Stochastic first passage time accelerated with CUDA

Vincenzo Pierro, Luigi Troiano, and Elena Mejuto
Dept. of Engineering, University of Sannio, Corso Garibaldi, 107, I-82100 Benevento, Italy

Giovanni Filatrella
Dept. of Sciences and Technologies, University of Sannio, Via Port'Arsa, 11, I-82100 Benevento, Italy

The numerical integration of stochastic trajectories to estimate the time to pass a threshold is an interesting physical quantity, for instance in Josephson junctions and atomic force microscopy, where the full trajectory is not accessible. We propose an algorithm suitable for efficient implementation on graphical processing unit in CUDA environment. The proposed approach for well balanced loads achieves almost perfect scaling with the number of available threads and processors, and allows an acceleration of about $400\times$ with a GPU GTX980 respect to standard multicore CPU. This method allows with off the shell GPU to challenge problems that are otherwise prohibitive, as thermal activation in slowly tilted potentials. In particular, we demonstrate that it is possible to simulate the switching currents distributions of Josephson junctions in the timescale of actual experiments.

I. INTRODUCTION

The passage across a threshold of a stochastic process, the First Passage Time (FPT) process [1], is a valuable tool to study theoretical properties of random systems (e.g. Kramers' rate theory [2, 3]). From the experimental point of view, in Josephson Junctions (JJ) superconducting devices [4] the measurement of the threshold for the escapes is the only possibility (or the simplest way) to gain information on the internal dynamics of systems as quantum devices [5, 6], threshold detectors [7–9] and arrays of JJ [10, 11] of great interest for high frequency, up to the THz region, local oscillators [12]. The threshold for escapes is an essential physical information also in diverse systems . Examples are the atomic force spectroscopy (AFS) [13], in some cases characterized by energy barrier similar to the washboard potential [14] and the fast reversal of nanoparticles in a thermal bath [15].

FPT process differs from standard stochastic evolution [16], for the physically interesting quantity is no more some dynamical observable (e.g., the position of a particle) as a function of time, but the (random) time to reach a given coordinate. Thus, while in the stochastic evolution one is interested in the probability distribution of the positions at a given time, in the analysis of passage times one deals with the probability distribution of the time necessary to reach a given position. This is necessary when there exist special points where something happens that can be recorded by the instruments, as it is the case in the JJ physics that we discuss in this work. This inversion of the roles of time and position calls for a different approach to parallel calculation, and prevents the straightforward application of already existing stochastic parallelization for the determination of FPTs. Computationally efficient methods are advantageous in JJ physics for a number of reasons. First, measurements are performed while changing the bias (that is, in the mechanical analogue, tilting the poten-

tial), and thus the resulting nonequilibrium problem requires numerical simulations. Second, the bias change occurs on the scale of the conventional "DC" electronics, typically in the scale of kHz, while the characteristic time scale for the Josephson dynamics is in the range $0.1 \div 1$ THz, thus realistic and accurate simulations require something like $10^9 \div 10^{12}$ integrations steps. Third, if JJ are to be used as signal detectors, the tests require extensive simulations to accurately retrieve the response to the binary hypothesis. Fourth, JJ can be employed to form larger structures, or arrays and extended or long junctions where vortices are nucleated [17], that evidently increase the computational burden as many JJs are employed. Fifth, as will be discussed in details in Sect. IV A, the estimate of the quasipotential [18, 19] requires simulations for extremely rare escape events. However, all the above problems amount to the collection of random exit times; If stochastic replicas can be efficiently assigned to different processor units one straightforwardly benefits of an accelerated integration. A possible method to achieve efficient distributed computation is the topic of this work.

The call for a specific algorithm for the numerical evaluation of FPTs can be summarized in a rough physical and intuitive way as follows. If the representative trajectory of a system is computed by a single processor (with sequential algorithms), in both stochastic evolution and random FPTs the processor just integrates the stochastic equation, and when the goal is reached (either the final position or the passage time, that are the desired information) the computation is arrested and a new calculation can begin. Thus, in single processor calculations one can use with little changes the algorithm performing the integration of stochastic equations also for the computation of the FPTs. Instead, for parallel computing the termination times on the processors are different, inasmuch the stochastic nature of the process entails different FPTs. It is the lack of synchrony between the processors that limits the

efficiency of parallel computation. Our purpose is to device an efficient parallel algorithms that assigns a new job to the processor as soon as it has terminated its task, and to show how this can be efficiently done on a GPU in a CUDA environment. There is a huge literature for parallel solution of stochastic differential equations where an arrival time is to be calculated, see for instance [20] and references therein. The main point is that the performances are strongly related to the communication time, and hence to the processor architecture. Therefore, the knowledge accumulated for traditional parallel supercomputing cannot be directly transferred to CUDA environment.

The work is organized as follows. In Sect. II we discuss some physical examples where an accelerated (and cheap) procedure might prove useful, with special reference in Sect. II A to the model equations for superconducting electronics. The details of the computational parallel solution (CUDA oriented) are given in Sect. III, and the results of test simulations, together with the observed performances, are collected in Sect. IV. The conclusions are in the last Sect. V.

II. THE PHYSICAL PROBLEM

To focus on a specific physical setting, we discuss a practical realization of the washboard potential, a superconductive JJ [4]. For this device, the FPT is of interest inasmuch it is associated to a measurable quantity. In JJ the trajectory it is not available because its measure is incompatible with the Heisenberg principle, for the intrinsic quantum nature of the representing coordinate. What is experimentally accessible is the voltage, that is proportional to the velocity, associated to finite and large jumps, for the time scale in the THz region makes it prohibitive to follow the voltage time evolution. The voltage jumps only occur when the phase passes the maximum of the metastable (washboard like) potential and enters the running regime, see Fig. 1. This exit across the separatrix is naturally formulated as a FPT on the top of the potential barrier and can be simulated with a GPU in CUDA environment. Besides the general interest in FPT, there is an additional reason to pay a special attention to JJ. The well established [21] noisy dynamics of Josephson junctions is of great contemporary interest [7, 22–24] in connection with macroscopic quantum tunneling [25–32]. The quantum behavior of JJ makes them a good candidate for the realization of a quantum bit; however, the quantum regime is only observed after the ordinary thermal activated regime has been tamed [33]. Thus, it is of interest to have precise numerical simulations of the thermal regime, simulations that are deeply connected with the theoretical predictions (and corrections) of Kramers' formula in the underdamped regime [3].

