

Numerical estimation of the area of the Mandelbrot set using quad tree tessellation and distance estimation

T. Förstemann*
(Dated: January 23, 2017)

An analytical expression of the area of the Mandelbrot set is unknown. There are different numerical approaches to estimate the area. A drawback of some approaches is that the numerical error of the area estimation is based on statistics and has limited theoretical significance. This approach is based on quad tree tessellations combined with distance estimation and provides precise lower and upper bounds of the area of Mandelbrot set. So far the most accurate numerical estimation of the precise lower and upper bounds by Y. Fisher and J. Hill (2001) on 100 standard single core PCs is based on a quad tree tessellation of depth 19. In this approach a corresponding tessellation of depth 26 is reached using Mathematica and OpenCL, i.e. about $(2^{26-19})^2 = 16\,000$ times larger tessellation. The difference between the upper and lower bound can be halved.

CONTENTS

I. Introduction	1
A. Problem Statement	1
B. Solution Statement	1
II. Materials and Methods	2
A. Hardware and Software	2
B. Methods	2
1. Exterior and Interior Distance Estimation	3
2. Quad Tree Tessellation	4
3. Performance Considerations	4
4. Outline of the Quad Tree Algorithm	6
5. Post Processing of the Tessellation Tree	7
III. Results	7
A. Tile counts at different depths	7
B. Area Estimate	7
C. Time Consumption	7
D. Parameters of depths 25 and 26	7
E. Comparison between $\alpha = 1$ and $\alpha = 2$	9
IV. Discussion	9
A. Soucre Code	12
1. GPU Code for the <i>decide</i> Procedure	12
2. Definition of Mathematica Functions	16
3. Initialisation of Working Files	18
4. Mathematica Code for the Main Loop	18
References	19

I. INTRODUCTION

A. Problem Statement

An analytical expression of the area of the Mandelbrot[1] is unknown. There are many numerical estimations of the area, e.g. [2], based on different methods. At present the most precise estimations are based on Monte Carlo integration[3][4]. A main drawback of Monte Carlo integrations is that they do not provide precise upper and lower bounds of the Mandelbrot set area. Their numerical error is just based on statistics and has limited theoretical significance.

At the moment the most accurate estimation of precise upper and lower bounds by Y. Fisher and J. Hill in about 2001 [5] is based on 100 standard single core PCs. The aim of this approach is to narrow the precise upper and lower bound of this numerical area estimation.

B. Solution Statement

This approach is based on quad tree tessellations combined with distance estimation and provides precise lower and upper bounds of the area of Mandelbrot set. The initial suggestion for this approach came from Prof. L. Bartholdi (Univerity of Göttingen) in 2015. This approach is based on the same method applied by Y. Fisher and J. Hill[5].

By now the approach by Y. Fisher and J. Hill is the most accurate numerical estimation of the precise lower and upper bounds and is based on a quad tree tessellation of depth 19. In this approach a tessellation of depth 26 is reached using Mathematica and OpenCL, i.e. about $(2^{26-19})^2 = 16\,000$ times larger tessellation.

In addition Y. Fisher and J. Hill supposes a factor 2 in the interior distance estimation method based on their numerical findings (refer to [5] section 5, there). Both cases, i.e. $\alpha = 1$ and $\alpha = 2$, are investigated and compared, here.

The GPU code runs in a Mathematica 10 environment on a standard PC with two dual core graphic cards.

* thorsten@foerstemann.name



FIG. 1. Standard PC with two graphic cards.

II. MATERIALS AND METHODS

A. Hardware and Software

A typical PC (refer to figure 1) is used for the numerical calculations. Specific details of the hardware and software configuration are:

- CPU: Intel Core i7 2600K (about \$500 full system)
- GPU: 2x Radeon HD 5970. Thus, there are 4 GPUs with 1600 stream processors (Cypress) each. Ebay prices: \$100 each.
- Power consumption under load: approx. 350 watts. Thus, about 300 kWh per month, i.e. about \$90 energy costs (Germany).
- Mathematica 10, Windows 7 (refer to figure 2)
- ATI driver Catalyst 11.2 with AMD Stream SDK 2.3 [7]



FIG. 2. Mathematica's OpenCL system information. Most important information: support for double precision.

B. Methods

In this section the applied method based on quad tree tessellation and distance estimation is illustrated. The implementation with OpenCL and Mathematica is detailed in the following sections.

Quad tree tessellations and distance estimations are quite established. The challenges of this approach are in particular performance considerations detailed in section II B 3.

1. Exterior and Interior Distance Estimation

The distance to the edge of the Mandelbrot set is needed in the quad tree tessellation, which is outlined in section II B 2.

There are two different algorithms to estimate the distance of a given point to the edge of the Mandelbrot set, e.g. [8]: The exterior distance estimator is suitable for points outside the Mandelbrot set and the interior distance estimator is suitable for points inside the Mandelbrot set.

The estimation of the exterior distance is based on iteration sequences generated by the standard escape time algorithm of the Mandelbrot set, e.g. [11]. For the interior distance estimation additional orbit detection is needed. Parameters of the implemented escape time algorithm with orbit detection are:

- The escape radius is set to 100 instead of 4 in order to gain a few additional iterations for a better convergence of the exterior distance estimation (refer to line 24 in section A 1). For a numerical evaluation of this inaccuracy refer to [5] appendix A, there.
- The maximum iteration is larger than $1024 + 20\,000 \cdot 2^9 = 10\,241\,024$. At tessellation depths 25 and 26 the limit is $1024 + 20\,000 \cdot 2^0 = 21\,024$. Specialities of depths 25 and 26 are discussed in section III D.
- Orbits of maximum cycle length $20\,000 \cdot 2^8 = 5\,120\,000$ can be detected if the cycle starts at the latest after $1024 + 20\,000 \cdot 2^8 = 5\,121\,024$ iterations. At tessellation depths 25 and 26 orbits of maximum cycle length $20\,000 \cdot 2^0 = 20\,000$ can be detected if the cycle starts at the latest after 1024 iterations.
- Cardioid and period-2 bulb testing are implemented (refer to e.g. [12]).

The interior and exterior distance estimator is part of the *decide* procedure, which is described in section II B 4. The OpenCL implementation of the *decide* procedure is outlined in section A 1. In order to implement these algorithms in OpenCL pseudo code resources are helpful, e.g. [9][10] for the exterior method and [10] for the interior method.

Due to Koebe quarter theorem[13] the factor 2 at line 229 in section A 1 should amount to 4 instead. In fact, [10] uses factor 4. On the other side Y. Fisher and J. Hill[5] suggests that factor 2 works quite well for the area estimation applied here. The factor 2, here, corresponds to the case $\alpha = 2$ in [5]. Accordingly, factor 4 corresponds to $\alpha = 1$ in [5]. In addition, using the actual implementation of the *MandelbrotDistance* function in Mathematica 10 yields the same results as this implementation with factor 2. More details are discussed in section III E.

Thus, we have the following two different interior distance estimators:

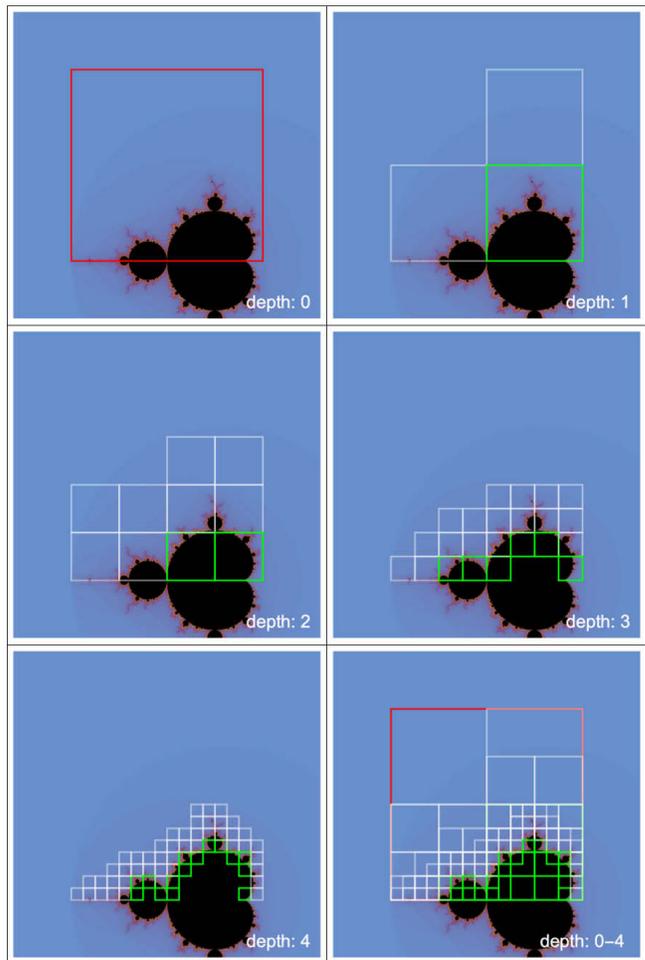


FIG. 3. The lowest five depths of the applied tessellations and a synopsis (bottom right), i.e. the tessellation tree. The initialization tile is marked red. Tiles that are fully outside or inside the Mandelbrot set are omitted (type 1 or 2, refer to section II B 2). Tiles with center outside the Mandelbrot set but covering parts of the Mandelbrot set are marked white (type 4). Tiles with center inside the Mandelbrot set but partly outside the Mandelbrot set are marked green (type 2). Observe that the synopsis of the different depths covers the hole area of the red tile at depth = 0. Interior distance estimation, here, is based on case $\alpha = 1$ (refer to section II B 1).

