Computer Physics Communications 277 (2022) 108366

Contents lists available at ScienceDirect

Computer Physics Communications

www.elsevier.com/locate/cpc

Enhancing higher-energy spectral resolution for electron particle simulations in air ☆,☆☆

Anthony Schmalzried [a,*], Alejandro Luque [a], Nikolai Lehtinen [b]

[a] Instituto de Astrofísica de Andalucía (IAA), CSIC, PO Box 3004, 18008 Granada, Spain
[b] Birkeland Center for Space Science, University of Bergen, PO Box 7803, N-5020 Bergen, Norway

ARTICLE INFO

Article history:
Received 28 March 2021
Received in revised form 30 September 2021
Accepted 31 March 2022
Available online 11 April 2022

Keywords:
Monte-Carlo
Roulette/splitting
Variance reduction
Super-particle compaction
Electron spectral resolution
Electron swarm
Target super-spectrum
Thermal runaway

ABSTRACT

In the presence of an electric field, electrons would theoretically accelerate asymptotically to relativistic energies. However, regular collisions with air molecules limit the increase in electron energy. The stochastic nature of collisions leaves a theoretical probability that an electron elude inelastic collisions thereby accumulating an atypically high energy. Such an electron, under specific criteria, could be called a "thermal" or "cold runaway". Depending on the electric field, the runaway probability might be too low to be computationally observed without resorting to Monte Carlo importance sampling. This article provides a method for fixing the spectral energy resolution of electrons through the combined methodology of Russian roulette and probabilistic splitting in order to render the study of runaway mechanism amenable to electron swarm simulations in various plasma physics applications.

Program summary

Program title: particle-energy-compaction.py
CPC Library link to program files: https://doi.org/10.17632/k2y2r73t69.1
Developer's repository link: https://osf.io/c6wyh
Licensing provisions: CC By 4.0
Programming language: Python 3
Nature of problem: Currently, electron thermal runaway simulations in electric discharges cannot properly resolve the electron energy-spectrum tail. Most codes apply a super-particle restriction algorithm without efficiently or systematically allocating more space for scarcer electrons located at higher-energies. This results in a poor estimation of the thermal runaway rates at electric fields below the critical runaway threshold.
Solution method: We provide a methodology based on the Monte-Carlo variance reduction techniques of Russian roulette and splitting, to allocate particles in different energy domains according to a given target super-particle energy density function. The algorithm converts an input set of super-particle energies and weights into another set that matches the desired spectral resolution. To avoid a sudden surge of super-particles in a low-populated spectral region, the algorithm takes also as input a minimum super-particle weight threshold relative to the physical number of particles in the simulation.
Additional comments including restrictions and unusual features: A good knowledge of the problem is required in order to design the target spectrum function. An inappropriate selection of the target spectrum can result in a decrease of spectral resolution or also a severe deterioration of the swarm properties due to large stochastic fluctuations. It is recommended that the users first devise the target spectrum based on the physical spectrum obtained from their simulations, and start designing from there.

1. Introduction

Monte-Carlo simulations for particle transport are now widespread in many different domains. They are used in nuclear reactor analysis [1], dose calculations [2], detector response [3], medical treatment [4] and diagnosis [5] (radiography, proton therapy, positron scan, etc.), computer vision (photon transport [6]) and

☆ The review of this paper was arranged by Prof. David W. Walker.
☆☆ This paper and its associated computer program are available via the Computer Physics Communications homepage on ScienceDirect (http://www.sciencedirect.com/science/journal/00104655).
* Corresponding author.
E-mail addresses: aschmalz@iaa.es (A. Schmalzried), aluque@iaa.es (A. Luque), nikolai.lehtinen@uib.no (N. Lehtinen).

https://doi.org/10.1016/j.cpc.2022.108366

countless other domains [7]. The relevance and efficiency of those simulations greatly depends on the underlying particle management algorithm. During the simulation, some particles might contribute little useful information while others enter into regions or regimes where the resolution (as a function of the number of particles) needs to be raised to a certain desired level of precision. To maintain physical coherency while reaching the specific objectives of the simulation, statistical weights [8, p. 141] are assigned to particles. Regions with high(low) resolution will typically contain many(few) particles of lower(higher) weights. Particles are curtailed by probabilistic disposal (known as "Russian roulette"), and are multiplied by splitting each into several copies of lower weights. A particle is determined to be split, curtailed or preserved through a so-called importance function (cf. [9] and Appendix A) that translates how relevant the particle is to the outcome of the simulation.

The scope of this paper is to study the proportion and distribution of higher-energy electrons in simulations of electron swarms in an electrified gas. The underlying motivation is to understand better the occurrence probability of cold-runaway electrons [10–14] in electric discharges. The energy state accessible to an electron can be intuitively apprehended with the dynamic friction curve [15] that represents its average energy or momentum loss through collisions. While at low (thermal) energies, this friction increases with energy, it peaks around 100 eV in air before decreasing again [16]. As a result, for a sufficiently high electric field, there exists an energy range where electrons are more likely to be driven into higher energies. This range along with its electrons are subsequently described as 'runaway'. Cold, or thermal, runaway represents the transition from thermal energies up to the runaway regime favoured by intense electric fields. Such a transition is conditioned by the probability that an electron endures less energy loss in random collisions than it gains by acceleration in the field between collisions. Thermal runaway plays an important role in discharge physics; it is thought to occur around leader or streamer tips of lightning discharges in thunderclouds with high activity and is likely related to the occurrence of Terrestrial Gamma-Ray Flashes [17], mostly observed from space [18]; and to X-ray bursts [19], experimentally observed in laboratory sparks [20,21].

The challenge behind simulating the region fostering thermal runaway is the high number of electrons involved in a discharge, growing exponentially in high-field regions. Active (free) electrons in the downstream region of a discharge can at times peak around 10^{15} over a region of 1 cm^3 [22]. The proportion of higher-energy electrons above the friction curve peak at about 100 eV over thermal ones can be extremely low: 1 in a billion or less. This does not, however, exclude the possibility of one electron out of 10^{15} reaching an exceptionally high energy [14]. In order to reduce noise in electron energy distributions to an arbitrarily determined resolution, we implemented an equivalent method to the importance sampling [23], used in Monte-Carlo particle transport codes. This is done through the use of a so-called "contrived spectrum" provided by the user that sets the desired relative abundance of particles in the simulation (denoted as super-particles) according to their energy. In this article, the suffix "super-" will be used for quantities that ignore the stochastic weight w of simulation particles when being calculated.

To our knowledge, although several codes implemented various approaches for preferential particle sampling, their methodology systematically relies upon a simulation-domain partition into two [11], three [24], M [13], etc. respectively and is in some cases limited to particle pair-coalescence [25,26] which either imposes a restriction upon the particle resolution or relaxes the control over the number of simulation particles. Our methodology takes a different approach which preserves flexibility on resolution while maintaining control over the super-particle number. This enhancement is due to: first, a shift from a discrete to a *continuous* partition of our simulation domain (contrived spectrum) and second, a replacement of particle coalescence by *curtailment*. A continuous partition allows for an explicit analytical formula for the target super-particle distribution, whereas discrete partitions always relied upon *relative* resolution levels of neighbouring cells.

We first present the fundamental methodology behind the particle weight adaptation. Then, we address the computational difficulty of the physical exponential growth of electrons in an intense electric field with a particle-number-limiting (curtailment) algorithm. We analyse and compare the performance of 5 implementations in which we let the super-particle number grow to a different factor of the target particle number in order to grasp the conditions yielding optimal usage of the algorithm. The comparison is done through a figure of merit (FoM) accounting for the time consumed and the variance reduction in the higher-energy electron statistics.

The algorithm we discuss below is implemented in the code provided together with this article. Further details are given in the `README.txt` file.

2. Particle energy contrived spectrum

In the following text, we will distinguish:

- N_s : number of (simulated or super-) particles comprised in the numerical simulation: those are the particles effectively stored in an array and tracked individually, each of them bearing a certain statistical weight w.
- N_s^* : predetermined value that fixes the *desired number* of super-particles to be held in the code. Sometimes, some leeway can be provided so that N_s can vary within a range $[N_s^-; N_s^+]$.
- N : number of (real or physical) particles represented by the simulation: this number is yielded by the sum of all statistical weights of the simulated particles.

Thus in practice we have:

$$N \equiv \sum_{i=1}^{N_s} w_i \,. \tag{1}$$

The operation of reallocating N_s particles to match their goal number N_s^* was previously implemented into various techniques including: adaptive particle management [27,13,25,26], resampling [28] or rezoning [29]. In general, *how many* particles ought to be discarded or added in the simulation will depend on the particle number deviation expressed as $\Delta N_s = N_s - N_s^*$; whereas the criteria for determining *which ones*, can be multiple: position, velocity (direction and/or energy). Those can be defined through collective quantities, such as: density, flux, temperature, etc. The processes of adding and discarding can eventually be combined into particle coalescence [30], where the properties of discarded particles are used to generate those to be added.

