

Fast Computation of Recurrences in Long Time Series

Tobias Rawald, Mike Sips, Norbert Marwan and Doris Dransch

Abstract We present an approach to recurrence quantification analysis (RQA) that allows to process very long time series fast. To do so, it utilizes the paradigm *Divide and Recombine*. We *divide* the underlying matrix of a recurrence plot (RP) into sub matrices. The processing of the sub matrices is distributed across multiple graphics processing unit (GPU) devices. GPU devices perform RQA computations very fast since they match the problem very well. The individual results of the sub matrices are *recombined* into a global RQA solution. To address the specific challenges of subdividing the recurrence matrix, we introduce means of synchronization as well as additional data structures. Outperforming existing implementations dramatically, our GPU implementation of RQA processes time series consisting of $N \approx 1,000,000$ data points in about 5 min.

1 Introduction

Many different systems show recurring behavior and its study has attracted attention in almost all scientific fields. The climate system can express recurring behavior due to Milankovich cycles [1], seasonal changes, El Niño/Southern Oscillation, etc. The study of these recurrences allows for a better understanding of the climate system;

T. Rawald (✉) · D. Dransch · M. Sips
German Research Center for GeoSciences GFZ,
Berlin, Germany
e-mail: trawald@gfz-potsdam.de

D. Dransch
e-mail: dransch@gfz-potsdam.de

M. Sips
e-mail: sips@gfz-potsdam.de

N. Marwan
Potsdam Institute for Climate Impact Research, Potsdam, Germany
e-mail: marwan@pik-potsdam.de

its past as well as its future [2–4]. The cardiorespiratory system is investigated by its recurrence properties to get insights in its mechanisms or to measure dysfunctions for diagnosing life threatening conditions [5–9]. Recurrence analysis is promising for investigating brain activity [10] and early detection of epileptic states [11]. Furthermore, it is applied to monitor engines and technological processes, like power generation gas turbines health, cutting processes or crack detection in materials [12–14].

Recurrence plots (RPs) and recurrence quantification analysis (RQA) are powerful methods for analyzing recurrences in measured time series [15]. Their application in many fields have proven their potential for various kinds of analyses [16]. A recurrence plot is a two-dimensional representation of a time series when a m -dimensional phase space trajectory recurs to former (or later) states. Recurrence of a state at time i at a different time j is captured within a two-dimensional squared matrix \mathbf{r} [15]:

$$r_{i,j} = \Theta(\varepsilon - \|x_i - x_j\|), \quad x_i \in \mathbb{R}^m, \quad i, j = 1 \dots N. \quad (1)$$

Both of its axes represent the set of states in temporal order. N is the number of considered states x_i (length of phase space trajectory). ε is a threshold distance, $\|\cdot\|$ a norm, and $\Theta(\cdot)$ the Heaviside function. A pair of states that fulfills the threshold condition is assigned with the value 1 (*recurrence point*), whereas a pair that is considered to be dissimilar is assigned with the value 0. Further details about the reconstruction of phase space vectors from a scalar time series, the recurrence parameters, as well as the typical visual characteristics of RPs can be found in [15].

Small scale structures in the RP, like diagonal lines, are used to define measures of complexity establishing the recurrence quantification analysis (RQA) [15, 17, 18]. As an example, we present the RQA measure *percent determinism* (*DET*):

$$DET = \frac{\sum_{l=d_{\min}}^N l H_D(l)}{\sum_{i,j=1}^N r_{i,j}}. \quad (2)$$

It is the fraction of recurrence points that form diagonal lines; $H_D(l)$ is the number of diagonal lines of exactly length l and d_{\min} is a minimal length necessary to be a diagonal line. This measure characterizes the deterministic nature of a dynamical system from a heuristic point of view (further discussions can be found in [15, 19]). Further measures quantify average line lengths or the complexity of the line length frequency distributions $H_D(l)$ (diagonal lines) and $H_V(l)$ (vertical lines).

