

# Runaway-electron formation and electron slide-away in an ITER post-disruption scenario

A Stahl<sup>1</sup>, O Embréus<sup>1</sup>, M Landreman<sup>2</sup>, G Papp<sup>3</sup> and T Fülöp<sup>1</sup>

<sup>1</sup> Department of Physics, Chalmers University of Technology, Göteborg, Sweden

<sup>2</sup> Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, Maryland 20742, USA

<sup>3</sup> Max Planck Institute for Plasma Physics, Garching, Germany

E-mail: [stahla@chalmers.se](mailto:stahla@chalmers.se)

**Abstract.** Mitigation of runaway electrons is one of the outstanding issues for the reliable operation of ITER and other large tokamaks, and accurate estimates for the expected runaway-electron energies and current are needed. Previously, linearized tools (which assume the runaway population to be small) have been used to study the runaway dynamics, but these tools are not valid in the cases of most interest, i.e. when the runaway population becomes substantial. We study runaway-electron formation in a post-disruption ITER plasma using the newly developed non-linear code **NORSE**, and describe a feedback mechanism by which a transition to electron slide-away can be induced at field strengths significantly lower than previously expected. If the electric field is actively imposed using the control system, the entire electron population is quickly converted to runaways in the scenario considered. We find the time until the feedback mechanism sets in to be highly dependent on the details of the mechanisms removing heat from the thermal electron population.

## 1. Introduction

Runaway electrons pose a severe threat to the safety and reliability of ITER and other high-plasma-current fusion devices [1]. The larger the runaway-electron population, the larger the threat to the integrity of the device. However, if the electron momentum-space distribution function becomes highly non-Maxwellian due to the presence of a high-energy tail of runaways, existing numerical tools employing linearized collision operators are no longer valid. The same is true if the electric field (even momentarily) becomes comparable to the Dreicer field [2].

We recently presented **NORSE** [3] – an efficient solver of the kinetic equation in a homogeneous plasma – which includes the full relativistic non-linear collision operator of Braams & Karney [4, 5]. **NORSE** – which will be discussed in Section 2 – is able to model Dreicer and hot-tail runaway generation in the presence of electric fields of arbitrary strength and synchrotron-radiation reaction: one of the most important energy-loss channels for runaways [6].

Since **NORSE** is able to treat highly distorted distributions, a range of new questions may be addressed. One issue of particular interest is: will non-linear phenomena accelerate or dampen the growth of runaways? Naturally, this is of great importance in view of ITER and other large tokamaks, as it potentially impacts the requirements on the disruption mitigation system (the design of which is currently being finalized) [1, 7]. In addition, the electric field is expected to reach values as high as 80-100 V/m during the current quench in ITER [8], and runaway



generation is likely to be strong enough for 60% or more of the plasma current to be converted to runaway current. In Section 3, we use NORSE to model the evolution of the electron population in a typical ITER post-disruption scenario.

If the electric field is strong enough, the net parallel force experienced by electrons due to the electric field and collisions becomes positive in the entire momentum space, leading to a phenomenon known as electron *slide-away*. This is expected to happen when  $E > 0.215E_D \equiv E_{SA}$ , where  $E_D = ne^3 \ln \Lambda / 4\pi\epsilon_0^2 T$  is the so-called Dreicer field [2];  $n$ ,  $T$  and  $-e$  are the electron number density, temperature and charge;  $\ln \Lambda$  is the Coulomb logarithm; and  $\epsilon_0$  is the vacuum permittivity. The associated surge in runaway current can have a large impact on the potential for material damage, as well as the subsequent evolution of the parallel electric field. The slide-away process cannot be consistently modelled using linear tools such as CODE [9, 10] or LUKE [11], which assume a Maxwellian background plasma and therefore require  $E \ll E_{SA}$ , as well as that the runaway fraction is small  $n_r/n \ll 1$ .

A strong electric field represents a source of energy that quickly heats up the electron distribution. This heating can induce a transition to the slide-away regime – even under a fixed applied electric field which is initially below the threshold  $E < E_{SA}$  – since the collisional friction is lower in a hotter distribution. As a consequence, the Dreicer field is also lower, making the effective normalized field  $E/E_{D,\text{eff}}$  higher for a given field strength. If the temperature increase is large enough, the slide-away regime is reached, which happens at a field of  $E/E_{D,\text{eff}} = 0.215$  in the case of a constant applied electric field, i.e. it coincides with the standard slide-away field at the effective temperature  $T_{\text{eff}}$  [3].