Our study focuses on electron swarms in electric fields above the conventional breakdown, implying a continual exponential growth of electrons. In this specific context, we will choose to speak of (electron) *compaction*; comparing the many super-electrons that will have to be discarded against the relatively few to be kept or added. The information of the original swarm is thus to be *compacted*. Moreover, whenever the topic should stress the super-particle number reduction (irrespectively of their individual importance), we will speak of *curtailment*; referring then, more specifically, to the procedure of discarding super-particles. This terminology is later illustrated in Fig. 6.

When compacting electrons, we restrict our considerations to their *kinetic energy* ε only, through the one-dimensional criterion of the electron density in energy space: $n_s(\varepsilon)$. This is an ideal quantity representing the number of particles per energy segment under the assumption that there are enough particles to be tallied in that segment. We obtain N_s by integration:

$$N_s(t) = \int_{\varepsilon=0}^{\varepsilon=\infty} n_s(\varepsilon, t) d\varepsilon \,. \tag{2}$$

Of course, in practice, this integration will be bounded by the maximum observed kinetic energy of an electron among the swarm: $\varepsilon^+ \equiv \max_i \varepsilon^{(i)};\ i \in \{1..N_s\}$. As the simulation progresses, all the quantities defined above are subjected to evolve with time t, which when not written specifically will be considered implicit throughout the article.

While N_s^* sets *how many* super-particles are to be held, to be able to determine *which* ones are worth simulating, we appeal to the concept of a contrived (or target) super-spectrum $n_s^*(\varepsilon)$. It can be interpreted as an abstract measure that sets the *desired* number of super-particles per energy segment: it is a *target super-electron density* in energy. Thus when $n_s^*(\varepsilon) < n_s(\varepsilon, t)$ (or .. > ..), all super-electrons whose energy lies "close to" ε, will have to be removed (or cloned) with an according probability $p(\varepsilon, t)$. As opposed to n_s, the contrived spectrum is the same throughout the whole simulation, and thus *independent* of time t. Along with n_s^* comes a contrived set of particle *target* weights w^* whose value can be deduced by the correspondence relationship between simulated and real particle distributions:

$$n(\varepsilon, t) = n_s(\varepsilon, t) \cdot w(\varepsilon, t) = n_s^*(\varepsilon) \cdot w^*(\varepsilon, t) \tag{3}$$

Although n_s^* is temporally and spatially identical throughout the whole simulation, the target weights w^* will have to adapt according to the evolution of the physical number of particles $N(t)$. We also observe that a real particle can be interpreted as a simulated particle of unitary statistical weight. Our notation is summarised below:

Definition: Electron ...	Total Number	Density in energy	Statistical weight
real (physical)	N	n	1
super (simulated)	N_s	n_s	w
target (contrived)	N_s^*	n_s^*	w^*

The distributions here are not normalised and their integrals yield respectively the real, the simulated and the target number of particles:

$$\int_0^{\varepsilon^+} n(\varepsilon, t)\, d\varepsilon = N(t) : \text{real} \,, \tag{4}$$

$$\int_0^{\varepsilon^+} n_s(\varepsilon, t)\, d\varepsilon = N_s(t) : \text{super} \,, \tag{5}$$

$$\int_0^{\varepsilon^+} n_s^*(\varepsilon)\, d\varepsilon = N_s^* : \text{target} \,. \tag{6}$$

As explained above, the conversion of the actual distribution n_s toward the contrived one n_s^* uses probabilistic discarding ($n_s > n_s^*$: particles in excess) and cloning ($n_s < n_s^*$: particles in dearth). This is known in Monte-Carlo techniques as Russian roulette (for discarding) and splitting (for cloning). To conserve *on average* the physical number of particles and their distribution in phase-space, the weights of the super-particles will have to be adapted according to the probability test that they will undergo [9]. The correspondence between densities and weights is expressed through (3) as:

$$\frac{n_s(\varepsilon)}{n_s^*(\varepsilon)} = \frac{w^*(\varepsilon)}{w(\varepsilon)} \,. \tag{7}$$

Graphically, left and right terms in (7) are seen on bottom and top graphs in Fig. 1, where the differences between n (solid), n_s, n_s^* (bins) and w vs. w^* (dots) can all be observed before and after the compaction of an overpopulated swarm. The distributions and weights obtained afterwards are expected to be as close as possible to their target (ideal) values: n_s^* and w^*.

2.1. Russian roulette

When $n_s(\varepsilon) > n_s^*(\varepsilon)$, some particles of energy close to ε must be discarded (curtailed) with a probability p such that $(1 - p(\varepsilon))n_s = n_s^*$. This relation [31, p. 128] can also be expressed with particle weights through (7):

$$(1-p)\frac{w^*}{w} = 1 \quad \text{and therefore:} \quad p = 1 - \frac{w}{w^*} \,. \tag{8}$$

The surviving particles are attributed the target weights w^* to bear the same statistical influence as the original particles with weights w (cf. Fig. 1-top under "Excess" mound).

2.2. Splitting

Conversely, $n_s(\varepsilon) < n_s^*(\varepsilon)$ (equivalently $w > w^*$) means that some super-particles around ε must be added to the simulation. Their *number* is determined by separating the density ratio into its integer k and fractional d components as follows [9]:

$$\frac{n_s^*(\varepsilon)}{n_s(\varepsilon)} = \frac{w}{w^*} = k + d \ : k \in \mathbb{N}, d \in [0; 1[\tag{9}$$

$$p(K = k) = 1 - d \qquad ; \qquad p(K = k+1) = d \tag{10}$$

The random variable K of splitting a particle of energy around ε into k or $k+1$ sub-particles (of weights w/k or $w/(k+1)$), has associated probabilities $1 - d$ and d respectively. This process is known as "sampled-splitting" [32]; the split and intact super-particles are seen on Fig. 1-top above the "Dearth" region, as a series of overlapping streaks taking different weights at a given energy. The streaks are individually associated to the progression of the integer k from 1 to its maximal value.

There exists an alternative splitting that is called "expected-value" splitting [32] that uses a single value weight outcome $w/(k+d)$ but conserves the total weight only on average. Although it was argued [33] to be preferable to the sampled-splitting presented in eq. (10), the comparison was conducted on the basis of the restrictions imposed by the importance function (cf. Appendix A). In our case, we opted for the sampled-splitting for not introducing fluctuations in the total weight after splitting.

2.3. Compaction as the combination of splitting and curtailing

Together, the operations of splitting and curtailing are complementary to each other:

- curtailing: all super-particles reach their targeted weights (or are discarded), but the total weight is not necessarily conserved (only on average);

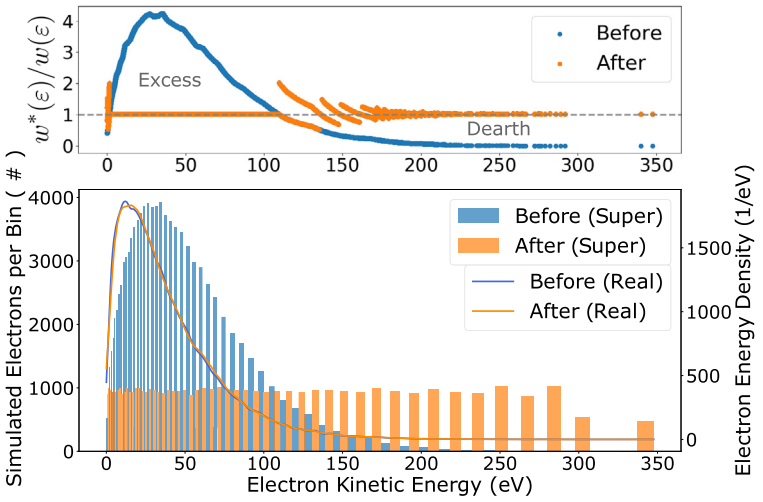

Fig. 1. Illustration of the compaction procedure applied on a swarm. (Top): weight ratios $w(\varepsilon)^*/w(\varepsilon)$ before and after compaction. All curtailed super-particles (in excess) have a weight raised to the target value. Only those super-particles that are split (in dearth) have their weights divided according to their offspring. The average value of split and intact super-particles weights is their target weight. (Bottom): Real distributions and simulated particle tallies before and after a compaction step. While the real density is fairly unchanged, the super-particles are reallocated so that all bins contain an equal amount of super-particles on average. The bins are determined according to Fig. 4. The fact that top and bottom graphs have the same shape is a direct implication of eq. (7).

- splitting: the total weight is exactly conserved but the super-particles do not necessarily attain their respective target weights (only on average).

To avoid having to reiterate accidentally (aka "thrashing" [34]) on particles that were already split once but whose new weights are under the target weight, some leeway is often provided around w^*: in literature this is known as the weight-window technique [35–37]. Particles whose weights lie within a window around w^* are exempted from the probability tests. A typical window has $w_s = 2w^*$ as the upper boundary and naturally $w_d = 2/3w^*$ as the lower boundary for particles who incidentally happened to be split into three. In our case however, we did not impose any weight-window, we used a minimal allowed weight ratio instead (see sec. 2.5 below).

2.4. Density estimation

Both relationships (4) and (7) rely on a continuous description of the quantities n, n_s and n_s^*. In practice however, the density estimation is subjected to the imperfect sampling of the energy space through $n_s(\varepsilon_i)w(\varepsilon_i),\ i \in \{1..N_s\}$ for each super-particle.