Finally, we wish to stress that there are other physical contexts where a similar setting is of practical interest. For instance, AFS is also due to the passage of the dip over the maximum of the atomic interaction potential [13, 14]. Interestingly, the model potential for AFS is often of the same type as the JJ, namely the washboard potential. The description of the stochastic motion in a washboard potential is the subject
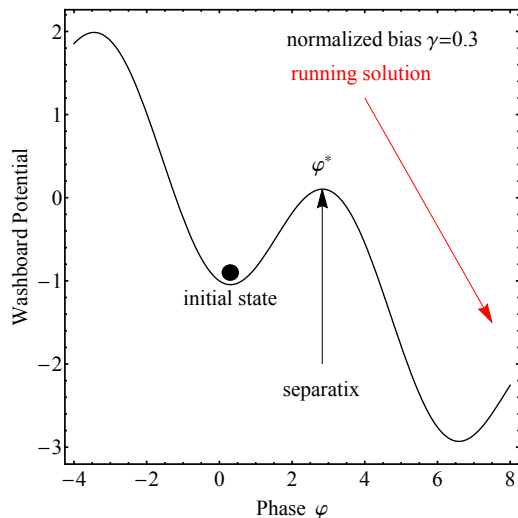

FIG. 1. (color online) Sketch of the washboard potential. The circle indicates the initial position of the virtual mechanical particle associated, for instance, to the quantum phase φ of a JJ. The φ^* value is the threshold indicating the separatrix. The red arrow shows the running state of the system after the escape over the barrier. Retrapping in subsequent minima never occurs in practice if the system is underdamped.

of the next Section.

A. The washboard model: basic equations

In this Section we describe the basic equations for the thermal stochastic washboard potential of Fig. 1.

To be specific, we derive the equations with reference to JJ. The usual model for a point-like JJ reads [4]:

$$\frac{C\hbar}{2e}\frac{d^2\varphi}{dt'^2} + \frac{1}{R}\frac{\hbar}{2e}\frac{d\varphi}{dt'} + I_0\sin\varphi = I_b + I(t') \qquad (1)$$

that includes inertia (determined by the capacitance C), dissipation (as governed by the dissipative element R) and fluctuations (the random current $I(t')$ supplied by the resistance), the nonlinear periodic term (the oscillating current of amplitude I_0), and a constant bias I_b. Here, as usual, $\hbar$ is the reduced Planck constant, and e is the electron charge. Fluctuations are assumed to be Gaussian with:

$$\langle I(t') \rangle = 0, \qquad (2)$$

$$\langle I(t')I(s') \rangle = \frac{2k_BT}{R}\delta(t'-s'), \qquad (3)$$

where k_B is the Boltzmann constant, T is the absolute temperature, $\delta(\cdot)$ the Dirac function, $\langle\cdot\rangle$ is the expectation operator. The normalized units [6] are the following. The current is normalized to the critical current I_0:

$$\gamma = \frac{I_b}{I_0}. \qquad (4)$$

The time is normalized to the quantity $\hbar/M$, where $M = C(\hbar/2e)^2$ is the effective mass of the junction. We introduce the time $\omega_N^{-1} = \hbar/E_C$, where E_C is the Coulomb energy $E_C = (2e)^2/C$, or the energy of the condenser charged by the elementary Cooper's pair. We have therefore:

$$t = t' \frac{E_C}{\hbar} = t' \frac{1}{C} \frac{(2e)^2}{\hbar} = t' \frac{\hbar}{M}. \tag{5}$$

Dissipation is given by the parallel resistor that in dimensionless units becomes

$$\beta = \frac{1}{R} \frac{\hbar}{(2e)^2}. \tag{6}$$

Also the fluctuating current is normalized to I_0:

$$\gamma_N = \frac{I(t')}{I_0}. \tag{7}$$

The normalized noise amplitude D reads

$$D = \beta \frac{k_B T}{E_C} \tag{8}$$

that is, the thermal energy scaled to the energy E_C. To evidentiate the fluctuation-dissipation theorem, one can also introduce the normalized dissipation β and the normalized temperature θ in the fluctuation correlator (3), we have:

$$\theta = \frac{k_B T}{E_C} \Rightarrow D = \beta\theta.$$

Using the above normalizations Eq. (1) becomes

$$\frac{d^2\varphi}{dt^2} + \beta\frac{d\varphi}{dt} = V_0\left(-\sin\varphi + \gamma + \gamma_N\right), \tag{9}$$

where $V_0 = E_J/E_C$ and $E_J = I_0\hbar/(2e)$ is the energy barrier of phase particle. The statistical properties of the Gaussian (thermal) noise normalized current $\gamma_N(t)$ are fully described by

$$\langle\gamma_N(t)\rangle = 0, \tag{10}$$
$$\langle\gamma_N(t)\gamma_N(t+\Delta t)\rangle = 2D\,\delta\left(\Delta t\right) = 2\beta\theta\,\delta\left(\Delta t\right). \tag{11}$$

Eqs. (9,10,11) constitute the physical setting in which we describe the method for accelerated computation of the FPTs.

B. The connection between escape times and switching currents probability distribution

The evolution of a JJ according to the Langevin Eqs. (9,10,11) is the basis to evaluate the exit time, that is the practically available physical quantity. Schematically, the trajectories that determine the random exit times are represented in Fig. 2. The physically interesting quantity is the time at which the phase φ hits the separatrix, because in the Josephson case this is the time at which a sizable voltage appears. This sudden passage from a static metastable solution to a running solution
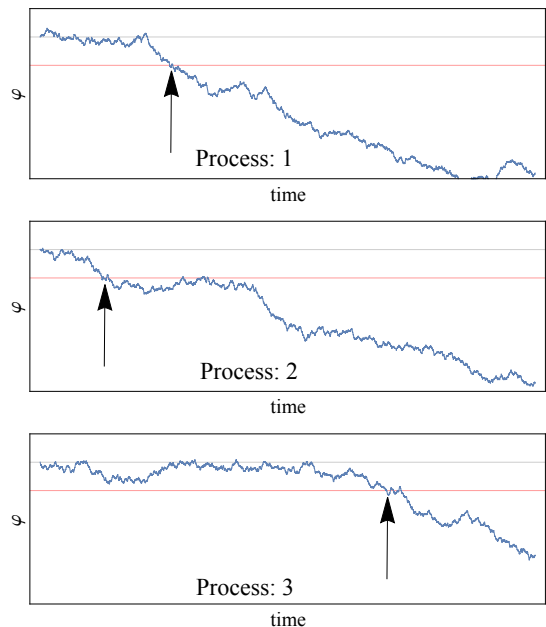

FIG. 2. (color online) Sketch of independent parallel stochastic evolutions, each belonging to a different process. The upper black horizontal line denotes the initial point, and the lower red line is the separatrix that defines the escapes. The claim that the FPT of a single realization has been computed occurs when the random trajectory touches the red line (i.e., the time instants determined by the black vertical arrows).

across the separatrix φ^*, occurs at random times under the influence of noise. Rather than the direct measurement of the escape time, it is convenient in experimental setup to ramp the bias current, and to read the current at which the passage occurs. Thus, the exit time (or the FPT) becomes a *Switching Current* (SC) if the reverse current probability is negligible [35], which is commonly the case for underamped systems. Also, there is a deep connection, for overdamped systems, between the MFPT and the mean transition time, that is the statistical quantity that corresponds to the switching currents in JJ. However, in the first place, from a computational point of view the algorithms for the MFPT and the mean transition time are essential identical. In the second place, the MFPT are more general and of larger applications in systems outside the JJ physics. We also notice that, if mean transition times are to be used, one can use a modified Kramers formula [36]. The histogram of the so defined SCs represents the most important physical feature, from which a number of properties are deduced. Computationally-wise, the evaluation of the FPT or the switching currents is a rather similar task, as will be shown in Sect. III B.