In practice, many processes act to remove heat from the plasma. In a cold post-disruption plasma, line radiation and bremsstrahlung from interactions with partially ionized impurities are important loss channels, as is radial heat transport. Including a heat sink in numerical simulation of such scenarios is therefore desirable, and the sink effectively acts to delay or prevent the transition to slide-away. In this paper, we demonstrate that the evolution of the runaway electron population – including the time to reach slide-away – is highly sensitive to the properties of the applied heat sink, making a detailed investigation of the various loss channels an area of interest for future work.

## 2. NORSE

We will use the newly developed fully relativistic non-linear tool NORSE [3] to study the dynamics of the electron population. NORSE, which is valid in spatially uniform plasmas, solves the kinetic equation

$$\frac{\partial f}{\partial t} - \frac{e\mathbf{E}}{m_e c} \cdot \frac{\partial f}{\partial \mathbf{p}} + \frac{\partial}{\partial \mathbf{p}} \cdot (\mathbf{F}_s f) = C_{ee}\{f\} + C_{ei}\{f\} + S, \quad (1)$$

where  $f$  is the electron distribution function,  $t$  is the time,  $m_e$  and  $\mathbf{p}$  are the electron rest mass and momentum,  $\mathbf{E}$  is the electric field,  $c$  is the speed of light,  $\mathbf{F}_s$  is the synchrotron-radiation-reaction force,  $C_{ee}$  is the relativistic non-linear electron-electron collision operator,  $C_{ei}$  is the electron-ion collision operator, and  $S$  represents heat and particle sources or sinks. For a detailed description of the various terms and operators, see Ref. [3]. For the remainder of this paper, we define the electric field such that electrons are accelerated in the positive  $p_{\parallel}$  direction.

In NORSE, the particle momentum  $\mathbf{p}$  is represented in terms of the magnitude of the normalized momentum  $p = \gamma v/c$  (where  $v$  is the velocity of the particle and  $\gamma$  is the relativistic mass factor) and the cosine of the pitch angle  $\xi = p_{\parallel}/p$ . The kinetic equation is discretized using finite differences in both  $p$  and  $\xi$ . A linearly implicit time-advancement scheme is used, where the five relativistic Braams-Karney potentials [4] – analogous to the Rosenbluth potentials in the non-relativistic case – are calculated explicitly from the known distribution. These are then used to construct the electron-electron collision operator  $C_{ee}$ , and the remainder of the kinetic equation is solved implicitly.

For the results presented in this paper, a non-uniform finite-difference grid was used to improve the computational efficiency. In the pitch-angle coordinate, the grid points were chosen with a dense spacing close to  $\xi = \pm 1$  – in particular at  $\xi = 1$  where the runaway tail forms – but with a sparser grid at intermediate  $\xi$ . In the  $p$ -direction, a grid mapping with a tanh step in spacing was used in order to produce a grid with dense spacing at low momenta (to accurately resolve the bulk dynamics), and larger spacing in the high-energy tail, where the scale length of variations in  $f$  is larger.

The runaway region in NORSE was determined by studying particle trajectories in phase space, neglecting momentum-space diffusion but including self-consistent collisional friction and synchrotron-radiation reaction [3]. The trajectory that terminates at  $\xi = 1$  and  $p = p_c$  marks the lower boundary of the runaway region, since particles that follow it neither end up in the bulk nor reach arbitrarily high energies. The critical momentum  $p_c$  is the momentum at which the balance of forces in the parallel direction ( $\xi = 1$ ) becomes positive (i.e. the lowest momentum at which the accelerating force of the electric field overcomes the collisional and synchrotron-radiation-reaction drag). If the balance of forces is positive for all  $p$  at  $\xi = 1$ , however, all electrons experience a net acceleration, and the population is in the slide-away regime.

