

# GPU accelerated Monte Carlo simulation of Brownian motors dynamics with CUDA<sup>☆</sup>



J. Spiechowicz, M. Kostur, L. Machura<sup>\*</sup>

Institute of Physics, University of Silesia, 40-007 Katowice, Poland

## ARTICLE INFO

### Article history:

Received 4 August 2014

Received in revised form

5 December 2014

Accepted 26 January 2015

Available online 11 February 2015

### Keywords:

Langevin equation

GPU

CUDA

Monte Carlo

Brownian motor

Gaussian noise

Poissonian noise

Dichotomous noise

## ABSTRACT

This work presents an updated and extended guide on methods of a proper acceleration of the Monte Carlo integration of stochastic differential equations with the commonly available NVIDIA Graphics Processing Units using the CUDA programming environment. We outline the general aspects of the scientific computing on graphics cards and demonstrate them with two models of a well known phenomenon of the noise induced transport of Brownian motors in periodic structures. As a source of fluctuations in the considered systems we selected the three most commonly occurring noises: the Gaussian white noise, the white Poissonian noise and the dichotomous process also known as a random telegraph signal. The detailed discussion on various aspects of the applied numerical schemes is also presented. The measured speedup can be of the astonishing order of about 3000 when compared to a typical CPU. This number significantly expands the range of problems solvable by use of stochastic simulations, allowing even an interactive research in some cases.

### Program Summary

*Program title:* Poisson, dich

*Catalogue identifier:* AEVP\_v1\_0

*Program summary URL:* [http://cpc.cs.qub.ac.uk/summaries/AEVP\\_v1\\_0.html](http://cpc.cs.qub.ac.uk/summaries/AEVP_v1_0.html)

*Program obtainable from:* CPC Program Library, Queen's University, Belfast, N. Ireland

*Licensing provisions:* GNU Lesser General Public License, version 3

*No. of lines in distributed program, including test data, etc.:* 3338

*No. of bytes in distributed program, including test data, etc.:* 45009

*Distribution format:* tar.gz

*Programming language:* CUDA C.

*Computer:* Any with CUDA-compliant GPU.

*Operating system:* No limits (tested on Linux).

*RAM:* Hundreds of megabytes for typical case

*Classification:* 4.3, 23.

*External routines:* The NVIDIA CUDA Random Number Generation library (cuRAND)

*Nature of problem:* Graphics processing unit accelerated numerical simulation of stochastic differential equation.

*Solution method:* The jump-adapted simplified weak order 2.0 predictor–corrector method is employed to integrate the Langevin equation of motion. Ensemble-averaged quantities of interest are obtained through averaging over multiple independent realizations of the system generated by means of the Monte Carlo method.

<sup>☆</sup> This paper and its associated computer program are available via the Computer Physics Communication homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

<sup>\*</sup> Corresponding author. Tel.: +48 32 359 2107; fax: +48 32 258 3653.

E-mail address: [lukasz.machura@us.edu.pl](mailto:lukasz.machura@us.edu.pl) (L. Machura).

*Unusual features:* The actual numerical simulation runs exclusively on the graphics processing unit using the CUDA environment. This allows for a speedup as large as about 3000 when compared to a typical CPU.

*Running time:* A few seconds

© 2015 Elsevier B.V. All rights reserved.

## 1. Introduction

The stochastic dynamics has been a very useful tool for analysis of problems in various areas of science, engineering and finance [1]. It takes into consideration the indispensable fluctuations of the system which often play a dominant role in explanation of the occurring phenomena. A remarkable example of this situation is Brownian particle dynamics [2–5].

One of the approaches that has been developed for analysis of dynamical systems subjected to noise is formulated in terms of a stochastic differential equation (SDE) [6] that specifies the evolution of stochastic process. Similarly to the case of ordinary differential equations, they are usually rather intractable analytically and explicit solutions are given only for very few cases [1]. However, from a computational point of view, stochastic differential equations can be implemented on modern personal computers without much effort. Moreover, when the dimensionality of the problem is greater than three it is often the only applicable numerical method. In practice, to obtain statistical quantities of interest one has to run a large number of realizations of stochastic differential equation and take average over all of these paths. Each of them is generated by means of Monte Carlo method [7]. This proceeding is simpler and often much faster than other approaches that have been developed so far for studying such systems. It is due to the recent evolution of computer architectures towards multiprocessor platforms.

In particular, the advent of Compute Unified Device Architecture (CUDA) [10], a general purpose parallel computing architecture of modern NVIDIA graphics processing units (GPUs) has taken the computational science into completely new level of possibilities [11–15]. Today, the latest commodity of NVIDIA GPUs are capable of performing trillions of floating point operations per second as a result of employing the power of massively parallel architecture with hundreds of cores on a single silicon chip. In order to take advantage of such hardware one has to carefully redesign the algorithms. The ideal situation occurs when the problem inherently decomposes into a large number of independent tasks whose results can be combined in a simple way, e.g. ensemble averaging in Monte Carlo simulation of stochastic differential equation [8,9]. It is an example of a so called “embarrassingly parallel problem” that may particularly benefit from a parallel workload. It has been a few years since the last article devoted to the subject of numerical solution of stochastic differential equations by harvesting the power dormant in the modern GPU was published [11] and therefore we decided to present an updated and extended step by step guide how to properly accelerate the stochastic dynamics simulations.


This paper is organized as follows: first, we briefly introduce the CUDA environment. Next, a short overview of the discussed Brownian dynamics models is presented. Then, the algorithms employed to numerically solve these problems are laid out. This is followed by the guide how to properly accelerate numerical solution of stochastic differential equations with CUDA. Next, programs validations and performance measurements on GPU and CPU hardware are presented, both in single and double precision. The final section contains summary and some conclusions. We also provide the source code of example programs which demonstrate the techniques described in the paper.

## 2. A primer on CUDA


The name CUDA often embodies not only the hardware architecture of the GPU but also the software which is used to access it. We will now briefly outline the architecture of CUDA compatible devices as well as the abstraction layer between the programmer and the hardware.

NVIDIA distinguishes CUDA GPUs by their Compute Capability [10]. Devices of a greater Compute Capability offer more advanced features and better performance. For example, GPUs with Compute Capability 1.3 and higher support double precision floating point arithmetic [10]. From the hardware viewpoint, CUDA is built on the concept of a streaming multiprocessor (MP). Such a multiprocessor is a SIMT (Single Instruction, Multiple Threads) unit formed of several scalar processors (SPs) which are capable of executing a certain job. MPs usually execute a whole warp of threads at once. The warp size specifies the number of threads within a warp. The time it takes to execute one warp depends on what kind of operation the threads need to perform. Moreover, each MP has four types of limited on-chip memory: a set of 32-bit registers, a shared memory block, a constant and texture cache. Since a scalar processor can only access its own registers they cannot be used to share the data between threads. However, the other types of memory are common to all SPs and therefore can be utilized to do it.