Density estimators can roughly be categorised into histogram binning [38] and kernel methods [39]. The first category uses a partition of the energy space into M intervals to tally the number of particles present in each energy segment $b_k(\Delta\varepsilon_k),\ k \in \{1..M\}$; also called a bin. This partition can be regular (bins of equal width $\Delta\varepsilon_k = const$) or irregular. The second method uses a function called a kernel $K(\varepsilon, \Delta\varepsilon)$ whose replica are centred on all data-points and summed at the points of interest to compute a density estimate. The kernel's integral is normalised to unity and its bandwidth $\Delta\varepsilon$ serves to construct an estimation of the particle density locally around ε.

Computing histograms is faster than recurring to kernel methods due to simplicity. However, the bins present a jagged (discontinuous) density profile that can require a smoothing procedure. Kernel density estimators are much less sensitive to small data variations and yield a smooth profile, depending on the kernel/bandwidth chosen. Unfortunately, their computational complexity is higher and so their use for a high number of particles N_s might be impractical.

In general, the formal expression (7) is replaced by its numerical approximation with the actual $i \in \{1..N_s\}$ super-particles:

$$n(\varepsilon) = n_s^*(\varepsilon)w^*(\varepsilon) \simeq \begin{cases} \sum_{j=1}^{n(b_k)} w_j/\Delta\varepsilon_k,\ \varepsilon_j \in b_k \ \leftarrow \text{histogram} & \text{(a)} \\ \sum_{i=1}^{N_s} \frac{w_i}{\Delta\varepsilon} K(\varepsilon_i - \varepsilon, \Delta\varepsilon) \ \leftarrow \text{kernel} & \text{(b)} \end{cases} \tag{11}$$

In the binned case, the sum is restricted only to those $n(b_k)$ particles in the k^{th} bin that contains ε; whereas in the kernel case, all particles are involved: their contribution being gauged by the kernel shape and bandwidth $\Delta\varepsilon$.

For each super-particle, the ratio w/w^* corresponds to the expected number of simulation particles on average after roulette or splitting has been applied. Since all tests are performed independently, this means that the sum of those ratios for all particles will give the average expected value for the total number of simulated particles remaining after compaction. Ideally, this number should converge towards N_s^* as the number of simulated particles increases. Nevertheless, there will always be a difference:

$$\sum_{j=1}^{N_s} \frac{w_j}{w^*(\varepsilon_j)} = \begin{cases} \sum_{j=1}^{N_s} w_j \frac{n_s^*(\varepsilon_j)\Delta\varepsilon_k}{\sum_{i=1}^{n(b_k(\varepsilon_j))} w_i} \\ \sum_{j=1}^{N_s} w_j \frac{n_s^*(\varepsilon_j)\Delta\varepsilon}{\sum_{i=1}^{N_s} w_i K(\varepsilon_i - \varepsilon_j, \Delta\varepsilon)} \end{cases} \overset{?}{\neq} N_s^* \tag{12}$$

In the limit $N_s \to \infty$ and $[M \to \infty$ or $\Delta\varepsilon \to 0]$, both the binning and the kernel approach the Dirac peak distribution and the above expression reduces correctly to the integral of $n_s^*(\varepsilon)$ which, if normalised according to (6), will yield the correct value of N_s^*.

In the meantime, if the bins or $K(\varepsilon, \Delta\varepsilon)$ are chosen such that there are no overlaps of particles when computing the density (as it is the case with Voronoi cells [40,41]), then the middle term in (12) reduces to a Riemann sum of n_s^* at bin centroids about the bin width; or at ε_j points about $\Delta\varepsilon$. When the bins centred on a particle (or the kernels) are not contiguous, the sampling of n_s^* presents holes, or in more formal words the support of $n(\varepsilon) : supp(n) = \{\varepsilon : n(\varepsilon) \neq 0\}$ is not contiguous which entails an inevitable overall under-estimation of the n_s^* integral:

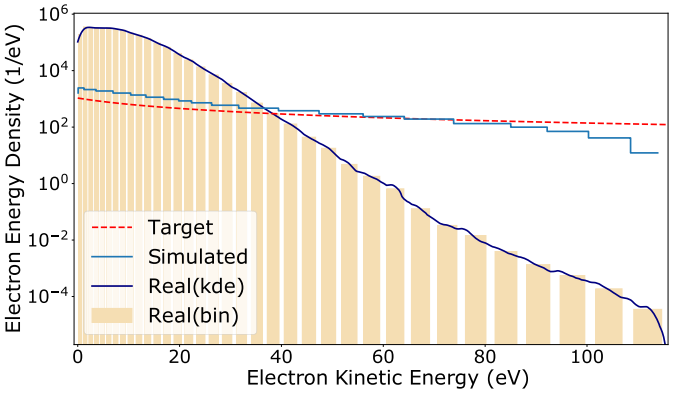

Fig. 2. Illustration of all types of electron spectra (real/super) and how they are calculated (sample taken after 25 ps for a uniform field of 10 MV/m from a swarm initially at thermal equilibrium). The target super-spectrum shape corresponds to (18) and tries to be matched by the actual simulated-particle spectrum. The real spectra are calculated both with a kernel density estimator (kde-solid) and a histogram (bins-columns), the latter is used for the simulated electron spectrum.

$$\int_{\varepsilon \in supp(n)} n_s^*(\varepsilon) d\varepsilon \lesssim N_s^* \tag{13}$$

Therefore, the most straightforward way to ensure that the resulting number of particles after compaction equals the target value *on average*, is to perform a normalising operation on the ensemble of $w_i/w^*(\varepsilon_i)$ values:

$$\text{Corrected weight ratios} \equiv \frac{w_i}{w^*(\varepsilon_i)} \times \frac{N_s^*}{\sum_i \frac{w_i}{w^*(\varepsilon_i)}} . \tag{14}$$

Those latter corrected ratios are the ones to be used in the discarding/splitting tests in (8) and (10).

An overview of histogram and kernel density estimators is presented in Fig. 2. Since the energy range can span over several orders of magnitude, a linear distribution of bins would present a maladapted precision at lower and higher energies (respectively too coarse and too fine). Instead, it seems adequate to adapt bins in function of the expected spectrum. Bins $(b_k(\varepsilon_k))$ can be obtained from a uniform distribution of the target super-spectrum cumulative integral, then reversed back into the energy space:

$$\frac{1}{M} = \int_{b_k(\varepsilon_k)} n_s^*(\varepsilon) d\varepsilon \; ; \; \forall b_k(\varepsilon_k), \; k \in \{1..M\} , \tag{15}$$

like it is depicted in Fig. 4 for an illustrative case with four bins ($M = 4$). This bin configuration is the one used in our simulations, with M determined by the Freedman-Diaconis rule [42] and n_s^* as defined in (18) of the next section.

2.5. Contrived spectrum shape

Since the function $n_s^*(\varepsilon)$ fixes how many particles should be allocated in each energy region, it thereby has a crucial impact on swarm fluctuations in the energy spectrum. As the simulation progresses, the maximal energy ε_+ attained by an electron either converges to a higher value or increases indefinitely in case of runaway. This is because our simulation does not include radiation braking and the electric field is spatially unbounded. For a fixed target number of electrons N_s^*, a minor ratio of thermal electrons over higher-energy ones may lead to large fluctuations (cf. Fig. 3). Consider a constant uniform contrived spectrum:

$$n_s^*(\varepsilon) = \frac{1}{\varepsilon_+} . \tag{16}$$

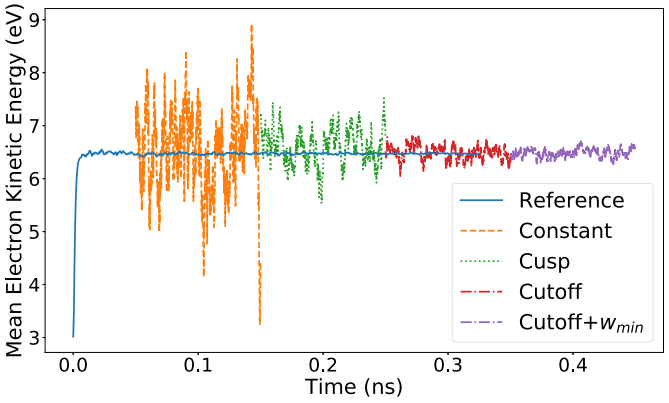

Fig. 3. Mean kinetic energy fluctuations depending on the contrived spectrum shape (displayed in consecutive time intervals for clarity). Shapes further different from the specific physical spectrum yield higher fluctuations of the mean value. Simulation parameters: 8000 simulated electrons at 7 MV/m. Spectrum shapes (eq.?) associated with each label in the figure and their fluctuations (∼?) in eV : Reference(∼ 0.02)=uniform-probability Russian roulette (no shape imposed), Constant(∼ 0.91)=eq. (16), Cusp(∼ 0.33)=eq. (17), Cutoff(∼ 0.14)=eq. (18), Cutoff(w_{min}^*)(∼ 0.09)=eq. (18) and $w_{min}^* = N/10^9$. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

For this distribution with a given energy threshold value ε_{th}, the fraction of super-electrons below ε_{th} is simply $\varepsilon_{th}/\varepsilon_+$. The risk of unphysical fluctuations comes from two sources. For N_s^* fixed, if ε_+ is not bounded, this ratio does not have a lower bound. If however ε_+ has an upper bound, the number of thermal electrons is nonetheless conditioned by N_s^*; which can be considerably low if the energy range is broad.

