Thus while the 50 TeV scenario does not result in detectable rates, the lower energy scenarios indicate that a good balance between energy and event yield can be achieved, particularly with the implementation of muon cooling.

These results demonstrate the potential of the proposed method to create a new energy frontier in neutrino physics, where controlled, high-energy neutrino–proton interactions can be measured.

While the results show the potential of the proposed FCC-like setup, the staggered beamline could also be a basis for future advancements to enable higher neutrino–proton energies and rates. Adaptations might for example use muon collider accelerators or some of their specific techniques to address issues like beam spot size, emittance growth, or target vulnerability to high currents.

6. Feasibility

The proposed method is extremely challenging from an engineering and accelerator design perspective, and very idealized assumptions were made in the simulation. Below are some of the issues and possible means of addressing them:

Pion beam focusing: In a typical neutrino accelerator experiment pions emerge from the production target and magnetic horn focusing system with a certain initial beam spot size (on the order of millimeters to centimeters [24]) and an inherent divergence due to their transverse momentum (p_T). In our scenario, we assume the pion beam can be focused to a small transverse size of

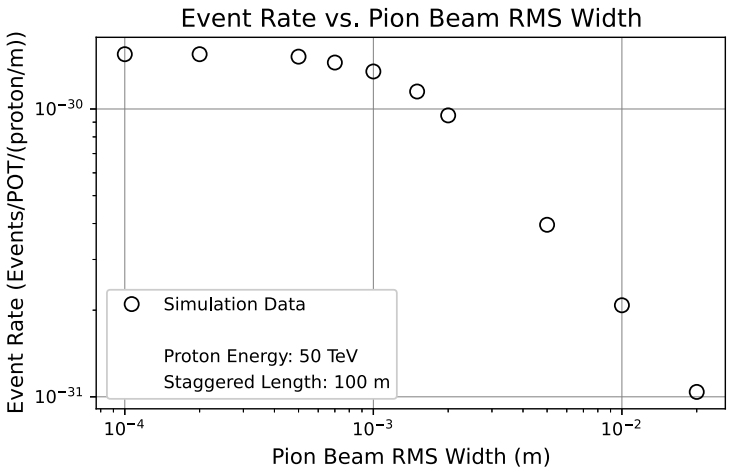

Fig. 5. Log–log plot of neutrino–proton event rates versus RMS width of the pion beam. Open circles represent simulation data.

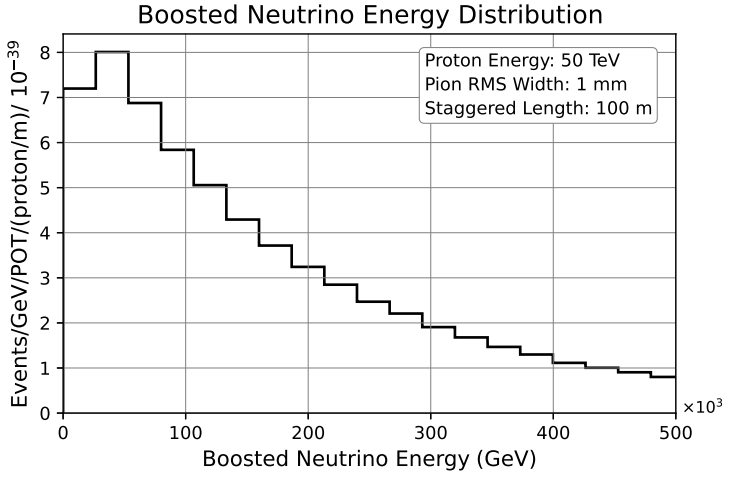

Fig. 6. Neutrino energy distribution in the proton beam rest frame.

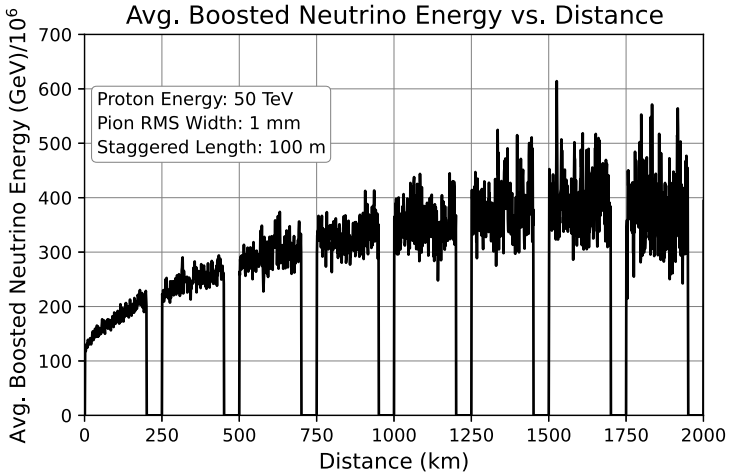

Fig. 7. Average neutrino energy in the proton rest frame as a function of the distance from the target (including passing through multiple 200 km straight sections and 50 km intervening arcs).