### 2.1. Heat sink

Including a heat sink (HS) in the numerical simulations is of great importance for accurate modelling of the distribution evolution during a disruption. The heat sink used in NORSE to remove heat from the thermal population has the form  $S = \partial/\partial\mathbf{p} \cdot (k_h \mathbf{S}_h f)$ , where  $\mathbf{S}_h(p)$  is an isotropic function of momentum (i.e.  $\mathbf{S}_h \parallel \mathbf{p}$ ) and  $k_h(t)$  is the magnitude of the source. The terms in the kinetic equation that affect the total energy content are the electric-field and synchrotron-radiation-reaction terms, however; when considering a subset  $\Omega$  of momentum space, collisions can also transfer energy in or out of  $\Omega$ , and a corresponding term must be included. The total energy change  $dW/dt$  in  $\Omega$  can thus be written as

$$\frac{dW}{dt} = m_e c^2 \int_{\Omega} d^3p (\gamma - 1) \left( -\frac{e\mathbf{E}}{m_e c} \cdot \frac{\partial f}{\partial \mathbf{p}} + \frac{\partial}{\partial \mathbf{p}} \cdot (\mathbf{F}_s f) - C\{f\} + k_h \frac{\partial}{\partial \mathbf{p}} \cdot (\mathbf{S}_h f) \right). \quad (2)$$

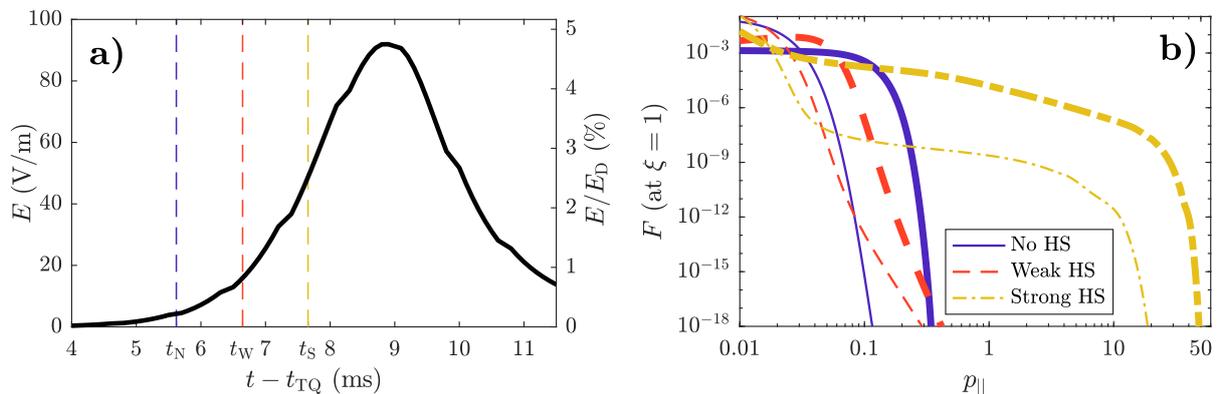
The magnitude  $k_h$  of the sink in each time step can be determined by requiring  $dW/dt = 0$ . In this work, we take  $\Omega$  to represent the thermal bulk of the distribution, which we define as all particles with  $v < 4v_{th,0}$ , with  $v_{th,0} = \sqrt{2T_0/m_e}$  the thermal speed at the initial temperature  $T_0$ .

In this study, the  $p$ -component of  $\mathbf{S}_h$  was chosen to have the shape of a Maxwellian at the desired temperature  $T$ . In practice, the momentum dependence of the sink will be more complicated and subject to the details of the particular physical processes at work. It will also likely have a limited energy-removal rate, dictated by for instance spatial gradients or impurity content, which could limit its efficiency in maintaining a given temperature. A detailed investigation of the characteristics of the sink is left for future work; the aim of this paper is to highlight the sensitivity of the runaway-electron evolution to the particulars of the sink, and for that purpose we will impose a limit on the energy-removal rate, as will be discussed in the next section.