In order to avoid this, several precautions can be made. First and simplest, n_s^* can be chosen such that its integral (6) for $\lim \varepsilon_+ \to \infty$ stays finite. The user can then check the lowest theoretical number of simulated electrons in the thermal range: $N_s^*(\varepsilon < \varepsilon_{th})$, where ε_{th} is taken as the thermal boundary. However, if this bounded constraint over the shape of n_s^* is not satisfactory, the code must itself set a limit to the maximal electron energy allowed in the simulation, to which the value ε_+ is lesser or equal. Setting such a limit can either be explicit (in which case all electrons exceeding this limit are simply non-probabilistically eliminated) or implicit (based on restricting the minimum particle weight w_{min}^*). In the latter case, w_{min}^* is set as a small fraction of the physical number of particles N. Physically, this corresponds to choosing the lowest particle resolution in the simulation. For example, fixing $w_{min}^* = N/10^9$ means that the minimally provisioned weight for a super-electron shall account for at least one-billionth of the whole swarm; the simulation will not be able to resolve particles with a lower scarcity. Although not absolute, such a relative weight boundary brings a probabilistic limit to the maximal energy observed through the combination of n and n_s^* as can be seen in eq. (A.2). Once these combinations of choices are made, N_s^* can be set accordingly. A few iterations of this process can be repeated if the N_s^* entailed is deemed too high.

To illustrate the implications of the contrived spectrum's shape and its boundedness, we provide a simulation with $N_s^* = 8000$ and show the difference in fluctuations of the electron swarm's mean kinetic energy in Fig. 3. The selected contrived spectra presented hereafter are compared to a simple uniformly random elimination of particles (labelled 'Reference' see also in sec. 3).

Besides the uniform spectrum in eq. (16), we included a cusp-shaped spectrum with an amplitude α of 1/2 defined as:

$$n_s^*(\varepsilon) = \frac{N_s^*}{\varepsilon_+}(1 + \alpha \cos(2\pi \varepsilon/\varepsilon_+)), \; \alpha = 0.5 . \tag{17}$$

Assuming that the natural spectrum shape beyond thermal energies of an electron swarm is exponential due to the impact ionisation cascade, we additionally devised an inverse power-law $n_s^*(\varepsilon)$ with a cutoff ε_0 at 15 eV, taken to represent the boundary for thermal energies (see results section):

$$n_s^*(\varepsilon) = \frac{N_s^*}{\varepsilon_0 \ln(1+\varepsilon_+/\varepsilon_0)} \frac{1}{1+\varepsilon/\varepsilon_0} \,. \tag{18}$$

For all spectra, $\varepsilon_+ = 500$ eV, and an additional implementation of the cutoff spectrum featured $w^*_{min} = N/10^9$ alone.

The robustness of our algorithm is underlined by the average physical statistics; identical and independent from the particular n_s^* shape chosen. The fluctuations of variable amplitudes demonstrate an inappropriate combination choice of the spectrum shape n_s^*, along with the allowed number of particles N_s^*, the warranted maximal energy ε_+ and the individual particle resolution w^*_{min}. Moreover, although we can observe that ε_+ and w^*_{min} play a milder role next to the n_s^* shape, the minimal weight introduces a better quenching of the simulated spectrum with superior control over fluctuations. An examination of the average energy for higher-energy electrons (above 15 eV or even 60 eV) reveals that fluctuations are roughly the same for both weight-limited and unlimited configurations.

2.6. Constraining the number of super-particles

In the previous sections, we developed a procedure to curtail a number N_s to N_s^* particles *on average*, while matching a given $n_s^*(\varepsilon)$ as closely as possible. This operation is however discrete and aside from it, the electron swarm will keep growing at its own exponential rate which depends on the applied electric field (higher than the conventional breakdown around 3 MV/m in standard air [43]). The curtailment must thus be operated on a regular basis. A simple procedure consists in setting a super-particle threshold N_s^+ or a timer of period τ. The operation triggers respectively every once $N_s \geq N_s^+$; or when the simulated time interval since the last operation Δt exceeds τ. If a steady and reliable measurement of the growth rate ν_s is available, then one can set:

$$N_s^+ = N_s^0 \exp(\nu_s \tau) \Leftrightarrow \tau = \frac{\ln(N_s^+/N_s^0)}{\nu_s} \,, \tag{19}$$

and the two procedures are equivalent on a large number of trials, provided ν_s is an unbiased estimator of the actual growth rate. Here, N_s^0 is the number of particles present when the time accumulator Δt is reinitialised. It can be thought at first that $N_s^0 = N_s^*$, but in practice, N_s^0 could also be a lower threshold $N_s^- < N_s^*$. Indeed, if N_s^* represents the *average* number of super-particles present in the simulation during a period Δt, then by definition:

$$N_s^* = \frac{1}{\Delta t} \int_0^{\Delta t} N_s^0 \exp(\nu_s t) dt = \frac{N_s^0}{\nu_s \Delta t} (\exp(\nu_s \Delta t) - 1) \,. \tag{20}$$

If we set the multiplication factor $\exp(\nu_s \Delta t) = m > 1$, then:

$$N_s^0 \equiv N_s^- = N_s^* \frac{\ln(m)}{m-1} < N_s^* \,. \tag{21}$$

As $\Delta t \to 0 : m \to 1$ so $N_s^- \to N_s^*$ and we maintain a tight control over the super-particles allowed in the code. On the other hand, this means that the compaction step will have to be performed more frequently. There should be some compromise linked to the simulation outcome between how drastically and frequently should the particle number be curtailed. In the next section, we discuss this issue, present assessment tools and illustrate the performance of different compaction configurations seen in table 1.

Table 1
Lower and upper super-particle number limits for five different compaction schemes. The limits are not applicable to the Reference simulation which simply uses uniformly random elimination of particles.

Label	:	N_s^-	,	N_s^+	
A "Fixed"	:	N_s^*	,	N_s^*	
B "Plus Quarter"	:	N_s^*	,	$1.25N_s^*$	
C "Doubled"	:	N_s^*	,	$2N_s^*$	
D "Halved"	:	$N_s^*/2$	,	N_s^*	
E "Averaged"	:	$N_s^* \ln 2$	,	$2N_s^*$	; $m = 2$
R "Reference"	:	$N_s = N_s^*$			

3. Assessing variance reduction schemes

To test the algorithm, we apply it to simulations of electrons swarms in homogeneous electric field in air. The ultimate goal behind limiting the number of super-particles is dual:

1. Reduce the variance σ^2 of the quantities computed in energy ranges of interest.
2. Reduce the computational demands expressed through the total simulation time T_{simu}.

These two objectives can be incorporated into a figure of merit (FoM) [44] and computed as:

$$\mathrm{FoM} = \frac{1}{T_{\mathrm{simu}} \sigma^2} \,. \tag{22}$$

This figure will enable us to make an assessment over the various possibilities in compacting N_s super-particles in the range defined by $N_s^- \leq N_s^* \leq N_s^+$.

In the following we will present five configurations all using the aforementioned "cutoff" super-spectrum (18) as their target, contrasted against a reference simulation (cf. Table 1)

The first "fixed" (A) scheme curtails whenever $N_s \neq N_s^*$, with the minimum temporal resolution permitted by the simulation time-step which is defined as the inverse of the maximal collision rate according to the null-collision [45] methodology.

With the growth rate depending on the electric field, there might be a large lapse of time between two curtailments for a value of electric field not much above the conventional breakdown. We explore this issue by repeating all 5 configurations at 5 MV/m and 10 MV/m with $N_s^* = 50000$ electrons in homogeneous air at atmospheric pressure. Additionally, since we are interested in the phenomenon of thermal runaway [12,13,16,14], we set a minimum relative weight of $w^*_{min} = N/10^{15}$ to reflect the order of magnitude of free electrons present in a typical streamer corona as already mentioned in the introduction. As a precaution, we also set $\varepsilon_+ = 1$ keV: a traditional limit [22,14] taken in studies about thermal runaway. The initial conditions of the swarm are already in thermal equilibrium. The cross-sections used for processing collisions were taken from Biagi's database [46]. The particle code used is described in more detail in a previous article [47].

