Letter

Universal performance bounds of restart

Dmitry Starkov [1,2] and Sergey Belan [1,3,*]

[1] *Landau Institute for Theoretical Physics, Russian Academy of Sciences, 1-A Akademika Semenova Av., 142432 Chernogolovka, Russia*
[2] *National Research University Higher School of Economics, Faculty of Mathematics, Usacheva 6, 119048 Moscow, Russia*
[3] *National Research University Higher School of Economics, Faculty of Physics, Myasnitskaya 20, 101000 Moscow, Russia*

(Received 13 December 2022; accepted 11 May 2023; published 5 June 2023)

As has long been known to computer scientists, the performance of probabilistic algorithms characterized by relatively large runtime fluctuations can be improved by applying a restart, i.e., episodic interruption of a randomized computational procedure followed by initialization of its new statistically independent realization. A similar effect of restart-induced process acceleration could potentially be possible in the context of enzymatic reactions, where dissociation of the enzyme-substrate intermediate corresponds to restarting the catalytic step of the reaction. To date, a significant number of analytical results have been obtained in physics and computer science regarding the effect of restart on the completion time statistics in various model problems, however, the fundamental limits of restart efficiency remain unknown. Here we derive a range of universal statistical inequalities that offer constraints on the effect that restart could impose on the completion time of a generic stochastic process. The corresponding bounds are expressed via simple statistical metrics of the original process such as harmonic mean h, median value m, and mode M, and, thus, are remarkably practical. We test our analytical predictions with multiple numerical examples, discuss implications arising from them and important avenues of future work.

DOI: 10.1103/PhysRevE.107.L062101

Introduction. Restarting was first proposed as a promising optimization tool of probabilistic algorithms in the early 1990s [1,2]. The restart-induced speed-up may seem counterintuitive at first glance, but the basic idea behind this technique is quite transparent: If the current realization of the randomized algorithm takes too long, it may be faster (on average) to retry attempt with a new random seed to avoid prolonged wandering in the region of the configuration space far from the actual solution. Since then, restarting has become a routine procedure used to hasten computational tasks whose run-time exhibits significant fluctuations [3–17]. In particular, the option of restart is built into state-of-art constraint satisfaction problem solvers [18–21] (see also [22,23] for other computer science applications).

A current wave of interest in this topic from the statistical physics community has been sparked by the work of Evans and Majumdar [24] who showed that stochastic (Poisson) restart expedites diffusion-mediated search. After that it has been demonstrated on a number of different examples that a specially selected restart frequency makes it possible to minimize the average time for completing random search tasks [25–50]. Also, recent development of a model-independent renewal approach, originally proposed for the purposes of describing the single enzyme kinetics [51], has provided a shortcut to the exact completion time statistics of an arbitrary stochastic process under an arbitrary restart protocol [52–54]. This fruitful approach helped to reveal unexpected universality in statistics of optimally restarted processes [53], to establish remarkably simple sufficient conditions for when restart is beneficial [53,55,56] and to rigorously quantify impact of restart on various statistical characteristics of random processes [57–60].

The natural course of development of the research field poses the following question to us: What, if any, are the fundamental limitations of the optimization via restart? The knowledge of the exact completion time distribution allows one to determine the optimal restart strategy (which may be to not restart at all) and the corresponding best possible performance for a given stochastic process. In practice, however, the complete statistics of the completion time is usually unavailable; see, e.g., [2,7,9,10,15,17,51]. The existing literature lacks understanding of how to evaluate the potential performance of restart based on some limited set of statistical characteristics of an observable process. In this paper we fill this gap by exploring how much restart can lower the expected completion time of a stochastic process with partially specified properties. More specifically, we present the universal lower bound on the mean completion time respected by any stochastic process under an arbitrary restart protocol. Besides, we construct the universal upper bound on the mean completion time of generic stochastic process at optimal restart conditions. Both types of probabilistic inequalities are generalized to the high order statistical moments of random completion time. Being formulated in terms of easily interpreted statistical characteristics of an unperturbed stochastic process, the resulting bounds provide valuable insight into what restart can and cannot do. As a useful corollary, our analysis provides a novel sufficient condition for restart to be beneficial which requires very little information on

*sergb27@yandex.ru

the statistics of the random process of interest. Finally, we propose a broad generalization of the well-known inequality offering constraint on the relative fluctuation of the expected completion time of optimally restarted processes.

Model formulation. Consider a generic stochastic process which ends after a random time T if allowed to take place without interruptions. Statistical properties of the variable T are described by the probability density $P(T)$ with a proper normalization $\int_0^\infty dTP(T) = 1$. This implies that the process terminates in finite time with probability 1. As discussed below, our key results remain unchanged or require a trivial modification when one introduces the nonzero probability of never stopping.

The restart protocol $\mathcal{R}$ is characterized by a (possibly infinite) sequence of inter-restart time intervals $\tau_1, \tau_2, \ldots$. If the process is completed prior to the first restart event, the story ends there. Otherwise, the process will start from scratch and begin anew. Next, the process may either complete prior to the second restart or not, with the same rules. This procedure repeats until the process finally reaches completion. Importantly, we assume that protocol $\mathcal{R}$ is uncoupled from the process internal dynamics: The restart decisions do not use information on the current internal state of the process.