## 3. Runaway generation in an ITER disruption

### 3.1. Post-disruption scenario

In this section, we use NORSE to model the evolution of the electron distribution during a typical ITER disruption. The electric field evolution (which is shown in Fig. 1a) and other parameters are taken from the ITER inductive scenario no. 2 in Ref. [8], but the temperature evolution has been simplified to facilitate the numerical calculation. We assume the electron population to be completely thermalized and use the final temperature  $T = T_0 = 10$  eV throughout our

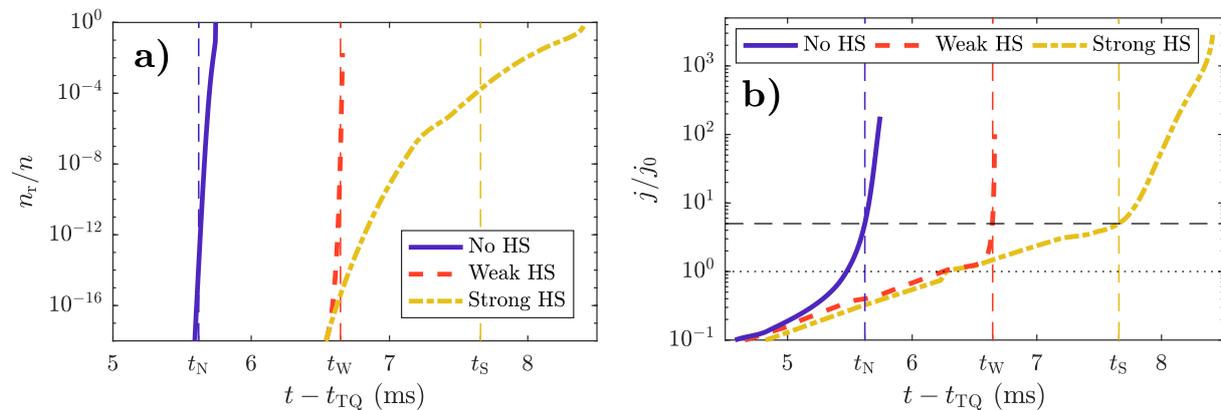


**Figure 1.** a) Electric field in V/m (left vertical axis) and normalized to the Dreicer field  $E_D$  at the temperature  $T_0$  and density  $n_0$  (right vertical axis), as a function of time after the thermal quench. b) Tail of the parallel electron distribution. Thin lines show  $f$  at  $t_N$  (no HS),  $t_W$  (weak HS) and  $t_S$  (strong HS), and thick lines show  $f$  immediately before the transition to slide-away.

simulation, together with the density  $n = n_0 = 7.1 \cdot 10^{19} \text{ m}^{-3}$ . This is likely to underestimate the runaway-electron generation in the early phase, in which the temperature is still dropping, however the chosen set-up is sufficient for our purposes. We use the magnetic field on axis ( $B = 5.3 \text{ T}$ ) and  $Z_{\text{eff}} = 1$ . The initial current density in the scenario is  $j_0 = 0.62 \text{ MA/m}^2$ , however our calculations start from a Maxwellian distribution and make no attempt to maintain the experimental current evolution explicitly.

To highlight the importance of the temperature evolution of the bulk, in Section 3.2 we will consider three scenarios: *no heat sink* (subscript N), *weak heat sink* (W) and *strong heat sink* (S). In the no-heat-sink scenario, all the energy supplied by the electric field will remain in the simulation, leading to rapid bulk heating; with the strong heat sink, a bulk temperature of  $T_0 = 10 \text{ eV}$  will be enforced in accordance with Eq. 2, i.e. any excess heat in the bulk region will be removed using a heat sink. In the intermediate case of a weak heat sink, the energy-removal rate of the heat sink will be restricted to  $0.5 \text{ MW/m}^3$ . This particular value has been chosen at will, but is meant to represent some inherent limitation in the physical processes responsible for the energy loss. In both the weak and strong cases, the heat sink will affect only the thermal population, allowing the supra-thermal tail to gain energy from the electric field. Physically, this corresponds to processes not included in the simulation (such as radial transport or radiative losses) which primarily affect the thermal population.