Finding a sampling strategy that would best exploit each scheme's advantages turned out a unrealisable task. Ideally, one would need to track the time step corresponding to the highest resolution in the energy range of interest. Since this moment depends on the swarm's growth rate conditioned in turn by the electric field, the sampling rate for each scheme would then have to be different, leading to an unequal number of samples for a fixed period of simulation T_{simu}. To preserve fairness, one would need to fill-in the missing samples with ones taken at other times. Otherwise, longer simulation times for schemes with a higher compaction ratio would be needed. Despite various trials, the conclusions about each scheme's performance remained unchanged

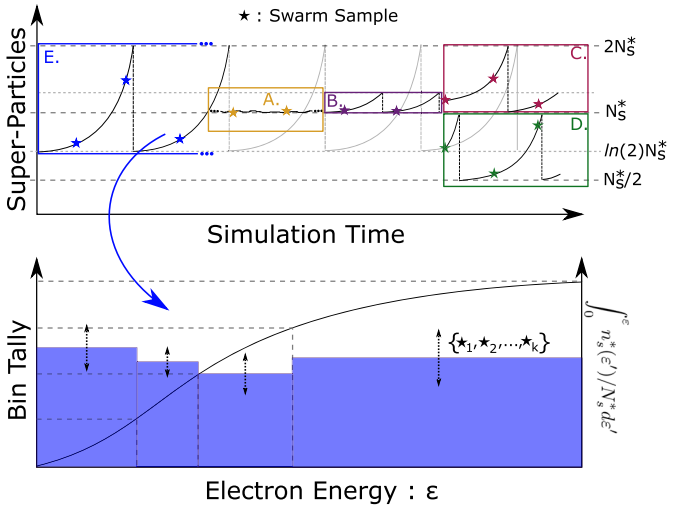

Fig. 4. Schematic view of all five compaction configurations. (Top): Evolution of simulated particles along simulation time sampled at a regular rate irrespective of their number N_s. Vertical dashed cuts represent compaction operations when $N_s > N_s^+$. (Bottom): Construction of an average histogram of all collected samples with bins uniformly delimited in the cumulative integral of the target super-spectrum (cf. eq. (15)); i.e. such that their expected number is equally distributed. The letters A–E represent the five compaction configuration.

due to a compensation of resolution gains by losses either in: longer simulation times, less samples or samples of lesser resolution. In the end for a fair comparison, all simulations were sampled at regular intervals of 10 ps during 1 ns, thus cumulating a total of 100 samples per run from which statistical analysis are conducted.

The FoM as defined above (22), applies to a macroscopic quantity whose variance is to be computed. Unfortunately, neither the notions of "higher-energy" nor even "runaway" electrons lend themselves to a very precise and universal mathematical definition. Although the runaway threshold has a theoretical value based upon the averaged electron friction (collisional) force equated by the electric driving force [10]; for the electric fields considered, this value exceeds the keV boundary [15] and thus, we are confronted to the risk of not having any runaways to tally at all (if the schemes proved themselves to be too weak). Instead, we decided to compare our simulations on the ground of three simple definitions of average energy: one is without restriction, the others only involve electrons above 15 eV and 60 eV respectively. The first boundary was chosen as a limit of the thermal neighbourhood, the other was chosen as a consequence that the maximal energy observed in the Reference simulations was around 60 eV.

A more qualitative comparison will be based on the expected variance reduction in the targeted energy ranges. Since the regions of interest are quantitatively incorporated in the contrived spectrum's shape, we also provide energy-density histograms averaged over those 100 samples, whose binning decomposition is based upon a uniform partition in the cumulative integral space of the contrived spectrum: $y = \int_0^{\varepsilon} n_s^*(x)/N_s^* dx$. In the ideal case of electrons distributed throughout the range according to n_s^*, each bin should contain an equal number of particles, a configuration that minimises the variance of the average particle number in each bin. The whole procedure is illustrated in Fig. 4.

3.1. Results

Table 2 summarises different macroscopic averages and their fluctuations (standard deviations) for each compaction scheme (A to E against the Reference) in the case of 10 MV/m applied electric field. As a prerequisite for assessing the effectiveness, we took the characteristic energy level $\varepsilon_{th} = 15$ eV, to delimit the thermal-energy neighbourhood. This choice was based upon the fact that the electron friction curve rises steeply from 15 eV in air owing to the onset of impact ionisation collisions. We surmise therefore that cutting off the contrived spectrum at 15 eV is a reasonable choice, if we are to assume that the electron distribution function drops beyond this neighbourhood.

Looking at the three first rows of "milestone" averages, we can observe how the resolution in energy degrades as the threshold of thermal electrons is passed when no splitting is applied. A second glance reveals however that the variance reduction at higher energies is *not necessarily* attained for every scheme. All C, D and E perform poorly in resolving the spectrum beyond thermal energies. It seems therefore that a certain compaction frequency is required to maintain a sample of electrons in higher-energy regions, lest the simulation resources be mostly spent on the thermal electrons. The percentage of super-electrons above 60 eV is also very low, at most 0.01% as opposed to approximately a quarter in the A and B schemes. As a result, the only reliable estimation of the physical proportion of electrons swarming above 60 eV under 10 MV/m is between 2 to 3 in one million. It must be noted that a physical resolution of $1/10^6$ above 60 eV implies that a super-electron at that energy should have a 10^6 times lesser weight than a thermal electron in order to provide a comparable signal-to-noise ratio in the spectrum. Such resolution is impractical to reach if one resorts to a discrete partitioning with *relative* weight ratios between neighbouring regions. For instance with a partition of 6, the minimal jump from one to the next region would be of a factor 10. On the other hand, increasing the number of bins leads naturally to our continuous description. This is why we did not include comparisons with discrete partitions (because they cannot attain comparable levels of resolution under practical considerations) nor pair-wise coalescing schemes (because they require an additional closest neighbour search).

The increased resolution in higher energies can be better grasped by looking at a comparison of average electron density (top) and super-electron tally histograms (bottom) in Fig. 5 constructed as explained in the previous section (cf. Fig. 4-bottom: $M = 100$). The bins are chosen such that the target super-spectrum would appear as a horizontal line (equipartition of super-electrons). From A to E, there is a clear ranking of closest to farthest from the target distribution. Additionally, one can observe the difference between the spectral sampling fluctuations. The difference between A and B is minute and implies that there could be an optimal compaction frequency. The smoothed distributions of A and B all along the spectrum contrast with the noise of the remaining schemes above 40 eV. Both A and B spectra extend beyond two hundred eV but only the portion up to 100 eV was presented for better clarity. Although it is true that C,D,E make a step towards broadening the energy-spectrum sampling boundary, the resolution obtained is not sufficient to make any physical induction about runaway or spectrum shape.

The two middle rows in Table 2 report a measure of duration in computer time and average number of particle collisions per unit of simulated time. Although we could argue that the first is machine dependant whereas collisions should be independent, it is also unjust to interpret less collisions as a good computational time-saving indicator. Indeed, the increased presence of high-energy electrons naturally entails a higher rate of collisions since the (inelastic) collision rate has a maximum around 100 eV in air [16]. Thus, the time spent in more collisions is not a defect in the application of the algorithm but rather an unavoidable effect coming along with the user's constraints. Timing performances were run separately on the same machine under identical conditions, to ensure the least possible influence on the timing scores.

Subsequently, the figures of merit for all three thresholds of average kinetic energy were reported below in Table 3 for both

Table 2
Comparison of 6 simulation setups all under a homogeneous electric field of 10 MV/m. Numbers presented are 100-sample averages (and their standard deviations: σ). The first column "R" (=Reference) does not use spectrum-based compaction; $*$ therefore $N \propto N_s$. The other schemes' nomenclature refers to the swarm's simulated particle growth allowance prior to curtailment (see section 3).

Scheme	R	(A)	(B)	(C)	(D)	(E)
Mean (eV)	8.03(0.027)	8.03(0.045)	8.03(0.043)	8.03(0.034)	8.03(0.05)	8.03(0.037)
Mean $>$ 15 eV (eV)	19.4(0.12)	19.4(0.057)	19.4(0.056)	19.4(0.071)	19.4(0.094)	19.4(0.077)
Mean $>$ 60 eV (eV)	64.3(4.5)	64.6(0.21)	64.6(0.26)	64.4(4.1)	64.8(4.2)	64.7(4.1)
$N_s > 60$ eV/N_s		0.27(0.0027)	0.24(0.034)	$10^{-4}(10^{-3})$	$< 10^{-4}$	$< 10^{-4}$
$(N > 60$ eV/$N)$ (x10^{-6})	2.4(7.4)*	2.5(0.21)	2.5(0.26)	1.6(6.4)	2.6(10)	4(10)
T_{simu} (s/ps)	0.78(0.11)	1.3(0.2)	1.1(0.27)	1.1(0.22)	0.5(0.1)	1(0.27)
T_{simu} (#coll/fs)	115(8.7)	330(11)	320(58)	180(32)	90(13)	160(37)
N_s^{del}/N_s (%)	0.14(0.017)	33(1.4)	41(2.4)	77(5)	74(6.2)	80(1.6)
N_s^{new}/N_s^* (%)	0	23(2.3)	29(3.1)	53(9.3)	48(12)	40(4.6)