Case $\alpha = 1$: This is a conservative, i.e. smaller, under estimate of the interior distance based on Koebe quarter theorem. In this case the factor 2 at line 229 in section A 1 should amount to 4.

Case $\alpha = 2$: This is a more progressive, i.e. larger, under estimate of the interior distance. This case corresponds with factor 2 at line 229 in section A 1. Y. Fisher and J. Hill[5] suppose this method based on their numerical findings. In addition, the actual implementation of the *MandelbrotDistance* function in Mathematica 10 uses this approach.

2. Quad Tree Tessellation

The lowest five depths of the applied tessellations are depicted in figure 3. Selected parts of the tessellation tree at larger depths are shown in figures 4 and 7.

Each tile of the tessellation tree can fully be specified by its center (x, y) and its side length s , since only quadratic tiles are allowed. The red initialization tile in figure 3 upper left has center $(-0.75, 1.25)$ and a side length of 2.5.

When depth is increased by one each tile is divided into four sub tiles ('quad'); or, in other words, the side length of each new tile is divided by two. This repeated segmentation is similar to repeated branching ('tree').

Together with the distance estimator (refer to IIB1) five different types of tiles can be distinguished:

Type 1: Full member tiles – The center of the tile is inside the Mandelbrot set and the interior distance of the center is larger than $s\sqrt{2}$, where s is the side length of the tile. These tiles are omitted in figures 3 and 4 because they need no further refinement in the following depths of the tessellation.

Type 2: Full outside tiles – The center of the tile is outside the Mandelbrot set and the exterior distance of the center is larger than $s\sqrt{2}$, where s is again the side length of the tile. These tiles are omitted in figures 3 and 4, too. These tiles also need no further refinement in the following depths of the tessellation.

Type 3: Partly member tiles – The center of the tile is inside the Mandelbrot set and the interior distance of the center is equal or smaller than $s\sqrt{2}$, where s is again the side length of the tile. These tiles are marked green in figures 3 and 4. These tiles will be segmented in the next depth of the tessellation.

Type 4: Partly outside tiles – The center of the tile is outside the Mandelbrot set and the exterior distance of the center is equal or smaller than $s\sqrt{2}$, where s is again the side length of the tile. These tiles are marked white in figures 3 and 4. These tiles will also be segmented in the next depth of the tessellation.

Type 5: Chaotic tiles – With regard to specific parameters of the distance calculation, i.e. the maximum iteration, it can not be decided whether the center of the tile is outside or inside the Mandelbrot set. Thus, the distance of the center can not be estimated. These tiles are marked red in figure 4 and will be segmented in the next depth of the tessellation. Observe that the initialization tile in figure 3 is marked red, too.

Thus, for a given tile, i.e. the center (x, y) and the side length s the *decide* procedure returns the respective type (1 – 5) of the tile. That's why the procedure is

called 'decide', here. The procedure is part of the quad tree algorithm, which is described in section IIB4. The OpenCL implementation of the procedure is outlined in section A1.

Another important procedure of the quad tree algorithm is the *expand* procedure. This procedure runs on the CPU and is implemented as a compiled function within Mathematica (refer to A2). Essentially, the *expand* procedure calculates the new tile parameters, i.e. the new centers (x, y) and the halved side lengths s of the new tiles in the next depth of the tessellation.

Basically both, the *decide* and the *expand* procedures, are applied alternating to the tiles. Refer to figure 5 for a flow chart of the quad tree algorithm.

The complete tessellation tree of the red initialization tile in figure 3 upper left consists of quadratic tiles of different sizes, i.e. tiles at different depths. With respect to figure 3 bottom right observe on the one hand that the tiles cover the hole initialization tile – but, on the other hand, the total area of the tiles is larger than the area of the initialization tile, because the sub tiles do partly overlap. Nevertheless, it is possible to calculate the exact parts of the area of the different tile types. The calculation of the different areas is outlined in section IIB5.

To sum this up the following nomenclature is used in the following sections:

Tile: Tiles are always quadratic and can be fully specified by their center (x, y) and by their side length s . There are five different types of tiles (see above).

Tessellation: Here, a tessellation is a set of tiles of the same size. Observe that a given tessellation typically does not gapless cover the area. Tessellations can be enumerated by an index called depth, here.

Tessellation tree: A tessellation tree is a set of successive tessellations – starting with the tessellation of depth = 0, that contains only one (initialization) tile – up to tessellation of depth = n . Observe that a given tessellation tree covers the complete area of the initialization tile, but the area of all tiles in the tessellation tree does not sum up to the area of the initialization tile, because the tiles from different tessellations do partly overlap.

3. Performance Considerations

With increasing depth the number of tiles is growing fast, i.e. factor 4. Actually, because the tiles of type 1 and 2 are removed from the quad tree the number of tiles grows just by a factor of 3.5 approximately. This rapidly growing number leads to a rapidly growing CPU time and a rapidly growing file size of the quad tree, which both have to be reasonably addressed.

The *decide* and the *expand* procedures can be applied in parallel to a large number of tiles. Thus, a OpenCL