III. GPU COMPUTATION WITH CUDA

We here demonstrate that it is possible to accelerate the computation of the FPTs, first describing the general struc-

ture in Sect. III A and the logic of the coding in Sect. III B. The last Sect. is devoted to the speed-up analysis.

A. GPU Architecture

As reference we use the NVIDIA Maxwell micro architecture used in the GEFORCE GTX980 graphics card based on GM204 GPU. This GPU, as shown in Fig.3, is composed of an array of 4 Graphics Processing Clusters (GPC), 16 Streaming Multiprocessors (SMM), and 4 memory controllers.

In particular the GTX980 GPU is a 393 mm^2 die made of 5.2 billion transistors at 28nm integration scale, equipped with 2048 cores running at 1126 MHz (boost clock at 1216 MHz), able to reach 4.612 TFLOPs. The on-board memory is 4096MB, with L2 Cache size of 2048KB and memory clock at 7010 MHz, able to reach a bandwidth of 224.3 GB/sec.

The GPU architecture is based on an array of Graphics Processing Clusters (GPCs). These high-level blocks include hardware resources to perform almost all graphics processing, including pixel, texture, raster, vertex and geometry operations. In practice, the blocks are self-contained GPUs whose work is scheduled and coordinated by the GigaThread Engine at higher level. The blocks have access to memory by dedicated Memory Controllers. In each GPC, the processing is further parallelized using four Maxwell Streaming Multiprocessors (SMMs).

As shown in Fig.4, each SMM features control logic, L2 cache memory and instruction cache memory, shared memory and four blocks. Each block has dedicated control logic and registers, that gain access to 32 CUDA cores organized in a grid. Each row of four cores has access to load/store units to calculate source and destination addresses, to move (in/out) the data at each address and to cache or DRAM. In addition, for each row there is a Special Function Unit (SFU) designed to implement instructions such as sin, cosine, reciprocal, and square root. The basic processing unit is a CUDA core (Fig.5). This computational unit is a simple scalar processor made of a fully pipelined integer arithmetic logic unit (ALU) and floating point unit (FPU), based on IEEE 754-2008 floating-point standard, which improves the single and double precision multiplication and multiply-addition. The ALU is a 32-bit unit, optimized to support 64-bit operations. Among, the instructions there are Boolean, shift, move, compare, convert, bit-field extract, bit-reverse insert, and population count.

The algorithm we here propose for the CUDA architecture allows to reach a massive parallelization of simulations, as those depicted in Fig.2. Escape times are collected in a floating-point array T_e (see the Algorithm 1) allocated to the GPU memory. The size of the array corresponds to the number of simulation runs one plans to execute. The code of each run is embedded in a *kernel* and sent to the GPU using the host interface. The CUDA programming model is based on kernel, that is the portion of code that can be processed in parallel by threads. The kernel is fetched by the GigaThread Engine and sent to GPC and SMM. Each run is assigned to a specific location in the array and scheduled to be executed as thread by a core. Threads are scheduled by each SMM in groups of
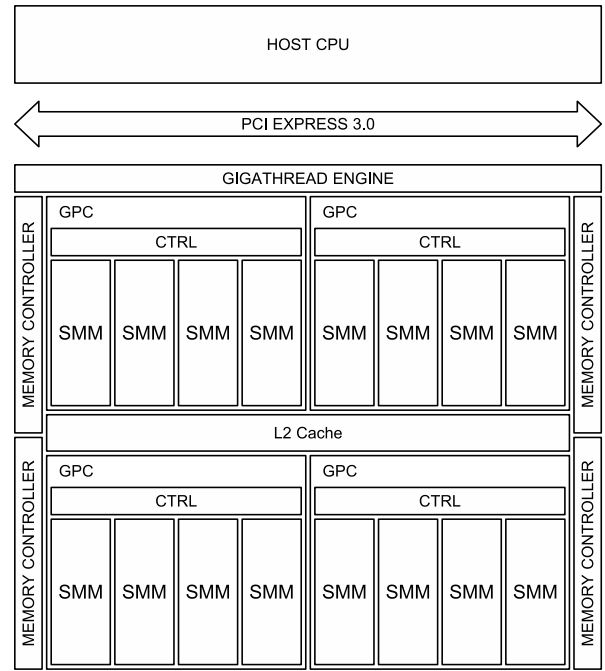

FIG. 3. Cores layout for the GPU routine. Here GPC stands for Graphic Processors Cluster and SMM for Maxwell Streaming Multiprocessors.

32 parallel threads called "warps". Each SMM features four warp schedulers and four dispatch units, so that four warps can be issued and executed concurrently. Each simulation run is executed independently from the other loads: when a core terminates a run, the next scheduled thread can be assigned to the core. It is therefore possible to exploit the hardware parallelism offered by the GPU almost to the full extent of its potentiality.
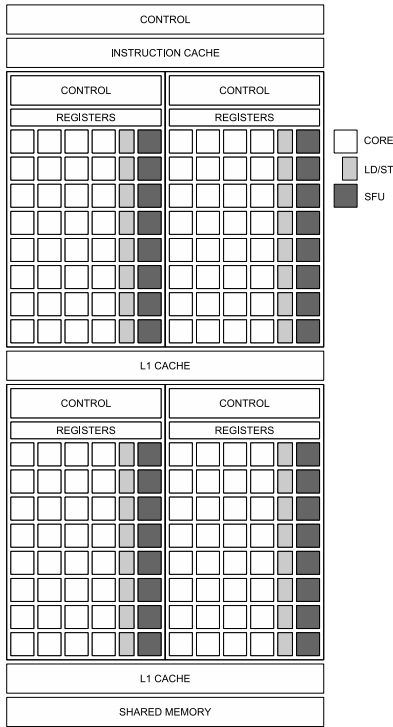
FIG. 4. Enlargment of the Maxwell Streaming Multiprocessors (SMM) blocks.

FIG. 5. Details of the CUDA core.