NORSE simulations of the evolution of the electron distribution function in the presence of the electric field in Fig. 1a were performed for the three different scenarios. The simulations were aborted when the runaway population reached  $n_r/n = 1$ ; i.e. a transition to the slide-away regime was observed. The simulation results can however only be considered characteristic of a natural ITER disruption for current densities comparable to, or somewhat larger than, the initial value  $j_0$ , since after that point the strong response of the inductive electric field to the increased local current would invalidate the  $E$ -field evolution used. We will therefore mark the time where the current density reaches  $j/j_0 > 5$  in all plots, and denote it with  $t_N$ ,  $t_W$  and  $t_S$ , respectively, for the no-sink, weak-sink and strong-sink scenarios. The distribution evolution at later times can only be considered accessible in scenarios where the loop voltage is actively sustained using the control system. Nevertheless, this regime will turn out to be of interest, since a non-linear feedback mechanism leading to a rapid transition to slide-away is observed.



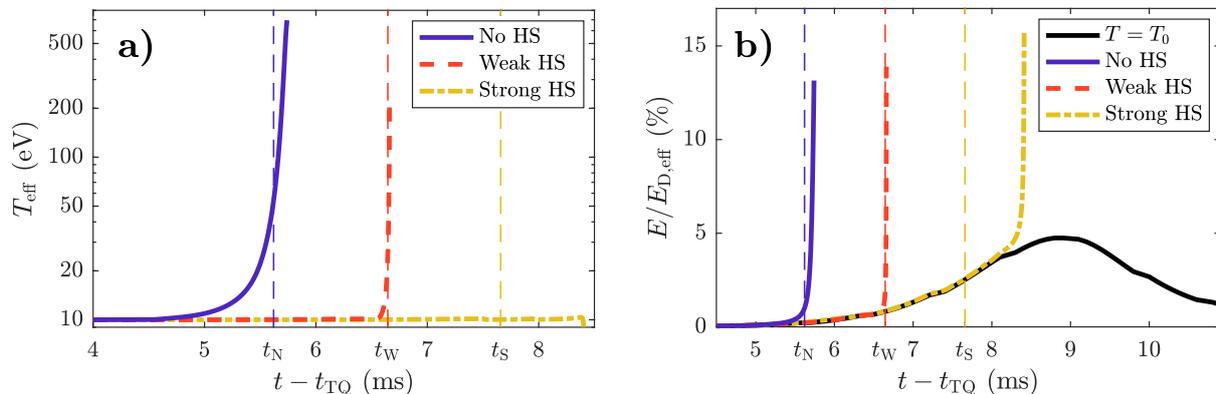
**Figure 2.** a) Runaway fraction and b) current density normalized to its initial value  $j_0$ , as a function of time after the thermal quench in the different heat-sink scenarios. The times  $t_N$ ,  $t_W$  and  $t_S$  (vertical thin dashed lines), mark the time where the current density reaches  $j/j_0 > 5$  for the no-heat-sink, weak-heat-sink and strong-heat-sink scenarios, respectively.

### 3.2. Evolution of the runaway-electron population

In Fig. 1b, the tails of the distributions in the parallel direction are shown at  $t_N$ ,  $t_W$  and  $t_S$  (thin lines), as well as at the final times (thick lines) in each scenario (just before the transition to slide-away is reached). In the figure, the distribution is normalized such that  $F = f/f_M(t=0, p=0)$ , where  $f_M$  is a Maxwell-Jüttner distribution, so that  $F$  initially takes the value unity at  $p = 0$ . The maximum achieved particle energies are highly dependent on whether a heat sink was applied or not; in the no-heat-sink and weak-heat-sink scenarios, the particles did not have time to reach relativistic energies, whereas in the strong heat sink case,  $p \approx 19$  and  $p \approx 44$  (corresponding to energies of roughly 9 and 22 MeV), were obtained at  $t_S$  and just before reaching slide-away, respectively. The reason for this is that, as we shall see, the current density growth and subsequent transition to slide-away in the latter case occur at much later times, and the runaways have time to gain more energy.

Figure 2a shows the evolution of the runaway fraction during the course of the simulation. In the no-heat-sink case, the runaway fraction increases sharply, but as shown in Fig. 1b, the runaways are all at low energy. The transition to slide-away happens already at  $t = 5.7$  ms; early on in the electric-field evolution (cf. Fig. 1). In the two scenarios employing a heat sink, the growth in runaway fraction occurs later, but in the weak-heat-sink case the growth rate is comparable to the case when a sink is absent once the process is initiated. In this case, the transition to slide-away happens at  $t = 6.7$  ms. With the strong heat sink, the runaway population grows steadily, eventually dominating the entire distribution at  $t = 8.4$  ms, however in this case the transition is gradual, rather than rapid. In all three scenarios, including the one with an ideal strong heat sink, the slide-away regime is thus reached even before the electric field (calculated assuming a linear treatment) has reached its peak.