Probably the most important feature of the CUDA architecture which often determines whether the program is optimized or not is the memory hierarchy [10]. The crucial difference between its types is access times - see a schematic representation on Fig. 1. What matters most is how close they are to the actual host memory. A rule of thumb is: the bigger the size of a certain kind of memory, the more time it takes to access it. The slowest type of memory is the global device memory of the GPU. It is the main memory of the graphic card that can be read and written from both the host and device. The name global suggests that it is available to all threads launched on the device and persists for the lifetime of the application. However, global memory accesses are latency expensive operations, taking several hundred clock cycles of the GPU to complete. The fastest kind of memory currently available on graphic cards is the previously mentioned set of registers. Generally, accessing a register consumes zero extra clock cycles per instruction, but delays may occur due to register read-after-write dependences and register memory bank conflicts. Maximum number of 32-bit registers per thread is presently limited to 255 [10]. If there are insufficient registers, data is stored in a local memory. However, here the adjective “local” refers to scope not physical locality. It means that its access time is comparable to the global memory and therefore data that cannot be stored in the set of registers should be loaded into the shared memory which



**Fig. 1.** The schematic representation of the GPU memory hierarchy. The above arrow points in the direction of ever-faster types of memory. Its width illustrates the available size of particular type of memory.



**Fig. 2.** Periodic structure in the form of *symmetric* piece-wise linear potential of period  $L = 2$ .

in terms of speed stands in between the global memory and the registers.

The abstraction layer between programmer and the device is built around the concept of a kernel. It is a part of the source code that is executed in parallel by the scalar processors. Kernels are invoked as bunches of threads arranged in one-, two- or three-dimensional blocks. Each block is assigned to a unique multiprocessor. They are further organized into a one- or two-dimensional grid. The composition of the grid and blocks is determined at the time of kernel invocation. Each parallel thread has its unique ID which can be calculated assuming knowledge of the grid and block structure. Threads that reside in a single multiprocessor or equivalently block can exchange data between themselves through the previously mentioned shared memory. Another important feature which is available for the block of threads is an execution synchronization.

The CUDA also provides a comprehensive development environment for C developers building GPU-accelerated applications. It includes a CUDA C compiler, math libraries, and tools for debugging and optimizing the performance of applications. CUDA C is a simple extension of the C programming language which includes several new keywords and expressions that make it possible to distinguish between host and device functions as well as transfer data between them.

### 3. Noise induced transport in periodic structures

We now present two models of Brownian particle dynamics which we concentrate upon later and are of particular interest in many disciplines of science. It is already a well known fact that one can obtain a measurable particle current in periodic structures without application of any external biasing force or field gradient [16–18]. The consideration of such unusual noise induced transport phenomenon originated from biology and stimulated many scientists working in diverse areas of physics, chemistry and biology itself. As a first approximation of real fluctuations with a short correlation time, Gaussian white noise has been predominantly used. Such additive white noise is not capable of inducing a finite current due to the laws of equilibrium dynamics [19]. However, there are other stochastic processes frequently met in physical, chemical or biological context which are able to do it due to their inherent statistical asymmetry [20,21].

A prominent example of such non-equilibrium fluctuations is a shot noise. Its realizations consist of sequences of very sharp pulses with random heights and randomly distributed times between subsequent pulses. A particular mathematical model of shot noise is formulated as white Poissonian noise [22–24]. It is a sequence of  $\delta$ -shaped pulses which occur at times forming Poissonian point process. The amplitudes of the pulses are independent random variables which are distributed according to a common probability density. Such white Poissonian noise commonly occurs in various micro-structures [25]. This non-equilibrium process is statistically asymmetric in the sense that its cumulants of odd order do not

vanish and as a consequence it may induce a directed transport [20,21].

#### 3.1. Model I

As the first model let us consider an over-damped Brownian particle which is driven by white Poissonian noise in a spatially periodic potential. The dynamics of the system is described by Langevin equation of the form [25]

$$\dot{x}(t) = -V'(x) + \eta(t) + \xi(t), \quad (1)$$

where  $V(x)$  is a periodic potential with period  $L$ , so  $V(x + L) = V(x)$ . It is assumed to be symmetric piece-wise linear potential, namely

$$V(x) = \begin{cases} 1 + x, & x \in [-1, 0] \bmod L \\ 1 - x, & x \in [0, 1] \bmod L. \end{cases} \quad (2)$$

The periodic potential  $V(x)$  is plotted on Fig. 2 for  $L = 2$ . The process  $\eta(t)$  is white Poissonian noise defined as

$$\eta(t) = \sum_{i=1}^{n(t)} z_i \delta(t - t_i) - a. \quad (3)$$

The random times  $t_i$  form a Poisson sequence, i.e., the probability that a sequence of  $k$  impulses occurs in the interval  $(0, t)$  is given by the Poisson distribution

$$\Pr\{n(t) = k\} = \frac{(\lambda t)^k}{k!} e^{-\lambda t}. \quad (4)$$

The parameter  $\lambda$  determines the average number of  $\delta$ -pulses per unit time. The amplitudes  $\{z_i\}$  are distributed according to exponential probability density

$$\rho(z) = \zeta^{-1} \theta(z) \exp(-z/\zeta), \quad (5)$$

where  $\theta(z)$  denotes the Heaviside step function. As a consequence the mean value  $\langle z \rangle = \zeta$ . The quantity

$$a = \lambda \langle z \rangle \quad (6)$$

describes according to (3) the negative valued bias of the white Poissonian noise realization between consecutive  $\delta$ -pulses. The process  $\eta(t)$  is therefore of zero mean and possesses auto-correlation function

$$\langle \eta(t) \rangle = 0, \quad \langle \eta(t) \eta(s) \rangle = 2\lambda \zeta^2 \delta(t - s). \quad (7)$$

The last equation defines the Poissonian noise intensity

$$D_p = \lambda \zeta^2. \quad (8)$$

Thermal equilibrium fluctuations due to the coupling of the particle with the environment are modeled as usual by unbiased Gaussian white noise

$$\langle \xi(t) \rangle = 0, \quad \langle \xi(t) \xi(s) \rangle = 2k_B T \delta(t - s), \quad (9)$$

where  $k_B$  is the Boltzmann constant and  $T$  is the temperature. The noise intensity factor

$$D_G = k_B T \quad (10)$$

follows from the fluctuation–dissipation theorem [26]. Furthermore, we assume that thermal fluctuations  $\xi(t)$  are uncorrelated with non-equilibrium noise  $\eta(t)$ , so

$$\langle \xi(t) \eta(s) \rangle = \langle \xi(t) \rangle \langle \eta(s) \rangle = 0. \quad (11)$$

### 3.2. Model II

Along with the white Poissonian noise, the dichotomous process, known also as the random telegraph signal [27–29], is another important type of a non-equilibrium noise that has wide applications in various fields of science. In particular, it may provide a realistic representation of an actual physical situation such as random transitions between meta-stable configurations in a two level system, the phenomenon which is a paradigm for numerous theories. Its realizations consist of randomly distributed jumps between the two states. Being important example of colored noise it mimics the effects of the finite correlation time of the real fluctuations [30].

