

5-5-2008

A Constant Entropy Increase Model for the Selection of Parallel Tempering Ensembles

Dubravko Sabo
University of Rhode Island

Markus Meuwly

David L. Freeman
University of Rhode Island, dfreeman@uri.edu

J. D. Doll

Follow this and additional works at: https://digitalcommons.uri.edu/chm_facpubs

Citation/Publisher Attribution

Sabo, D., Meuwly, M., Freeman, D. L. & Doll, J. D. (2008). A Constant Entropy Increase Model for the Selection of Parallel Tempering Ensembles. *Journal of Chemical Physics*, 128(17), 174109. doi: 10.1063/1.2907846

Available at: <http://dx.doi.org/10.1063/1.2907846>

This Article is brought to you by the University of Rhode Island. It has been accepted for inclusion in Chemistry Faculty Publications by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons-group@uri.edu. For permission to reuse copyrighted content, contact the author directly.

A Constant Entropy Increase Model for the Selection of Parallel Tempering Ensembles

Publisher Statement

© 2008 American Institute of Physics.

Terms of Use

All rights reserved under copyright.

A constant entropy increase model for the selection of parallel tempering ensembles

Dubravko Sabo,^{1,2,3} Markus Meuwly,⁴ David L. Freeman,² and J. D. Doll^{1,a)}

¹*Department of Chemistry, Brown University, Providence, Rhode Island 02912, USA*

²*Department of Chemistry, University of Rhode Island, Kingston, Rhode Island 02881, USA*

³*Computational Bioscience, Sandia National Laboratories, Albuquerque, New Mexico, 87185, USA*

⁴*Department of Chemistry, University of Basel, Klingelbergstrasse 80, CH-4056 Basel, Switzerland*

(Received 11 December 2007; accepted 20 March 2008; published online 5 May 2008)

The present paper explores a simple approach to the question of parallel tempering temperature selection. We argue that to optimize the performance of parallel tempering it is reasonable to require that the *increase* in entropy between successive temperatures be uniform over the entire ensemble. An estimate of the system's heat capacity, obtained either from experiment, a preliminary simulation, or a suitable physical model, thus provides a means for generating the desired tempering ensemble. Applications to the two-dimensional Ising problem indicate that the resulting method is effective, simple to implement, and robust with respect to its sensitivity to the quality of the underlying heat capacity model. © 2008 American Institute of Physics. [DOI: [10.1063/1.2907846](https://doi.org/10.1063/1.2907846)]

I. INTRODUCTION

Stochastic quadrature methods are valuable tools in the study of many-body systems. In particular, they offer a general approach to broad classes of both classical¹ and quantum-mechanical problems.² A common practical issue associated with the application of these methods is coping with the “sparse” probability distributions that accompany activated or “rare-event” processes. Unless special care is exercised, the random walk procedures typically used to implement such stochastic approaches can become trapped in the myriad of the system's local potential minima or “inherent structures,” thereby producing an inefficient, or even improper sampling of the underlying probability distribution.

Parallel tempering methods provide a general and effective technique for dealing with sparse sampling issues.^{3–6} The essence of these procedures is to create an ensemble of replicas of the system of interest corresponding to a range of one or more control parameters and to utilize the resulting ensemble to improve the sampling. For example, if the tempering control parameter involved is the system's temperature, the strategy is to use information from the ensemble's high-temperature members, where activation barriers are more easily surmounted, to improve the sampling at lower temperatures, where barrier crossings are otherwise exponentially suppressed. In practice, this is accomplished by introducing attempted exchanges of configurations between ensemble members. For these attempted moves to be useful, the temperature range of the ensemble must be sufficiently large that its higher energy members can surmount relevant activation barriers and the spacing must be sufficiently small that exchange attempts between the various temperatures are statistically significant.