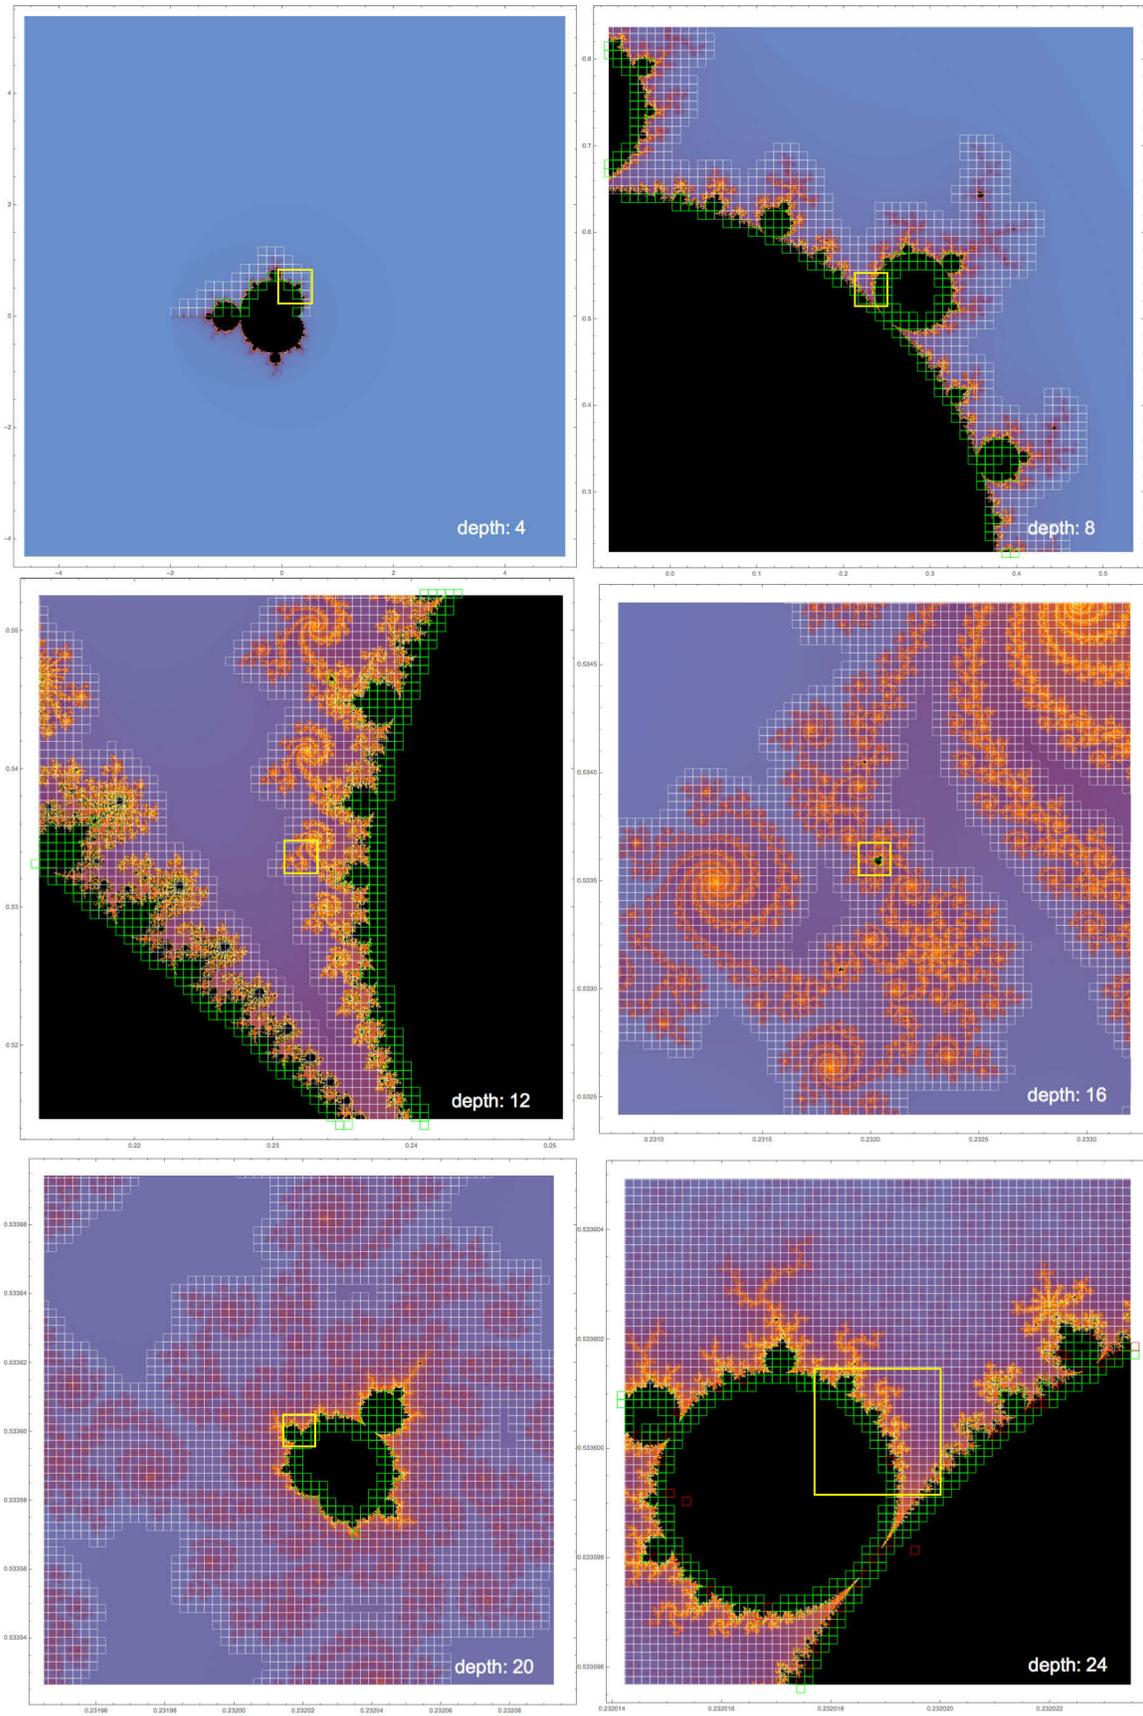


FIG. 4. Selected parts of the tessellation tree at different depths with $\alpha = 1$. Zoomed regions are marked yellow in lower depths. For details refer to caption of figure 3 and section IIB2. The yellow square at depth 24 corresponds to figure 7.

implementation is handy. Here, the *decide* procedure is implemented in OpenCL (refer to A 1) and the *expand* procedure is implemented as a compiled Mathematica function with compilation target C and allowed parallelization (refer to A 2).

In order to handle a large number of tiles at a given depth of tessellation, the list of tiles is partitioned into lists of less than 100 000 tiles. Each list is stored in a separate file. For each tile the center (x, y) and the type is stored. The side length s need not to be stored because it is constant at a given depth. A typical file in text format has a file size of about 7 MB. Compression reduces the file size to about 0.7 MB. The last tessellation that is stored on the hard disc has depth 23. It consists of about 2.3 million files and thus has a total file size of about 1.6 TB of already compressed data. Compression consumes about 30% of CPU time. For more details about timings refer to III C.

Since the tessellation of depth 23 is the last one saved on the hard disc, the tiles of the tessellation at depth 24 are counted but their coordinates are not saved to disc. In order to count the tiles of the tessellations 25 and 26 the tiles of tessellation 23 have to be expanded twice and thrice before counting (refer to the dashed arrow in figure 5 and following section II B 4).

In order to reduce CPU time the maximum iteration of tessellations 25 and 26 is reduced from 10 241 024 to 21 024 (refer to section II B 1). Since the maximum iteration mainly affects the lower bound of the area estimation the upper bound can still be lowered even with this reduced maximum iteration. Specialities of depths 25 and 26 are discussed in section III D.

4. Outline of the Quad Tree Algorithm

A flow chart of the quad tree algorithm is depicted in figure 5. The algorithm starts with the initialization tessellation at depth = 0 consisting of one tile with center $(-0.75, 1.25)$ and a side length of 2.5 (refer to figure 3 upper left). The type of the initialization tile is set to 5. Thus, it is marked red in figure 3. Refer to section A 3 for the definition of the initialization tile in the program code.

In the next step the *separate* procedure is called. Tiles of type 1 and 2 of the current tessellation are removed and counted. The results are stored in separate 'result' files. The coordinates of the other tiles are stored in 'quad tree' files. If the number of tiles is larger than 100 000 the tiles stored in two 'quad tree' files. Thus, the file size of the 'quad tree' files is limited to about 0.7 MB. The 'quad tree' files are compressed using Mathematica's build-in *Compress* function. Refer to section A 2 for the program code. The separation and the export/import is combined in a function called 'saveDecide'.

To sum this up, there are two important file types:

Quad tree files: These files contain the data needed for representing the tessellation tree. To limit their size

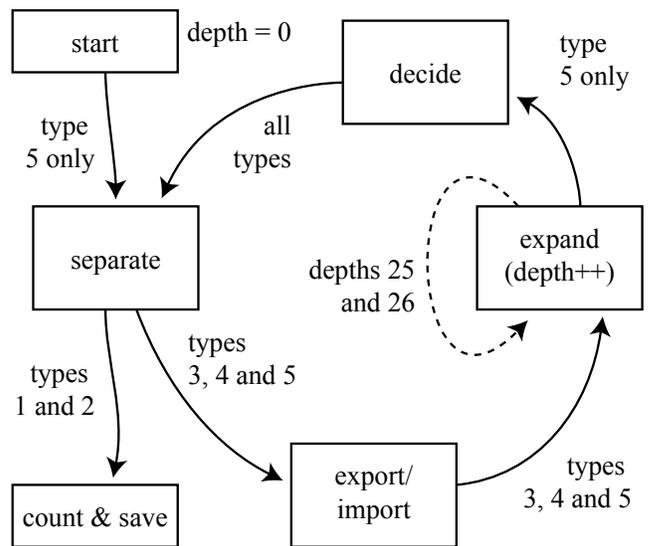


FIG. 5. The algorithm of the quad tree tessellation. The quad tree algorithm is outlined in section II B 4.

to about 0.7 MB the maximum number of tiles in on file is limited to 100 000. The files are compressed using Mathematica's *Compress* function. Only tiles of type 3, 4 and 5 are stored in these files. So, in order to store the hole tessellation at depth = 23 for example, 2.5 million files and thus a total file size of about 1.6 TB are needed.