Algorithm 1 Kernel used for the simulations of the escape time)

1: **procedure** ESCAPETIME(V_0, β, γ, D, h)
  V_0, normalized potential barrier
  β, friction coefficient
  γ, normalized bias current
  D, Gaussian noise amplitude
  h, Euler time step

2:    $pos \leftarrow cuda.grid(1)$
3:    **if** $pos < sizeof(T_e)$ **then**
4:        $\sigma \leftarrow \sqrt{2D \cdot h}$
5:        $\varphi_0 \leftarrow \arcsin \gamma$
6:        $th \leftarrow \pi - \varphi_0$
7:        $\dot{\varphi}_0 \leftarrow 0$, initial phase velocity
8:        $k \leftarrow 1$
9:        **while** $k \leq klim$ **do**
10:            $r \leftarrow nrand()$
11:            $\varphi_k \leftarrow \varphi_{k-1} + \dot{\varphi}_{k-1} \cdot h$
12:            $\dot{\varphi}_k \leftarrow \dot{\varphi}_{k-1} + (-\beta\dot{\varphi}_{k-1} - V_0 \sin \varphi_{k-1} + V_0\gamma) \cdot h + \sigma \cdot r$
13:            **if** $\varphi_k > th$ **then**
14:                $T_e[pos] \leftarrow k \cdot h$
15:                **return**
16:            **end if**
17:            $k \leftarrow k + 1$
18:        **end while**
19:        $T_e[pos] \leftarrow -1$
20:    **end if**

21: **end procedure**

B. Programming

The simulation is coded as kernel in the environment CUDA. The pseudo-code of the kernel for the computation of the exit time distribution is outlined in Algorithm 1. The Algorithm is implemented by the ESCAPETIME procedure that accepts five arguments as input. The symbols and the physical meaning of the first four parameters have been introduced in Eq. (9). The last parameter, the time step h of the stochastic integration method is used in an Euler scheme [34], lines 11, 12. The simulations are retrieved following the scheme outlined in Figs. 2 and 6, by integration of the stochastic differential equation (9) and checking for the threshold crossing, see line 13. The sequence of operation in the escape time Algorithm 1 is as follows

1. Each processor device begins the trajectory integration (see Fig. 2) checking at each time step if a threshold crossing (i.e., a FPT) has occurred

2. When the FPT has been encountered the integration ends on a processor device and the device communicates the escape time to the host processor (see Fig. 6)

3. The host device assigns a new job to the processor device that has just completed the task

In Algorithm 1 the function $nrand$ at line 10 generates a Gaussian random variable with zero average and unit vari-

ance. This Gaussian random generator is based on the Box-Muller transform [37] applied to pseudo-random number obtained with the Mersenne Twister generator [38]. The corresponding CUDA core pseudocode is outlined in Algorithm 2. The computation is based on an internal subsequence that is initialized according to the MTRG-INIT procedure. This subsequence is named mt and it is made of 624 32-bits integers locally allocated to the device (line 2) and stored at L1 cache as shown in Fig.4. The seed, used as the first element of the sequence (line 3), is a unique value assigned to each thread. This initial value is externally provided from the host code and passed to the kernel. To generate the seed we generally use the current time in milliseconds.

Algorithm 2 Marsenne Twister Random Generator

1: **procedure** MTRG-INIT(seed)
$\quad$ $seed$, random sequence initial value

2: $\quad$ $mt \leftarrow cuda.local.array(624, dtype = int32)$
3: $\quad$ $mt[0] \leftarrow seed$
4: $\quad$ **for** $i = 1..623$ **do**
5: $\quad\quad$ $m = (mt[i-1] \wedge (mt[i-1] >> 30)) + i)$
6: $\quad\quad$ $mt[i] = (1812433253 \cdot m) \wedge \texttt{0xFFFFFFFF}$
7: $\quad$ **end for**
8: $\quad$ **return** mt

9: **end procedure**

10: **procedure** MTRG-EXTRACT(r, mt)
$\quad$ r, sequence index
$\quad$ mt, internal sub-sequence

11: $\quad$ $ri \leftarrow r \bmod 624$
12: $\quad$ **if** $ri = 0$ **then**
13: $\quad\quad$ **for** $i = 0..623$ **do**
14: $\quad\quad\quad$ $y \leftarrow mt[i] \wedge \texttt{0x80000000}$
15: $\quad\quad\quad$ $y \leftarrow y + (mt[(i+1) \bmod 624] \wedge \texttt{0x7fffffff}))$
16: $\quad\quad\quad$ $mt[i] \leftarrow mt[(i+397) \bmod 624] \oplus (y >> 1)$
17: $\quad\quad\quad$ **if** y is odd **then**
18: $\quad\quad\quad\quad$ $mt[i] \leftarrow mt[i] \oplus \texttt{0x9908b0df}$
19: $\quad\quad\quad$ **end if**
20: $\quad\quad$ **end for**
21: $\quad$ **end if**
22: $\quad$ $y \leftarrow mt[ri]$
23: $\quad$ $y \leftarrow y \oplus (y >> 11) \wedge \texttt{0xFFFFFFFF}$
24: $\quad$ $y \leftarrow y \oplus (y << 7) \wedge \texttt{0x9D2C5680}$
25: $\quad$ $y \leftarrow y \oplus (y << 15) \wedge \texttt{0xEFC60000}$
26: $\quad$ $y \leftarrow y \oplus (y >> 18) \wedge \texttt{0xFFFFFFFF}$
27: $\quad$ **return** y

28: **end procedure**

The seed used by MTRG-INIT is obtained by adding the thread identification number to initial value. With this expedient each thread makes use of a new seed based on the assigned simulation run number. The scheme allows each thread to keep a unique independent pseudorandom integer sequence that is controlled by the initial seed. From the value $mt[0]$, all the following values of the subsequence mt are obtained by iteration (lines 4-7). Once completed, the subsequence mt is returned as a result of the initialization (line 8) and used

by MTRG-EXTRACT according to the sequence index r and mt (line 10). The procedure works by blocks of 624 values. Which value of the block to use is simply indexed by a modulo operation (line 11). At the beginning, and at each time they are employed (line 12 of Algorithm 1), a new block of values is computed from the previous one (lines 13-20). Then, the value $mt[ri]$ is transformed (lines 22-26 of Algorithm 2) and returned (line 27).

Algorithm 3 Kernel used for simulations of the switching currents

---

1: **procedure** SWITCHINGCURRENT($\delta\gamma, V_0, \beta, D, h$)
$\quad$ $\delta\gamma$, normalized ramp step
$\quad$ V_0, normalized potential barrier
$\quad$ β, friction coefficient
$\quad$ D, Gaussian noise amplitude
$\quad$ h, Euler time step

2: $\quad$ $\gamma_0 \leftarrow 0$, initial normalized bias current
3: $\quad$ $pos \leftarrow cuda.grid(1)$
4: $\quad$ **if** $pos < sizeof(\Gamma_e)$ **then**
5: $\quad\quad$ $\sigma \leftarrow \sqrt{2D \cdot h}$
6: $\quad\quad$ $\varphi_0 \leftarrow \arcsin \gamma$
7: $\quad\quad$ $th \leftarrow \pi - \phi_0$
8: $\quad\quad$ $\dot{\varphi}_0 \leftarrow 0$, initial phase velocity
9: $\quad\quad$ $k \leftarrow 1$
10: $\quad\quad$ $\gamma_k \leftarrow \gamma_0$
11: $\quad\quad$ **while** $\gamma_k \leq 1$ **do**
12: $\quad\quad\quad$ $r \leftarrow nrand()$
13: $\quad\quad\quad$ $\gamma_k \leftarrow \gamma_0 + \delta\gamma \cdot k$
14: $\quad\quad\quad$ $\varphi_k \leftarrow \varphi_{k-1} + \dot{\varphi}_{k-1} \cdot h$
15: $\quad\quad\quad$ $\dot{\varphi}_k \leftarrow \dot{\varphi}_{k-1} + (-\beta\dot{\varphi}_{k-1} - V_0 \sin \varphi_{k-1} + V_0\gamma_k) \cdot h + \sigma \cdot r$
16: $\quad\quad\quad$ **if** $\varphi_k > th$ **then**
17: $\quad\quad\quad\quad$ $\Gamma_e[pos] \leftarrow \gamma_k$
18: $\quad\quad\quad\quad$ **return**
19: $\quad\quad\quad$ **end if**
20: $\quad\quad\quad$ $k \leftarrow k + 1$
21: $\quad\quad$ **end while**
22: $\quad\quad$ $\Gamma_e[pos] \leftarrow 1$
23: $\quad$ **end if**