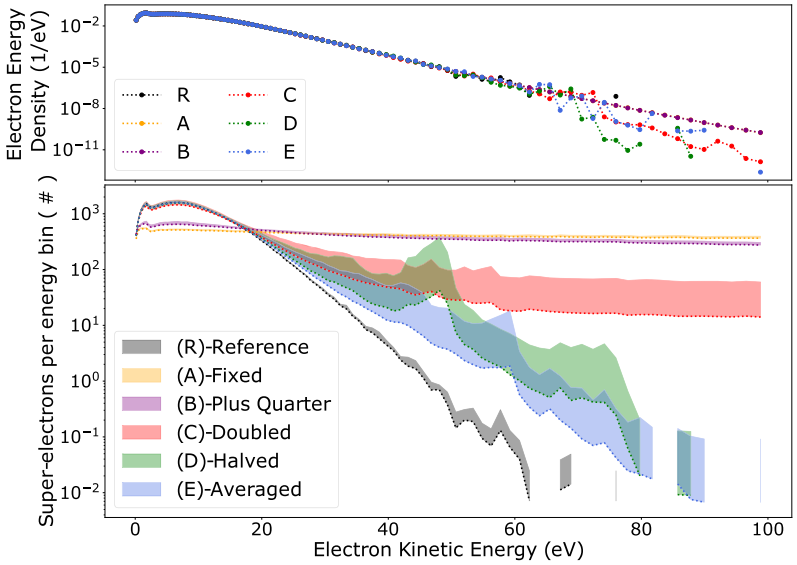

Fig. 5. Energy spectrum comparison of all five configurations and the Reference simulations defined in section 3. Top: (Physical) electron energy density. Bottom: Super-electron tally in 100 bins defined uniformly from the cumulatively integrated target super-spectrum (as in Fig. 4 and eq. (15) with $M = 100$). Shaded areas represent the errors calculated as the standard deviations of the 100 bin-tallies. An ideally resolved spectrum would tally equally in each bin and thus present a horizontal line. The x-axis is shared for both figures.

fields: 10 MV/m and 5 MV/m (rows shaded in grey). Inevitably, the fluctuations about the average kinetic energy are, on the long run, higher when compaction is applied more frequently. Nonetheless, the aim of increasing resolution of electron energy density at higher energies (above 15 or 60 eV) is only obtained with the schemes A and B. Furthermore, the parting of A and B scores at 5 MV/m indicates that the lower the electric field, the more demanding will be the goal of resolving higher-energy regions: compaction should be applied more often. There quickly arises a compromise among the resolutions desired above various energy thresholds. While compacting every time-step amounts to an unnecessarily frequent manipulation of the electron swarm, waiting a certain factor growth is not universally effective and is much affected by the actual electric field applied.

This loss of efficiency can be seen on the cyclical "strain" applied on the simulated swarm after every compaction step. If the interval is too long, the departure of $n_s \nsim n_s^*$ will require a higher rearrangement to be made among the swarm. This can be measured by looking at the last two rows of Table 2 featuring the proportion of super-electrons to be eliminated N_s^{del}/N_s and those to be added N_s^{add} prior to (N_s) and posterior to (N_s^*) compaction. Those proportions are laid out graphically in Fig. 6. Evidently, the higher N_s^+/N_s^-, the greater the information loss from the original swarm per curtailment operation; an information that can be compacted into a fraction as little as 20% of its original population in the case of the "averaged" (E) scheme. The colour coding in Fig. 6 could help clarifying the terminology used; arguably one could distinguish the action of "curtailing" N_s down to N_s^- super-particles (height difference between red and blue column tops); and the action of "compacting" the information of N_s into N_s^* (red column top over green). Here, the "Reference" scheme is purely curtailing, whereas A to E are also adding some particles split from previously existing ones. Thus, compaction is the combined action of curtailing and splitting (as we had proposed in sec. 2.3).

An immediate consequence of this higher-energy resolution boost is a surge in super-particle secondary ionisation right after the compaction seen on Fig. 7. Without an uninterrupted supply of higher-energy super-particles, the swarm will gradually evolve back to a "super-less" distribution (each super-particle weight converging back to a uniform value: the physical limit). Both figures (6 and 7) indicate that there is an increasing risk of introducing amplified sporadic stochastic fluctuations as more growth is allowed to the swarm between two compaction operations. Several samples taken from D and E *right after compaction* presented a severe deviation (accounting for the irregularity observed in Fig. 5-bottom around 50 eV), albeit non-biased on average, from the actual mean kinetic energy. This deviation is gradually damped as the swarm grows without curtailment. The low absolute value of N_s^- in D and the higher ratio of N_s^+/N_s^- for E are favouring factors for this deviation.

Table 3
Figures of merit for both 5 MV/m (rows shaded in gray) and 10 MV/m (non-shaded rows) simulations for energy averages of selected electron populations. They are obtained from the fluctuations (σ) on the first three rows in Table 2 and the machine time (first row with T_{simu}) according to equation (22). At 5 MV/m there were no electrons observed above 60 eV except for the first scheme A, which makes it the only successful case in estimating the swarm spectrum at higher energies.

Focus population	Reference	Fixed	Plus quarter	Doubled	Halved	Averaged
All	1734	369	473	776	753	723
	3361	641	2411	2110	2076	2442
Any $>$ 15 eV	90	237	279	178	210	169
	103	750	145	145	122	135
Any $>$ 60 eV	0.06	18	13	0.05	0.11	0.06
	-	49	-	-	-	-

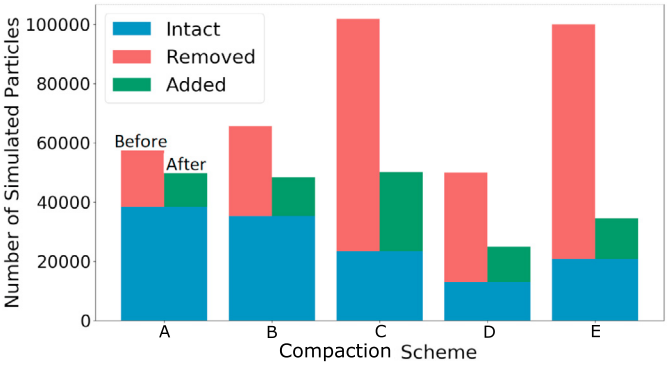

Fig. 6. Comparison of super-electron balances performed during a compaction operation. The original swarm begins at the top of the columns in red and is rearranged into a resolution-boosted swarm at the top of the columns in green. Particles curtailed (sec. 2.1) are represented in red; those that were left intact, in blue; and those that are mere copied splinters (sec. 2.2) of intact particles, in green. The height of blue columns should be taken thus as containing all the residual information from the original swarm. Graphically, "curtailment" corresponds to removing the red columns and leaving out the blue; to be contrasted with "compaction" that rearranges the number of electrons in a swarm seen as column heights from left (before) to right (after).

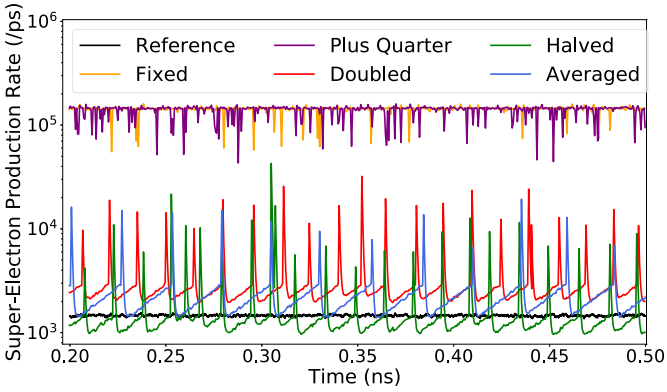

Fig. 7. Super-Electron production rate for all five schemes. The discontinuous peaks all correspond to the curtailment steps which means there is an immediate secondary super-electron production surge after compaction. This surge reduces in amplitude as the curtailment frequency increases, due to an overall background rise in secondary super-ionisation.

While Fig. 7 demarcates a clear dividing line between successful and unsuccessful schemes for enhancing higher-energy regions, the observation on Fig. 5-bottom seems to imply a gradual improvement from E to A, with C scoring somewhere in between. The transition in success is probabilistic; after 1.3 ns of simulation, the C scheme featured an electron whose maximal energy (100 eV) was such that after compaction, the secondary production rate was significantly boosted to shorten the pseudo-periodical compaction time-lapse (construed as the peaks in Fig. 7). Once triggered, this mechanism reinforces itself; since higher-energy electrons promote a faster super-electron growth, this growth, in turn, increases the compaction frequency and thus stabilises the supply in higher-energy super-electrons. From 1.5 ns, the C scheme was able to resolve energies up to several hundreds of eV at 10 MV/m as A and B alike. The density plot in Fig. 5 shows that higher-energy electrons could be ultimately resolved. Nevertheless, the FoM only takes the fluctuations into account, which remain high for C since more than half of the simulation time was spent without enhanced spectral resolution. This implies that the maintenance of higher-energy electrons in swarm simulations is a *state* and not a *property* of the scheme chosen. The difference lies in the *probability* (almost certain for A and B, low for C and even lower for D and E) with which the scheme can foster higher-energy electrons.

Overall, the efficiency of a correct implementation of contrived spectral compaction is demonstrated by considering how little FoM (less than one order of magnitude) is lost on the total energy average compared to how much is preserved at higher energies (increase of 4 orders of magnitude above 60 eV or simply by having any data at all for low electric fields).