The time complexity of basic RQA measures is $O(N^2)$, where N denotes the number of data points. This property hampers an efficient computation for very long data. Furthermore, current implementations are limited by the memory these tools can manage. The CRP Toolbox for MATLAB [20] is limited to $N < 10,000$ data points when calculating the entire RP; for standard PC configurations even less, i.e., $N < 5,000$ data points. The RQA software by Webber [21] is capable of processing only up to $N = 5,000$ data points.

We present an algorithm based on the concept of *Divide and Recombine* (D&R) that allows RQA of very long time series in Sect. 2 and evaluate our algorithm in Sect. 3. In Sect. 4 we highlight the utilization of our approach for a concrete application example taken from the climate research domain.

2 Our Approach

2.1 Divide and Recombine

D&R is a very general approach to address large computational problems. The basic idea is to divide a data set into small sub sets allowing the fast computation of analytical results of the subsets. The intermediate results of the sub sets are recombined into a global solution.

The main application of D&R as presented by Guha et al. in [22] is to enable the analysis of big data sets. They use the *MapReduce* [23] framework to distribute the data and analytical computation between several computing nodes. In contrast, the central challenge in the context of RQA is not the amount of data itself but rather that the determination of RQA measures is compute intensive. To meet this challenge, we utilize the paradigm D&R as follows.

We divide the underlying matrix of a RP, the so called *recurrence matrix* [see Eq. 1] into small sub matrices (*Divide*). For each sub matrix, we compute RQA measures in a massively parallel manner on GPU devices. This includes especially the detection of diagonal and vertical lines. The key issue of the divide step is distributing RQA computations of sub matrices between multiple GPU devices. In a final step, the individual results of the sub matrices are recombined into a global RQA solution (*Recombine*).

Having distributed the computational load to several GPU devices using D&R, we further reduce the runtime by exploiting the parallel processing capabilities of a GPU device itself. We subdivide the processing within a sub matrix into a set of subtasks that can be processed concurrently. The underlying workflow of our approach is summarized in Fig. 1.

An important challenge with D&R for RQA is that diagonal and vertical lines may spread over multiple sub matrices (see Fig. 2). To compute a valid global frequency distribution of diagonal and vertical line lengths, we introduce the *carryover buffer*. A carryover buffer is a data structure that allows to share information about the length of diagonal or vertical lines that exceed a sub matrix. In the following, we describe how our approach addresses the challenges of computing valid global frequency distributions of diagonal and vertical line lengths (see Sect. 1).




Fig. 1 Our D&R approach. **a** Given a recurrence matrix. **b** We divide the recurrence matrix into a set of sub matrices. **c** We distribute sub matrices to several GPU devices. For each sub matrix, we compute the frequency distributions of *diagonal* and *vertical* lines. The computation of the frequency distributions of a single sub matrix is done in a massively parallel manner on a GPU device by identifying independent sub tasks; depicted as *dotted arrows*. **d** The individual results are recombined into a global RQA solution

2.2 Detection of Vertical and Diagonal Lines

To determine vertical lines within the i th column of the recurrence matrix \mathbf{r} , the computation starts at element $r_{i,0}$, representing its first element. The index j is increased until the first recurrence point (see Sect. 1) has been found. Assuming the element $r_{i,j}$ is a recurrence point, the counter representing the length of the current vertical line is increased by 1. This counter is initially set to 0. If $r_{i,j+1}$ is also a recurrence point, the counter is increased again. If $r_{i,j+1}$ is not a recurrence point, the vertical line stops at $r_{i,j}$. We then update the frequency distribution of vertical line lengths $H_V(l)$, reset the line length counter to 0 and continue the detection of vertical lines at $r_{i,j+2}$. Note, each column of the recurrence matrix has a separate line length counter attached.