In the simplest case of strictly periodic protocol, which is of particular methodological importance as explained below, the process is restarted whenever τ units of time pass. The expected value of the random completion time T_τ of the process subject to such a restart procedure can be obtained by averaging of an appropriate renewal equation. The result is given in the following expression:

$$\langle T_\tau\rangle = \frac{\int_0^\tau P(T)T\,dT + \tau\int_\tau^\infty P(T)dT}{\int_0^\tau P(T)dT}, \tag{1}$$

which, thus, relates the expectation of the random completion time in the presence of periodic restart to the statistics of the "bare" (i.e., restart-free) process.

Restart performance is limited to a quarter of the harmonic mean completion time. First of all, we seek to derive inequality of the form $\langle T_\mathcal{R}\rangle \geqslant C\mathcal{T}$, where $T_\mathcal{R}$ denotes the random completion time of the generic stochastic process under arbitrary restart protocol $\mathcal{R}$, $\mathcal{T}$ is expressed through some simple statistical characteristics of the original process (such as statistical moments, quantiles, or mode of the probability density $P(T)$), and C is the universal positive constant which depends neither on the specific form of $P(T)$ nor on the particular restart schedule $\mathcal{R}$.

Previous works have shown the importance of relative fluctuation σ/μ, where $\mu = \langle T\rangle$ and $\sigma = \sqrt{\langle T^2\rangle - \langle T\rangle^2}$, for the analysis of the potential response of the stochastic process to restart. Namely, the inequality $\sigma/\mu > 1$ represents a sufficient condition for the existence of a restart protocol that reduces the expected completion time [52,53,56]. Given this result, let us first find out if knowledge of the mean value μ and the standard deviation σ allows one to write a lower bound on the average performance of restart. Consider, probability density $P(T) = p\delta(T - t_1) + (1 - p)\delta(T - t_2)$, where $0 \leqslant t_1 \leqslant t_2$ and $0 \leqslant p \leqslant 1$. Putting $t_2 = \frac{\mu^2+\sigma^2}{\mu}$, $p = \frac{\sigma^2}{\mu^2+\sigma^2}$, $\tau = t_1 + 0$, and $t_1 \to 0$, one immediately obtains from Eq. (1) $\langle T_\tau\rangle = t_1/p \to 0$. We see that for the fixed values of μ and σ, the completion time $\langle T_\tau\rangle$ can be arbitrarily small. Therefore, the pair (μ, σ) does not produce any nontrivial lower bound.

Our derivation of the desired lower bound limit is based on the special properties of the periodic restart strategy. As shown by Luby *et al.* [2] for a discrete time case and generalized to continuous settings by Lorenz [17] (see also [61] for simpler and even more general proof), if you found a value $\tau_* \geqslant 0$ (probably $\tau_* = +\infty$) such that $\langle T_{\tau_*}\rangle \leqslant \langle T_\tau\rangle$ for any $\tau \geqslant 0$, then $\langle T_{\tau_*}\rangle \leqslant \langle T_\mathcal{R}\rangle$ for all $\mathcal{R}$. In other words, optimally tuned periodic restart beats any other restart strategy. In addition, the same authors have proved (see also [61]) that the mean performance of an optimal periodic restart obeys the condition,

$$\langle T_{\tau_*}\rangle \geqslant \frac{1}{4}\min_\tau \frac{\tau}{Pr[T \leqslant \tau]}. \tag{2}$$

Let us show that using Eq. (2) together with the optimal property of the periodic restart strategy leads to a simple performance bound of restart. Namely, applying Markov's inequality [62] to the variable $\omega = 1/T$ we find $\Pr[T \leqslant \tau] = \Pr[\omega \geqslant \frac{1}{\tau}] \leqslant \tau\langle\omega\rangle = \tau\langle\frac{1}{T}\rangle$. Next, taking into account Eq. (2) one obtains $\langle T_{\tau_*}\rangle \geqslant \frac{1}{4}h$, where $h = \langle T^{-1}\rangle^{-1}$ is the harmonic mean completion time of the original process. And finally, since $\langle T_\mathcal{R}\rangle \geqslant \langle T_{\tau_*}\rangle$ for any $\mathcal{R}$, this yields

$$\langle T_\mathcal{R}\rangle \geqslant \tfrac{1}{4}h. \tag{3}$$

No constraints have been imposed on the form of $P(T)$, and, therefore, Eq. (3) is universally valid for any setting. What is more, this estimate remains valid also when stochastic process may have nonzero probability of never ending, if h is always understood as the harmonic mean completion time of the halting trials [63].

Particular case of smooth unimodal distribution. A somewhat less general, but still informative, result can be obtained if we assume that the completion time distribution $P(T)$ is smooth and exhibits single local maximum at some nonzero value of T. This class of probability densities covers, in particular, a vast number of random search models; see, e.g., Refs. [24,27–29,33,36,38,42–44,48]. The efficiency of any restart protocol in this case satisfies the inequality,