Figure 2b shows the evolution of the current density. It indicates that the rapid increase in the runaway fraction is correlated with a similar increase in the current density, although in the no-sink case, the growth rate is somewhat smaller. Again, the growth in the strong-heat-sink case is gradual, rather than explosive. Note that in the no and weak heat-sink cases, the runaway fraction is still negligibly small at the start of the rapid transition to slide-away. The transition is thus not a non-linear phenomenon triggered by the size of the runaway population; it starts in a regime where linearized tools are normally expected to be valid, and before the current density becomes significantly larger than its initial value.



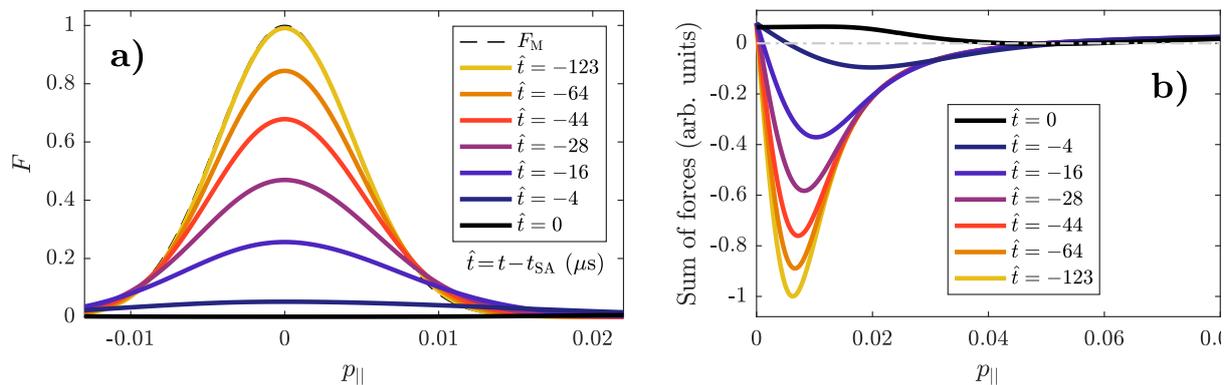
**Figure 3.** a) Effective temperature of the bulk population and b) corresponding effective normalized electric-field strength (also taking changes to the bulk density into account) as functions of time, in the different scenarios. The black solid line in panel b) corresponds to the normalized field with  $T = T_0$  and  $n = n_0$  (cf. Fig. 1).

The explanation can be found by examining the thermal population. By comparing the energy moment of the bulk of the distribution ( $W_{\Omega}$ ) with that of a relativistic Maxwellian ( $W_{\text{M}}(T)$ ),

$$W_{\Omega} = m_e c^2 \int_{\Omega} d^3 p (\gamma - 1) f = W_{\text{M}} = \frac{m_e c^2 n}{\Theta_{\text{eff}} K_2(1/\Theta_{\text{eff}})} \int_0^{p_{\text{max},\Omega}} dp p^2 (\gamma - 1) \exp\left(-\frac{\gamma}{\Theta_{\text{eff}}}\right), \quad (3)$$

an effective temperature  $T_{\text{eff}}$  for a given distribution  $f$  can be determined by solving for  $\Theta_{\text{eff}} = T_{\text{eff}}/m_e c^2$ . In the above equation,  $p_{\text{max},\Omega}$  is the upper boundary in  $p$  of the bulk region in momentum space, and  $K_2(x)$  is the modified Bessel function of the second kind (and order two). The effective temperature  $T_{\text{eff}}$  is plotted in Fig. 3a as a function of time for all three scenarios. In the no-heat sink case, it increases by roughly two orders of magnitude during the simulation, i.e. when energy is not actively removed from the system, the Ohmic heating is sufficient to heat the plasma to a temperature of about 700 eV before the onset of slide-away, or 55 eV if the electric field is not artificially sustained. A similar (albeit weaker) tendency is seen in the weak-heat-sink case, where the temperature increases to  $T_{\text{eff}} \approx 210$  eV in the phase leading up to the slide-away transition, or 25 eV in a non-driven case. This heating is a consequence of the imposed limited maximum energy-removal rate of the heat sink in the weak case, since the temperature is efficiently kept constant in the beginning of the simulation, where  $(dW/dt)_{\text{HS}} < 0.5 \text{ MW/m}^3$ . The strong heat sink manages to keep  $T_{\text{eff}} - T_0$  to within a few tenths of an eV during the entire simulation, corresponding to a source with unlimited (or at least higher than required) maximal energy-removal rate.