24: **end procedure**

In conclusion the code MTRG-EXTRACT returns a random integer uniformly distributed in the range of 32-bit integers. This integer with a suitable linear normalization generate an uniform deviate in the range $[0, 1]$. The uniform deviate is used in the external function $nrand$ with the Box-Muller transform to generate Gaussian pseudo random numbers with zero average and unit variance.

The proposed Algorithm 1 is useful for the computation of escape time distribution by means of direct event simulations. In threshold device based on JJ it is experimentally simpler to retrieve the SCs distribution [39]. The proposed code is straightforwardly changed to compute the SCs instead of the exit time, as per the SWITCHINGCURRENT that is only slightly different from the escape time process (see Algorithm 3). During the computation, the bias current is increased from zero to 1 (that in our normalized units, see Eq.(4), corre-
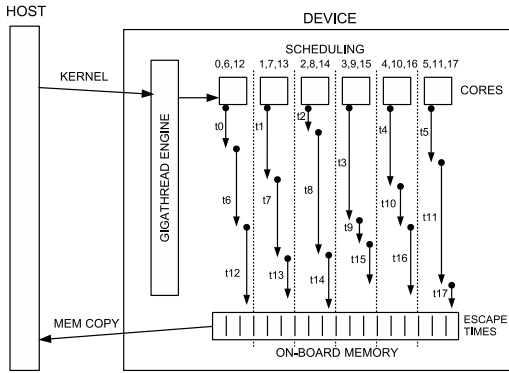
FIG. 6. Sketch of the streaming process, the GigaThread Engine that dictates the task scheduling. The processor computes an element (e.g. $t0$, $t6$, ...) of the escape time sequence and eventually stores the result in the array T_e. The order of the single tasks termination does not correspond to the array index order of T_e.

FIG. 7. Sketch of the execution time (in seconds) as a function of the number of realizations N_r with different numbers of active threads per block N_{tpb}.

sponds to the maximum superconducting current that can flow through the JJ). The numerical method to increase the bias current at each time step, however, requires a particular care. In fact the GPUs normally work in single precision. When this is the case, the smallest bias increment is 10^{-7}, that is the reciprocal of the largest integer in single precision. Consequently, for the extremely long simulations reported in the experiments demand that the bias current is not incremented with a single precision integer loop. The difficulty can be overcome with a tuple of n 32-bit integers.

C. Parallelism and Speedup tests

To investigate the scaling with the number of processing units, we have studied the execution time as a function of the number of realizations N_r for three different numbers (denoted as N_{tpb}) of *threads per block* (TPB). This is shown in Fig. 7. From the Figure it is evident that for a given number of processors the execution time linearly scales with the number of realizations, as expected for serial computation. The advantage is in the slope, that decreases when the number of units is increased.

The advantage of a large number of processing units is shown in more details in Figs. 8(a,b). In (a) we display the execution time as a function of the TPB for different numbers of realizations N_r, from $\simeq 10^3$ to $\simeq 10^5$. The figure essentially confirms the linear behavior of Fig. 7: the increased number of processor units proportionally scales with the required CPU time only above a certain number of threads. This number, that depends upon the number of realizations, represents the point at which the additional processors do not contribute. In Fig. 8(b) we compare two technologies (Python Numba vs CUDA Fortran solution) to implement the same algorithm. It is evident that the overall behavior is the same for both technologies, albeit the CUDA Fortran is much more efficient, of about a factor 30. This is to be expected, as Numba is based on CUDA Toolkit version 7.5 which does not provide a native

implementation of the random generator library, made available in version 8 as CURAND_DEVICE library and used by the CUDA Fortran compiler. Thus, the execution of MTRG-EXTRACT in Algorithm 2 significantly impacts on the simulation kernel if the number of operations required by the random generator (MTRG and Box-Muller) is comparable to the number of operations for the integration. In addition, CUDA Fortran uses compiler optimizations that are not available to the LLVM compiler infrastructure on which Numba is based. In both figures we observe that time required by simulations decreases linearly with respect to number of active TPB. A larger number of threads improves the GPU occupancy, up to the limit of 512 TPB.

The speedup efficiency, defined as the ratio S/N_{tpb}, is displayed in Fig. 9. In the efficiency the speedup S is defined as the ratio between the average execution time for the sequential and the parallel algorithm (see Eq.(A.5) and the Appendix A for the details of the model and notation). The data show that the asymptotic efficiency converges to 1 as the number of realizations increases. For a constant number of realizations N_r, the efficiency makes worse increasing the number of threads per block (i.e., reducing the computation burden per thread N_L). This effect is manly due to the fluctuations in the execution time of each thread, that emerge in the small sample limit when the load is distributed over too many threads.

In the inset of Fig. 9 we show a 3σ error limits due to the statistical variability of the speedup S for the particular case $N_{tpb} = 20$. We remind that the statistical model employed in the Appendix neglects the short time behavior of the system. However, the qualitative asymptotic behavior of the speedup, and more specifically of the statistics concerning the execution time, is mainly due to the exponential tail of the random variable distribution.
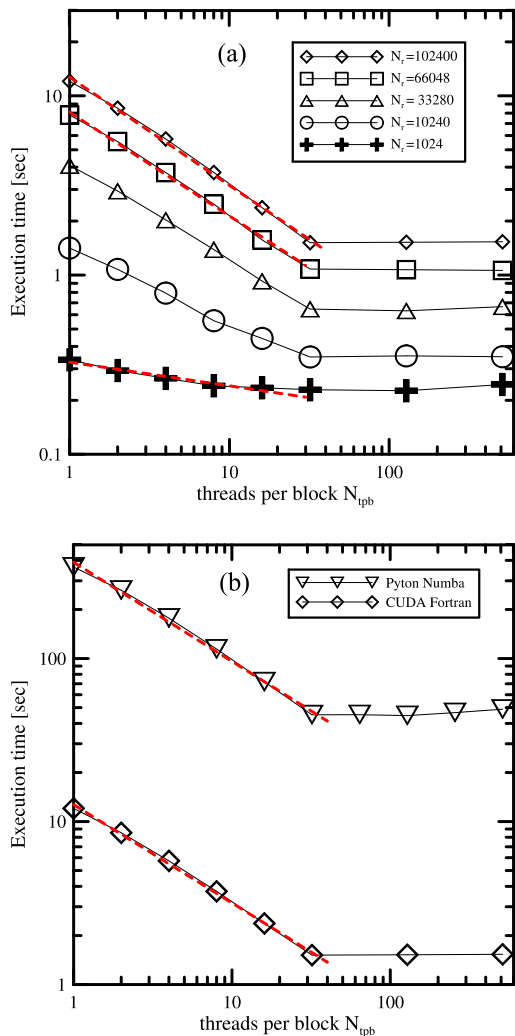
FIG. 9. Speedup efficiency S/N_{tpb} as a function of the number of realizations (on a log scale), as per Appendix A. The inset shows the speedup S (solid line) for the particular case $N_{tpb} = 20$, together with the $\pm 3\sigma$ uncertainty (dots) due to statistical variability.