$$\langle T_\mathcal{R}\rangle \geqslant \tfrac{1}{4}M, \tag{4}$$

where $M = \mathrm{argmax}_T P(T) > 0$ is the mode of the probability distribution $P(T)$, i.e., the value of the random completion time T that occurs most frequently. To prove this result let us introduce $\tau_0 \equiv \mathrm{argmin}_\tau \frac{\tau}{\Pr[T\leqslant\tau]}$. Clearly, assumption $M > 0$ implies that $\tau_0 > 0$. Since the smooth function $f(\tau) = \frac{\tau}{\Pr[T\leqslant\tau]}$ attains its minimal value at $\tau = \tau_0$, one obtains $df(\tau_0)/d\tau = 0$ or, equivalently, $P(\tau_0)\tau_0 = \int_0^{\tau_0} P(T)dT$. Next, as the unimodal function $P(T)$ is nondecreasing on the interval form 0 to M, this extrema condition implies the inequality $\tau_0 \geqslant M$ and, therefore, $\frac{\tau_0}{\Pr[T\leqslant\tau_0]} \geqslant \frac{M}{\Pr[T\leqslant\tau_0]} \geqslant M$. Together with Eq. (2) this yields inequality $\langle T_{\tau_*}\rangle \geqslant \frac{1}{4}M$. Recalling that $\langle T_\mathcal{R}\rangle \geqslant \langle T_{\tau_*}\rangle$ for all $\mathcal{R}$, we then obtain Eq. (4). Note also, that if the probability distribution $P(T)$ has multiple local maxima, then $\langle T_\mathcal{R}\rangle \geqslant \frac{1}{4}M^{\min}$, where $M^{\min}$ is the leftmost mode. Moreover, similarly to Eq. (3), inequality (4) holds even for potentially nonstopping processes, with the obvious caveat that M should now be considered as the most frequent completion time of halting trials.

The twice median sets upper bound on the optimized mean completion time. Having considered the lower bounds, we now turn to the opposite question. How good is the best restart strategy? In other words, we wish to construct an inequality of the form $\langle T_{\tau_*}\rangle \leqslant C\mathcal{T}$, where the time scale $\mathcal{T}$ is determined by the properties of the original stochastic process, and C is the universal positive constant which depends neither on specific form of $P(T)$ nor on the optimal restart period τ_*.

It is easy to understand that the upper bound limit on optimal performance cannot be expressed via the harmonic mean h or the mode M. Indeed, for the half-normal distribution $P(T) = \sqrt{\frac{2}{\pi\sigma^2}}e^{-\frac{T^2}{2\sigma^2}}$ one has $\tau_* = +\infty$, so that $\langle T_{\tau_*}\rangle = \langle T\rangle > 0$, whereas $h = M = 0$. Therefore, inequalities of the form $\langle T_{\tau_*}\rangle \leqslant C_1 h$ and $\langle T_{\tau_*}\rangle \leqslant C_2 M$, where C_1 and C_2 are positive constants, cannot be universally valid.

The desired universal upper bound can be expressed in terms of the median completion time m of the original process obeying by definition the equation $\Pr[T \leqslant m] = 1/2$. Indeed, taking into account that $\langle T_{\tau_*}\rangle \leqslant \langle T_\tau\rangle$ for any $\tau \geqslant 0$ together with the inequality $\langle T_\tau\rangle \leqslant \frac{\tau}{Pr[T\leqslant\tau]}$, which straightforwardly follows from Eq. (1), we find $\langle T_{\tau_*}\rangle \leqslant \frac{\tau}{\Pr[T\leqslant\tau]}$. Substituting m for τ in the last inequality one obtains

$$\langle T_{\tau_*}\rangle \leqslant 2m. \tag{5}$$

Thus, no matter how heavy the tails of $P(T)$ are, in the presence of an optimally tuned periodic restart, the average completion time does not exceed twice the median of the unperturbed process. More generally, if the process has nonzero probability of never halting, we arrive at the estimate $\langle T_{\tau_*}\rangle \leqslant 2m_s/q$, where m_s denotes the median completion times of the halting trials, whereas q is the probability that process ends for a finite time [61].

Importantly, the bound dictated by Eq. (5) is sharp: For any given m there is a probability density $P(T)$ which is characterized by median value m and for which $\langle T_{\tau_*}\rangle = 2m$. Indeed, for $P(T) = \frac{1}{2}\delta(T-t) + \frac{1}{2}\delta(T-3t)$ one obtains $m = t$ and $\langle T_{\tau_*}\rangle = 2t$, where $\tau_* = t$. Note also that Eqs. (3) and (5) do not contradict each other since the relation $h \leqslant 2m$ is always valid as shown in [61]. Also, Eq. (4) is in accord with Eq. (5) since $M \leqslant 2m$ for any continuous unimodal probability distribution (see [61]).

Given this result, it is natural to ask if the median value can be used to construct the bottom bound of restart performance in the spirit of Eqs. (3) and (4). The answer is no. A simple counterexample demonstrating that the inequality $\langle T_\mathcal{R}\rangle \geqslant Cm$, where C is the universal nonzero constant, cannot be valid is given by the Weibull distribution $P(T) = \frac{k}{\lambda^k}T^{k-1}e^{-(\frac{T}{\lambda})^k}$ with $0 < k < 1$ for which $\langle T_{\tau_*}\rangle = 0$, where $\tau_* \to 0$, and $m > 0$.

Beyond the mean performance. Inequality constraints given by Eqs. (3)–(5) can be generalized to higher order statistical moments of random completion time. First of all, since $\sqrt[k]{\langle T_\mathcal{R}^k\rangle} \geqslant \langle T_\mathcal{R}\rangle$ for any natural k due to Jensen's inequality [62], we immediately find from Eq. (3) that $\sqrt[k]{\langle T_\mathcal{R}^k\rangle} \geqslant \frac{1}{4}h$ for a generic stochastic process under an arbitrary restart protocol. Likewise, Eq. (4) trivially entails inequality $\sqrt[k]{\langle T_\mathcal{R}^k\rangle} \geqslant \frac{1}{4}M$ which is valid in the case of unimodal completion time distribution.