Result files: For each 'quad tree' file a corresponding 'result' file is written to disc. These files contain only the frequencies of the different tile types. Thus, they have very small file size and do not need further compression. The precise numbers of different tile types at different depths of the tessellation tree can be derived from these files (refer to section II B 5).

When all 'quad tree' and 'result' files of the actual tessellation are written to hard disc, the 'quad tree' files are successively imported again. Then the *expand* procedure is applied to each tile of the actual tessellation. Each tile is divided into four sub tiles with new centers (x, y) and halved side lengths. The type of the new tiles is set to 5. The parameter depth is increased by one. Refer to section A 2 for the program code.

An important feature of the *expand* procedure is that it can be called repeatedly, as indicated in by the dashed arrow in figure 5. This feature is used to reach the tessellations of depth 25 and 26 without need for saving tessellations higher than 23 on hard disc (refer to section II B 3).

In the next step the types of the new tiles have to be determined. The *decide* procedure is applied in parallel to each tile. Thus, the *decide* procedure is implemented in OpenCL (refer to A 1). Finally all tiles are associated to the five types and they are again transferred to the

separate procedure. A new cycle of the quad tree algorithm starts.

The quad tree algorithm ends at a given depth with one or many 'result' files for each depth containing the counts of the different tile types. These files need further handling which is detailed in section II B 5.

5. Post Processing of the Tessellation Tree

'Result' files contain the frequencies of the different tile types (refer to section II B 4). There is at least one 'result' file for each depth of the tessellation tree. If there are more than one 'result' file the counts in these files have to be totalized, resulting in one 'total result' file for each depth of the tessellation tree.

Since every tessellation consists of tiles only of type 3, 4 and 5 of the previous tessellation, the counts for the tiles of type 1 and 2 of the actual tessellation have to be reconstructed from previous tessellations. For example, the single tile of type 2 at depth 1 adds $(2^{26-1})^2 = 33\,554\,432^2 = 1\,125\,899\,906\,842\,624$ tiles of type 2 to the tessellation at depth 26.

The reconstructed counts for the different tile types at different depths are listed in table I. Observe, that the tile counts are given for both cases $\alpha = 1$ (conservative) and $\alpha = 2$ (progressive) of the interior distance estimator (refer to section II B 1).

III. RESULTS

A. Tile counts at different depths

The evaluated tile counts for depths $d = 0..26$ are listed in table I. For each depth the tile counts of both cases $\alpha = 1$ (conservative) and $\alpha = 2$ (progressive) of the interior distance estimator are given. Refer to section II B 1 for details of the applied interior distance estimator.

The individual counts are calculated as described in section II B 5. Precise upper and lower bounds of the area of the Mandelbrot set can be derived from the tile counts with $\alpha = 1$ (conservative) (refer to section III B and table II).

Observe, that some tile counts at depths 25 and 26 change discontinuously. The reason is that the parameters of the distance estimation (refer to section II B 1) are changed, due to performance considerations (refer to section II B 3). Specialities of depths 25 and 26 are discussed in section III D.

B. Area Estimate

Based on the tile counts in table I precise lower and upper bounds of the area of the Mandelbrot set can be calculated. The lower bound corresponds to the tile count of type 2 tiles. The upper bound corresponds to the total

tile counts of type 2 and hybrid tiles. For the definition of hybrid tiles refer to table I and section II B 2. The calculation of the precise lower and upper bounds is based on tile counts resulting from the conservative interior distance estimator ($\alpha = 1$).

Calculated lower and upper bounds of the area at different tessellation depths are listed in table II and depicted in figure 6.

The current precise lower and upper bounds can be read of table II: the area of the Mandelbrot set is between 1.506 40 (refer to row 26'(1)) and 1.531 21 (refer to row 26'(1)). Corresponding bounds from Y. Fisher and J. Hill amount to 1.502 97 and 1.570 13 [5]. The difference of 0.063 16 between the bounds from Y. Fisher and J. Hill can approximately be more than halved by this approach.

The bounds stated by Y. Fisher and J. Hill are quite close to the calculated bounds of depth 19, here. Curiously Y. Fisher and J. Hill call this level 17. The small differences of the bounds may result from different parameters of the exterior/interior distance estimation (refer to section II B 1) and from the parameters of the initialization tile (refer to section II B 2).

C. Time Consumption

In table II the CPU (and GPU) times for calculations with $\alpha = 1$ are listed for each depth, i.e. the listed timings are not accumulated. The values at depths 24, 25 and 26 are estimated because of discontinued computations. With $\alpha = 2$ the times are approximately 2% smaller.

In the range from 18 to 23 the CPU time is increasing by factor 3 with increasing depth. Thus, at depth 24 a CPU time of about 400 h is expected. Since no quad tree files (refer to section II B 4) are written to disc at depth 24, the resulting CPU time is about 30% smaller (refer to section II B 3).

Based on extrapolation the timings for depth 25 and 26 add up to 1200 h (50 days) and 3600 h (150 days). In order to keep the CPU time below one month depths 25 and 26 are calculated with altered parameters of the interior distance estimator (refer to section III D) resulting in about 6× smaller timings.

D. Parameters of depths 25 and 26

In order to significantly reduce CPU time (refer sections II B 3 and III C) the maximum iteration is considerably reduced (refer to section II B 1) at depths 25 and 26. To indicate this, these depths are typically primed, i.e. 25' and 26', in tables and figures.

Figure 7 contrasts the resulting tessellations based on the standard maximum iteration (left column) with the corresponding tessellations based on the reduced maximum iteration (right column) at depth 26 exemplarily. The yellow marked zoom region in figure 4 corresponds

TABLE I. Tile counts at different depths for both cases $\alpha = 1$ (conservative) and $\alpha = 2$ (progressive) of the interior distance estimator (refer to section IIB 1). For more details about the different tile types refer to enumeration in section IIB 2. Colors are used in figures 3 and 4. Depths 25 and 26 are indicated by primes because they are calculated with different parameters. Thus, some of the tile counts change discontinuously. For more details refer to section III A.