A. Arrhenius plots

First, to validate the algorithms we have investigated the case of a constant bias JJ subject to different noise intensity, see Fig. 10. In the Figure the error bar is not visible, for the speed of the algorithm has allowed to collect a large number of realizations ($N_r = 5120$). This high accuracy allows to distinguish between the Arrhenius behavior

$$\langle \tau \rangle \propto \exp \frac{\Delta U}{\Delta \theta} \tag{12}$$

and the detailed Kramers rate, obtained taking into account the prefactor corrections, for the special case of moderately underdamped systems. The model fits in this case with prefactor becomes [1]:

$$\langle \tau \rangle \propto \frac{1}{\theta} \exp \frac{\Delta U}{\Delta \theta}. \tag{13}$$

The accurate simulations of the mean FPT as a function of the inverse of the temperature is a task useful in several applications, most importantly to retrieve the so-called quasipotential (or pseudopotential) for non-Hamiltonian systems [18, 19]. When a bona fide potential does not exist or is not explicitly known, one can estimate an effective quasipotential from numerical simulations as those of Fig. 10, reversing the logic of Eq. (13), i.e., assuming:

$$\Delta U \equiv \frac{\log (\theta \langle \tau \rangle)}{\Delta \theta}. \tag{14}$$

The resulting method requires to simulate the system in the low noise limits, when the relations (12,13) are strictly valid [18, 19]. By the same token, in this limit simulations are extremely long, and hence the call for fast, parallel simulations. In this work we retrieve the potential ΔU form the slope of the

FIG. 8. (color online) (a) Sketch of the execution time (in seconds) as a function of the number of TPB for different number of realizations N_r reported in the legend. The dashed red curves give the power law interpolation. For $N_r = 1024$ the scaling coefficient is 0.13, it converges for larger values of N_r to the values ≈ 0.6. (b) Python Numba vs CUDA Fortran execution time (in seconds) as function of TPB. We note the same power law scaling (i.e. ≈ 0.6), as a function of N_{tpb} in both practical implementations.

IV. TESTS ON PHYSICAL SYSTEMS

At the end of the previous Section we have investigated the execution time statistics of the Algorithms 1 and 3. In the following we collect numerical results of the benchmark physical system, that is of superconducting JJs. We do so to validate the whole method, for the CUDA arithmetic is poorer than the ordinary CPU arithmetic, in that it is single precision and with a lower level of fidelity [40].
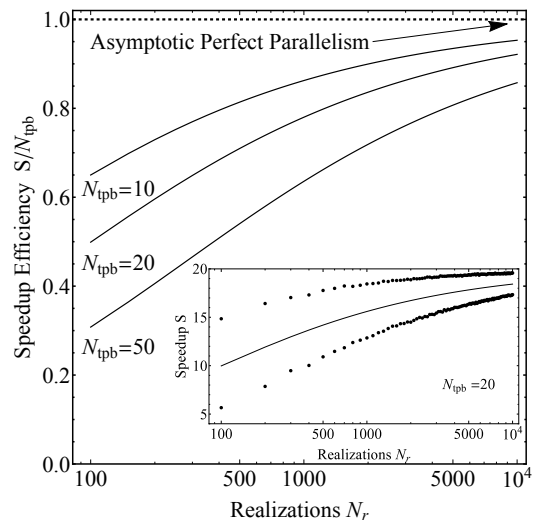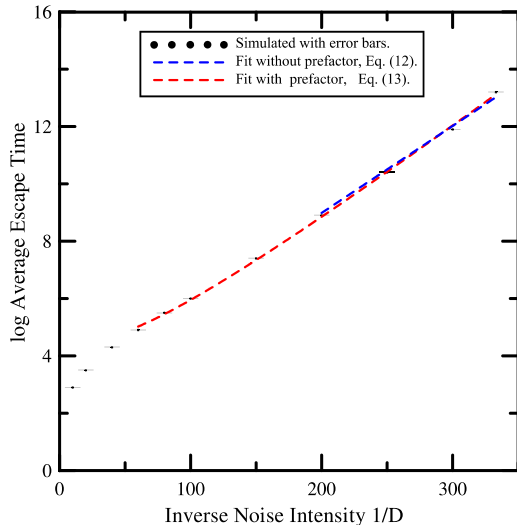
FIG. 10. (color online) Log of the average FPT (black dots) as a function of the inverse normalized temperature $D = \beta\theta$, compared to the theory, Eq. (12) (short dashed blue) and Eq. (13) (long dashed red). Parameters of the simulations are: $\gamma = 0.5$, $\alpha = 0.05$, $N_r = 5120$, integration step $h = 0.004$ The theoretical value for the energy barrier reads $\Delta U = 0.0342$, while the numerical evaluation reads $\Delta\tilde{U} \simeq 0.0303$ for fitting with Eq.(12) and $\Delta\tilde{U} \simeq 0.0359$ for fitting with Eq.(13). The accuracy of the numerical method is 11% and 4%, respectively.

Arrhenius plots of Fig. 10 of a system such as the Josephson potential associated to Eq.(9). In this case, being the potential known, we can use the analytic result to ensure that numerical simulations are reliable. From the data of Fig. 10 we conclude that Kramers fitting offers a good estimate of the potential barrier ΔU, with a relative error of few percents.

B. Escapes and switching current distributions

As a benchmark for the simulations we also employ the switching current distribution of the Josephson junctions. We have preliminarily checked that CUDA simulations are statistically reliable (see Sect. IV A). The accuracy is further confirmed by the comparison of simulations with ordinary double precision CPU arithmetic and the GPU arithmetic. For the sake of comparison, we have employed PGI Suite Compilers version 17.4, community edition, with Cuda Toolkit 8.0 for both the CPU-only and the GPU version of the code. Routines have therefore been encoded without external libraries, such as those provided by the Intel Math Kernel Library, that includes optimized implementation of Mersenne Twister random number generator. The comparison is shown in Fig. 11(a) as the escape times cumulative distributions computed for the same parameters for both GPU and CPU code.

The difference between the two distributions (shown in the inset) is small and lies within the statistical fluctuations for this type of stochastic simulations.