4. Discussion and conclusion

The use of super-particles in simulations translates as a reallocation of resolution in certain regions of space, energy, momentum, time or any other desired property. This reallocation can be implemented by several means, including particle coalescence, resampling or probabilistic discarding/splitting. Using the latter method, we presented here a specific application to electron swarms in standard air under high electric fields above the conventional breakdown ($\sim$3 MV/m). The motivation was to investigate the small probability of occurrence of thermal runaway, where amongst a very large swarm of electrons $\gtrsim 10^{15}$, a few are probabilistically able to accelerate up to higher kinetic energies of several hundred eV or more. Runaway *precursors* ($\gtrsim$60 eV) were differentiated by thermal electrons ($\lesssim$15 eV) solely based on their energy. The double-sided goal of both hosting such a large swarm and favouring the conditions for observing thermal runaway, was embodied by a so-called contrived target super-spectrum for the energy distribution of super-particles. Its norm ensured a limit on the number of super-particles, while its shape arranged their allocation in energy-space.

When compacting a swarm, the loss in thermal-energy resolution in steady state is observed as an increase in the fluctuations about the mean energy. The two causal factors were identified as the particular shape chosen for the *target super-spectrum* $n_s^*(\varepsilon)$ with its norm N_s^*, maximal boundary ε_+ and possibly a minimal allowed weight w_{min}^*; and the *frequency* with which the compaction of a non-optimal to an optimally resolved spectrum is applied. The study about the influence of frequency was divided into 5 different implementations labelled from A to E. While n_s^* limits the expected number of super-particles in the thermal domain, tuning the curtailment frequency amounts to choosing between a

repeated series of continual albeit small perturbations (A and B schemes) as opposed to a sporadic occurrence of considerable perturbations to the swarm (C, D and especially E).

Depending on the electric field, the compaction frequency might be ineffective, for a lower electric field will be less able to sustain high-energy electrons and thus will require a more frequent resupply of those. Although B performed best at 10 MV/m, its effectiveness (FoM) was greatly reduced under 5 MV/m and overrun by scheme A. By imposing a minimal target weight w^*_{min}/N relative to the physical number of particles (N), the success of observing an electron above a certain energy threshold can be directly attributed to the physical probability occurrence in a heavily populated swarm. This minimal weight thus acts as a tunable resolution parameter (lower value leads to higher resolution).

Overall, in order to fulfil the goal for resolving an electron swarm at higher energies, it is best to carefully choose a target super-spectrum shape that closely matches its physically expected form and at the same time raises the occurrence probability of rare super-electrons according to the total number of simulation particles allowed in the code and the target energy domain to be reached. Additionally, the compaction algorithm should ideally be applied according to a periodicity based on the relaxation time of the swarm for the given electric field. Without a reliable model, this time can be estimated as the inverse of the swarm mean kinetic energy growth rate (onset of "Reference" curve on Fig. 3) starting from room-temperature initial conditions (i.e. virtually equal to 0 eV). While a discrete partition of the spectrum requires to define many bins to be able to reach adequate resolution levels with restricted flexibility, the present continuous description of the target super-spectrum provides a systematic way for applying preferential sampling.

The present algorithm could be extended to the application of Relativistic Runaway Electron Avalanches that perhaps requires another suitable contrived spectrum shape that would reach up to few hundreds of MeV. Moreover, future improvements could also seek to incorporate *spatial* dependence when the electric field is highly inhomogeneous as at the tip of a streamer, and smooth the compaction procedure into a continuous process rather than an abrupt and discrete rectification. With those improvements, the algorithm could be used to better characterise electron surfing on negative streamer fronts and subsequent thermal runaway rates depending on the applied external electric field.

This work was supported by the European Research Council (ERC) under the European Union H2020 programme/ERC grant agreement 681257. The authors acknowledge financial support from the State Agency for Research of the Spanish MCIU through the "Center of Excellence Severo Ochoa" award for the Instituto de Astrofísica de Andalucía (SEV-2017-0709), and from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement n. 320839 and the Research Council of Norway under contracts 208028/F50, 216872/F50 and 223252/F50 (CoE). We comply with the new FAIR data access policy. The code for compacting according to the contrived spectrum methodology is freely accessible through the Open Science Frame at https://osf.io/c6wyh [48].

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.

Appendix A. Relation to the importance function

Our method using the contrived spectrum as a way to discriminate particles according to their energy can be equivalently described through the so-called importance function $I(\bar{r})$ widely used in Monte-Carlo particle literature [49,50, s. 7.5.1]. This function maps points $\bar{r} \in \mathbb{R}^d$ in the particle phase-space into a real positive value that appraises how contributive a particle of properties $\bar{r}$ is to the simulation output. Thus, when a particle moves from a position $\bar{r}_i$ to $\bar{r}_j$ with a substantial difference in importance, the test proceeds as follows:

$$q = \frac{I(\bar{r}_j)}{I(\bar{r}_i)}, \quad \begin{cases} q < 1 : p_r = 1 - q \ \text{ roulette test} \\ q > 1 : p_s(\lceil q \rceil) = 1 - (q - \lfloor q \rfloor) \ \text{ split test} \end{cases} \tag{A.1}$$

We can see here that there is a clear correspondence between the importance function and the particle target weight. If we assumed that all weights match exactly the target weights assigned, then $I(\bar{r}) \propto 1/w^*(\bar{r})$. Evidently, this is not always the case which is why the weight-window extension is added to the particle management algorithm to offer a better control over the absolute weight of a particle and its fluctuations. In terms of importance function, we replace the target weight by isolating it in (3):

$$w^*(\varepsilon) = \frac{n(\varepsilon)}{n^*_s(\varepsilon)} \tag{A.2}$$

$$I(\varepsilon) \propto \frac{n^*_s(\varepsilon)}{n(\varepsilon)} \tag{A.3}$$

This is a well-known relation between the importance function and the contrived distribution n^*_s with a certain coefficient of proportionality [9]. Thus, imposing either a target distribution or an importance function leads to an equivalent setting for the simulation, provided a reliable estimate of the real distribution $n(\varepsilon)$ is available for any given energy ε region. However, the importance function methodology has several limitations making it inappropriate for simulations of electric discharges:

1. *Population control*: the electron proliferation through impact ionisation in intense electric fields (above conventional breakdown) is exponential. While the importance function enables only a relative control on the particle weights, combined with the weight-window, this control can become absolute. However, the total number of super-particles allowed in the simulation is not formally controlled. We thus sought a method that maintains this number bounded.
2. *Homogeneous domain*: as opposed to calculations in nuclear engineering where the position-energy space is discretised according to the geometry of the problem [3], we simulate a uniform gap filled with air. Therefore, there was no need for using a fixed grid-partition of the phase-space and computing importance jumps from region to region as through equation (A.1). Spatial homogeneity justifies also why we focused only on electron energy.
3. *Unspecific optimisation*: unlike when trying to maximise the response to a detector [51], our approach seeks to highlight and study the *general* characteristics of thermal runaway for which there is no specific energy range whose sampling needed to be levelled; we took instead a holistic approach to resolving all higher-energy regions above the thermal domain.
4. *Adjustable Spectral Resolution*: as a transient phenomenon, thermal runaway does not present a steady-state that can be precalculated and used as an input into (A.3) to help determine what importance function ought to be used. There is no *a priori* knowledge of the physical spectrum $n(\varepsilon, t)$, and thus we found it more physically intuitive to think in terms of particle allocation per energy bin (=contrived spectrum) instead of the more abstract relative particle importance according to its energy (=importance function).

For these four reasons - one to avoid computational overload and three demanding flexibility - we took the path of setting the discarding/splitting probabilities with the help of the contrived spectrum method. With a proper normalisation relationship (6), the average number of particles allowed in the simulation after compaction should match N_s^*.

References

[1] P.D. Soran, D.C. McKeon, T.E. Booth, IEEE Trans. Nucl. Sci. 37 (2) (1990) 936–942.

[2] E.A. Ainsbury, D. Samaga, S. Della Monaca, M. Marrale, C. Bassinet, C.I. Burbidge, V. Correcher, M. Discher, J. Eakins, P. Fattibene, N. Güçlü, M. Higueras, E. Lund, N. Maltar-Strmečki, S. McKeever, C.L. Rääf, S. Sholom, I. Veronese, A. Wieser, C. Woda, F. Trompier, Radiat. Prot. Dosim. 178 (4) (2017) 382–404, https://doi.org/10.1093/rpd/ncx125, https://academic.oup.com/rpd/article-pdf/178/4/382/24274696/ncx125.pdf.

[3] J. Wagner, D. Peplow, T. Evans, Nucl. Technol. 168 (2009) 799, https://doi.org/10.13182/NT09-A9309.

[4] P. Andreo, Phys. Med. Biol. 36 (7) (1991) 861–920, https://doi.org/10.1088/0031-9155/36/7/001.

[5] S. Staelens, I. Buvat, in: P. Verdonck (Ed.), Advances in Biomedical Engineering, Elsevier, Amsterdam, 2009, pp. 177–209, http://www.sciencedirect.com/science/article/pii/B9780444530752000058.

[6] J. Vorba, J. Křivánek, ACM Trans. Graph. 35 (4) (2016) 42, https://doi.org/10.1145/2897824.2925912.

[7] W.K. Chan (Ed.), Theory and Applications of Monte Carlo Simulations, InTc, 2013.

[8] A. Haghighat, in: Monte Carlo Methods for Particle Transport, CRC Press, 2020, pp. 141–142.