The efficiency of the parallel tempering method depends

on the details of the control parameter selection. A number of techniques for guiding this selection are available and have been reviewed elsewhere.^{6,7} One fruitful line of the development is represented by the works of Kofke,⁸ Kone and Kofke,⁹ and Predescu *et al.*,¹⁰ who have explored connections between the success of tempering exchanges and the heat capacity of the system being studied. More recently, Predescu *et al.*¹¹ have explored ensemble selection schemes based on alternative performance metrics such as the “effective fraction.” Trebst and co-workers^{12–16} have examined the optimization of generalized ensembles from the point of view of stochastic flows in the control parameter space, work that has subsequently been unified and extended by Nadler and Hansmann.^{17–19}

The present paper examines a simple, intuitive approach to the question of parallel tempering temperature selection. As in the works of Kofke,⁸ Kone and Kofke,⁹ and Predescu *et al.*,¹⁰ the system's heat capacity plays a central role. After introducing the basic idea in Sec. II, we examine its relationship to existing methods. In Sec. III, we compare the performance of various temperature selection techniques when applied to a common problem.

II. FORMAL DEVELOPMENTS

In its most basic form, conventional parallel tempering involves the creation of an ensemble of equilibrium simulations of the physical system in question that covers a range of temperatures, denoted here by $\{T_m\}$, $m=1, M+1$. Sampling moves involve a mixture of conventional Metropolis displacements at the individual temperatures as well as attempted exchanges of information between different data streams. One has great freedom in selecting both the range of temperatures as well as the details of the exchange attempts. Indeed, it is guidance concerning the selection of such procedures that we seek.

^{a)}Electronic mail: jimmie_doll@brown.edu.

Kofke,⁸ Kone and Kofke,⁹ emphasize the notion that the probability of accepting tempering exchange attempts between different ensemble members is related to the entropy differences involved. To optimize the performance of parallel tempering, we suggest that it is reasonable to supplement this idea and to require that the *increase* in entropy (or “information”) between successive temperatures be *uniform* over the entire ensemble. That is, if ΔS is the overall entropy difference between the minimum and maximum temperatures of the ensemble, T_1 and T_{M+1} (presumed known), the suggestion is to require the entropy increase between each successive pair of temperatures, T_m and T_{m+1} to be a constant, $\Delta S/M$, for all m . Explicitly, assuming a constant volume,

$$\int_{T_m}^{T_{m+1}} dT \frac{C_v(T)}{T} = \frac{\Delta S}{M}. \quad (1)$$

This choice of temperatures has the desirable feature that each successive temperature stream in the ensemble bears the task of producing a fixed amount of “new” information, a seemingly equitable way of “sharing” the computational workload between the ensemble members. It is clear that this constant entropy increase (CEI) approach can be implemented for general, temperature-dependent heat capacities. Given the overall temperature range, the number of ensemble replicas to be used, and an estimate of the system’s heat capacity (obtained either from experiment, a preliminary simulation, or a suitable physical model), Eq. (1) provides a means for generating the desired tempering ensemble. Specifically, one first integrates the model heat capacity from the minimum to maximum temperatures, T_1 and T_{M+1} , to obtain the overall entropy change ΔS , and then solves for the successive temperatures T_2, T_3, \dots, T_M that satisfy Eq. (1). Depending on one’s wishes, this approach can be implemented either as a single-step or as an iterative procedure in which successive estimates of the system’s heat capacity are utilized to refine the quality of the ensemble selection.

Beyond the question of “how” to select the ensemble temperatures, one must also consider the question of “how many” temperatures to include. A reasonable estimate can be obtained by requiring that entropy difference between successive ensemble members be less than or equal to the natural entropy fluctuations induced by energy fluctuations at those individual temperatures. Such a requirement assures at least a minimal similarity between the configurations sampled at the two different temperatures, and thus, a reasonable chance that tempering exchange attempts between those temperatures will be successful. Recalling that the increase in entropy with respect to energy at fixed volume is given by

$$\left(\frac{\partial S}{\partial E} \right)_v = \frac{1}{T}, \quad (2)$$

and denoting the relevant energy fluctuations at a given temperature as σ_E , a reasonable choice of M is thus to require that the inequality

$$\frac{\Delta S}{M} \leq \frac{\sigma_E}{T} \quad (3)$$

be satisfied at each of the temperatures in the ensemble. Expressing the energy fluctuations in terms of the system’s (generally temperature dependent) heat capacity, Eq. (3) becomes

$$\frac{\Delta S}{M} \leq (k_B C_v)^{1/2}. \quad (4)$$

For a given value of ΔS , we thus want to select a value of M that is large enough to assure that Eq. (4) holds at each of the temperatures in the ensemble.