Subdividing \mathbf{r} , its columns are split into a number of parts belonging to different sub matrices. We introduce the *vertical carryover buffer* to address this challenge. For each column of \mathbf{r} it stores the length of a vertical line that exceeds the horizontal border of a sub matrix.

Figure 3 compares the detection of vertical lines using (a) the recurrence matrix as a whole with applying (b) the vertical carryover buffer to the set of sub matrices. The states of the carryover buffer element corresponding to the column containing the vertical line after processing a particular sub matrix are shown above the recurrence




Fig. 2 Challenges of divide and recombine for RQA. The single *diagonal* and *vertical line* in (a) are distributed between several sub matrices in (b). **a** Full recurrence plot. **b** Divided recurrence plot

plot. If a vertical line reaches the last element of a column of a sub matrix, the carryover buffer element stores its current length. Otherwise its value is 0.

The value of the carryover buffer element is used as input for processing parts of the i th column of \mathbf{r} that belong to the adjacent sub matrix. To compute valid results, applying the carryover buffer requires a particular order of processing the sub matrices. We present the *vertical execution order rule* that reflects this dependency.

A sub matrix of \mathbf{r} is referred to as $S_{g,h}$. The sub matrix representing the bottom left corner of the recurrence matrix has the indices $g = 0$ and $h = 0$.

Vertical Execution Order Rule Suppose a sub matrix $S_{g,h}$ with g being the horizontal index and h being the vertical index. All sub matrices $S_{g,n}$ with $0 \leq n < h$ have to be processed before $S_{g,h}$ is processed.

Sub matrices which do not share any element of the carryover buffer can be processed concurrently. This allows us to compute the local frequency distribution of vertical line lengths of multiple sub matrices at the same time.

This concept can easily be adapted to the detection of diagonal lines, including the use of a carryover buffer and a particular order of execution concerning the set of sub matrices. The major difference is that a diagonal line may transcend not only the horizontal but also the vertical borders of sub matrices. Furthermore, the size of the carryover buffer is equivalent to the number of diagonals of \mathbf{r} .

Figure 4 illustrates the detection of diagonal lines. For the purpose of demonstration, the RP contains only a single diagonal line of length 6. Since diagonal lines may cross horizontal as well as vertical sub matrix borders, we define a *diagonal execution order rule* that reflects this property.

Diagonal Execution Order Rule Suppose a sub matrix $S_{g,h}$ with g being the horizontal index and h being the vertical index. All sub matrices $S_{m,n}$ that fulfill either