depth	side length $s = 2^d$	total tiles $s^2 = T_1 + T_2 + T$	full outside tiles type 2 T_2	full inside tiles type 1 T_1	hybrid tiles $T = T_3 + T_4 + T_5$	partly outside type 3, white T_3	p.inside t.4, green T_4	chaotic tiles type 5, red T_5
0(1)	1	1	0	0	1	0	0	1
0(2)	1	1	0	0	1	0	0	1
1(1)	2	4	1	0	3	2	1	0
1(2)	2	4	1	0	3	2	1	0
2(1)	4	16	6	0	10	8	2	0
2(2)	4	16	6	0	10	8	2	0
3(1)	8	64	39	2	23	17	6	0
3(2)	8	64	39	2	23	17	6	0
4(1)	16	256	177	18	61	46	15	0
4(2)	16	256	177	18	61	46	15	0
5(1)	32	1024	771	92	161	129	32	0
5(2)	32	1024	771	92	161	129	32	0
6(1)	64	4096	3238	407	451	359	92	0
6(2)	64	4096	3238	413	445	359	86	0
7(1)	128	16384	13370	1750	1264	1022	241	1
7(2)	128	16384	13370	1780	1234	1022	211	1
8(1)	256	65536	54564	7316	3656	3083	573	0
8(2)	256	65536	54564	7443	3529	3083	446	0
9(1)	512	262144	221242	30137	10765	9304	1455	6
9(2)	512	262144	221242	30504	10398	9301	1091	6
10(1)	1024	1048576	893222	122698	32656	28982	3656	18
10(2)	1024	1048576	893222	123750	31604	28974	2612	18
11(1)	2048	4194304	3597092	496280	100932	91681	9188	63
11(2)	2048	4194304	3597092	499164	98048	91640	6345	63
12(1)	4096	16777216	14460134	1999241	317841	294959	22593	289
12(2)	4096	16777216	14460134	2006829	310253	294792	15177	284
13(1)	8192	67108864	58059613	8032011	1017240	960764	55339	1137
13(2)	8192	67108864	58059613	8051978	997273	960037	36108	1128
14(1)	16384	268435456	232915056	32214438	3305962	3166121	135418	4423
14(2)	16384	268435456	232915056	32266676	3253724	3163158	86202	4364
15(1)	32768	1073741824	933790101	129073027	10878696	10535474	325653	17569
15(2)	32768	1073741824	933789804	129209229	10742791	10523680	201755	17356
16(1)	65536	4294967296	3741951161	516822508	36193627	35351140	771867	70620
16(2)	65536	4294967296	3741947638	517181676	35837982	35305699	462813	69470
17(1)	131072	17179869184	14989686577	2068585656	121596951	119522811	1792260	281880
17(2)	131072	17179869184	14989661153	2069534138	120673893	119351062	1055891	266940
18(1)	262144	68719476736	60029958552	8277406710	412111474	406879291	4121677	1110506
18(2)	262144	68719476736	60029809562	8280019251	409647923	406235644	2390547	1021732
19(1)	524288	274877906944	240353790780	33116792335	1407323829	1393553259	9414774	4355796
19(2)	524288	274877906944	240353007132	33124439264	1400460548	1391156484	5368652	3935412
20(1)	1048576	1099511627776	962189238785	132483851445	4838537546	4800134425	21256443	17146678
20(2)	1048576	1099511627776	962185387560	132507822680	4818417536	4791239255	11885269	15293012
21(1)	2097152	4398046511104	3851336461567	529973966852	16736082685	16621019575	47327542	67735568
21(2)	2097152	4398046511104	3851318353176	530054346629	16673811299	16588078940	25880154	59852205
22(1)	4194304	17592186044416	15413996505530	2119984208854	58205330032	57833365121	103565735	268399176
22(2)	4194304	17592186044416	15413913972934	2120269727493	58002343989	57711539338	55236852	235567799
23(1)	8388608	70368744177664	61685165726455	8480136960072	203441491137	202153609362	221774194	1066107581
23(2)	8388608	70368744177664	61684798155308	8481196327648	202749694708	201703361795	115099822	931233091
24(1)	16777216	281474976710656	246839623842876	33920994359344	714358508436	709653138630	462828303	4242541503
24(2)	16777216	281474982710656	246838016125231	33925044558813	711922026612	707994674519	233811346	3693540747
25'(1)	33554432	1125899906842624	988090988392369	135684048606756	2522499667947	2174847868839	23704025	347628095083
25'(2)	33554432	1125899906842624	987688207138272	135699211786362	2512487917990	2171104588977	23459772	341359869241
26'(1)	67108864	4503599627370496	3956445430108888	542736526509859	8938453210521	7547894457325	47400846	1390511352350
26'(2)	67108864	4503599627370496	3951902496891266	542796894067968	8900236411262	7534750938783	46894494	1365438577985

to the upper row of figure 7. This row shows that the reduced maximum iteration leads to red tiles inside the Mandelbrot set. The middle row of figure 7 is a zoom view from the 'spike' region in the far left of the Mandelbrot set. This middle row shows that the reduced maximum iteration does not affect white tiles outside the Mandelbrot set. The lowest row of figure 7 is a zoom view from the boundary of the main cardioid of the Mandelbrot set. This lowest row shows the effect of the implemented cardioid testing (refer to section IIB 1).

Thus, the reduction of the maximum iteration mainly affects the lower bound of the area estimation, because the new red tiles appear mainly inside the Mandelbrot set. Since at higher depths the main reduction of the area bounds is gained from the upper bound, the reduction of the maximum iteration still helps to reduce the area bounds.

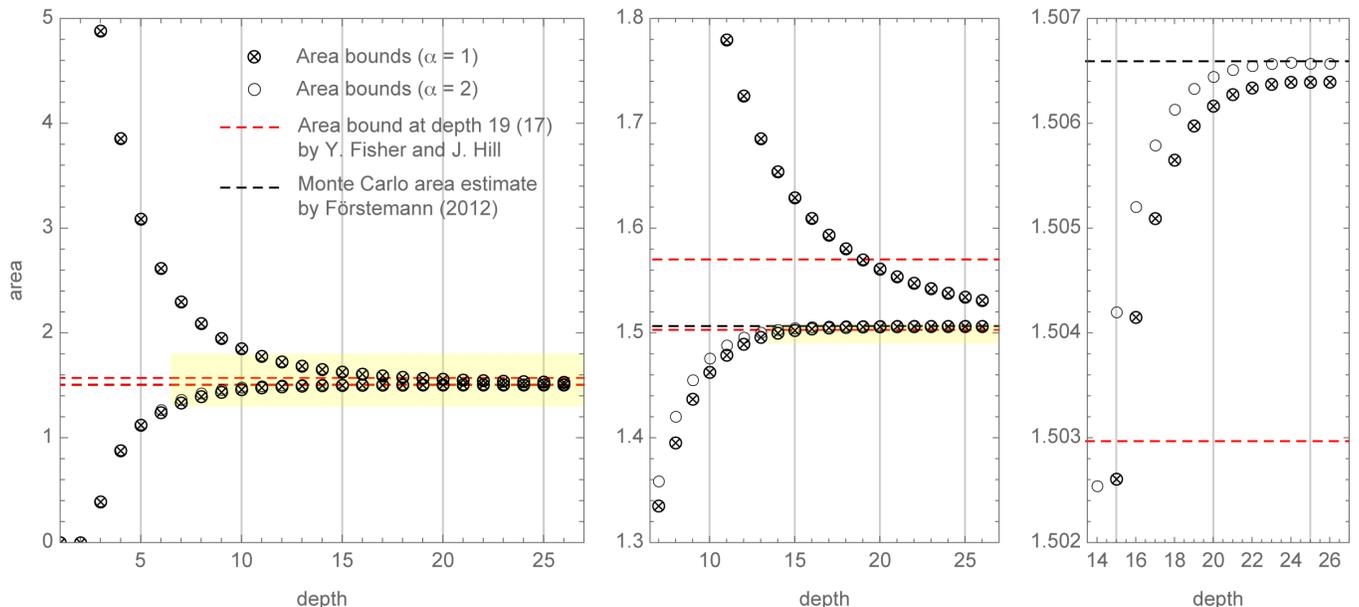


FIG. 6. The lower and upper bounds of the Mandelbrot area as a function of depth. Bounds are given for $\alpha = 1$ and $\alpha = 2$. The final area bound by Y. Fisher and J. Hill is marked by a dashed red line (refer to III B). The Monte Carlo area estimate by Förstemann[4] is marked by a dashed black line. The middle and the right figures are blow-ups – zoomed regions are marked yellow. This figure is inspired by the similar figure 2 in Y. Fisher and J. Hill [5].

TABLE II. Precise lower and upper bounds of the area of the Mandelbrot set at different tessellation depths d with $\alpha = 1$. In addition the difference between the upper and the lower bound and the CPU/GPU time needed are listed.

$d(\alpha)$	lower bound	upper bound	difference	CPU time
0(1)	0.000 00	12.500 00	12.500 00	< 0.01 h
1(1)	0.000 00	9.375 00	9.375 00	< 0.01 h
2(1)	0.000 00	7.812 50	7.812 50	< 0.01 h
3(1)	0.390 63	4.882 81	4.492 19	< 0.01 h
4(1)	0.878 91	3.857 42	2.978 52	< 0.01 h
5(1)	1.123 05	3.088 38	1.965 33	< 0.01 h
6(1)	1.242 07	2.618 41	1.376 34	< 0.01 h
7(1)	1.335 14	2.299 50	0.964 35	< 0.01 h
8(1)	1.395 42	2.092 74	0.697 32	< 0.01 h
9(1)	1.437 04	1.950 36	0.513 31	< 0.01 h
10(1)	1.462 67	1.851 96	0.389 29	< 0.01 h
11(1)	1.479 03	1.779 83	0.300 80	0.01 h
12(1)	1.489 55	1.726 36	0.236 81	0.02 h
13(1)	1.496 08	1.685 55	0.189 48	0.03 h
14(1)	1.500 10	1.654 05	0.153 95	0.04 h
15(1)	1.502 61	1.629 25	0.126 64	0.05 h
16(1)	1.504 15	1.609 49	0.105 34	0.09 h
17(1)	1.505 09	1.593 57	0.088 47	0.21 h
18(1)	1.505 65	1.580 61	0.074 96	0.55 h
19(1)	1.505 98	1.569 98	0.064 00	1.53 h
20(1)	1.506 17	1.561 17	0.055 01	4.47 h
21(1)	1.506 28	1.553 84	0.047 57	13.38 h
22(1)	1.506 34	1.547 70	0.041 36	41.83 h
23(1)	1.506 37	1.542 51	0.036 14	135.65 h
24(1)	1.506 39	1.538 12	0.031 72	~ 275 h
25'(1)	1.506 40	1.534 40	0.028 00	~ 170 h
26'(1)	1.506 40	1.531 21	0.024 81	~ 550 h

E. Comparison between $\alpha = 1$ and $\alpha = 2$

The parameter α and the two different cases $\alpha = 1$ and $\alpha = 2$ are introduced in section II B 1. In figure 8 the resulting tessellations based case $\alpha = 1$ (left column) and case $\alpha = 2$ (right column) are exemplarily shown. The zoomed regions are the same as in figure 7 (refer to section III D). Both cases are calculated with standard maximum iteration (refer to section III D).