The CEI approach provides a general means by which we can select the parallel tempering ensemble. Under particular circumstances, these general results simplify significantly. For example, if the heat capacity is relatively constant over a particular temperature interval, Eq. (1) suggests that the differences between successive temperatures be taken as

$$\Delta T_{m,m+1} = T_m \left(\frac{k_B}{C_v} \right)^{1/2}, \quad (5)$$

as suggested by Nymeyer *et al.*²⁰ on the basis of somewhat different arguments. If, in addition, the heat capacity is constant over the entire interval, Eq. (5) produces a geometric distribution for the temperatures in the parallel tempering ensemble.

III. NUMERICAL EXAMPLE

In this section, we consider the application of the CEI approach to a specific example, the zero-field, two-dimensional (2D) Ising model.²¹ We choose this model both to illustrate the use of the present technique on a well-characterized problem and to facilitate the comparison of the CEI method’s performance to that of other, recently published approaches.^{14,16} With the latter goal in mind, we utilize a tempering ensemble that consists of 21 temperatures covering the reduced temperature range of $T=(0.1,10.0)$ throughout.

The exact, analytical heat capacity of the 2D Ising system²¹ is plotted in Fig. 1 as a function of T . When used in the CEI model, this heat capacity produces the 21 member parallel tempering ensemble depicted in Fig. 2. In this figure, the temperature index $N(T)$, which is an integer index ranging from 1 (lowest temperature in the ensemble) to 21 (highest temperature in the ensemble), is plotted as a function of T . As can be seen by comparing Figs. 1 and 2, within the CEI model, a maximum in the heat capacity leads naturally to a “clustering” of the parallel tempering members in the corresponding temperature range. We note also that the strong temperature dependence of the heat capacity for the 2D Ising problem produces an ensemble that is far from geometric in nature.

The system’s exact, analytical heat capacity will typically not be available for general applications. Consequently, it is important to assess the CEI method’s sensitivity to the quality of the heat capacity model used to drive the approach. Does the approach retain its utility when a heat ca-




FIG. 1. Shown is the exact, analytical heat capacity for the 2D Ising model as a function of temperature.

capacity obtained from approximate and/or numerical means is used as the basis for ensemble selection? In Fig. 3, we show the exact 2D Ising model heat capacity of Fig. 1 (dashed line) along with the approximations numerically obtained from Monte Carlo simulations that utilize 20×20 (solid circles) and 4×4 (open circles) periodic lattice grids, respectively. The CEI tempering ensembles generated by using the various estimates of the system's heat capacity are presented in detail in Table I. The various CEI ensembles are labeled according to the heat capacity model used to produce them. For example, the CEI-S(4) ensemble is obtained by using the




FIG. 3. Heat capacities for a 2D Ising model as a function of lattice size (L). The figure shows the analytical results for the infinite lattice (dashed line), for a 20×20 grid (solid circles), and for a 4×4 grid (open circles). These heat capacity results in conjunction with the CEI method generate the parallel temperature ensembles listed in Table I.

numerical 4×4 grid estimate of the system's heat capacity to drive the CEI approach, the CEI-S(20) using the numerical 20×20 grid estimate, and so on. The tempering ensemble produced by the techniques of Ref. 16 is also shown. Finally, Table I contains an *ad hoc* tempering ensemble designed by a knowledgeable (but anonymous!) researcher familiar with the replica exchange methods when asked to generate a




FIG. 2. Shown are the temperatures in a parallel tempering ensemble selected on the basis of the constant entropy increase model and the heat capacity results of Fig. 1. For this example, a total of 21 temperatures in the interval $(0.1, 10.0)$ are used. Note the clustering of temperatures in the vicinity of the heat capacity maximum. The dashed line is intended as a guide for the eye and has no physical significance.