The significance of the observed bulk heating is its influence on the Dreicer field  $E_{\text{D}}$ , which (apart from the weak dependence on  $\ln \Lambda$ ) is inversely proportional to  $T$ . For a given electric-field strength, the normalized field  $E/E_{\text{D}}$  thus increases as the bulk heats up. The effective normalized electric field is shown in Fig. 3b, indicating that the rapid growth in the runaway fraction in Fig. 2a is correlated with a sudden increase in the normalized electric field in the no-heat-sink and weak-heat-sink cases. In the strong-heat-sink case, the increase in normalized field is not caused by the temperature, which is kept constant during the entire simulation, but by the decrease in the bulk density as the runaway population becomes substantial. As can be seen from the yellow dash-dotted line in the figure, this has a similar effect as a temperature increase, since  $E_{\text{D}} \sim n_{\text{bulk}}$ . The effective  $E/E_{\text{D}}$  starts to deviate from the baseline value (black solid line)



**Figure 4.** a) Evolution of the bulk of the electron distribution and b) balance of forces, in the direction parallel to the electric field, just before the transition to slide-away in the weak-heat-sink scenario.  $\hat{t}$  is the time relative to the transition to slide-away in  $\mu\text{s}$ .

already when  $n_{\text{bulk}}/n_0 \approx 0.97$ . The feedback process is thus initiated when the runaway fraction is just 3%; a regime where linear tools are expected to be valid.

The two effects of increasing  $T_{\text{eff}}$  and decreasing  $n_{\text{bulk}}$  lead to a positive feedback mechanism which is responsible for the rapid growth in runaway fraction and current density seen in the no-heat-sink and weak-heat-sink cases. Once the bulk temperature has increased enough (or a high enough runaway tail is produced, as seen in the strong-heat-sink case), the normalized electric field becomes strong enough to cause a depletion of the bulk through primary runaway generation. The reduced bulk density in turn leads to a more efficient heating of the remaining bulk particles. Both of these effects contribute to a reduction in the Dreicer field  $E_D$ , and a corresponding reduced collisional friction on the bulk electrons, which makes runaway acceleration easier. This further increases the rate of bulk depletion, and so on. Eventually the friction becomes low enough that the parallel balance of forces becomes positive everywhere, marking the transition to the slide-away regime. At this point, the bulk of the distribution can no longer be well described by a Maxwellian and the positive feedback mechanism makes the transition possible even though  $E/E_{D,\text{eff}} < 0.215$  (in this case at around  $E/E_{D,\text{eff}} \approx 0.15$ ).

The bulk depletion and associated change in the parallel force balance is shown in Fig. 4a and Fig. 4b, respectively, in the phase leading up to the transition to slide-away in the weak-heat-sink case. Initially, the force balance is positive in the tail – i.e. particles there experience a net acceleration – while it is negative in the bulk, meaning particles are slowed down by collisions. As the feedback process starts, the minimum in the sum of forces becomes gradually less pronounced, in tandem with the depletion of the bulk population, up until the point where the sum of forces becomes positive everywhere, and the slide-away regime is reached. This highly non-linear process cannot be accurately captured by a linear model.