Figure 8 suggests that the progressive interior distance estimator with $\alpha = 2$ yields reasonable results. The CPU time benefit of the progressive estimator of about 2% (refer to section III C) is rather small in the calculations applied here.

IV. DISCUSSION

The area bounds of the Mandelbrot set are quite slowly converging, as already mentioned by Y. Fisher and J. Hill [5]. In this approach a $16\,000\times$ larger tessellation just leads to halved area bounds.

The tile counts listed in table I are predestined for further fractal dimension analysis [15] based on box counting dimension or Minkowski-Bouligand dimension [16]. In addition, since the interior distance estimator needs orbit detection, lower area bounds of hyperbolic components [17] can be derived from numerical data gathered here. These and further investigations shall be subject to future projects.

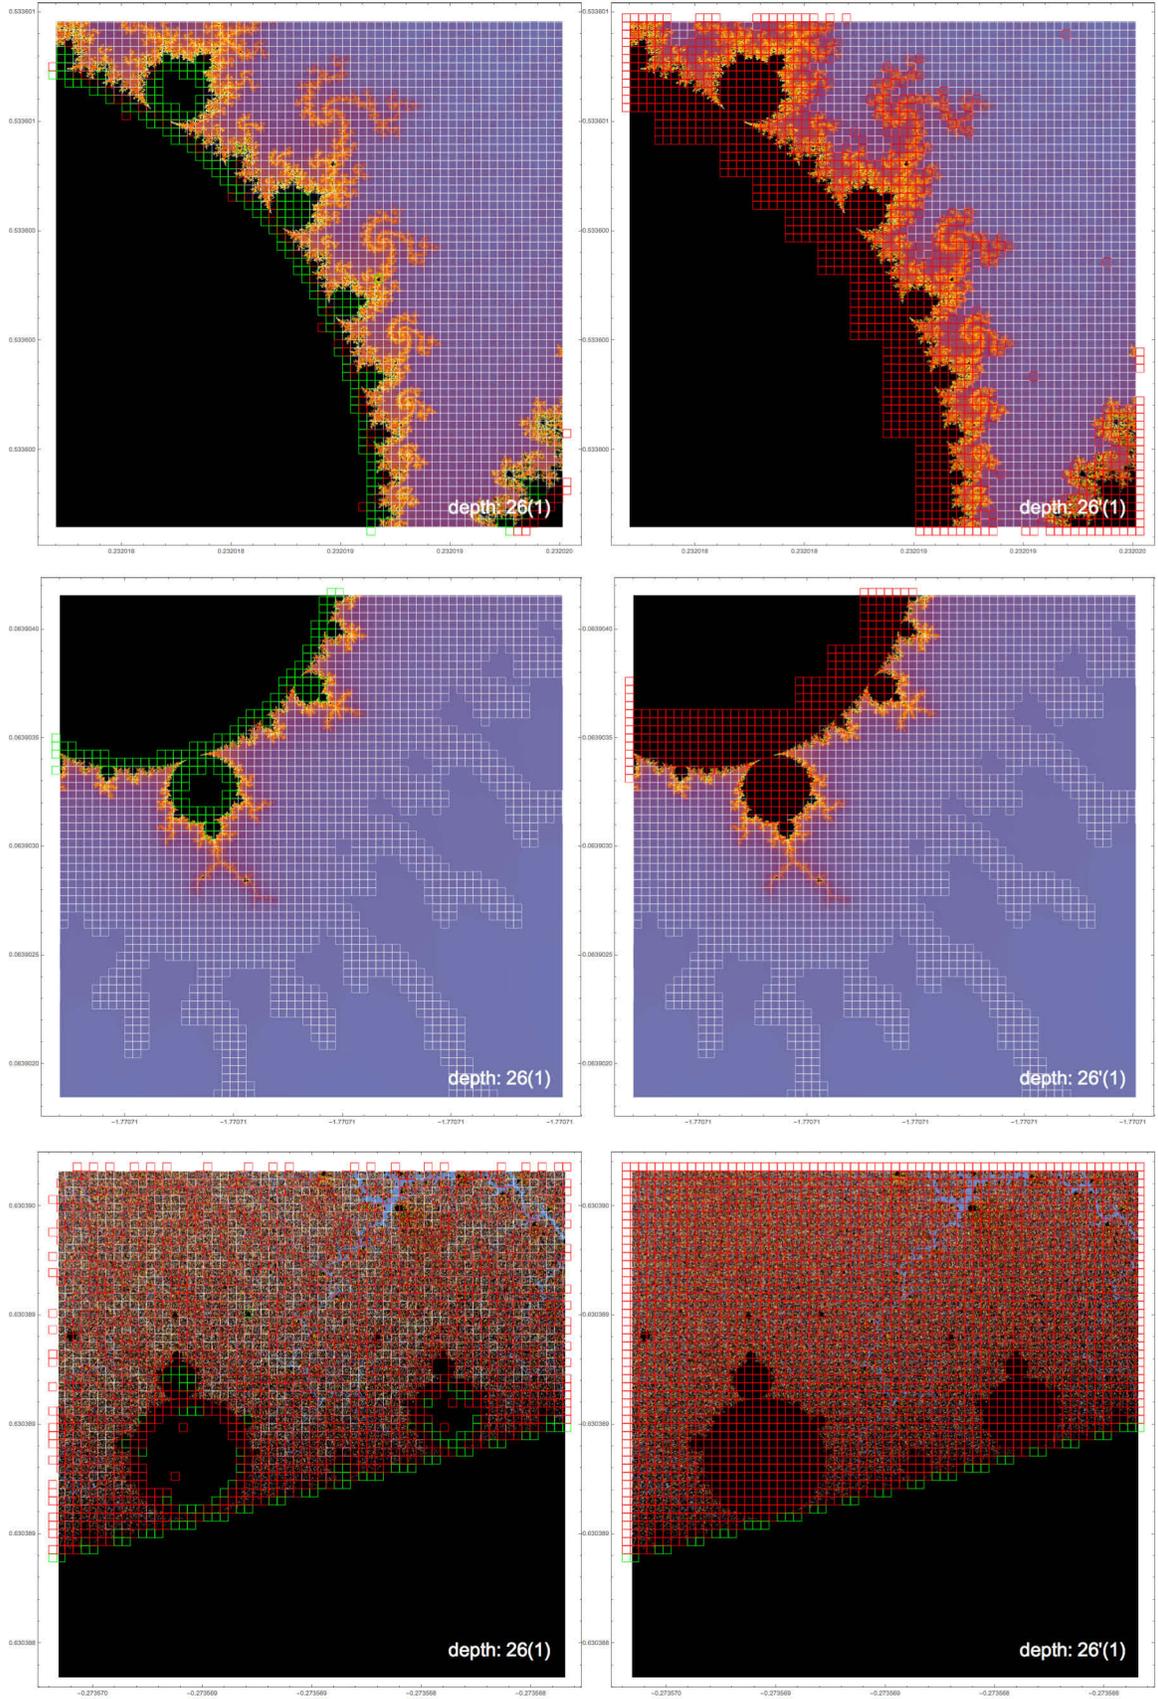


FIG. 7. Comparison of tessellations based on the standard maximum iteration (left column) with the corresponding tessellations based on the reduced maximum iteration (right column) for three different regions at depth 26 with $\alpha = 1$. For discussion refer to section III D.