TABLE I. Shown are the temperatures of the various parallel tempering ensembles used for the 2D Ising problem. The results shown correspond to the CEI-S(exact), CEI-S(20), CEI-S(4) ensembles, the ensemble of Katzgraber *et al.* [Ref. 16], and the *ad hoc* ensemble described in the text.

CEI-S(exact)	CEI-S(20)	CEI-S(4)	Ref. 16	<i>ad hoc</i>
0.1000	0.1000	0.1000	0.100 000	0.100
1.4688	1.4635	1.4310	1.394 883	0.807
1.6866	1.6820	1.6460	1.668 180	1.250
1.8373	1.8332	1.7980	1.871 952	1.514
1.9538	1.9496	1.9220	2.025 904	1.650
2.0478	2.0433	2.0310	2.134 071	1.850
2.1250	2.1201	2.1320	2.205 451	2.000
2.1879	2.1836	2.2270	2.255 943	2.100
2.2374	2.2386	2.3210	2.297 688	2.221
2.2702	2.2883	2.4150	2.335 762	2.300
2.3064	2.3365	2.5130	2.375 000	2.500
2.3603	2.3896	2.6160	2.419 257	2.700
2.4321	2.4518	2.7290	2.474 513	2.929
2.5265	2.5341	2.8550	2.552 927	3.200
2.6473	2.6464	3.0010	2.686 760	3.636
2.8115	2.8019	3.1780	2.900 401	4.343
3.0372	3.0220	3.4060	3.228 705	5.000
3.3708	3.3483	3.7270	3.754 932	5.757
3.9414	3.8847	4.2390	4.617 088	7.171
5.1247	4.9906	5.2970	6.226 545	8.586
10.000	10.000	10.000	10.000 00	10.000




FIG. 4. A visual depiction of the various parallel tempering ensembles listed in detail in Table I. Each horizontal line of dots indicates the temperatures for that ensemble. Successive rows in the figure correspond to (starting from the bottom of the figure) the CEI-S(exact), CEI-S(20), CEI-S(4), Ref. 16, and *ad hoc* ensembles of Table I, respectively.

21-temperature ensemble for the current problem. The purpose of including such an *ad hoc* ensemble is to obtain a rough sense of the gain in performance of the various systematic approaches to ensemble selection compared to that achievable using purely intuitive selection approaches. To facilitate a visual comparison of the various ensembles, the temperature points in the neighborhood of the heat capacity maximum for the various ensembles are graphically displayed in Fig. 4.

Figure 5 displays the calculated results for a typical equilibrium property, the heat capacity, obtained from parallel tempering simulations by using the various ensembles of Table I. Each simulation depicted involves 300 loops of 400 000 Monte Carlo moves per loop for each of the 21 members of the tempering ensemble. Each of these 400 000 moves consists, in turn, of one attempted exchange of configurations between adjacent, randomly chosen ensemble temperatures and a conventional Metropolis move for a single, randomly chosen spin in each of the remaining 19 ensemble members.

All tempering ensembles used in Fig. 5 produce accurate estimates of the Ising model's heat capacity. As indicated in Fig. 6, however, they do so with differing efficiency. In Fig. 6, we plot the standard deviations of the heat capacity estimates of the 2D Ising model obtained by using the various tempering ensembles with a fixed number of Monte Carlo points. The results shown in Fig. 6 are single standard deviation estimates obtained during the simulations corresponding to the results in Fig. 5. Clearly, the sampling errors vary depending on the choice of tempering ensemble. Owing to the similarity of the heat capacities upon which they are based, the performance of the CEI-S(exact) and CEI-S(20) ensembles are essentially indistinguishable. Interestingly, both of these CEI results are also effectively identical in