The second model we analyze is that of an over-damped Brownian particle which is subjected to the dichotomous noise in a periodic structure. Consequently, the dynamics of the system is given by the following equation [27]

$$\dot{x}(t) = -V'(x) + \Gamma(t) + \xi(t). \quad (12)$$

Here  $V(x)$  and  $\xi(t)$  are the previously defined the symmetric piecewise linear potential and the Gaussian white noise, respectively. The process  $\Gamma(t)$  represents non-equilibrium fluctuations modeled by the dichotomous noise

$$\Gamma(t) = \{-F_a, F_b\}, \quad F_a, F_b > 0, \quad (13)$$

with the transition probabilities

$$\Pr(-F_a \rightarrow F_b) = \mu_a = \frac{1}{\tau_a}, \quad \Pr(F_b \rightarrow -F_a) = \mu_b = \frac{1}{\tau_b}, \quad (14)$$

where  $\tau_a$  and  $\tau_b$  are the mean waiting times in states  $F_a$  and  $F_b$ , respectively. The process  $\Gamma(t)$  possesses the auto-correlation

$$\langle \Gamma(t)\Gamma(s) \rangle = \frac{D_D}{\tau} e^{-|t-s|/\tau}, \quad (15)$$

with the intensity  $D_D$  and the correlation time  $\tau$  determined by

$$D_D = \frac{\mu_a \mu_b (F_a + F_b)^2}{(\mu_a + \mu_b)^3}, \quad \frac{1}{\tau} = \frac{1}{\tau_a} + \frac{1}{\tau_b} = \mu_a + \mu_b. \quad (16)$$

We further assume that  $F_a \mu_b = F_b \mu_a$ . Then

$$\langle \Gamma(t) \rangle = 0, \quad D_D = F_a F_b \tau. \quad (17)$$

If  $F_a = F_b$ , the dichotomous noise is symmetric. Otherwise it is asymmetric. As usual, we assume that  $\Gamma(t)$  is not correlated with  $\xi(t)$ .

The most prominent transport quantity for the above systems is the average velocity  $\langle v \rangle$ . In the asymptotic long time limit it can be defined as

$$\langle v \rangle = \lim_{t \rightarrow \infty} \frac{\langle x(t) \rangle}{t}. \quad (18)$$

The corresponding stationary probability density  $p(y)$  is expressed in the following way

$$p(y) = \lim_{t \rightarrow \infty} \langle \delta(y - x(t)) \rangle. \quad (19)$$

Generally, this steady state carries a probability current  $J$ . It is significant that both considered systems can actually be studied analytically in the limiting case of no coupling between the particle and the environment, so when  $D_C = k_B T = 0$ . The stationary probability densities and currents satisfy the following differential equations which were obtained in [20,31] for the first and the second model, respectively:

$$J_I = \zeta[-V'(x) - \lambda \zeta]p'(x) + [-V'(x) - \zeta V''(x)]p(x), \quad (20a)$$

$$J_{II} = -\frac{D_{eff}(x)}{1 - \tau V''(x)}p'(x) - \frac{V'(x) + D'_{eff}(x)}{1 - \tau V''(x)}p(x), \quad (20b)$$

where the prime denotes a derivative with respect to  $x$  and the effective diffusion function  $D_{eff}(x) = \tau[F_a + V'(x)][F_b - V'(x)]$ . The value of  $J$  is obtained by the periodic boundary condition  $p(-L/2) = p(L/2)$  and normalization  $\int_{-L/2}^{L/2} dx p(x) = 1$ . We remark that for the piece-wise linear potential the current can be calculated explicitly. For instance, when  $L = 2$  it reads

$$J_I = \frac{1}{4D_p} \frac{\alpha_+ - \alpha_-}{(\alpha_+ - 1)(\alpha_- - 1)}, \quad (21a)$$

$$J_{II} = \frac{1}{4D_D} \frac{\beta_+ - \beta_-}{(\beta_+ - 1)(\beta_- - 1)}, \quad (21b)$$

with

$$\alpha_{\pm} = \exp\left(\frac{a}{D_p(a \mp 1)}\right) \quad (22)$$

$$\beta_{\pm} = \exp\left(\frac{1}{(F_a \pm 1)(F_b \mp 1)\tau}\right).$$

We end this section with the presentation of the relationship between the stationary average velocity  $\langle v \rangle$  which can be obtained from estimate of the moment  $\langle x(t) \rangle$  by using (18) and the steady state current  $J$

$$\langle v \rangle = \langle \dot{x} \rangle = \langle -V'(x) \rangle = -\int_0^L dx V'(x)p(x) = LJ \quad (23)$$

where the expression for  $V'(x)p(x)$  has been inserted from (20) and periodicity of  $p(x)$  and  $V(x)$  has been utilized. We stress that this relation is not so simple for the case of non-stationary states.

### 4. Numerical solution of stochastic differential equations

This section contains a detailed presentation of algorithms employed to solve the models given by (1) and (12). Since stochastic differential equations describing real problems usually cannot be tackled analytically, direct numerical simulations have to be carried to obtain quantities of interest. There is an abundance of methods for solving stochastic differential equations. At the present moment two books written by Platen et al. [8,9] are the most comprehensive references in this field. Most numerical algorithms use a stochastic Taylor expansion technique based on a discrete time approximation. They are classified as either strong or weak schemes. Methods that generate good approximation of a probability distribution and its moments are considered weak schemes. In the remaining part of this paper we present only such class of algorithms.