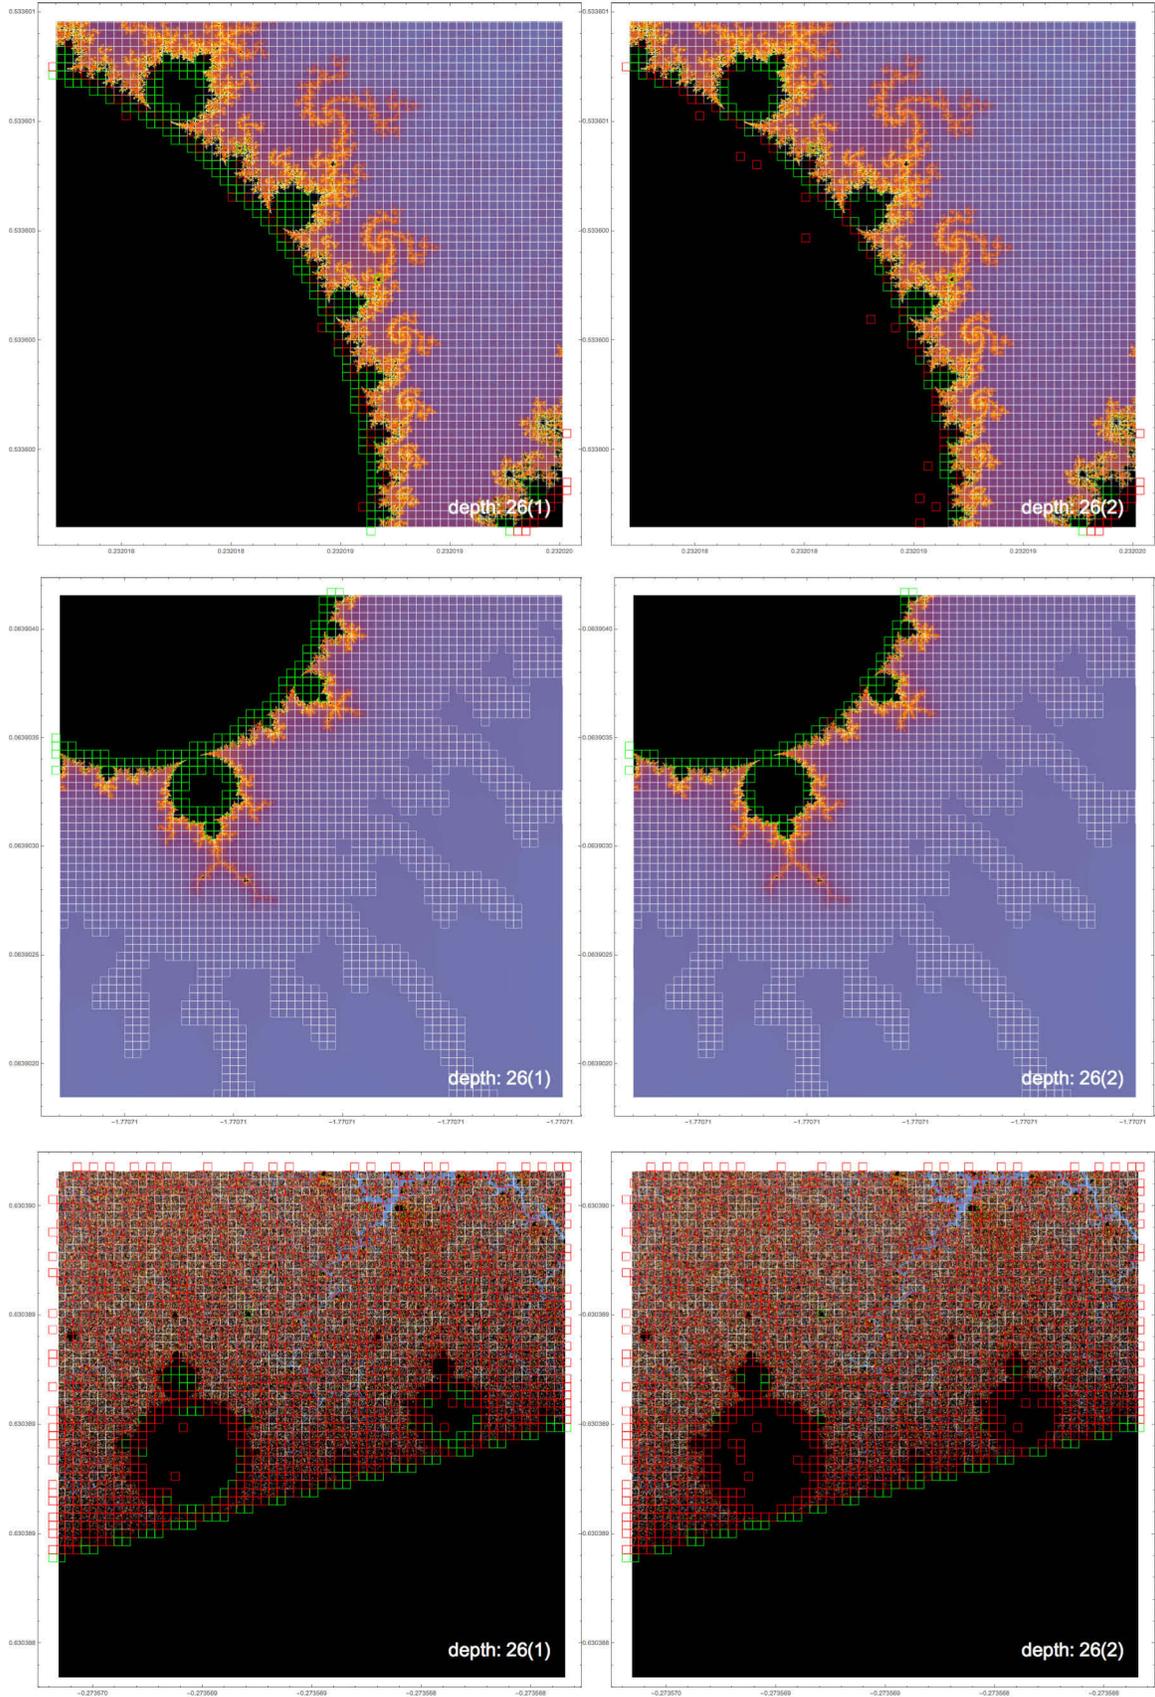


FIG. 8. Comparison of resulting tessellations based on the conservative interior distance estimator with $\alpha = 1$ (left column) and the corresponding tessellations based on the progressive interior distance estimator with $\alpha = 2$ (right column). For discussion refer to section III E.

Appendix A: Soucre Code

1. GPU Code for the *decide* Procedure

This is the GPU code of the *decide* procedure. The code is stored as a simple string. Essentially an exterior and interior distance estimator is implemented in OpenCL.

```

1 (* openCL source code *)
2 clSourceDecide="
3 #ifdef USING_DOUBLE_PRECISIONQ
4 #pragma OPENCL EXTENSION cl_khr_fp64 : enable
5 #pragma OPENCL EXTENSION cl_amd_fp64 : enable
6 #endif /* USING_DOUBLE_PRECISIONQ */
7 __kernel void rawDecide(
8   __global double * a, // real part of tile center
9   __global double * b, // real part of tile center
10  __global double * c, // result flag
11  __global double * d, // result cycle count
12  mint startIterations, // maximum start iterations, parameter
13  mint maxDepth, // maximum depth, parameter
14  double boxSize, // side length of box, parameter
15  mint n) {
16   int iGID = get_global_id(0);
17   if (iGID < n) {
18
19     // prolog: init variables
20     double cx = a[iGID]; // real part of tile center
21     double cy = b[iGID]; // real part of tile center
22     double zx = a[iGID];
23     double zy = b[iGID];
24     double escapeRadius2 = 10000.; // square of escape radius 100, empiric paramter
25     long maxCycleCount = 20000; // maximum length of detected cycles, empiric paramter
26     double q = 0.0; // interim value
27     bool refuge = false; // cx ,cy is refugee
28     bool member = false; // cx ,cy is member
29     bool card = false; // cx ,cy is member of cardioid
30     bool circ = false; // cx ,cy is member of circle
31     long iterationCount = 0; // iteration counter
32     long iterationDepthLimit = 0; // iteration limit post iteration
33     long cycleCount = 0;
34     long depth = 0;
35     double tx = 0.0; // interim value
36     double ty = 0.0; // interim value
37     bool card1 = false; bool card2 = false;
38     bool card3 = false; bool card4 = false;
39     bool circ1 = false; bool circ2 = false;
40     bool circ3 = false; bool circ4 = false;
41     double z2 = 0.; double absz = 0.;
42     double z2x = 0.; double z2y = 0.;
43     double dzx = 1.; double dzy = 0.;
44     double zCyclEx = 0.; double zCyclEy = 0.;
45     double D1r = 1.; double D1i = 0.;
46     double D2r = 0.; double D2i = 0.;
47     double D3r = 0.; double D3i = 0.;
48     double D4r = 0.; double D4i = 0.;
49     double D1rT = 0.; double D1iT = 0.;
50     double D2rT = 0.; double D2iT = 0.;
51     double D3rT = 0.; double D3iT = 0.;
52     double D4rT = 0.; double D4iT = 0.;
53     double valA = 0.; double valB = 0.;
54     double distance = 5.; // default value
55     long i = 0;
56     double result = 0.;
57
58     // cardioid and circle test
59
60     if (!refuge) {
61       q = pown((cx - 0.25), 2) + cy * cy;
62       card = q * (q + (cx - 0.25)) < 0.25 * cy * cy;