#### 4. Discussion and conclusions

In this paper we have examined the evolution of the electron distribution and runaway generation in an ITER-like post-disruption scenario where the electric fields reach values as high as 90 V/m. With the help of the newly developed tool NORSE, which is a relativistic non-linear solver for the electron momentum-space distribution function, we have shown that the slide-away regime, i.e. a net parallel acceleration of electrons in all of momentum space, is reached in this scenario, provided the electric field evolution used is artificially enforced by the control system. In the stage leading up to the transition, a positive feedback mechanism sets in by means of which the bulk quickly gets depleted by primary runaway generation which reduces the friction on the

thermal population, leading to further bulk depletion, until the point where the slide-away is reached. This process can be initiated at significantly weaker fields than the slide-away field  $E_{SA}$  expected from linear theory.

The time to a transition to slide-away is highly dependent on the ability of loss processes to remove heat from the thermal electron population, but even with an ideal sink (the strong-heat-sink case in Section 3), complete runaway generation was seen 8.4 ms after the thermal quench. These results were obtained without taking avalanche or hot-tail runaway generation into account, which would only lead to more prominent runaway growth.

Also in the case of a disruption where the electric field is not artificially sustained, strong bulk heating leads to a rapid growth in the runaway fraction because of the increase in the normalized electric field  $E/E_D$ , but the current density becomes large enough to significantly affect the electric field evolution (suppressing the growth of  $E$ ) before the slide-away regime is reached. This is observed in the absence of a heat sink, as well as with a heat sink with a limited maximum energy-removal rate (in this case  $0.5 \text{ MW/m}^3$ ). If the efficiency of the heat sink is not limited, the runaway fraction grows more slowly and the runaways have time to reach significantly higher energies before the electric field becomes affected by the growing current density. The severity of disruptions in ITER could thus be greatly affected by the properties of the heat sinks present in the plasma.

The feedback mechanism described in this paper has important consequences for the understanding of runaway-electron dynamics. With the entire electron population experiencing a net accelerating force at much weaker electric fields than previously expected, very large runaway-electron current generation is likely. This would impact the subsequent electric-field evolution, leading to a reduction in field strength and duration which could occur at relatively early times if the heat sink has a limited energy removal rate. Therefore, it is difficult to determine the magnitude of the effect on the current evolution and post-quench dynamics without a self-consistent calculation of the electron distribution and the electric field. Nevertheless, this paper shows that feedback effects play an important role in post-disruption runaway dynamics, and that the details of the heat-loss channels may have a big impact on what strength and duration of electric field can be tolerated before the positive feedback, and possible subsequent transition to slide-away, is induced.

## Acknowledgments

This work was supported by the Swedish Research Council (Dnr. 2014-5510), the European Research Council (ERC-2014-CoG grant 647121), and the Knut and Alice Wallenberg Foundation. M.L. was supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science, under Award Numbers DE-FG02-93ER54197 and DE-FC02-08ER54964.

## References

- [1] Hollmann E M, Aleynikov P B, Fülöp T, Humphreys D A, Izzo V A, Lehnen M, Lukash V E, Papp G, Pautasso G, Saint-Laurent F and Snipes J A 2015 *Physics of Plasmas* **22** 021802
- [2] Dreicer H 1959 *Physical Review* **115**(2) 238–249
- [3] Stahl A, Landreman M, Embréus O and Fülöp T 2016 *Submitted to Computer Physics Communications* URL <http://arxiv.org/abs/1608.02742>
- [4] Braams B J and Karney C F F 1987 *Physical Review Letters* **59** 1817–1820
- [5] Braams B J and Karney C F F 1989 *Physics of Fluids B: Plasma Physics* **1** 1355–1368
- [6] Stahl A, Hirvijoki E, Decker J, Embréus O and Fülöp T 2015 *Physical Review Letters* **114** 115002
- [7] Boozer A H 2015 *Physics of Plasmas* **22** 032504
- [8] Smith H, Helander P, Eriksson L G, Anderson D, Lisak M and Andersson F 2006 *Physics of Plasmas* **13** 102502
- [9] Landreman M, Stahl A and Fülöp T 2014 *Computer Physics Communications* **185** 847 – 855 ISSN 0010-4655
- [10] Stahl A, Embréus O, Papp G, Landreman M and Fülöp T 2016 *Nuclear Fusion* **56** 112009
- [11] Decker J and Peysson Y 2004 DKE: A fast numerical solver for the 3D drift kinetic equation Tech. Rep. EUR-CEA-FC-1736 Euratom-CEA