In this section, we will consider for simplicity only one equation which in the limiting cases reduces to the previously formulated models, namely

$$\dot{x}(t) = -V'(x) + \eta(t) + \Gamma(t) + \xi(t). \quad (24)$$

Putting  $D_D = 0$ , one gets (1). Similarly, in the absence of  $\eta(t)$  ( $D_p = 0$ ) Eq. (12) is restored. To obtain the average velocity  $\langle v \rangle$  multiple realizations of the above system have to be simulated and then according to (18) an ensemble average must be performed. The major disadvantage of this procedure is that a large number of sample paths must be generated to get statistically reliable results. The statistical error due to finite sampling is proportional to  $1/\sqrt{N}$ , where  $N$  is the number of sample paths. However, there is also another type of error, called the systematic error, which emerges from the finite size of the time step  $\Delta t$ . The smaller the time step is, the larger number of sample trajectories must be simulated to get reliable results. In [9] it was pointed out that the following rule makes sense under suitable conditions: with a scheme of weak order of convergence  $\beta$  it is sensible to increase  $N$

at the order  $\Delta t^{-2\beta}$ . Apart from the above discussed shortcomings of computer simulation of stochastic differential equations another common challenge is to choose algorithms which are numerically stable for the task at hand. The precise definition of the numerical stability depends on the context but is typically related to the long term accuracy of an algorithm when applied to a given dynamics. Numerical errors such as round-off and truncation errors are often unavoidable during practical simulations. Therefore, it is important to choose a numerical scheme which does not propagate uncontrolled approximation errors. In particular, if quantities of interest concern the asymptotic long time limit features of the system such as the stationary average velocity  $\langle v \rangle$  numerical stability is an extremely important problem to obtain reliable results. In practice, the basic question on numerical stability should be almost always answered first when deciding which numerical scheme to use. Higher order of convergence is a secondary issue.

**Algorithm 1** The weak order 2.0 predictor–corrector method to integrate  $\dot{x}(t) = -V'(x) + \eta(t) + \Gamma(t) + \xi(t)$ .

---

```

1:  $F_1 \leftarrow -V'(x_n)$ 
2:  $\bar{x}_{n+1} \leftarrow x_n + F_1 \Delta t + \Delta \xi_n$ 
3:  $F_2 \leftarrow -V'(\bar{x}_{n+1})$ 
4: Predictor:  $\bar{x}_{n+1} \leftarrow x_n + \frac{1}{2}(F_1 + F_2) \Delta t + \Delta \xi_n$ 
5:  $F_2 \leftarrow -V'(\bar{x}_{n+1})$ 
6: Corrector:  $x_{n+1} \leftarrow x_n + \frac{1}{2}(F_1 + F_2) \Delta t + \Delta \xi_n + \Delta \eta_n + \Delta \Gamma_n$ 

```

---

Taking all this into account, we decided to employ the weak order 2.0 predictor–corrector method [9] to simulate the stochastic dynamics given by (24). According to the most comprehensive reference [9] it remains accurate, even when using relatively large time step sizes. Furthermore, this method is only slightly computationally more expensive than corresponding explicit schemes, as it uses the same random variables. Predictor–corrector algorithm is similar to implicit methods but does not require the solution of an algebraic equation at each step. Therefore it offers good numerical stability properties which it inherits from the implicit counterparts of its corrector. Moreover, the stated order of convergence is always defined for a worst case scenario, and higher order is often possible for specific classes of stochastic differential equations, e.g. those with additive noise as it is in (24). In the above algorithm one first computes at each step the predicted approximate value  $\bar{x}_{n+1}$  and afterward the corrected value  $x_{n+1}$ . The difference  $\chi = \bar{x}_{n+1} - x_{n+1}$  between these values provides information about the local approximation error. This quantity can be utilized to introduce a dynamic time step size control during the calculations to increase the efficiency of simulation.

**Algorithm 2** The simplified scheme to generate the Gaussian increment  $\Delta \xi_n$ .

---

```

1:  $G = \sqrt{2D_C}$ 
2:  $\varphi \leftarrow U(0, 1)$ 
3: if  $\varphi \leq 1/6$  then
4:    $\Delta \xi_n \leftarrow -G\sqrt{3\Delta t}$ 
5: else if  $\varphi > 1/6 \wedge \varphi \leq 1/3$  then
6:    $\Delta \xi_n \leftarrow G\sqrt{3\Delta t}$ 
7: else
8:    $\Delta \xi_n \leftarrow 0$ 
9: end if

```

---

For weak convergence we only need to approximate the probability distribution induced by the process  $x(t)$ . Therefore, we can replace the Gaussian increments  $\Delta \xi_n$  appearing in Algorithm 1 by other random variables with similar moment properties.

We obtain a simpler scheme by choosing more easily generated random numbers, e.g.

$$\Pr\{\Delta \xi_n = \pm G\sqrt{3\Delta t}\} = \frac{1}{6}, \quad \Pr\{\Delta \xi_n = 0\} = \frac{2}{3}, \quad (25)$$

where  $G = \sqrt{2D_C}$  is thermal noise intensity factor. It leads to the simplified weak order 2.0 predictor–corrector method [9]. This variant of the algorithm can be tens of times faster than the typical version based on the Gaussian random variables. The increase in efficiency reflects the fact that the simplified scheme requires only generation of uniformly distributed random number and therefore is computationally less expensive than its Gaussian counterpart. The loss in accuracy, which is due to the use of multi-point random variables, is for typical parameters below a few percent in terms of relative errors. However, in some particular cases the accuracy of simplified scheme can be superior to that built on Gaussian random numbers. Moreover, in general, one can say that the simplified algorithm often increases numerical stability [9].

**Algorithm 3** The algorithm to calculate the Poissonian increment  $\Delta \eta_n$ .

---

```

1:  $a \leftarrow \sqrt{\lambda D_P}$ 
2: if  $s \leq 0$  then
3:    $\zeta \leftarrow \sqrt{D_P/\lambda}$ 
4:    $\varphi \leftarrow U(0, 1)$ 
5:    $s \leftarrow [-(1/\lambda) \ln(1 - \varphi)]/\Delta t$ 
6:    $\varphi \leftarrow U(0, 1)$ 
7:    $z \leftarrow -\zeta \ln(1 - \varphi)$ 
8:    $\Delta \eta_n \leftarrow z - a\Delta t$ 
9: else
10:   $s \leftarrow s - 1$ 
11:   $\Delta \eta_n \leftarrow -a\Delta t$ 
12: end if