Having tested the code, we now turn our attention to some

physically interesting results. The high speed simulations allow to investigate the adiabatic approximation used in [39] that, neglecting non-equilibrium corrections to the Kramers' theory, gives a formula for the SC cumulative distribution only valid in the limit of vanishingly small bias ramp speed (for possible generalizations of the adiabatic approach, however overdamped, see [41] . This formula can be compared with the numerical result of Fig. 11(b), that displays the computed cumulative distribution function of the JJ switching currents with the same parameters of the experiments reported in Ref. [42]. In particular, to check the theoretical approximations we have performed simulations with a bias current step that in normalized units correspond to about 10^{-12}. It is shown that simulated and theoretical SC cumulative distributions agree within 1.5%, as displayed in the inset of Fig. 11(b). This computation has required about $1.5 \times 10^5 sec \simeq 2d$ on the GPU device. Assuming an acceleration of $\simeq 400$ we estimate an execution time of about $5 \times 10^8 sec \simeq 700d$ on an ordinary CPU (e.g., Intel Core $i7 - 6700$). The direct numerical simulations of the experiments in the timescale of the laboratory electronics is therefore very demanding, for the different time scales between the JJ internal dynamics and the driving external electronics (see Sect. II A).

V. CONCLUSIONS

We have challenged the problem of an effective and fast algorithm for the mean first passage times with CUDA. This is specially useful in the GPU simulations of escape times of Josephson junctions, where the intrinsic timescale of the system is extremely fast (up to the THz magnitude), while the electronic to control the escapes from the zero voltage state is much slower (typically around or below 1KHz).

We have exploited the CUDA environment to compute in parallel several different realizations of the escape process as schematically shown in Figs. 2 and 6. Each parallel thread uses a pseudorandom number generator with a different seed.
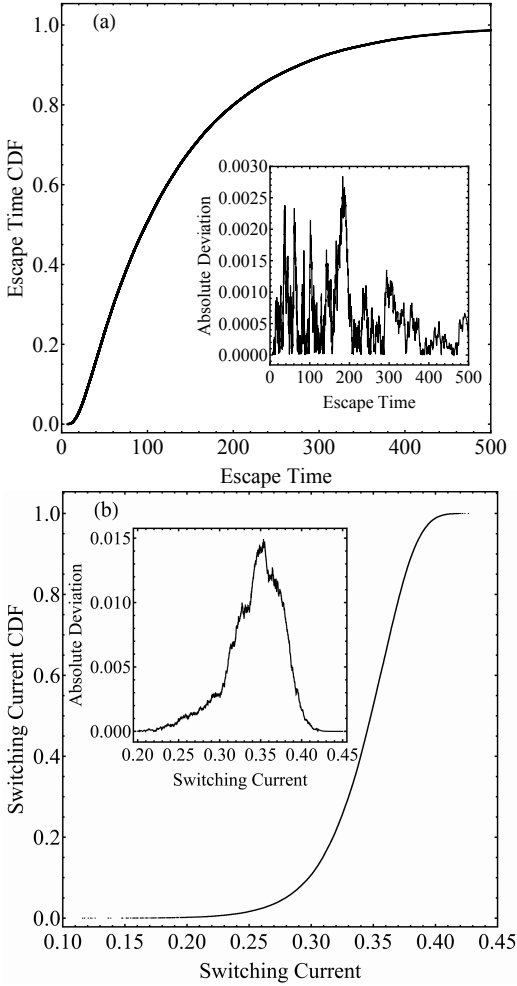
FIG. 11. Cumulative distributions of the escape time and switching current, and deviations from other reference models. (a) Escape time cumulative distribution through Monte Carlo solution of Eq.(9). The numerical simulations are performed with CUDA Fortran following Algorithm 1 and with ordinary GNU compiler in standard double precision arithmetic. On the scale of the figure the two numerical solutions are the same, therefore the difference between the two results is displayed in the inset. (b) Numerical simulations and theoretical predictions for the switching current cumulative distribution. The numerical simulations have been performed with CUDA Fortran following Algorithm 3, while the theoretical predictions are based on Kramers adiabatic rate [39]. On the scale of the figure the two results are the same, therefore the difference is displayed in the inset. Parameters of the simulations, reported in physical units, are the same of a realistic device tested in Ref. [42]: $R = 250\Omega$, $C = 88 \times 10^{-15}F$, $I_c = 0.748 \times 10^{-6}A$, $T = 1.2K$, and the sweep rate is $200Hz$. The SC empirical CDF is computed using $N_r = 30\,000$ realizations. The initial bias in each realizations is $\gamma = 0$. The integration step $h = 10^{-4}$ corresponds to about 2×10^{-4} of the natural oscillating period, that is given by the dimensionless plasma frequency $V_0^{1/2}$.

The efficiency of the proposed algorithm reaches accelerations of about 400 respect to standard Intel Core $i7 - 6700$ host processors. The scaling of the performances with the rising of the number of threads has ben elucidated in Figs. 7 and 8. From the data, it is clear that: i) the choice of random

generator is crucial, it should be quick and low memory demanding; ii) the load for each thread should be high enough to avoid that fluctuations dominate and reduce the efficiency of paralllelism ; iii) in the running regime (that, however, is not essential for switching current distributions) the argument of the sin, namely φ, grows to a level where the single precision could be not enough . With this care the abovementioned efficiency applied to the benchmark Josephson junctions proved fast enough to challenge realistic experiments [42]. The numerical experiments have reproduced practical set-ups without extrapolation formulas and could be useful, for instance, to investigate the thermal-quantum transition process. In the same field we suggest that CUDA environment might be also relevant for arrays of Josephson junctions, possibly coupled to a resonator [43] or for voltage standards [44].

APPENDIX A. PARALLELISM STATISTICAL MODEL

We propose a simple probabilistic model for the parallel computation of escape times with GPU in CUDA environment that is at the basis of the performances of Fig. 9.

The execution time of Algorithm 1 for each parallel process is directly proportional to the random escape time. The overall average execution time random variable $\langle E_t \rangle$ is determined by the sum of the individual escapes t_i:

$$E_t = \sum_{i=1}^{N_r} t_i. \tag{A.1}$$

The variable t_i can be approximated with an exponential random deviate (although this is not exact, for the presence of an inertial short time prior to escapes over the barrier) identically distributed with an appropriated time scaling. The random variable (A.1) is distributed as the sum of iid exponentials' and the resulting distribution is an $Erlang(\cdot)$ function [45]

$$E_t \sim Erlang(N_r, 1). \tag{A.2}$$

Defining $N_L = N_r/N_{tpb}$ the number of runs assigned to each parallel thread the execution time random variable of the parallel algorithm is

$$E_p = \max(E_t^{(1)}, E_t^{(2)}, ..., E_t^{(N_{tpb})}). \tag{A.3}$$

We define the parallel execution time as the average $\langle E_p \rangle$. From (A.3) it follows that the CDF of the random variable E_p reads

$$CDF[E_p, x] = CDF[Erlang(N_L, 1), x]^{N_{tpb}} \tag{A.4}$$

by Eq. (A.4) is possible to compute in closed form all the moments, in particular the parallel execution time $\langle E_p \rangle$. If one defines the average speedup as the ratio

$$S = \frac{\langle E_t \rangle}{\langle E_p \rangle}, \qquad (A.5)$$

a perfect parallel algorithm entails: $S = N_{tpb}$. To display a quantity that is independent of the time scaling used in (A.1)

it is convenient to define the speedup efficiency as the ratio S/N_{tpb}. The theoretical estimate (A.5) for the speedup is the basis for the results displayed in Fig. 9 and discussed in Sect. III C.