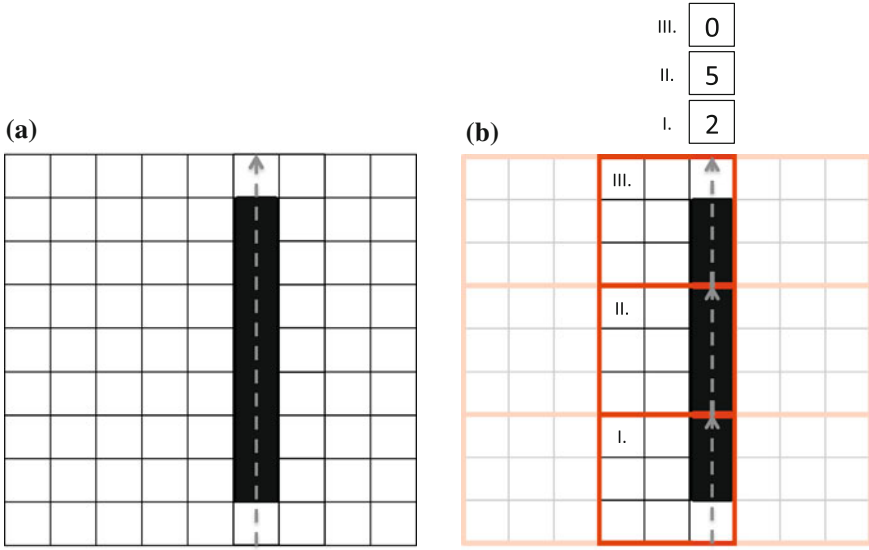


Fig. 3 Detection of vertical lines. Detecting the *vertical line* in (a) is straight-forward. Performing the processing on multiple sub matrices in (b) requires to preserve the execution order $I. \rightarrow II. \rightarrow III.$ The intermediate states of the carryover buffer element corresponding to the column after processing each sub matrix (2, 5 and 0) are depicted above the recurrence plot. **a** Full recurrence matrix. **b** Multiple sub matrices




Fig. 4 Detection of diagonal lines. Given a full recurrence matrix in (a), the detection of the *diagonal line* is straight-forward. By dividing the recurrence matrix in (b), the *diagonal line* stretches over 5 sub matrices. To preserve the validity of the detection result, they must be processed in the order $I. \rightarrow II. \rightarrow III. \rightarrow IV. \rightarrow V.$ The intermediate states of the corresponding carryover buffer element is depicted on the right next to the recurrence plot. **a** Full recurrence matrix. **b** Multiple sub matrices

the condition $(0 \leq m \leq g) \wedge (0 \leq n < h)$ or $(0 \leq m < g) \wedge (0 \leq n \leq h)$ have to be processed before $S_{g,h}$ is processed.

3 Performance Evaluation

To evaluate our approach, we compare the performance characteristics of three RQA implementations. We contribute an implementation of our D&R approach using version 1.1 of the *OpenCL* framework for parallel programming of heterogeneous systems [24]. To fulfill the performance requirements, we use the low-level programming language C++ for implementing the host program.

The hardware setup of the experiment consists of an off-the-shelf desktop workstation, containing an Intel i5-3570 quad-core CPU at up to 3.80 GHz and 16 GB of main memory. It also includes a NVIDIA GeForce GTX 690 that provides two GPU processors running at up to 1.019 GHz; each of them is supplied with 2 GB of memory. In the context of heterogeneous computing, each GPU processor is treated as a separate computing device. The workstation runs on a 64-bit version of *OpenSuse* 12.1 with version 4.2.1 of *CUDA*.

To reduce the impact of outliers, we repeat each experiment five times. Measuring the runtime, we rely on the *chrono* package of the *Boost* library [25] with an accuracy up to one millisecond.

As stated in [26], it is important to compare optimized code that is running on the GPU to optimized CPU code. For this reason, we compare the massively parallel GPU implementation approach to two non-distributed, non-D&R CPU implementations. As a baseline we refer to a single-threaded C++ implementation. Additionally, we employ a parallel C++ implementation that is extended with *OpenMP* statements. It executes the detection of diagonal and vertical lines using multiple CPU threads.

Table 1 Runtime for RQA calculation for time series of varying length

Length	OpenCL (1 × GPU) (s)	OpenCL (2 × GPU) (s)	OpenMP (CPU)	Single thread (CPU)
20,000	0.8	0.8	1.1 s	5.6 s
40,000	1.8	1.7	4.4 s	22.3 s
60,000	3.1	2.5	10.0 s	50.2 s
80,000	4.6	3.4	17.7 s	1 min 29.2 s
100,000	6.3	4.5	27.6 s	2 min 19.5 s
120,000	8.4	5.6	39.9	3 min 21.0 s
140,000	10.7	6.7	54.3 s	4 min 33.2 s
160,000	13.3	8.2	1 min 10.9 s	5 min 56.8 s
180,000	16.2	9.7	1 min 29.7 s	7 min 30.9 s
200,000	19.3	11.4	1 min 50.7 s	9 min 16.8 s

Parameters are embedding dimension $m = 2$ and embedding delay $\tau = 2$




Fig. 5 Runtime comparison. Our GPU implementation (OpenCL) of D&R outperforms both the single-threaded (C++) and multi-threaded (OpenMP) non-D&R CPU implementations. Balancing the work between two GPU processors, the runtime can be reduced additionally up to 40%

Table 1 and Fig. 5 compare the runtime of the three implementations based on a time series capturing the Sinus wave. We vary the length of the time series between 20,000 and 200,000 data points with a step size of 20,000. For each experiment we use the same parameters for reconstructing the states from the time series (see Sect. 1). Furthermore, we set the size of the sub matrices processed by a GPU processor to $20,000 \times 20,000$ elements.