[9] A. Haghighat, J.C. Wagner, Prog. Nucl. Energy 42 (1) (2003) 25–53, https://doi.org/10.1016/S0149-1970(02)00002-1, http://www.sciencedirect.com/science/article/pii/S0149197002000021.

[10] A. Gurevich, Sov. Phys. JETP 12 (5) (1961) 904–912.

[11] G.D. Moss, V.P. Pasko, N. Liu, G. Veronis, J. Geophys. Res. Space Phys. 111 (A2) (2006) A02307, https://doi.org/10.1029/2005JA011350, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2005JA011350.

[12] C. Li, U. Ebert, W. Hundsdorfer, J. Phys. D, Appl. Phys. 42 (20) (2009) 202003, https://doi.org/10.1088/0022-3727/42/20/202003.

[13] O. Chanrion, T. Neubert, J. Geophys. Res. Space Phys. 115 (A6) (2010) A00E32, https://doi.org/10.1029/2009JA014774, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2009JA014774.

[14] G. Diniz, C. Rutjes, U. Ebert, I.S. Ferreira, J. Geophys. Res., Atmos. 124 (1) (2019) 189–198, https://doi.org/10.1029/2018JD029178.

[15] K.I. Bakhov, L.P. Babich, I.M. Kutsyk, IEEE Trans. Plasma Sci. 28 (4) (2000) 1254–1262, https://doi.org/10.1109/27.893314.

[16] N.G. Lehtinen, N. Østgaard, J. Geophys. Res., Atmos. 123 (13) (2018) 6935–6953, https://doi.org/10.1029/2018JD028646, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JD028646.

[17] J.R. Dwyer, D.M. Smith, S.A. Cummer, Space Sci. Rev. 173 (1) (2012) 133–196, https://doi.org/10.1007/s11214-012-9894-0.

[18] N. Østgaard, T. Neubert, V. Reglero, K. Ullaland, S. Yang, G. Genov, M. Marisaldi, A. Mezentsev, P. Kochkin, N. Lehtinen, D. Sarria, B.H. Qureshi, A. Solberg, C. Maiorana, K. Albrechtsen, C. Budtz-Jørgensen, I. Kuvvetli, F. Christiansen, O. Chanrion, M. Heumesser, J. Navarro-Gonzalez, P. Connell, C. Eyles, H. Christian, S. Al-nussirat, J. Geophys. Res., Atmos. 124 (24) (2019) 14024–14036, https://doi.org/10.1029/2019JD031214, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019JD031214.

[19] J.R. Dwyer, H.K. Rassoul, M. Al-Dayeh, L. Caraway, A. Chrest, B. Wright, E. Kozak, J. Jerauld, M.A. Uman, V.A. Rakov, D.M. Jordan, K.J. Rambo, Geophys. Res. Lett. 32 (1) (2005) L01803, https://doi.org/10.1029/2004GL021782, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004GL021782.

[20] P.O. Kochkin, A.P.J. van Deursen, U. Ebert, J. Phys. D, Appl. Phys. 48 (2) (2014) 025205, https://doi.org/10.1088/0022-3727/48/2/025205.

[21] J.R. Dwyer, H.K. Rassoul, Z. Saleh, M.A. Uman, J. Jerauld, J.A. Plumer, Geophysical Research Letters 32 (20) (2005) L20809, https://doi.org/10.1029/2005GL024027, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2005GL024027.

[22] O. Chanrion, Z. Bonaventura, D. Çinar, A. Bourdon, T. Neubert, Environ. Res. Lett. 9 (5) (2014) 055003, https://doi.org/10.1088/1748-9326/9/5/055003.

[23] M.H. Kalos, P.A. Whitlock, Monte Carlo Methods, Wiley, 2008.

[24] E. Kunhardt, Y. Tzeng, J. Comput. Phys. 67 (2) (1986) 279–289, https://doi.org/10.1016/0021-9991(86)90263-9, https://www.sciencedirect.com/science/article/pii/0021999186902639.

[25] J. Teunissen, U. Ebert, J. Comput. Phys. 259 (2014) 318–330, https://doi.org/10.1016/j.jcp.2013.12.005, http://www.sciencedirect.com/science/article/pii/S0021999113008048.

[26] A. Sun, M.M. Becker, D. Loffhagen, Comput. Phys. Commun. 206 (2016) 35–44, https://doi.org/10.1016/j.cpc.2016.05.003, http://www.sciencedirect.com/science/article/pii/S0010465516301199.

[27] D. Welch, T. Genoni, R. Clark, D. Rose, J. Comput. Phys. 227 (1) (2007) 143–155, https://doi.org/10.1016/j.jcp.2007.07.015, http://www.sciencedirect.com/science/article/pii/S0021999107003245.

[28] D. Faghihi, V. Carey, C. Michoski, R. Hager, S. Janhunen, C. Chang, R. Moser, J. Comput. Phys. 409 (2020) 109317, https://doi.org/10.1016/j.jcp.2020.109317, http://www.sciencedirect.com/science/article/pii/S0021999120300917.

[29] G. Lapenta, J. Comput. Phys. 181 (1) (2002) 317–337, https://doi.org/10.1006/jcph.2002.7126, http://www.sciencedirect.com/science/article/pii/S0021999102971263.

[30] F. Assous, T.P. Dulimbert, J. Segré, J. Comput. Phys. 187 (2) (2003) 550–571, https://doi.org/10.1016/S0021-9991(03)00124-4, http://www.sciencedirect.com/science/article/pii/S0021999103001244.

[31] O.N. Vassiliev, Monte Carlo Methods for Radiation Transport: Fundamentals and Advanced Topics, 1st edition, Biological and Medical Physics, Biomedical Engineering, Springer, Cham, 2017.

[32] T. Booth, Sample problem for variance reduction in mcnp, Tech. Rep. lA-10363-MS Los Alamos National Laboratory, 10 1985.

[33] T.E. Booth, Nucl. Sci. Eng. 89 (4) (1985) 305–309, https://doi.org/10.13182/NSE85-A18622.

[34] T. Booth, Genesis of the weight window and the weight window generator in mcnp - a personal history, Tech. Rep., Los Alamos, lA-UR-06-5807, 08 2006.

[35] P. Saidi, M. Sadeghi, C. Tenreiro, in: Variance Reduction of Monte Carlo Simulation in Nuclear Engineering Field, InTech, 2013, pp. 154–173, Ch. 7.

[36] T.E. Booth, J.S. Hendricks, Importance estimation in forward Monte Carlo calculations, Nucl. Technol./Fusion; (United States) 5:1.

[37] M.A. Cooper, E.W. Larsen, Nucl. Sci. Eng. 137 (1) (2001) 1–13, https://doi.org/10.13182/NSE00-34.

[38] K.H. Knuth, Digit. Signal Process. 95 (2019) 102581, https://doi.org/10.1016/j.dsp.2019.102581, http://www.sciencedirect.com/science/article/pii/S1051200419301277.

[39] Stanislaw Weglarczyk, ITM Web Conf. 23 (2018) 00037, https://doi.org/10.1051/itmconf/20182300037.

[40] P.T. Luu, T. Tückmantel, A. Pukhov, Comput. Phys. Commun. 202 (2016) 165–174, https://doi.org/10.1016/j.cpc.2016.01.009, http://www.sciencedirect.com/science/article/pii/S001046551600028X.

[41] F. Pelupessy, W. Schaap, R. van de Weygaert, Astron. Astrophys. 403 (2) (2003) 389–398, https://doi.org/10.1051/0004-6361:20030314.

[42] D. Freedman, P. Diaconis, Z. Wahrscheinlichkeitstheor. Verw. Geb. 57 (4) (1981) 453–476, https://doi.org/10.1007/BF01025868.

[43] J.R. Dwyer, Phys. Plasmas 14 (4) (2007) 042901, https://doi.org/10.1063/1.2709652.

[44] G. Gualdrini, P. Ferrari, Radiat. Prot. Dosim. 146 (4) (2011) 425–433, https://doi.org/10.1093/rpd/ncr240, https://academic.oup.com/rpd/article-pdf/146/4/425/4540866/ncr240.pdf.

[45] K. Koura, Phys. Fluids 29 (11) (1986) 3509–3511, https://doi.org/10.1063/1.865826, https://aip.scitation.org/doi/abs/10.1063/1.865826.

[46] S.F. Biagi, Fortran program MagBoltz cross-sections, www.lxcat.net/Biagi, 2015.

[47] A. Schmalzried, A. Luque, J. Geophys. Res., Atmos. 125 (10) (2020) e2019JD031564, https://doi.org/10.1029/2019JD031564, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019JD031564.

[48] A. Schmalzried, Contrived energy-spectrum compaction algorithm code demonstration, https://osf.io/c6wyh, version 1 (Feb. 2021).

[49] A. Haghighat, Monte Carlo Methods for Particle Transport, CRC Press, 2020.

[50] R.R. Coveyou, V.R. Cain, K.J. Yost, Nucl. Sci. Eng. 27 (2) (1967) 219–234, https://doi.org/10.13182/NSE67-A18262.

[51] J. Wagner, E. Blakeman, D. Peplow, in: International Conference on Mathematics, Computational Methods & Reactor Physics, M&C 2009, 2009.