```

---

To incorporate the white Poissonian noise (3) into the presented scheme we first note that each sample realization of the process  $x(t)$  jumps by a random step from the prescribed distribution only at the Poisson time  $t_i$ . During the times between random kicks  $t_{i-1} < t < t_i$  the stochastic dynamics given by (24) reduces to

$$\dot{x}(t) = -V'(x) - a + \Gamma(t) + \xi(t). \quad (26)$$

Then it is sufficient to integrate it by employing Algorithm 1 with  $\Delta \eta_n = -a\Delta t$ . The difference between subsequent Poissonian times  $s_i = t_i - t_{i-1}$  is exponentially distributed

$$\psi_\lambda(s) = \lambda \theta(s) e^{-\lambda s}, \quad (27)$$

hence it can be obtained by the standard transformation  $s = -\ln(u)/\lambda$  of the random number  $u$  uniformly distributed on  $(0, 1)$ . In our case the random amplitudes  $\{z_i\}$  are also exponential random variates so they can be generated in the similar way. At each Poisson time  $t_i$  the kick gives rise to a finite jump  $z_i$  in  $x(t)$ . Then the increment is given as  $\Delta \eta_n = z_i - a\Delta t$ . An obvious shortcoming of this method is the fact that it requires variable time steps which must be adapted to each sequence of Poisson times  $t_i$ . It needs to be emphasized that presented jump-adapted simplified predictor–corrector approximation [9] becomes computationally demanding when the frequency  $\lambda$  of the Poissonian kicks is high, as then the integration time step must be reduced. However, a jump-adapted discretization allows for use of various schemes for the pure drift and diffusion part of the dynamics between subsequent Poisson times. In particular, higher order jump-adapted weak schemes avoid multiple stochastic integrals that involve the white Poissonian noise. For this reason they are relatively simple and can be implemented without much effort. There exist numerical methods that use a constant time step [32,33,9] but usually

their weak order of convergence is lower. Higher order regular schemes are possible to derive, however they are really complicated even for the simplest stochastic differential equations [9].

**Algorithm 4** The algorithm to generate the dichotomous increment  $\Delta\Gamma_n$ .

---

```

1: if  $s \leq 0$  then
2:   if  $d = 0$  then
3:      $d \leftarrow 1$ 
4:      $\varphi \leftarrow U(0, 1)$ 
5:      $s \leftarrow [-(1/\mu_b) \ln(1 - \varphi)/\Delta t]$ 
6:      $\Delta\Gamma_n \leftarrow F_b \Delta t$ 
7:   else
8:      $d \leftarrow 0$ 
9:      $\varphi \leftarrow U(0, 1)$ 
10:     $s \leftarrow [-(1/\mu_a) \ln(1 - \varphi)/\Delta t]$ 
11:     $\Delta\Gamma_n \leftarrow -F_a \Delta t$ 
12:   end if
13: else
14:    $s \leftarrow s - 1$ 
15:   if  $d = 0$  then
16:      $\Delta\Gamma_n \leftarrow -F_a \Delta t$ 
17:   else
18:      $\Delta\Gamma_n \leftarrow F_b \Delta t$ 
19:   end if
20: end if

```

---

Basically, the same reasoning can be repeated to include the dichotomous noise in the employed numerical algorithm, see [34]. For definiteness we consider only a case of the  $-F_a$  state. The sojourn times for staying in this state without flipping is governed by the exponential distribution

$$\psi_{\mu_a}(s) = \mu_a \theta(s) e^{-\mu_a s}. \quad (28)$$

During this period Eq. (24) is reduced to

$$\dot{x}(t) = -V'(x) + \eta(t) - F_a + \xi(t) \quad (29)$$

and can easily be integrated by using the previously outlined algorithms with  $\Delta\Gamma_n = -F_a \Delta t$ . For the next interval sampled from the above distribution with the transition rate  $\mu_b$  the dichotomous noise stays in the  $F_b$  state. Although the presented procedure is simple due to the structure of the dichotomous noise it should be emphasized that this method causes serious problems when is employed to stiff systems with the noise of a short correlation time. It can be remedied by use of the regular scheme which has a constant time step [35]. However, such algorithm is based on the Euler method and therefore is of lower order of convergence.

## 5. Best practices guide

This manual is a guide how to properly accelerate numerical solution of stochastic differential equations by harvesting the power dormant in modern NVIDIA GPU. It presents established parallelization and optimization techniques which help to obtain maximum performance from CUDA. Performance optimization revolves around three basic strategies [10]:

1. Optimize memory usage to achieve maximum memory throughput;
2. Maximize parallel execution to achieve maximum utilization;
3. Optimize instruction usage to achieve maximum instruction throughput.

Based on the recommendations contained in [10] we now will detail on each of these levels of optimization and comment how it is realized in the code supplemented to this paper.

**Algorithm 5** A CUDA kernel to simulate the models given by (1) and (12).

---

```

1: local  $i \leftarrow blockDim.x \cdot blockIdx.x + threadIdx.x$ 
2: load  $x_i$ , system parameters, the noise and RNG state from global memory and store them locally in the set of registers
3: for  $j = 1$  to samples do
4:   advance  $x_i$  by  $\Delta t$  using Algorithm 1
5: end for
6: save  $x_i$ , the noise and RNG state to global memory

```

---

### 5.1. Memory optimization

Memory optimizations are the most important area for performance. Its hierarchy described in Section 2 instantly suggests a strategy to follow which can be summarized as: move as much data as possible to the fastest kind of memory and keep it there as long as possible.

- *Minimize data transfer between the host and the device.*

Theoretically, a single CUDA kernel can be responsible for the entire numerical simulation. This results in just two data transfers, one before and after the simulation. The intermediate data structures are created in device memory, operated by the device and destroyed without ever being touched by the host. However, GPUs installed in desktop computers often have run time limit on kernels. In order to avoid the possible exceeding of this barrier it is reasonable to split the main kernel into smaller parts, each calculating *samples* number of time steps in a single invocation.

- *Make global memory accesses coalesced whenever possible.*

Besides organization of threads into blocks, there is also further arrangement into warps. These are the actual number of threads that gets calculated in SIMT. Global memory accesses by threads of a warp are coalesced by the device into as few as one transaction when certain requirements are met. In order to attain good bandwidth utilization we employed a natural parallelism – each realization of the stochastic process is calculated in a separate thread. Moreover, we run the simulation kernel in grid of one-dimensional blocks with typically 256 threads for each one of them. The latest commodity of NVIDIA GPUs have warp size set to 32 so such a fine-grained parallelism ensures that all global memory accesses are coalesced.

- *Minimize data transfer between global memory and the device.*

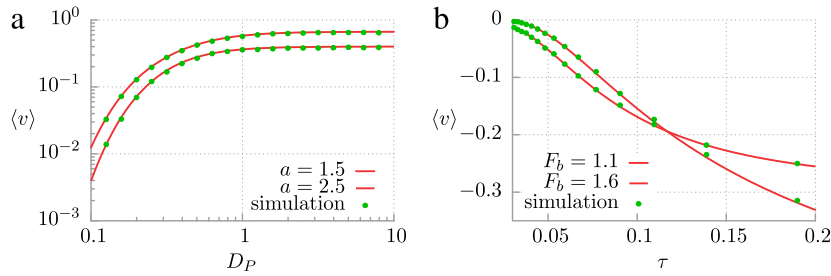
This goal is achieved by maximizing utilization of on-chip memory. In particular, it is recommended to use the fastest type of memory, i.e. the set of registers. At the beginning of each kernel invocation essential variables are loaded from global memory to registers and stay there for all calculations. The reverse operation is performed at the end of kernel execution.