about 1 mm and maintained at roughly that size throughout the decay channel. This might be achievable with a strong focusing lattice along the decay tunnel. High-energy pions experience less fractional deflection from a given p_T kick (since p_T/p_l is

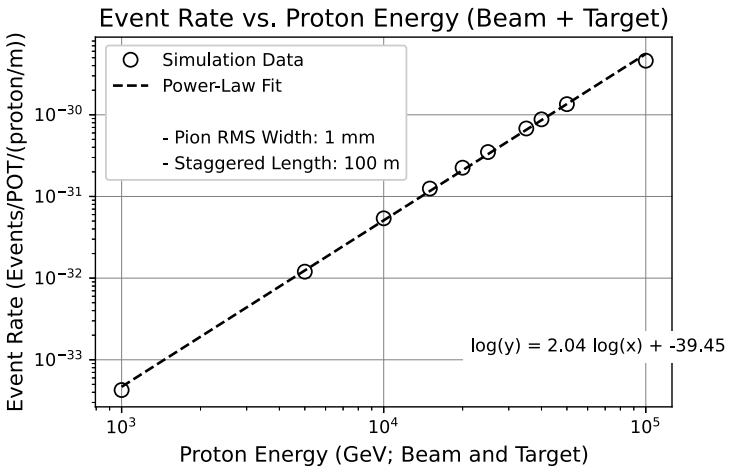

Fig. 8. Log–log plot of neutrino–proton event rates versus proton energy (for protons impinging upon the proton target to produce pions, and for protons in the beamline). Open circles represent simulation data, and the dashed line shows a linear fit in log–log space with the equation displayed.

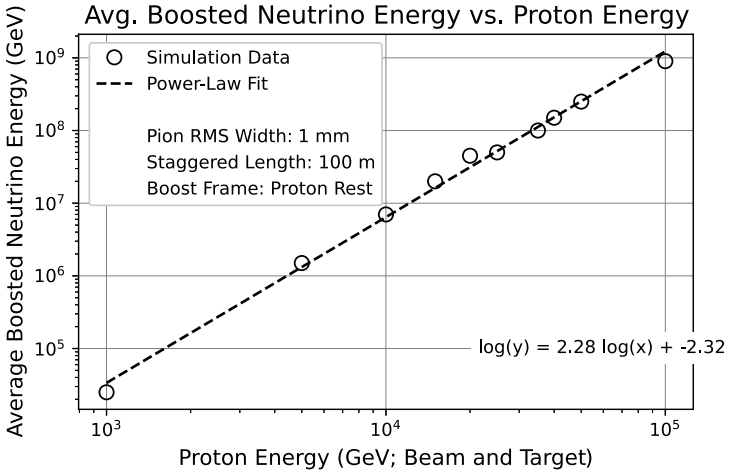

Fig. 9. Log–log plot of the average boosted neutrino energy versus proton energy (for protons impinging upon the proton target to produce pions, and for protons in the beamline). Open circles represent simulation data, and the dashed line shows a linear fit in log–log space with the equation displayed.

Table 2
Neutrino–proton events per year using parameters based on the FCC-hh. Proton line density and POT/year values are scaled as noted in the table headers.

Proton beam parameters					Interaction rates and energies		
Proton energy (TeV)	Proton beam current (A)	Turnaround time (h)	Proton line density ($\times 10^{10}$ protons/m)	POT/year ($\times 10^{18}$)	ν_μ–p ev/year	Cooled $\bar{\nu}_\mu$–p ev/year	(Anti)Neutrino E (pr. rest frame) (PeV)
50	0.5	1.8	1.0	5.1	0.1	0.5	250
20	20.	1.0	42.	370.	33.	130	45
10	312.	0.75	650.	7600.	1500.	6000	7

smaller at higher p_l), so focusing could be more manageable at multi-TeV momenta.