A similar extension of Eq. (5) is more tricky. It turns out that the statistical moments of the optimal completion time T_{τ_*} satisfy the inequality,

$$\sqrt[k]{\langle T_{\tau_*}^k\rangle} \leqslant 2\sqrt[k]{k!}m. \tag{6}$$

To prove Eq. (6) let us assume that the process, which is being restarted periodically in an optimal way, becomes subject to an additional restart protocol $\mathcal{R}_\Gamma$ characterized by random restart intervals $\tau_1, \tau_2, \ldots$ independently sampled from Gamma distribution $\rho(\tau) = \frac{\beta^k}{\Gamma(k)}\tau^{k-1}e^{-\beta\tau}$ with shape parameter k and infinitesimally small rate parameter β. In [61] we show that this produces a differential correction $\langle T_{\tau_*+\mathcal{R}_\Gamma}\rangle - \langle T_{\tau_*}\rangle \approx \frac{1}{k!}(\langle T_{\tau_*}\rangle\langle T_{\tau_*}^k\rangle - \frac{1}{k+1}\langle T_{\tau_*}^{k+1}\rangle)\beta^k$ to the mean completion time attained by the optimal periodic restart. Because of the dominance of a periodic restart over other restart strategies, one can be sure that this difference is positive, and therefore

$$\frac{\langle T_{\tau_*}^k\rangle}{\langle T_{\tau_*}\rangle^k} \leqslant k!. \tag{7}$$

Together with Eq. (5) this yields Eq. (6).

Numerical examples. For the sake of illustration we explored several probability distributions $P(T)$, whose response to restart has been extensively discussed in the physical and computer science literature: first-passage time densities for one-dimensional [24,43] and two-dimensional [27,41,43] diffusion processes, first-passage time density for one-dimensional drift-diffusion process [36], first-passage time density for one-dimensional diffusion in logarithmic potential [37], log-normal distribution [15,17,22,51,55], hyper-Erlang distribution [5,23,52,64], hyper-exponential distribution [5,23,52], and Pareto distribution [6,15,17]. The numerical parameters associated with the distributions are given in [61]. Numerical data for the average completion time at the optimally chosen frequency of periodic restart are summarized in the diagrams presented in Fig. 1. We see that all points fall into the region determined by Eqs. (3)–(5) together with the inequalities $h \leqslant 2m$ and $M \leqslant 2m$.

Importantly, theoretical analyses presented above does not answer the question of whether the bounds determined by Eqs. (3) and (4) are sharp. Thus, one may expect, that the constant 1/4 entering the right-hand sides of these equations can potentially be replaced by one larger. Numerical data presented in Fig. 1 allow us to argue that the unknown best possible constants C_1 and C_2 in the inequality constraints $\langle T_\mathcal{R}\rangle \leqslant C_1 h$ and $\langle T_\mathcal{R}\rangle \leqslant C_2 M$ representing the exact lower bounds on $\langle T_\mathcal{R}\rangle$ lie in the ranges $1/4 \leqslant C_1 < 5/4$ and $1/4 \leqslant C_2 < 5/3$.

Corollary 1: Novel criterion of restart efficiency. A notable implication of the upper bound dictated by Eq. (5) is a previously unknown sufficient condition of when restart is helpful in facilitating the process completion. Namely, as follows from Eq. (5), inequality $\mu/m > 2$, where $\mu = \langle T\rangle$, guarantees that there exists finite restart period decreasing the expected completion time. What is particularly interesting is that this simple inequality makes it possible to capture the benefit of restarting in those cases when analysis of the relative fluctuation cannot. In Fig. 2 we compare the applicability of two criteria, $\mu/m > 2$ and $\sigma/\mu > 1$, using the mix of two