### 5.2. Execution configuration optimizations

To maximize utilization the program should be structured in such a way that it exposes as much parallelism as possible and keep the device busy most of the time.

- *Maintain sufficient occupancy of the device.*

Occupancy is the ratio of the number of active warps per multiprocessor to the maximum number of possible active warps. It should be maintained at the appropriate level because thread instructions are executed sequentially in CUDA and therefore executing other warps when one warp is busy is the only way to hide latency. In order to increase the number of active threads in a straightforward manner we could divide the planned computational work between more threads. However, the application of this idea would unnecessarily complicate the code. Therefore we decided not to do it. The reason for this



**Fig. 3.** Comparison between the simulated (points) and theoretical (curves) asymptotic long time stationary average velocity  $\langle v \rangle$  of the Brownian particle described by the models (1) and (12). Panel (a): the average velocity  $\langle v \rangle$  versus the white Poissonian noise intensity  $D_P$  for the two values of the asymmetry parameter  $a = \sqrt{\lambda D_P} = 1.5$  (solid line) and  $a = 2.5$  (dashed line). Panel (b): the dependence of the same transport characteristic as on (a) on the dichotomous noise correlation time  $\tau$  is depicted for  $F_b = 1.1$  (solid line),  $F_b = 1.6$  (dashed line) and fixed  $F_a = 3$ . In both panels the Gaussian noise intensity  $D_G$  is set to zero.

lies in the peculiarities of the models (1) and (12) where one is usually interested in calculating certain quantities for the whole range of the system parameters. It is much more reasonable to structure the code in the clever way, such that the stochastic dynamics is solved for multiple values of the system parameters in a single kernel call. This procedure allows for very efficient exploration of the parameter space.

- **Thread and block heuristics.**

The number of threads per block should be a multiple of 32, as this provides the optimal computing efficiency and facilitates coalescing.

### 5.3. Instruction optimization

To maximize instruction throughput the program should mainly make as large as possible the usage of simple arithmetic instructions which often have their own hardware accelerated implementation on CUDA GPU.

- **Use single-precision arithmetic whenever it is possible.**

On currently available GPUs the single precision operations are about an order of magnitude faster than the double precision ones. This fact is often considered as significant limitation for numerical calculation. However, the performance gap is becoming smaller for the newer devices. Anyhow, we have found that for the presented models it is not a problem and the use of single precision arithmetic does not significantly impact obtained results. To increase the precision without sacrificing performance, we introduced the folding of  $x$  variable. As only  $x \bmod L$  is relevant for the system dynamics, the device is instructed to operate only on  $x \bmod L$ , simultaneously keeping track of the rest of  $x$  in another variable. This kind of folding creates the possibility of avoiding the numerically pathological situations in which a small quantity is added to a much larger one. In single precision, such an operation is problematic if one of the quantities is more than 7 orders of magnitude larger than the other one.

- **Use the fast math library whenever speed trumps precision.**

CUDA GPUs provide an alternative hardware implementation of various transcendental functions such as trigonometric functions, exponent, and logarithm. In particular, they are hardware-accelerated by the Special Functions Unit of the scalar microprocessor when their arguments do not fall within a narrow range specified for each of them. This leads to an increase of performance at the cost of precision. However, we have found that use of these functions does not degrade the results of the simulation if the above precautions are taken.

Of course, apart from the above listed optimizations there is still some room for a further performance improvement. In particular, it includes utilization of the parallel algorithms, e.g. a sum reduction needed in the ensemble averaging. However, from our experience it follows that such an optimization matters significantly

only for a case of a very large number of particles. Another possible improvement is connected with an avoidance of an instruction branching and a divergence within the same warp. Any flow control instruction (if, switch, do, for, while) can significantly affect the instruction throughput by causing threads of the same warp to diverge; that is, to follow different execution paths. It is generally not an easy task to avoid such problems in designing code which is supposed to be as generic as possible, especially for a programming language like CUDA C. Nevertheless this can be remedied by use of template-based Run-Time Code Generation (RTCG) [15] which is utilized to automatically generate optimized code for a specific system of stochastic differential equations. Moreover, it can also provide some isolation from low-level hardware details of rapidly evolving GPUs, greatly saving time and increasing the productivity of a programmer. Finally, while the computational power available in a single GPU is impressive, one may naturally run a distributed simulation using multiple GPUs which can be physically located in a single computer or in multiple machines connected over a network [15].

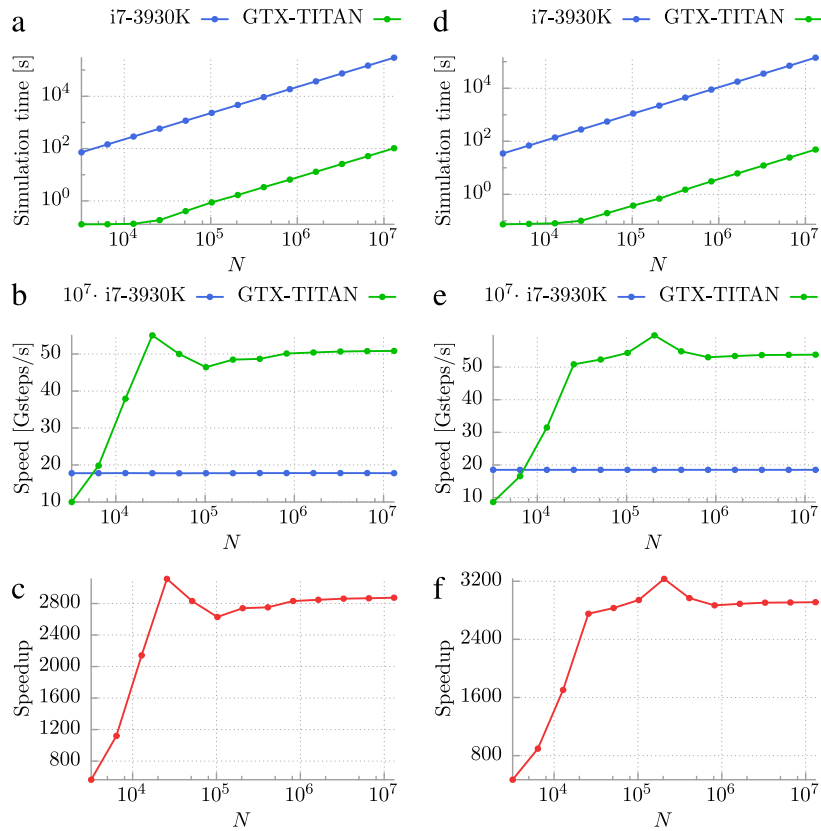
## 6. Validation

In order to validate our implementation of the presented algorithms, we performed simulations for the two previously discussed models and compared our results with the known analytical counterparts.