FIG. 5. Shown are the heat capacities computed using the various parallel tempering ensembles listed in Table I. The results shown correspond to the Katzgraber *et al.* (Ref. 16) (solid circles), CEI-S(exact) (open squares), CEI-S(20) (up triangles), CEI-S(4) (down triangles), and *ad hoc* (crosses) ensembles, respectively. It should be noted that all ensembles produce the correct results, although with differing levels of efficiency. It should also be noted that the error bars for these results are too small to be seen in this figure (cf. Fig. 6).

performance to the ensemble produced by the optimization methods discussed in Ref. 16. Finally, we note that the sampling errors for the CEI-S(4) ensemble, while somewhat less efficient than the CEI-S(exact) or CEI-S(20) results, are notably better than those generated by the *ad hoc* ensemble of




FIG. 6. Shown are the statistical errors for heat capacities shown in Fig. 5. The error bars for the various ensembles, too small to be seen in Fig. 5, are designated with the same symbols used in that figure. Errors are single standard deviation estimates obtained during the simulations used in the heat capacity simulations. All simulations utilized a common number of Monte Carlo points (see text for details). The ensembles are labeled by symbols specified in Fig. 5. The connecting lines are intended merely to guide the eye and have no physical significance.

TABLE II. Shown for the various parallel tempering ensembles listed in Table I is a measure of likelihood that a configuration successfully transits the entire ensemble via tempering exchanges without first returning to its starting temperature. The relevant performance index N_0 is defined in greater details in the text. For comparison, the results obtained by using both a 21-member geometric and equally spaced tempering ensemble are also given.

Method	$N_0/100$
Reference 16	2.2
CEI-S(exact)	2.2
CEI-S(20)	2.2
CEI-S(4)	2.7
<i>ad hoc</i>	4.1
Geometric	142
Equally spaced	153

Table I. This finding suggests that the CEI method can be of utility when used with even an approximate model of the heat capacity.

Although the standard deviation in a computed property of interest is perhaps the most direct measure of the quality of the underlying tempering ensemble, we present in Table II a second metric. We list there the likelihood that a configuration, starting at either of the ensemble's temperature extremes, successfully transits the entire ensemble via tempering exchanges without first returning to its starting temperature. As argued by Predescu *et al.*,¹⁰ this effective fraction is a useful metric for judging the quality of the sampling efficiency. The results in Table II are presented as an index N_0 , defined as the number of such starts required, on average, to produce a single, successful, direct transit. We see, for example, from Table II that, on average, roughly one in 2200 configurations "launched" from the $T=0.1$ (or $T=10.0$) data stream via a tempering exchange successfully reaches the $T=10.0$ (or $T=0.1$) data stream directly (i.e., without first returning to the data stream corresponding to its initial temperature). We see that the ranking of the various ensembles when judged by this measure mirrors that found in Fig. 6 on the basis of standard deviations. Although not shown, we note that the ranking of the ensembles obtained in Fig. 6 and Table II agrees with that based on an eigenvalue analysis of the relevant transfer matrix that underlies the tempering exchange process.

IV. SUMMARY

In summary, we have introduced a intuitively appealing criterion for selecting parallel tempering ensembles. The essential feature of this approach is to require that increases in entropy between adjacent members of the tempering ensemble be constant across the entire ensemble. The resulting constant entropy increase or CEI approach is initiated with an estimate of the system's heat capacity obtained from experiment, preliminary simulation, or a suitable physical model. Applications to the 2D Ising problem indicate that this method is effective, simple to implement, and relatively robust with respect to its sensitivity to the quality of the underlying heat capacity model.