In all cases, our OpenCL implementation outperforms the multi-threaded OpenMP implementation (up to a factor >5) and the single-threaded C++ implementation (up to a factor >28) using only one GPU processor. Using both GPU processors available, the runtime can be reduced additionally up to 40%.

4 Application to Climate Data

In the following we will investigate the hourly temperature dynamics by RQA, which will be applied on a measurement record of hourly air temperature in Potsdam. This record is one of the longest, non-interrupted, hourly climate records in the world. In our analysis it is covering the period from 1893 until 2011 (although the measurements are still ongoing), resulting in 1,043,112 data points (see Fig. 6a). For the period between 1893 and 1974, the warming trend of the annual mean temperature

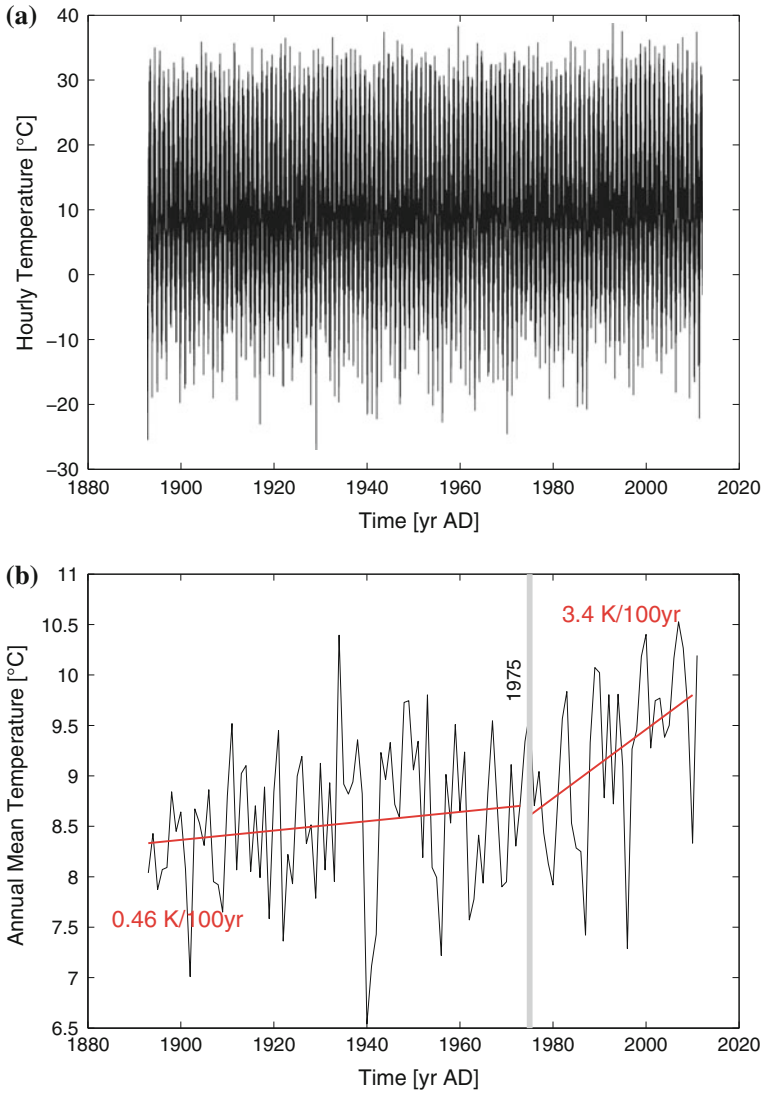


Fig. 6 **a** Hourly temperature in Potsdam. **b** Annual mean temperature in Potsdam and warming trend for the periods 1893–1974 and 1975–2011

was 0.46 K per century, but after 1974 the trend rose to 3.4 K per century (see Fig. 6b). In the following we will consider the full time period as well as the two periods 1893–1974 (718,776 data points) and 1975–2011 (324,336 data points) separately.