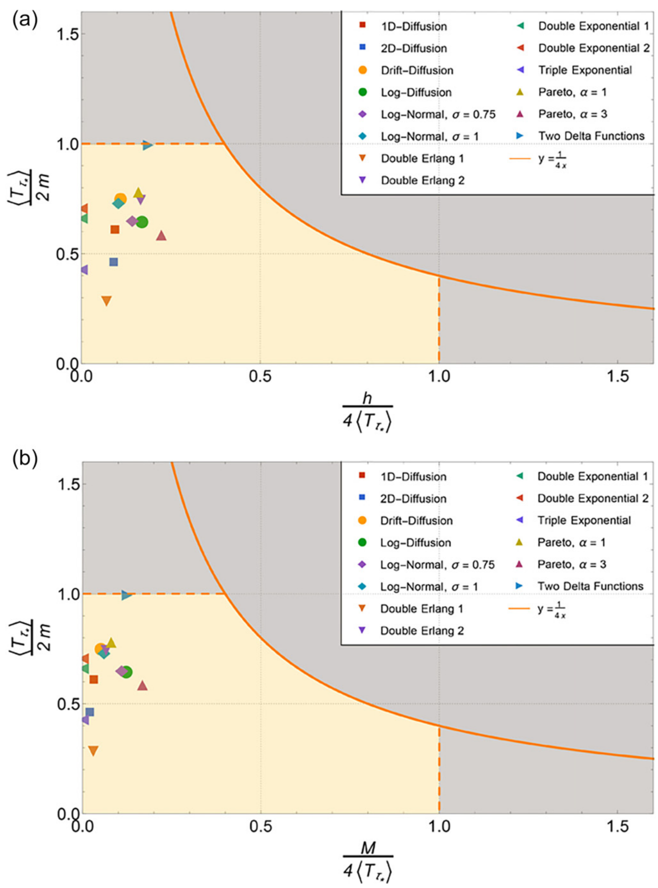

FIG. 1. (a) Diagram in the $x-y$ plane, where $x = h/(4\langle T_{\tau_*}\rangle)$ and $y = \langle T_{\tau_*}\rangle/(2m)$, depicting the relationship between mean completion time of an optimally restarted process $\langle T_{\tau_*}\rangle$, the median m, and the harmonic mean h of the original process. (b) Diagram in the $x-y$ plane, where $x = M/(4\langle T_{\tau_*}\rangle)$ and $y = \langle T_{\tau_*}\rangle/(2m)$, depicting the relationship between mean completion time $\langle T_{\tau_*}\rangle$ of an optimally restarted process, the median m, and the mode M of the restart-free process. As indicated in the legend, each point in these diagrams corresponds to a specified completion time distribution $P(T)$. In accordance with Eqs. (3)–(5) and probabilistic inequalities $h \leqslant 2m$ and $M \leqslant 2m$ all points belong to the light orange region determined by the conditions $0 \leqslant x \leqslant 1, 0 \leqslant y \leqslant 1, y \leqslant 1/(4x)$.

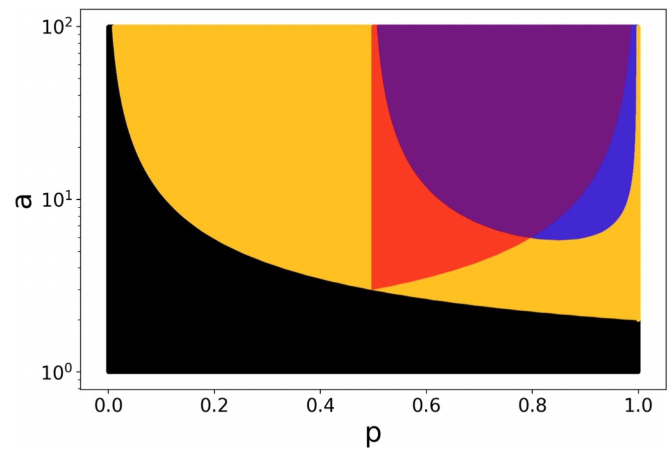

FIG. 2. A diagram of restart efficiency for the random process with completion time probability density $P(T) = p\,\delta(T - t_1) + (1 - p)\ \delta(T - t_2)$ in the plane of dimensionless parameters p and $a = t_2/t_1$. The black area corresponds to the case when restart is not efficient. For the purple area of the diagram restart is efficient and both criteria, $\sigma > \mu$ and $\mu > 2m$, are fulfilled. The region in blue corresponds to the scenario when restart efficiency is captured only by the inequality $\sigma > \mu$, while the region in red corresponds to the scenario when only the condition $\mu > 2m$ is satisfied. Finally, neither of the two sufficient conditions is met for the orange area of the diagram, but restart is useful nevertheless.

delta functions as a model distribution. Clearly, there is a region of parameters, where the relative fluctuation is less than unity, $\sigma/\mu < 1$, while $\mu/m > 2$. More generally, exploiting the well-known probabilistic inequality $|\mu - m| \leqslant \sigma$ [65–67], it is easy to see that the condition $\mu/m > 2$ implies that $\sigma/\mu > 1/2$. What is more, the numerical constant $1/2$ cannot be replaced by a larger one: In [61] we construct an example of probability density $P(T)$ with $\sigma/\mu \to 1/2$ and $\mu/m > 2$.

At the same time it should be noted that the opposite scenario, i.e., $\sigma/\mu > 1$ and $\mu/m < 2$ is also possible (see Fig. 2). These observations suggest that, since both conditions, $\sigma/\mu > 1$ and $\mu/m > 2$, are sufficient, but by no means necessary, in practice they should be used together to compensate (at least partially) each other's shortcomings.

Corollary 2: Generalized fluctuation relation for optimal restart. Let us also note that Eq. (7), which we have derived for auxiliary purposes, is interesting in itself as it represents a higher-order generalization of the well-known inequality constraint $\langle T^2_{\tau_*}\rangle/\langle T_{\tau_*}\rangle^2 \leqslant 2$ first derived by Pal and Reuveni [54]. Interestingly, for $k = 3$, we find from Eq. (7) an estimate $\langle T^3_{\tau_*}\rangle/\langle T_{\tau_*}\rangle^3 \leqslant 6$. These observations indicate that the optimal periodic protocol is generally characterized by smaller randomness of the process durations in comparison with optimal Poisson restart for which one has $\langle T^2_{\tau_*}\rangle/\langle T_{\tau_*}\rangle^2 = 2$ [53] and $\langle T^3_{r_*}\rangle/\langle T_{r_*}\rangle^3 \geqslant 6$ [61] provided $0 < r_* < +\infty$, where r_* is the optimal restart rate. It would be also interesting to compare the effects of periodic and stochastic restart strategies on randomness of completion times in terms of information-theoretic metrics [59,60].