The focus of the present discussion is on techniques designed to deal with the application of stochastic quadrature to problems whose underlying probability distributions are sparse. Such problems present well-known sampling difficulties, difficulties for which parallel tempering methods have proved to be a valuable tool. Based on the results presented here, CEI techniques appear to offer a simple means for improving the efficiency of such tempering approaches. It is important to emphasize that caution must always be exercised when dealing with sparse distributions. As discussed elsewhere,²² for example, tempering methods alone may be inadequate for the treatment of systems that exhibit low-temperature, solid-solid transitions. The CEI method presented here is thus intended as a way to improve the efficiency of tempering approaches in situations where they are appropriate, as opposed to a means for extending their intrinsic applicability.

ACKNOWLEDGMENTS

The present authors would like to thank the authors of Ref. 16 for graciously providing the detailed temperature distributions for their parallel tempering ensemble. This work has been supported by the National Science Foundation (Grant No. CHE0554922) and by the Swiss National Science Foundation (to M.M.).

- ¹M. Kalos and P. Whitlock, *Monte Carlo Methods* (Wiley, New York, 1986).
- ²*Quantum Monte Carlo Methods in Physics and Chemistry*, NATO Advanced Studies Institute, Series, Series C: Mathematical and Physical Sciences, edited by M. P. Nightingale and C. J. Umrigar (Kluwer, Dordrecht, 1999), Vol. X.
- ³C. J. Geyer, *Computing Science and Statistics Proceedings of the 23rd Symposium on the Interface* (American Statistical Association, New York, 1991), p. 156.
- ⁴C. J. Geyer and Elizabeth A. Thompson, *J. Am. Stat. Assoc.* **90**, 909 (1995).
- ⁵K. Hukushima and K. Nemoto, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- ⁶U. H. E. Hansmann, *Chem. Phys. Lett.* **281**, 140 (1997).
- ⁷D. J. Earl and M. W. Deem, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
- ⁸D. A. Kofke, *J. Chem. Phys.* **117**, 6911 (2002); **120**, 10852 (2004).
- ⁹A. Kone and D. A. Kofke, *J. Chem. Phys.* **122**, 206101 (2005).
- ¹⁰C. Predescu, M. Predescu, and C. V. Ciobanu, *J. Chem. Phys.* **120**, 4119 (2004).
- ¹¹C. Predescu, M. Predescu, and C. V. Ciobanu, *J. Phys. Chem. B* **109**, 4189 (2005).
- ¹²S. Trebst, D. A. Huse, and M. Troyer, *Phys. Rev. E* **70**, 046701 (2004).
- ¹³S. Trebst, M. Troyer, and U. H. E. Hansmann, *J. Chem. Phys.* **124**, 174903 (2006).
- ¹⁴H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, e-print, Archive/condmat/0602085.
- ¹⁵S. Trebst, D. A. Huse, E. Gull, H. G. Katzgraber, U. H. E. Hansmann, and M. Troyer, in *Developments in Computer Simulation Studies in Condensed Matter Physics XIX*, Springer Proceedings in Physics, Vol. 115, edited by D. P. Landau, S. P. Lewis, and H.-B. Schüttler (Springer, Berlin, 2007).
- ¹⁶H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, *J. Stat. Mech.: Theory Exp.* 2006, P03018.
- ¹⁷W. Nadler and U. H. E. Hansmann, *Phys. Rev. E* **75**, 026109 (2007).
- ¹⁸W. Nadler and U. H. E. Hansmann, *Phys. Rev. E* **76**, 065701R (2007).
- ¹⁹W. Nadler and U. H. E. Hansmann, *Phys. Rev. E* **76**, 057102 (2007).
- ²⁰H. Nymeyer, S. Gnanakaran, and Angel E. Garcia, *Methods Enzymol.* **383**, 119 (2004).
- ²¹K. Huang, *Statistical Mechanics* (Wiley, New York, 1963), Chap. 17.
- ²²V. A. Sharapov, D. Meluzzi, and V. A. Mandelshtam, *Phys. Rev. Lett.* **98**, 105701 (2007).

The Journal of Chemical Physics is copyrighted by the American Institute of Physics (AIP). Redistribution of journal material is subject to the AIP online journal license and/or AIP copyright. For more information, see <http://ojps.aip.org/jcpo/jcpcr/jsp>