REFERENCES

[1] H. Risken, *The Fokker-Planck Equation.* (Springer-Verlag, 1984)

[2] E. Pollak and R. Ianconescu, J. Phys. Chem. A, **120**, 3155 (2016).

[3] J. J. Mazo, F. Naranjo, and D. Zueco, Phys. Rev. B **82**, 094505 (2010).

[4] A. Barone, G. Paternò, *"Physics and Applications of the Josephson Effect"*, (John Wiley & Sons, 1982).

[5] M. H. Devoret, *"Quantum Fluctuations in Electrical Circuits"*, S. Reynaud, E. Giacobino and J. Zinn-Justin, eds., Les Houches, Session LXIII, 1995, Elsevier (1997).

[6] V. Pierro and G. Filatrella, Phys. Rev. A **94**, 042116 (2016).

[7] J. P. Pekola , Phys. Rev. Lett. **93**, 206601 (2004).

[8] G. Filatrella and V. Pierro, Phys. Rev. E **82**, 046712 (2010).

[9] P. Addesso, G. Filatrella, and V. Pierro, Phys. Rev. E **85**, 016708 (2012).

[10] B. Gross, J. Yuan, D. Y. An, M. Y. Li, N. Kinev, X. J. Zhou, M. Ji, Y. Huang, T. Hatano, R. G. Mints, V. P. Koshelets, P. H. Wu, H. B. Wang, D. Koelle, and R. Kleiner Phys. Rev. B **88**, 014524 (2013).

[11] Yu. M. Shukrinov, I. R. Rahmonov, K. V. Kulikov, A. E. Botha, A. Plecenik, P. Seidel and W. Nawrocki, Supercond. Sci. Technol. **30**, 024006 (2017).

[12] M. A. Galin, A. M. Klushin, V. V. Kurin, S. V. Seliverstov, M. I. Finkel, G- N- Goltsman, F. Müller, T. Scheller and A. D. Semenov, Supercond. Sci. Technol. **28**, 055002 (2015).

[13] L. B. Freund, Proc. Natl. Acad. Sci. U.S.A. **106**, 8818 (2009).

[14] J. J. Mazo, F. Naranjo, and D. Zueco, J. of Chem. Phys. **138**, 104105 (2013).

[15] A. A. Smirnov, and A. L. Pankratov, Phys. Rev. B **82**, 132405 (2010).

[16] M. Januszewski, M. Kostur, Comput. Phys. Commun. **181**, 183 (2010).

[17] K. G. Fedorov and A. L. Pankratov, Phys. Rev. Lett. **103**, 260601 (2009) .

[18] R. Graham and T. Tél, Phys. Rev. A **31**, 1109 (1985).

[19] R.L. Kautz, J. Appl. Phys. **76**, 5538 (1994).

[20] S.S. Artemiev and V.D. Korneev, Numer. Analys. Appl. **4**, 1 (2011).

[21] M. Büttiker, E.P. Harris, and R. Landauer, Phys. Rev. B **28**, 1268 (1983).

[22] G. Augello, D. Valenti, A.L. Pankratov, and B. Spagnolo, Eur. Phys. B **70**, 145 (2009).

[23] N. Grønbech-Jensen, M. G. Castellano, F. Chiarello, M. Cirillo, C. Cosmelli, L. V. Filippenko, R. Russo, and G. Torrioli, Phys. Rev. Lett. **93**, 107002 (2004).

[24] N. Grønbech-Jensen and M. Cirillo, Phys. Rev. Lett. **95**, 067001 (2005).

[25] J. M. Martinis, M. H. Devoret, and J. Clarke, Phys. Rev. B **35**, 4682 (1987).

[26] A. Shnirman, E. Ben-Jacob, and B. Malomed, Phys. Rev. B, **56**, 14677 (1997).

[27] J. M. Martinis, S. Nam, J. Aumentado, and C. Urbina, Phys. Rev. Lett. **89**, 117901 (2002).

[28] A. Wallraff, A. Lukashenko, J. Lisenfeld, A. Kemp, M. V. Fistul, Y. Koval, and A. V. Ustinov, Nature **425**, 155 (2003).

[29] A. N. Price, A. Kemp, D.R. Gulevich, F.V. Kusmartsev, and A.V. Ustinov, Phys. Rev. B, **81**, 014506 (2010).

[30] U. C. Coskun, M. Brenner, T. Hymel, V. Vakaryuk, A. Levchenko, and A. Bezryadin, Phys. Rev. Lett. **108**, 097003 (2012).

[31] D. Massarotti, A. Pal, G. Rotoli, L. Longobardi, M.G. Blamire, and F. Tafuri, Nature Comm. **6**, 7376 (2015).

[32] Y. Makhlin, G. Schön, and A. Shnirman, Rev. Mod. Phys. **73**, 357 (2001).

[33] J. A. Blackburn, M. Cirillo, N. Grnbech-Jensen, Physics Reports **611**, 1 (2016).

[34] P. E. Kloeden, E. Platen and H. Schurz, *Numerical Solution of SDE Through Computer Experiments* (Springer-Verlag, 1994).

[35] A.L. Pankratov Phys. Lett. A **234**, 329 (1997).

[36] A.N. Malakhov , A.L. Pankratov, Physica C **269**, 46 (1996).

[37] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes* (Cambridge University Press, 1995), Vol. I, Ch. 19.

[38] M. Matsumoto, and T. Nishimura, ACM Trans. Model. Comput. Simul. **8**, 3 (1998).

[39] T. A. Fulton and L. N. Dunkleberger, Phys Rev B **9**, 4760 (1974).

[40] http://docs.nvidia.com/cuda/floating-point/index.html, consulted on January 8th, 2018.

[41] A. L. Pankratov and M. Salerno, Phys Rev E **61**, 1206 (2000); Phys. Lett. A **273**, 162 (2000).

[42] Shao-Xiong Li, Wei Qiu, Siyuan Han, Y. F. Wei, X. B. Zhu, C. Z. Gu, S. P. Zhao, and H. B. Wang, Phys. Rev. Lett. **99**, 037002 (2007).

[43] R. Yamapi and G. Filatrella, Phys. Rev. E **89**, 052905 (2014).

[44] R. Behr, O. Kieler, J. Kohlmann, F. Müller, and L. Palafox, Meas. Sci. Technol. **23**, 124002 (2012).

[45] M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions*, 4th Ed. (John Wiley & Sons, 2011) Ch. 15, p. 84.