Conclusion and outlook. Revealing explicit performance bounds is crucial in many areas of science and engineering. For example, establishment of the Carnot cycle efficiency [68] played a fundamental role for the development of combustion engines and thermal power plants as it sets a bound on the efficiency of any thermodynamic heat engine. Similarly, Shannon's limit of information capacity [69] has become a guiding principle in the design of communication systems. Although optimization via restart is widely used in the practice of computer programming and represents an active field of academic research in physics, the question of performance limits of this control tool has not been addressed thus far. In this study, we expressed these limits in terms of simple statistical metrics that can be easily estimated based on finite samples of the process completion time.

Importantly, the presented analysis was grounded on the assumption of instantaneous restart events. Although this scenario covers an overwhelming majority of the model problems considered in the literature, it should be kept in mind that in real-life settings restart may be accompanied by some time penalty [51,53,56,70–76]. Say, in the context of single molecule enzyme kinetics, where restart occurs naturally by virtue of intermediate dissociation, some time is required for the enzyme which unbinds from its substrate to find a new one in the surrounding solution [51,53,71]. Similarly, restart of the computer program typically involves a time overhead. Also, models with noninstantaneous restarts provide more realistic pictures of colloidal particle diffusion with resetting [76]. How does accounting for delays modify the bounds constructed here? The straightforward generalization of arguments leading to Eq. (5) brings us to the following simple result (see [61]): $\langle T_{\tau_*}\rangle \leqslant 2m + 2\langle T_{\text{on}}\rangle$, where $\langle T_{\text{on}}\rangle$ is the expectation of the generally distributed time penalty which collectively accounts for any delays that may arise prior to the completion attempt. A similar generalization of the lower bounds given by Eqs. (3) and (4) is an important (and apparently sophisticated) task for future research.

Acknowledgment. The work was supported by the Russian Science Foundation, Project No. 22-72-10052.

[1] H. Alt, L. Guibas, K. Mehlhorn, R. Karp, and A. Wigderson, A method for obtaining randomized algorithms with small tail probabilities, Technical Report TR- 91-057 (International Computer Science Institute, Berkeley, 1991).

[2] M. Luby, A. Sinclair, and D. Zuckerman, Optimal speedup of Las Vegas algorithms, Inf. Proc. Lett. **47**, 173 (1993).

[3] C. P. Gomes, B. Selman, and N. Crato, Heavy-tailed distributions in combinatorial search, in *International Conference on Principles and Practice of Constraint Programming* (Springer, Berlin/Heidelberg, 1997), pp. 121–135.

[4] A. Montanari and R. Zecchina, Optimizing Searches via Rare Events, Phys. Rev. Lett. **88**, 178701 (2002).

[5] A. P. Van Moorsel and K. Wolter, Analysis and algorithms for restart, in *First International Conference on the Quantitative Evaluation of Systems, 2004*, QEST 2004 Proceedings (IEEE, New York, 2004), pp. 195–204.

[6] H. Wu, Randomization and restart strategies, Master's thesis, University of Waterloo, 2006.

[7] H. Wu and P. V. Beek, On universal restart strategies for backtracking search, in *International Conference on Principles and Practice of Constraint Programming* (Springer, Berlin/Heidelberg, 2007), pp. 681–695.

[8] J. Huang, The effect of restarts on the efficiency of clause learning, in *IJCAI* (Morgan Kaufmann Publishers, Cambridge, 2007), Vol. 7, pp. 2318–2323.

[9] M. Streeter, D. Golovin, and S. F. Smith, Restart schedules for ensembles of problem instances, in *AAAI 2007* (AAAI Publications, Washington, DC, 2007), pp. 1204–1210.

[10] M. Gagliolo and J. Schmidhuber, Learning restart strategies, in *IJCAI 2007, Hyderabad, India*, edited by M. Veloso (AAAI Press, 2007), Vol. 1, pp. 792–797.

[11] K. Wolter, Stochastic models for restart, rejuvenation and checkpointing. Technical report, Habilitation Thesis, Humboldt-University, Institut Informatik, Berlin, 2008.

[12] N. A. Wedge and M. S. Branicky, On heavy-tailed runtimes and restarts in rapidly-exploring random trees, in *Twenty-third AAAI Conference on Artificial Intelligence* (AAAI, Washington, DC, 2008), pp. 127–133.

[13] A. Biere, M. Heule, and H. van Maaren, *Handbook of Satisfiability* (IOS Press, Amsterdam, 2009).

[14] V. Roulet and A. d'Aspremont, Sharpness, restart, and acceleration, SIAM J. Optim. **30**, 262 (2020).

[15] J. H. Lorenz, Runtime distributions and criteria for restarts, in *International Conference on Current Trends in Theory and Practice of Informatics* (Edizioni della Normale, Cham, 2018), pp. 493–507.

[16] J. H. Lorenz and J. Nickerl, The potential of restarts for ProbSAT, in *Computer Aided Systems Theory-EUROCAST 2019: 17th International Conference, Las Palmas de Gran Canaria, Spain, February 17–22, 2019, Revised Selected Papers, Part I 17* (Springer International Publishing, 2020), pp. 352–360.