To study the short-term dynamics, we remove the annual trend (seasonal cycle) from the data by phase averaging, resulting in an anomaly temperature record. We use a time delay embedding of dimension $m = 5$ and delay $\tau = 3$, which have

Table 2 RQA results for the full time series of hourly temperature anomaly data of Potsdam as well for the two periods 1893–1974 and 1975–2011

Measure	1893–2011	1893–1974	1975–2011
RR	0.12	0.12	0.13
DET	0.94	0.94	0.94
L	8.4	8.4	8.6
LAM	0.96	0.97	0.96

been found by false nearest neighbors approach for finding m [27] as well as a combined autocorrelation and visual recurrence plot inspection approach for finding an optimal τ [28]. We calculate the four RQA measures (1) recurrence rate RR , (2) determinism DET , (3) average diagonal line length L , and (4) laminarity LAM [15] for a recurrence threshold of $\varepsilon = 1$ (Euclidean norm). These measures reflect different aspects of the short-term dynamics, e.g., predictability. We find that all four measures do not remarkably change for the full period and the sub periods 1893–1974 and 1975–2011 (see Table 2). This result suggests that, in contrast to the longer time-scales, the short-term dynamics, and, thus, the short-term weather predictability, has not (yet) changed due to the climate change.

The calculation of these RQA measures benefits highly from our D&R approach, which makes the calculations possible for these long time series. We apply the same implementations as well as the same experimental environment as described in Sect. 3. Using the single-threaded common RQA software, it takes over six hours to calculate the RQA for the full time period, whereas the OpenMP implementation still needs over one and a half hour. The OpenCL implementation allows to reduce the runtime to about five minutes, using two GPU processors (see Table 3, Fig. 7). This significant runtime improvement of RQA will allow comprehensive investigations of big data collections of weather data, consisting of thousands of time series similar to the present Potsdam temperature record.

Despite its performance improvements, it is of great importance that the GPU implementation computes correct RQA results. Table 4 compares a selection of RQA measures computed by the different implementations for the period between 1893 and 2011. It shows that the GPU implementation calculates the same results as the

Table 3 Runtime for RQA calculation for the full time series of hourly temperature anomaly data of Potsdam as well for the two periods 1893–1974 and 1975–2011

Computation schema	1893–1974	1975–2011	1893–2011
OpenCL (1 × GPU)	4 min 40 s	57 s	10 min 11 s
OpenCL (2 × GPU)	2 min 25 s	30 s	5 min 10 s
OpenMP (CPU)	44 min 38 s	9 min 5 s	1 h 33 min 58 s
Single thread (CPU)	2 h 59 min 25 s	36 min 35 s	6 h 18 min 4 s




Fig. 7 Runtime for RQA calculation for the full time series of hourly temperature anomaly data of Potsdam as well for the two periods 1893–1974 and 1975–2011

Table 4 Results for the RQA measures recurrence rate RR , determinism DET , average diagonal line length L , and laminarity LAM computed by the different implementations for the period 1893–2011

Measure	OpenCL (1 × GPU)	OpenCL (2 × GPU)	OpenMP (CPU)	Single thread (CPU)
RR	0.12	0.12	0.12	0.12
DET	0.94	0.94	0.94	0.94
L	8.4	8.4	8.4	8.4
LAM	0.96	0.96	0.96	0.96

single-threaded C++ and the multi-threaded OpenMP implementation. In addition, distributing the computations between two GPU processors does not influence the results.