The tests were run in a single precision using the CUDA Toolkit 5.5 on a 64-bit Linux system. Enabling the full floating-point optimizations does not affect the results and therefore we set the `-O3 --use_fast_math` options of the NVIDIA compiler. In the case of the first model (1) we carried out the simulations of 1024 Brownian particles, each lasting of  $10^3 \cdot (1/\lambda)$  dimensionless units of time. Here, we note that  $1/\lambda$  is the natural characteristic time scale for the studied system. It is reasonable not to specify the time step directly. Instead, it should be dynamically adapted to the chosen characteristic time scale. We provided the `spp = 400` parameter which stands for the *steps per period*. This number together with the spiking rate of the white Poissonian noise  $\lambda$  was used to calculate the suitable step size  $\Delta t = (1/\lambda)/400$ . Each kernel invocation advanced the simulation by `samples = 4000` time steps. We estimated the asymptotic long time stationary average velocity  $\langle v \rangle$  from the results of numerical calculations in the following way

$$\langle v \rangle \approx \langle v \rangle_{sim} = \left\langle \frac{x(t_f) - x(t_i)}{t_f - t_i} \right\rangle, \quad (30)$$

where  $\langle \cdot \rangle$  stands for the ensemble average. The time average is performed over a large, but finite time span  $t_f - t_i$ . In order to discard the influence of the initial transient effects before the  $\langle v \rangle$  is estimated we turned 10% of the  $t_f$  down, i.e. set  $t_i = t_f/10$ . The simulations show the perfect agreement with the analytical predictions published by Łuczka et al. [20] (see Fig. 3(a)).



**Fig. 4.** Single precision performance estimates for the programs *Poisson* and *dich* as a function of the number of particles  $N$ . Left and right panels correspond to the first and the second model, respectively. The top plots present the total simulation time in seconds. The middle ones illustrate the speed estimation (31). The bottom panels depict the GPU speedup (32).

For the second model (12), we simulated the same number of Brownian particles. However, for the dichotomous dynamics there are two natural characteristic time scales, which correspond to the mean waiting time in the first  $\tau_a = 1/\mu_a$  and the second  $\tau_b = 1/\mu_b$  state of the two-state noise. In all calculations we took as a reference the smaller one. Each trajectory evolved until  $10^3$  dimensionless characteristic time scales of the system. The *spp* parameter was fixed to 200 and each kernel call calculated *samples* = 2000 time steps. Similarly to the previous case initial 10% of the simulation time was dropped down for the estimation of  $\langle v \rangle$  due to the possible transients effects. The comparison between the simulated and theoretical results presented in [31] is illustrated in Fig. 3(b). Again, perfect agreement is found.

### 7. Performance evaluation

In this section the performance evaluation of the discussed numerical solution of the models (1) and (12) is presented and compared to the corresponding CPU C implementation. The latter was obtained from the original CUDA source code with the minimum changes possible and the proper replacement of all of the parallel parts.

GCC 4.8.2 was used to compile the CPU version. Here, we present benchmarks only for the two limiting cases: highly optimized single precision version (`nvcc -O3 --use_fast_math, gcc -O3 -ffast-math`) and the double precision counterpart with the default compiler flags. All tests were conducted on the same 64-bit Linux machine using the following hardware:

- CPU: Intel Core i7-3930K @ 3.2 GHz and 32 GB RAM. A single core was used for the simulation.
- GPU: NVIDIA GeForce GTX-TITAN BE installed in a system with the above described CPU.

Figs. 4 and 5 illustrate the performance of the discussed numerical solutions as a function of the total number of the Brownian particles  $N$  for the single and double precision variant, respectively. The performance is defined by means of three measures: a simulation time given in seconds, a speed defined as

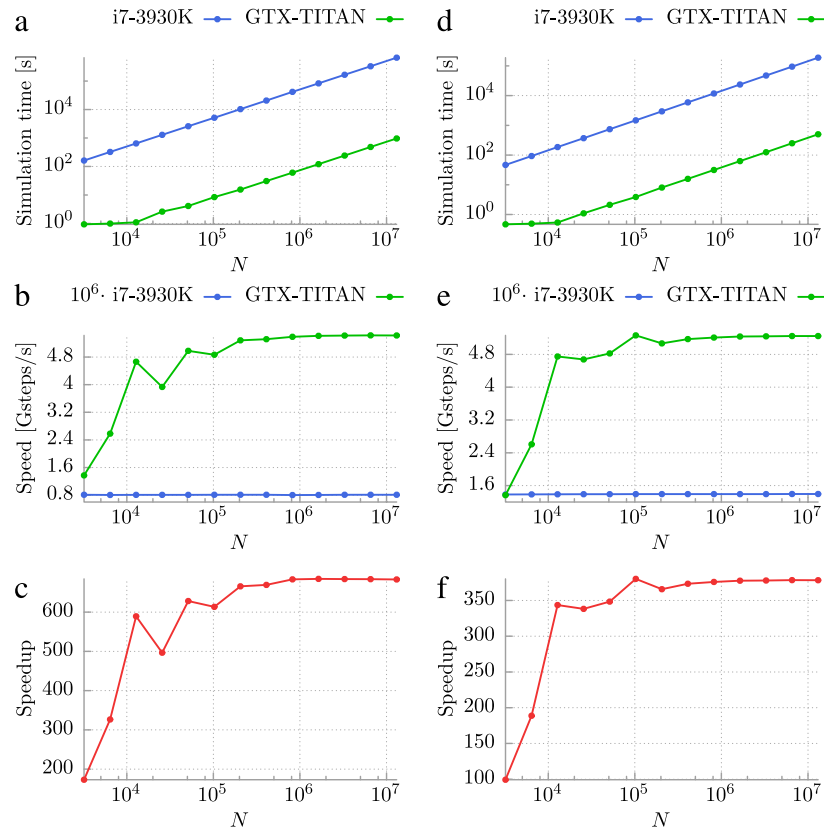
$$\text{Speed} = 10^9 \cdot \frac{\text{total number of time steps}}{\text{simulation time}} \quad (31)$$

and finally an actual speedup expressed in the following way

$$\text{Speedup} = \frac{\text{CPU simulation time}}{\text{GPU simulation time}} \quad (32)$$

Our tests show that in the case of the highly optimized single precision arithmetic the speedups of the astonishing order of about 3000 are possible for the models (1) and (12). Even when the entire simulation is carried by the double precision variables it is still feasible to achieve an amazing performance gain of the order of about 700 and 400 as compared to the CPU implementation. However, this performance profit is clearly dependent on the number of Brownian particles used in the simulation. Panels (c) and (f) of Figs. 4 and 5 depict that initially it increases with the number of simulated trajectories and then saturates around  $N = 10^5$ . This point corresponds to the situation when the computational resources of the GPU are fully utilized. Moreover, for each program there is an ideal number of Brownian particles when the measured speedup is largest. It is because for these points the device occupancy is reasonable high and at the same time the so called register pressure is optimal. Register pressure occurs when there are not enough registers available for a given task. Although each MP contains thousands of 32-bit registers these are partitioned among concurrent threads and therefore one should keep the pressure at the appropriate level [36].