[17] J. H. Lorenz, Restart strategies in a continuous setting, Theory Comput. Syst. **65**, 1143 (2021).

[18] C. Schulte, G. Tack, and M. Z. Lagerkvist, Modeling and Programming with GECODE (2010).

[19] A. Cire, S. Kadioglu, and M. Sellmann, Parallel restarted search, in *Proceedings of the AAAI Conference on Artificial Intelligence* (AAAI, Washington, DC, 2014), Vol. 28, No. 1.

[20] R. Amadini, M. Gabbrielli, and J. Mauro, SUNNY-CP and the MiniZinc challenge, Theory Pract. Logic Prog. **18**, 81 (2018).

[21] M. Wallace, Search control in MiniZinc, in *Building Decision Support Systems* (Springer, Cham, 2020), pp. 165–175.

[22] M. Schroeder and L. Buro, Does the restart method work? Preliminary results on efficiency improvements for interactions of web-agents, in *Proceedings of the Workshop on Infrastructure for Agents, MAS, and Scalable MAS at the Conference Autonomous Agents* (Springer, Berlin/Heidelberg, 2001).

[23] A. P. Van Moorsel and K. Wolter, Analysis of restart mechanisms in software systems, IEEE Transactions on Software Engineering **32**, 547 (2006).

[24] M. R. Evans and S. N. Majumdar, Diffusion with Stochastic Resetting, Phys. Rev. Lett. **106**, 160601 (2011).

[25] M. R. Evans and S. N. Majumdar, Diffusion with optimal resetting, J. Phys. A: Math. Theor. **44**, 435001 (2011).

[26] J. Whitehouse, M. R. Evans, and S. N. Majumdar, Effect of partial absorption on diffusion with resetting, Phys. Rev. E **87**, 022118 (2013).

[27] M. R. Evans and S. N. Majumdar, Diffusion with resetting in arbitrary spatial dimension, J. Phys. A: Math. Theor. **47**, 285001 (2014).

[28] Ł. Kuśmierz, S. N. Majumdar, S. Sabhapandit, and G. Schehr, First Order Transition for the Optimal Search Time of Lévy Flights with Resetting, Phys. Rev. Lett. **113**, 220602 (2014).

[29] Ł. Kuśmierz and E. Gudowska-Nowak, Optimal first-arrival times in Lévy flights with resetting, Phys. Rev. E **92**, 052127 (2015).

[30] A. Pal, A. Kundu, and M. R. Evans, Diffusion under time-dependent resetting, J. Phys. A: Math. Theor. **49**, 225001 (2016).

[31] S. Eule and J. J. Metzger, Non-equilibrium steady states of stochastic processes with intermittent resetting, New J. Phys. **18**, 033006 (2016).