5 Conclusion

We present an approach based on the paradigm of D&R that allows to perform RQA on very long time series ($>1,000,000$ data points) efficiently. By splitting the underlying matrix of a RP into a set of sub matrices, we are able to distribute the computational load between several GPU devices; each sub matrix is processed individually. We address the problem of diagonal and vertical lines transcending the borders of multiple sub matrices by introducing the diagonal and vertical carryover

buffer as global data structures. To preserve the correctness of the RQA measures, we provide specific execution order rules for processing the individual sub matrices.

In comparison to a parallel CPU implementation using OpenMP, we can improve the runtime performance significantly up to factor >9 , using two GPU processors. Our experiments have shown that GPU devices are well suited to compute basic RQA measures.

The application of the proposed RQA implementation to a specific problem from climate research has demonstrated its potential for an efficient recurrence study and will allow future RQA investigations of very long time series.

Acknowledgments We would like to thank T. Nocke and F.-W. Gerstengarbe for fruitful discussions. We acknowledge support from the Potsdam Research Cluster for Georisk Analysis, Environmental Change and Sustainability (PROGRESS, support code 03IS2191B).

References

1. Muller, R.A., MacDonald, G.J.: *Ice Ages and Astronomical Causes*. Springer, New York (2002). (Springer Praxis Books/Environmental Sciences)
2. Ponyavin, D.I.: *Sol. Phys.* **224**(1–2), 465 (2005). doi:[10.1007/s11207-005-4979-5](https://doi.org/10.1007/s11207-005-4979-5)
3. Donges, J.F., Donner, R.V., Trauth, M.H., Marwan, N., Schellnhuber, H.J., Kurths, J.: *Proc. Nat. Acad. Sci.* **108**(51), 20422 (2011). doi:[10.1073/pnas.1117052108](https://doi.org/10.1073/pnas.1117052108)
4. Goswami, B., Marwan, N., Feulner, G., Kurths, J.: *Eur. Phys. J. Special Topics* **222**, 861 (2013). doi:[10.1140/epjst/e2013-01889-8](https://doi.org/10.1140/epjst/e2013-01889-8)
5. Zbilut, J.P., Koebbe, M., Loeb, H., Mayer-Kress, G.: pp. 263–266. IEEE Computer Society Press (1990), doi:[10.1109/CIC.1990.144211](https://doi.org/10.1109/CIC.1990.144211)
6. Porta, A., Baselli, G., Montano, N., Gneccchi-Ruscione, T., Lombardi, F., Malliani, A., Cerutti, S.: *Biol. Cybern.* **75**(2), 163 (1996)
7. Marwan, N., Wessel, N., Meyerfeldt, U., Schirdewan, A., Kurths, J.: *Phys. Rev. E* **66**(2), 026702 (2002). doi:[10.1103/PhysRevE.66.026702](https://doi.org/10.1103/PhysRevE.66.026702)
8. Van Leeuwen, P., Geue, D., Thiel, M., Cysarz, D., Lange, S., Romano, M.C., Wessel, N., Kurths, J., Grönemeyer, D.H.W.: *Proc. Nat. Acad. Sci.* **106**(33), 13661 (2009). doi:[10.1073/pnas.0901049106](https://doi.org/10.1073/pnas.0901049106)
9. Marwan, N., Zou, Y., Wessel, N., Riedl, M., Kurths, J.: *Philos. Trans. R. Soc. A* **371**(1997), 20110624 (2013). doi:[10.1098/rsta.2011.0624](https://doi.org/10.1098/rsta.2011.0624)
10. Carrubba, S., Frilot II, C., Chesson Jr, A.L., Marino, A.A.: *Med. Eng. Phys.* **32**(8), 898 (2010). doi:[10.1016/j.medengphy.2010.06.006](https://doi.org/10.1016/j.medengphy.2010.06.006)
11. Acharya, U.R., Sree, S.V., Chattopadhyay, S., Yu, W., Ang, P.C.A.: *Int. J. Neural Syst.* **21**(3), 199 (2011). doi:[10.1142/S0129065711002808](https://doi.org/10.1142/S0129065711002808)
12. Bassily, H., Wagner: **10**, 629–635 (2008)
13. Litak, G., Sen, A.K., Syta, A.: *Chaos, Solitons Fractals* **41**(4), 2115 (2009). doi:[10.1016/j.chaos.2008.08.018](https://doi.org/10.1016/j.chaos.2008.08.018)
14. Iwaniec, J., Uhl, T., Staszewski, W.J., Klepka, A.: *Nonlinear Dyn.* **70**(1), 125 (2012). doi:[10.1007/s11071-012-0436-9](https://doi.org/10.1007/s11071-012-0436-9)
15. Marwan, N., Romano, M.C., Thiel, M., Kurths, J.: *Phys. Rep.* **438**(5–6), 237 (2007). doi:[10.1016/j.physrep.2006.11.001](https://doi.org/10.1016/j.physrep.2006.11.001)
16. Marwan, N.: *Eur. Phys. J. Special Topics* **164**(1), 3 (2008). doi:[10.1140/epjst/e2008-00829-1](https://doi.org/10.1140/epjst/e2008-00829-1)
17. Zbilut, J.P., Webber Jr, C.L.: *Phys. Lett. A* **171**(3–4), 199 (1992). doi:[10.1016/0375-9601\(92\)90426-M](https://doi.org/10.1016/0375-9601(92)90426-M)
18. Webber Jr, C.L., Zbilut, J.P.: *J. Appl. Physiol.* **76**(2), 965 (1994)