However, the large momentum spread would lead to chromatic effects, as lower-energy pions will diverge more between focusing elements than higher-energy ones. An achromatic design (perhaps pairing focusing and defocusing sections) might help, but beyond a certain momentum spread, it may not be possible to focus all pions equally well.

One possible mitigation is to keep only pions within a certain momentum range. Very low-energy pions (which would decay early anyway) and extremely high-energy pions (which might be too rigid to bend sufficiently in the staggered sections) could be removed. By reducing momentum spread, the focusing and bending challenges should become more tractable. Still, developing a lattice that can sustain a 1 mm beam over tens of kilometers will likely require new advances in magnet technology and beam dynamics.

A related issue concerns muon cooling. While the MICE collaboration demonstrated the principle of ionization cooling, it was performed at lower energies than those needed for our proposed concept.

Staggered bending sections: Each staggered segment end requires high-field bending magnets to offset the pion and proton beams

in one direction and then return them. Pion beams, in particular, are challenging due to their momentum spread. Achromatic bending cells, with dipoles and quadrupoles arranged so that different momenta have converging trajectories, are needed to prevent emittance growth. Short, distributed bending, with many bending magnets to form a smooth trajectory rather than a sharp kink, might alleviate synchrotron radiation. A similar approach might be used to at least partially address the bending around the 50 km synchrotron arc, placing numerous achromatic bending sections with high-field magnets and applying extensive chromatic corrections.

If the planned FCC 16 Tesla dipole magnets are used to bend the proton beam in a staggered fashion, then for a 20 TeV proton a likely sufficiently large offset of about 0.3 m occurs in a 100 m beamline length ($= 0.3 * B * length^2/8 * momentum$). The pion beam, due to lower energy, would need fewer such dipoles to achieve the same offset in the same distance. The remaining length could be occupied by quadrupole and other focusing magnets, as well as chicanes for the beams to not cross each other. As mentioned in Section 3, this distance is not accounted for in the current simulation.

Target and energies: The high intensities of the proton beams in the considered scenarios would destroy a stationary solid target. To address this, a possible solution could involve a moving or flowing target, such as liquid mercury [25]. The flow of liquid mercury would dissipate the extreme heat generated by the beam and lessen the damage from pressure waves caused by the beam pulses. Currents as high as those required for lower proton energies (e.g., 312 A at 10 TeV) are prohibitively large and unlikely to be feasible with near-future technology.

Detection: Detecting neutrino–proton events and determining their kinematical properties in this setup is challenging, with interactions occurring along the full length of the staggered parts of the proton beamline.

Several properties of the CC DIS neutrino proton events point to a possible detection method. The outgoing muons have energies in the TeV range, as well as very small angles relative to the proton beam. The DIS events also produce narrow high energy jets in the opposite direction.

These properties could be exploited by placing muon detector planes at fixed intervals around the proton beamline. The detector planes would be circular, surrounding the central magnet yoke, and providing 2D information on muon track positions. A torroidal magnetic field placed between sets of detector planes would provide momentum information through measurements of the sagita. The detector planes would also measure the angular spread and multiplicity of the jet.

Fig. 10 illustrates the proposed detector setup.

Preliminary Pythia simulations indicate the mean angle of the outgoing muon to be about 0.15 rad, with about 80% falling within 0.25 rad. Thus a 40 m separation between sets of detector planes with a 10 m diameter would enable the detection of most muons, while a 20 m separation would provide a significantly more accurate momentum determination through the long arm between the two sets of detector planes. The large number of tracks emerging from the collision should also enable the position determination of the interaction point.

The main background source are muons decaying from the pions in the pion decay pipe. They should be easily identified in our proposed setup as they are isolated, and have no accompanying jet in the opposite side. Additionally, these backgrounds can be rejected by requiring that the reconstructed interaction point (IP) lies within the proton beamline, and not the pion decay pipe. Multiple scattering (MS) should not be a limiting factor for beamline discrimination. For a baseline of about 100 m, an iron yoke of 0.5 m, and a detector plane radius of 10 m, the MS smearing [26] of the IP position is $\approx$ 10 cm in the z direction (compared to 100 m staggered length), and $\approx$ 1 cm radially (compared to 30 cm separation between beamlines).