[32] A. Nagar and S. Gupta, Diffusion with stochastic resetting at power-law times, Phys. Rev. E **93**, 060102(R) (2016).
[33] M. R. Evans and S. N. Majumdar, Run and tumble particle under resetting: A renewal approach, J. Phys. A: Math. Theor. **51**, 475003 (2018).
[34] A. Chechkin and I. Sokolov, Random Search with Resetting: A Unified Renewal Approach, Phys. Rev. Lett. **121**, 050601 (2018).
[35] Ł. Kuśmierz and T. Toyoizumi, Robust random search with scale-free stochastic resetting, Phys. Rev. E **100**, 032110 (2019).
[36] S. Ray, D. Mondal, and S. Reuveni, Péclet number governs transition to acceleratory restart in drift-diffusion, J. Phys. A: Math. Theor. **52**, 255002 (2019).
[37] S. Ray, and S. Reuveni, Diffusion with resetting in a logarithmic potential, J. Chem. Phys. **152**, 234110 (2020).
[38] P. Singh, Random acceleration process under stochastic resetting, J. Phys. A: Math. Theor. **53**, 405005 (2020).
[39] S. Ahmad, and D. Das, Role of dimensions in first passage of a diffusing particle under stochastic resetting and attractive bias, Phys. Rev. E **102**, 032145 (2020).
[40] M. Radice, One-dimensional telegraphic process with noninstantaneous stochastic resetting, Phys. Rev. E **104**, 044126 (2021).
[41] F. Faisant, B. Besga, A. Petrosyan, S. Ciliberto, and S. N. Majumdar, Optimal mean first-passage time of a Brownian searcher with resetting in one and two dimensions: experiments, theory and numerical tests, J. Stat. Mech.: Theory Exp. (2021) 113203.
[42] I. Abdoli and A. Sharma, Stochastic resetting of active Brownian particles with Lorentz force, Soft Matter **17**, 1307 (2021).
[43] M. R. Evans, S. N. Majumdar, and G. Schehr, Stochastic resetting and applications, J. Phys. A: Math. Theor. **53**, 193001 (2020).
[44] I. Santra, U. Basu, and S. Sabhapandit, Run-and-tumble particles in two dimensions under stochastic resetting conditions, J. Stat. Mech. (2020) 113206.
[45] G. R. Calvert and M. R. Evans, Searching for clusters of targets under stochastic resetting, Eur. Phys. J. B **94**, 228 (2021).
[46] G. Mercado-Vásquez and D. Boyer, Search of stochastically gated targets with diffusive particles under resetting, J. Phys. A: Math. Theor. **54**, 444002 (2021).
[47] O. L. Bonomo and A. Pal, First passage under restart for discrete space and time: Application to one-dimensional confined lattice random walks, Phys. Rev. E **103**, 052129 (2021).
[48] G. Tucci, A. Gambassi, S. Majumdar, and G. Schehr, First-passage time of run-and-tumble particles with noninstantaneous resetting, Phys. Rev. E **106**, 044127 (2022).
[49] H. Chen and F. Huang, First passage of a diffusing particle under stochastic resetting in bounded domains with spherical symmetry, Phys. Rev. E **105**, 034109 (2022).
[50] S. Ahmad, K. Rijal, and D. Das, First passage in the presence of stochastic resetting and a potential barrier, Phys. Rev. E **105**, 044134 (2022).
[51] S. Reuveni, M. Urbakh, and J. Klafter, Role of substrate unbinding in Michaelis-Menten enzymatic reactions, Proc. Natl. Acad. Sci. USA **111**, 4391 (2014).
[52] T. Rotbart, S. Reuveni, and M. Urbakh, Michaelis-Menten reaction scheme as a unified approach towards the optimal restart problem, Phys. Rev. E **92**, 060101(R) (2015).
[53] S. Reuveni, Optimal Stochastic Restart Renders Fluctuations in First Passage Times Universal, Phys. Rev. Lett. **116**, 170601 (2016).
[54] A. Pal and S. Reuveni, First Passage under Restart, Phys. Rev. Lett. **118**, 030603 (2017).
[55] I. Eliazar and S. Reuveni, Mean-performance of sharp restart: II. Inequality roadmap, J. Phys. A: Math. Theor. **54**, 355001 (2021).
[56] A. Pal, S. Kostinski, and S. Reuveni, The inspection paradox in stochastic resetting, J. Phys. A: Math. Theor. **55**, 021001 (2022).
[57] S. Belan, Restart Could Optimize the Probability of Success in a Bernoulli Trial, Phys. Rev. Lett. **120**, 080601 (2018).
[58] S. Belan, Median and mode in first passage under restart, Phys. Rev. Res. **2**, 013243 (2020).
[59] I. Eliazar and S. Reuveni, Entropy of sharp restart, J. Phys. A: Math. Theor. **56**, 024002 (2023).
[60] I. Eliazar and S. Reuveni, Diversity of sharp restart, J. Phys. A: Math. Theor. **56**, 024003 (2023).
[61] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.107.L062101 for derivation of Eq. (1), proof of global dominance of strictly regular restart, derivation of Eqs. (2) and (7), details of distributions used to generate Fig. 1, a check on the consistency of Eqs. (4) and (5), generalization of Eq. (5) accounting noninstantaneous restarts and nonzero probability of nonhalting, an example of probability density $P(T)$ with $\sigma/\mu = 1/2$ and $\mu/m > 2$, and analysis of completion time fluctuations under optimally tuned Poisson restart.
[62] R. G. Gallager, *Stochastic Processes: Theory for Applications* (Cambridge University Press, Cambridge, 2013).
[63] Noteworthy, once the probability $q = Pr[T < +\infty]$ of halting for a given stochastic process is known, one can improve the estimate given by Eq. (3) as follows: $\langle T_{\mathcal{R}}\rangle \geqslant \frac{1}{4}h/q$.
[64] A. Pal and V. V. Prasad, Landau-like expansion for phase transitions in stochastic resetting, Phys. Rev. Res. **1**, 032001(R) (2019).
[65] H. Hotelling and L. M. Solomons, The limits of a measure of skewness, Ann. Math. Statist. **3**, 141 (1932).
[66] C. A. O'cinneide, The mean is within one standard deviation of any median, Am. Stat. **44**, 292 (1990).
[67] C. Mallows, Another comment on O'Cinneide, Am. Statis. **45**, 257 (1991).
[68] S. Carnot, Reflections on the motive power of fire, and on machines fitted to develop that power, Paris: Bachelier **108**, 108 (1824).
[69] C. E. Shannon, Coding theorems for a discrete source with a fidelity criterion, IRE Nat. Conv. Rec. **4**, 1 (1959).
[70] M. R. Evans and S. N. Majumdar, Effects of refractory period on stochastic resetting, J. Phys. A: Math. Theor. **52**, 01LT01 (2019).
[71] T. Robin, S. Reuveni, and M. Urbakh, Single-molecule theory of enzymatic inhibition, Nat. Commun. **9**, 779 (2018).
[72] S. Ahmad, I. Nayak, A. Bansal, A. Nandi, and D. Das, First passage of a particle in a potential under stochastic resetting: A vanishing transition of optimal resetting rate, Phys. Rev. E **99**, 022130 (2019).

[73] A. Pal, Ł. Kuśmierz, and S. Reuveni, Time-dependent density of diffusion with stochastic resetting is invariant to return speed, Phys. Rev. E **100**, 040101 (2019).

[74] A. S. Bodrova and I. M. Sokolov, Resetting processes with noninstantaneous return, Phys. Rev. E **101**, 052130 (2020).

[75] A. S. Bodrova and I. M. Sokolov, Brownian motion under noninstantaneous resetting in higher dimensions, Phys. Rev. E **102**, 032129 (2020).

[76] O. Tal-Friedman, A. Pal, A. Sekhon, S. Reuveni, and Y. Roichman, Experimental realization of diffusion with stochastic resetting, J. Phys. Chem. Lett. **11**, 7350 (2020).