19. Marwan, N., Schinkel, S., Kurths, J.: *Europhys. Lett.* **101**, 20007 (2013). doi:[10.1209/0295-5075/101/20007](https://doi.org/10.1209/0295-5075/101/20007)
20. Marwan, N.: CRP Toolbox 5.17 (2013). Platform independent (for Matlab). <http://tocsy.pik-potsdam.de/CRPtoolbox>
21. Webber, C.L. Jr.: RQA Software 14.1 (2013). Only for DOS. <http://homepages.luc.edu/~decwebber>
22. Guha, S., Hafen, R., Rounds, J., Xia, J., Li, J., Xi, B., Cleveland, W.S.: *Stat.* **1**(1), 53 (2012). doi:[10.1002/sta4.7](https://doi.org/10.1002/sta4.7). <http://dx.doi.org/10.1002/sta4.7>
23. Dean, J., Ghemawat, S.: *Commun. ACM* **51**(1), 107 (2008)
24. OpenCL 1.1 Specification (2010). <http://www.khronos.org/registry/cl/specs/opencl-1.1.pdf>
25. Dawes, B., Abrahams, D., Rivera, R.: Boost—C++ libraries (2013). <http://www.boost.org>
26. Lee, V.W., Kim, C., Chhugani, J., Deisher, M., Kim, D., Nguyen, A.D., Satish, N., Smelyanskiy, M., Chennupati, S., Hammarlund, P., Singhal, R., Dubey, P.: In: *Proceedings of the 37th Annual International Symposium on Computer Architecture, ISCA'10*, pp. 451–460 (2010)
27. Kennel, M.B., Brown, R., Abarbanel, H.D.I.: *Phys. Rev. A* **45**(6), 3403 (1992). doi:[10.1103/PhysRevA.45.3403](https://doi.org/10.1103/PhysRevA.45.3403)
28. Marwan, N.: *Int. J. Bifurcat. Chaos* **21**(4), 1003 (2011). doi:[10.1142/S0218127411029008](https://doi.org/10.1142/S0218127411029008)