We are investigating for an upcoming publication the efficiency of the muon and jet reconstruction, accuracy in determining the kinematical properties of the event, as well as background rejection efficiency. Simulations will explore, among other, the effects of the dimensions of the circular detector planes, their separation, accuracy of the position determination, strength of the magnetic field, multiple scattering effects. The type of detector best suited for the purpose will depend on the required accuracy - if it is of the order of 100 μm or below, silicon pixel or micromegas could be a possible solution. A full reconstruction of jets requires high-granularity calorimeters. However, we will also simulate how much information of the jet can be determined by the charged track multiplicities and their angular spread alone, to determine if position detectors are sufficient to extract the relevant kinematic data of the event.

Cost and complexity: The scale of the proposed setup, with hundreds of kilometers of staggered tunnels, large and numerous detectors, likely thousands of magnets, a high-power target, and the requirement for very large currents, is beyond any current accelerator project and would depend on advances in accelerator technology and significant funding and energy commitments. This type of facility would combine elements of fixed-target neutrino production, which are known at moderate scales (meters to 100 m decay tunnels), with collider-scale beams and machine lengths (tens of kilometers), and with the need for extreme precision alignment (mm-level over km distances). Proton and muon colliders would likely need to be combined. Each aspect has been separately achieved to some extent, but not all together. Addressing these challenges would require a multi-stage R&D program, starting with scaled-down experiments at existing facilities to test the staggered beamline concept.

7. Conclusion

This paper presents a novel concept for a neutrino–proton collider experiment that could achieve interaction energies on the order of tens of PeV, significantly beyond the range of any existing or planned accelerator neutrino experiment.

Using a staggered beamline configuration and advanced beam-handling techniques such as muon cooling, the proposed method has the potential to reach the ultra-high-energy regime of astrophysical neutrinos. Simulations based on the FCC-hh parameters demonstrate that this approach could yield observable event rates in this energy range, with the potential for hundreds to thousands of interactions per year. Combining elements of a muon collider, in particular, could quadruple event rates.

Such a capability would provide measurements needed for models used in astrophysical neutrino observatories like IceCube and for exploring physics beyond the Standard Model at ultra-high energies. It would essentially extend the energy range of neutrino–proton fixed-target experiments by many orders of magnitude. Energies inaccessible elsewhere could be accessed in a controlled manner with rates enabling PDF fits.

However, significant engineering challenges need to be addressed. Several requirements, such as the ability to maintain a narrow pion beam through hundreds of kilometers, are far beyond the current state of the art. Scaled down tests are needed before deciding whether the proposed method, or its possible variations, could become a viable technique to generate neutrino–proton events at the PeV scale and beyond.

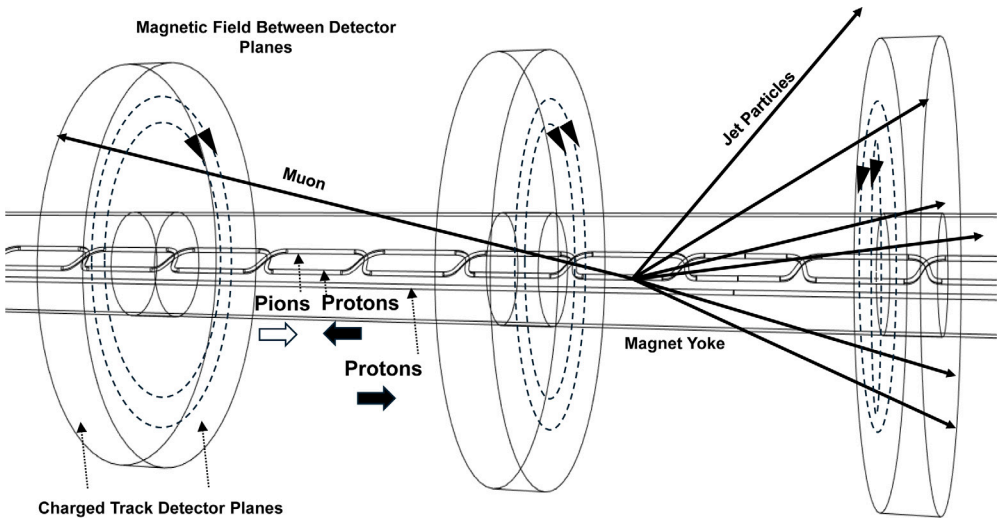

Fig. 10. Schematic of the proposed detection system for the CC DIS muons, showing three sets of circular detector planes with a magnetic field in between them, surrounding the magnet yoke. Within the yoke are the staggered pion decay tunnel and the proton beamline, as well as the proton beamline used for the fixed target collisions. Also shown are the muon and jet tracks from an interaction point in the pion decay tunnel.