**Fig. 5.** Double precision performance estimates for the programs *Poisson* and *dich* as a function of the number of particles  $N$ . For the rest of the details see Fig. 4.

Overall, by comprehensive profiling of our codes with NVIDIA Visual Profiler we found that the limiting factor of the simulation kernel execution is the GPU instruction throughput. However, there is not much room for improvement in this field because wherever it is possible we used alternative hardware version of transcendental functions. Therefore we conclude that our implementation of the task at hand is very close to optimum. Since the GPU that was used to benchmark our code is neither the most advanced nor the fastest NVIDIA GPU available on the market when this paper is written, one can basically expect even better speedups. Moreover, there is still possibility to run a distributed simulation using multiple GPUs which for convenience can also be located in a single machine. In such situation, achieving the performance gain of unbelievable order  $10^4$  would be just a piece of cake.

## 8. Conclusions

In this paper we have presented an updated and extended step by step guide on how to numerically solve the stochastic differential equations with both Gaussian and non-Gaussian fluctuations, especially in the form of the white Poissonian noise and the dichotomous process. In particular, we have properly harvested the power dormant in the latest commodity of Graphics Processing Units and demonstrated the possibility of achieving speedups of the order as large as about 3000 when compared to the same implementation executed by the single CPU core. This is thanks to the inherent parallel nature of the Monte Carlo integration of the stochastic differential equations. However, in order to achieve such astonishing performance gain one has to carefully redesign the known algorithms. Hopefully, with the help of our work this task will become much easier.

Finally, it is certain that the potential of the General Purpose computing on Graphics Processing Units still awaits to be fully

revealed. Nevertheless, even now it is safe to say that it opens a completely new chapter in the history of high performance computing.

## Acknowledgment

This work was supported in part by the MNiSW program Diamond Grant (J.S.).

## References

- [1] C.W. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences* (Springer Series in Synergetics), fourth ed., Springer, 2009.
- [2] P. Reimann, *Phys. Rep.* 361 (2002) 57.
- [3] P. Hänggi, F. Marchesoni, *Rev. Modern Phys.* 81 (2009) 387.
- [4] F. Jülicher, A. Ajdari, J. Prost, *Rev. Modern Phys.* 69 (1997) 1269.
- [5] E.R. Kay, A.D. Leigh, F. Zerbetto, *Angew. Chem. Int. Ed.* 46 (2007) 72.
- [6] P. Hänggi, H. Thomas, *Phys. Rep.* 88 (1982) 207.
- [7] K. Binder, D.W. Heermann, *Monte Carlo Simulation in Statistical Physics: An Introduction* (Graduate Texts in Physics), fifth ed., Springer, 2010.
- [8] P. Kloeden, E. Platen, *Numerical Solution of Stochastic Differential Equations* (Stochastic Modelling and Applied Probability), Springer, 1992.
- [9] E. Platen, N. Bruti-Liberati, *Numerical Solution of Stochastic Differential Equations with Jumps in Finance* (Stochastic Modelling and Applied Probability), Springer, 2010.
- [10] NVIDIA CUDA programming guide, <http://docs.nvidia.com/cuda/cuda-c-programming-guide/>.
- [11] M. Januszewski, M. Kostur, *Comp. Phys. Commun.* 181 (2009) 183.
- [12] A. Seibert, S. Denisov, A. Ponomarev, P. Hänggi, *Chaos* 21 (2011) 043123.
- [13] A. Barros, E. Garcia, R. Silva, *Comp. Phys. Commun.* 182 (2011) 989.
- [14] A.Y. Polyakov, T. Lyuty, S. Denisov, V.V. Reva, P. Hänggi, *Comp. Phys. Commun.* 184 (2013) 1483.
- [15] M. Januszewski, M. Kostur, [arXiv:1311.2404](https://arxiv.org/abs/1311.2404).
- [16] P. Hänggi, R. Bartussek, in: J. Parisi, S.C. Müller, W. Zimmermann (Eds.), *Nonlinear Physics of Complex Systems – Current Status and Future Trends*, Springer, 1996, pp. 294–308.
- [17] R.D. Astumian, P. Hänggi, *Phys. Today* 55 (2002) 33.
- [18] P. Hänggi, F. Marchesoni, F. Nori, *Ann. Phys.* 14 (2005) 51.
- [19] H. Risken, *The Fokker–Planck Equation: Methods of Solution and Applications* (Springer Series in Synergetics), second ed., Springer, 1989.
- [20] J. Łuczka, R. Bartussek, P. Hänggi, *Europhys. Lett.* 31 (1995) 431.

- [21] P. Hänggi, R. Bartussek, P. Talkner, J. Łuczka, *Europhys. Lett.* 35 (1996) 315.
- [22] J. Spiechowicz, J. Łuczka, P. Hänggi, *J. Stat. Mech* (2013) P02044.
- [23] J. Spiechowicz, J. Łuczka, arXiv:1406.2958.
- [24] J. Spiechowicz, P. Hänggi, J. Łuczka, *Phys. Rev. E* 90 (2014) 032104.
- [25] T. Czernik, J. Kula, J. Łuczka, P. Hänggi, *Phys. Rev. E* 55 (1997) 4057.
- [26] J. Łuczka, *Physica A* 274 (1999) 200.
- [27] J. Kula, T. Czernik, J. Łuczka, *Phys. Rev. Lett.* 80 (1998) 1377.
- [28] J. Kula, M. Kostur, J. Łuczka, *Chem. Phys.* 235 (1998) 27.
- [29] M. Kostur, J. Łuczka, *Phys. Rev. E* 63 (2001) 021101.
- [30] J. Łuczka, *Chaos* 15 (2005) 026107.
- [31] J. Kula, T. Czernik, J. Łuczka, *Phys. Lett. A* 214 (1996) 14.
- [32] C. Kim, E. Lee, P. Hänggi, P. Talkner, *Phys. Rev. E* 76 (2007) 011109.
- [33] M. Grigoriu, *Phys. Rev. E* 80 (2009) 026704.
- [34] V. Palleschi, *Phys. Lett. A* 128 (1988) 318.
- [35] C. Kim, E. Lee, P. Talkner, *Phys. Rev. E* 73 (2006) 026101.
- [36] C. Nugteren, B. Mesman, H. Corporaal, Presented at ODES-8: Proceedings of the 8th Workshop on Optimizations for DSP and Embedded Systems at CGO'10, 2010.