CRediT authorship contribution statement

Tome Anticic: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

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

Acknowledgments

This work was supported by the European Union's H2020 Spreading Excellence and Widening Participation - ERA Chair, grant agreement No 669014, and by the Ruder Boskovic Institute, Croatia.

Data availability

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

References

[1] L.A. Anchordoqui, et al., Cosmic neutrino pevatrons: a brand new pathway to astronomy, astrophysics, and particle physics, J. High Energy Astrophys. 1–2 (2014) 1–30.
[2] K. Murase, R. Laha, S. Ando, M. Ahlers, Testing the dark matter scenario for PeV neutrinos observed in IceCube, Phys. Rev. Lett. 115 (2015) 071301.
[3] J. Salvado, O. Mena, S. Palomares-Ruiz, N. Rius, Non-standard interactions with high-energy atmospheric neutrinos at IceCube, J. High Energy Phys. 2017 (2017) 141.
[4] Q. Wu, X. Xu, High-energy and ultra-high-energy neutrinos from primordial black holes, J. Cosmol. Astropart. Phys. (2025) 059.
[5] J. Salvado, O. Mena, S. Palomares-Ruiz, N. Rius, Neutrino interactions at ultrahigh energies, Phys. Rev. D 58 (1998) 093009.
[6] A. Kusenko, T.J. Weiler, Interactions of high energy neutrinos, Phys. Rev. Lett. 88 (2002) 161101.
[7] A.V. Kisselev, V.A. Petrov, Gravireggeons in extra dimensions and interaction of ultra-high energy cosmic neutrinos with nucleons, Eur. Phys. J. C 36 (2004) 103.
[8] V.B. Valera, M. Bustamante, C. Glaser, The ultra-high-energy neutrino-nucleon cross section: measurement forecasts for an era of cosmic eev-neutrino discovery, J. High Energy Phys. (2022) 105.
[9] FASER Collaboration, First direct observation of collider neutrinos with FASER at the LHC, Phys. Rev. Lett. 131 (2023) 031801.
[10] M. Aartsen, et al., Evidence for high-energy extraterrestrial neutrinos at the IceCube detector, Science 342 (2013) 1242856.
[11] M. Aartsen, et al., First observation of pev-energy neutrinos with IceCube, Phys. Rev. Lett. 111 (2013) 021103.
[12] M. Aartsen, et al., Observation of high-energy astrophysical neutrinos in three years of IceCube data, Phys. Rev. Lett. 113 (2014) 101101.
[13] S. Aiello, et al., Observation of an ultra-high-energy cosmic neutrino with KM3net, Nature 638 (2025) in press.
[14] S.E. Kopp, Accelerator neutrino beams, Phys. Rep. 439 (2007) 101.
[15] A. Abada, et al., FCC-hh: the hadron collider – future circular collider conceptual design report, Eur. Phys. J. Spec. Top. 228 (2019) 755–1107.
[16] V. Bertone, R. Gauld, J. Rojo, Neutrino telescopes as QCD microscopes, J. High Energy Phys. (2019) 217.
[17] R. Gandhi, et al., Neutrino interactions at ultrahigh-energies, Phys. Rev. D 58 (1998).
[18] G. Alves, M. Hostert, M. Pospelov, Neutron portal to ultra-high-energy neutrinos, 2025.
[19] M. Collaboration, Demonstration of cooling by the muon ionization cooling experiment, Nature 578 (2020) 53–59.
[20] C. Accettura, et al., Towards a muon collider, Eur. Phys. J. C 83 (2023) 755.
[21] M. Benedikt, et al., FCC-hh hadron collider – parameter scenarios and staging options, in: Proc. IPAC 2015, Richmond, VA, USA, 2015.
[22] T. Sjöstrand, et al., An introduction to PYTHIA 8.2, Comput. Phys. Comm. 191 (2015) 159–177.
[23] R. Brun, F. Rademakers, An object oriented data analysis framework, Nucl. Inst. Meth. Phys. Res. A 389 (1997) 81–86.
[24] J. Strait, et al., Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 3: Long-Baseline Neutrino Facility for DUNE, Tech. rep., 2016.
[25] H.G. Kirk, et al., The MERIT high-power target experiment at the CERN PS, in: Proc. 11th European Particle Accelerator Conference (EPAC08), Genoa, Italy, 2008.
[26] S. Navas, et al., Particle Data Group Collaboration Collaboration, Review of particle physics, Phys. Rev. D 110 (3) (2024) 030001.