```

```

63     circ = pown((cx + 1.0), 2) + cy * cy < 0.0625;
64     member = (card || circ);
65 }
66
67 // pre iteration
68
69 for ( iterationCount = 0;
70     (!member) && (!refuge) && (iterationCount < startIterations);
71     iterationCount++) {
72     z2 = pown(zx, 2) - pown(zy, 2) + cx;
73     z2y = 2.0 * zx * zy + cy;
74     z2x = z2;
75     z2 = 2.0 * (zx * dzx - zy * dzy) + 1;
76     dzy = 2.0 * (zx * dzy + zy * dzx);
77     dzx = z2;
78     zx = z2x;
79     zy = z2y;
80     refuge = (pown(zx, 2) + pown(zy, 2)) > escapeRadius2;
81 }
82
83 // post iteration and cycle detection
84
85 for (depth = 1;
86     (!refuge) && (!member) && (depth < maxDepth + 1);
87     depth++) {
88     iterationDepthLimit = startIterations * pown(2., depth);
89     zCyclex = zx;
90     zCycley = zy;
91     for (iterationCount = iterationCount;
92         (!refuge) && (!member) && (iterationCount < iterationDepthLimit);
93         iterationCount++) {
94         z2 = pown(zx, 2) - pown(zy, 2) + cx;
95         z2y = 2.0 * zx * zy + cy;
96         z2x = z2;
97         z2 = 2.0 * (zx * dzx - zy * dzy) + 1;
98         dzy = 2.0 * (zx * dzy + zy * dzx);
99         dzx = z2;
100        zx = z2x;
101        zy = z2y;
102        refuge = (pown(zx, 2) + pown(zy, 2)) > escapeRadius2;
103        member = (zx == zCyclex) && (zy == zCycley);
104    }
105 }
106
107 depth = depth - 1;
108
109 // case 1: chaotic tile center
110 if (!(iterationCount < startIterations * pown( 2., depth) ) ) {
111     // distance calculation needless
112     distance = 0.0; } // dist = 0 => chaotic point
113
114 // case 2: tile center is refuge (exterior distance estimator)
115 if (refuge) {
116     // distance calculation
117     absz = sqrt(zx * zx + zy * zy);
118     distance = 0.25 * 2. * log(absz) * absz / sqrt(dzx * dzx + dzy * dzy);
119     cycleCount = iterationCount; // return iteration count instead of cycle count
120 }
121
122 // case 3: tile center is member
123 if (member) {
124
125     // case 3a: member of cardioid
126     if (card) {
127         tx = cx + boxSize * 0.5;
128         ty = cy + boxSize * 0.5;
129         q = pown((tx - 0.25), 2) + ty * ty;
130         card1 = q * (q + (tx - 0.25)) < 0.25 * ty * ty;
131         tx = cx - boxSize * 0.5;
132         ty = cy + boxSize * 0.5;

```

```

133 q = pown((tx - 0.25), 2) + ty * ty;
134 card2 = q * (q + (tx - 0.25)) < 0.25 * ty * ty;
135 tx = cx - boxSize * 0.5;
136 ty = cy - boxSize * 0.5;
137 q = pown((tx - 0.25), 2) + ty * ty;
138 card3 = q * (q + (tx - 0.25)) < 0.25 * ty * ty;
139 tx = cx + boxSize * 0.5;
140 ty = cy - boxSize * 0.5;
141 q = pown((tx - 0.25), 2) + ty * ty;
142 card4 = q * (q + (tx - 0.25)) < 0.25 * ty * ty;
143 if (card1 && card2 && card3 && card4) {
144     distance = -boxSize; // bigger than box
145     cycleCount = 1; // cycle count of cardioid
146 }
147 else {
148     distance = -boxSize * 0.5; // smaller than box
149     cycleCount = 1; // cycle count of cardioid
150 }
151 }
152
153 // case 3b: member of circle
154 if (circ) {
155     tx = cx + boxSize * 0.5;
156     ty = cy + boxSize * 0.5;
157     circ1 = pown((tx + 1.0), 2) + ty * ty < 0.0625;
158     tx = cx - boxSize * 0.5;
159     ty = cy + boxSize * 0.5;
160     circ2 = pown((tx + 1.0), 2) + ty * ty < 0.0625;
161     tx = cx - boxSize * 0.5;
162     ty = cy - boxSize * 0.5;
163     circ3 = pown((tx + 1.0), 2) + ty * ty < 0.0625;
164     tx = cx + boxSize * 0.5;
165     ty = cy - boxSize * 0.5;
166     circ4 = pown((tx + 1.0), 2) + ty * ty < 0.0625;
167     if (circ1 && circ2 && circ3 && circ4) {
168         distance = -boxSize; // bigger than box
169         cycleCount = 2; // cycle count of circle
170     }
171     else {
172         distance = -boxSize * 0.5; // smaller than box
173         cycleCount = 2; // cycle count of circle
174     }
175 }
176
177 // case 3c: not member of circle nor member of cardioid
178 if (!circ && !card) {
179     // determine cycle length
180     member = false;
181     for (iterationCount = iterationCount, cycleCount = 0;
182         (!member) && (cycleCount < maxCycleCount + 2);
183         iterationCount++, cycleCount++) {
184         z2 = zx * zx - zy * zy + cx;
185         z2y = 2.0 * zx * zy + cy;
186         z2x = z2;
187         z2 = 2.0 * (zx * dzx - zy * dzy) + 1;
188         dzy = 2.0 * (zx * dzy + zy * dzx);
189         dzx = z2;
190         zx = z2x;
191         zy = z2y;
192         refuge = (pown(zx, 2) + pown(zy, 2)) > escapeRadius2;
193         member = (zx == zCyclex) && (zy == zCycley);
194     }
195
196     // too long cycle (chaotic) case: cycle count greater than maxCycleCount
197     if (cycleCount > maxCycleCount) { distance = 0.; } // dist = 0 => chaotic point
198
199     // go on case 3c: distance calculation (interior distance estimator)
200     else {
201         zx = zCyclex;
202         zy = zCycley;

```

```

203     for ( i = 1;
204           i <= cycleCount;
205           i++) {
206         D1rT = 2. *(zx * D1r - zy * D1i);
207         D1iT = 2. *(zy * D1r + zx * D1i);
208         D2rT = 2. *(zx * D2r - zy * D2i) + 1.;
209         D2iT = 2. *(zy * D2r + zx * D2i);
210         D3rT = 2. *((zx * D3r - zy * D3i) + (D1r * D1r - D1i * D1i));
211         D3iT = 2. *((zx * D3i + zy * D3r) + (2. * D1r * D1i));
212         D4rT = 2. *((zx * D4r - zy * D4i) + (D1r * D2r - D1i * D2i));
213         D4iT = 2. *((zx * D4i + zy * D4r) + (D1r * D2i + D1i * D2r));
214         D1r = D1rT; D1i = D1iT;
215         D2r = D2rT; D2i = D2iT;
216         D3r = D3rT; D3i = D3iT;
217         D4r = D4rT; D4i = D4iT;
218         z2 = zx * zx - zy * zy + cx;
219         zy = 2.0 * zx * zy + cy;
220         zx = z2;
221     }
222     valA = 1. - (D1r*D1r + D1i*D1i);
223     valB = (1 - D1r) * (1 - D1r) + D1i * D1i;
224     D1rT = (D2r * (1 - D1r) - D2i * D1i) / valB;
225     D1iT = (D2i * (1 - D1r) + D2r * D1i) / valB;
226     D2rT = D4r + (D3r * D1rT - D3i * D1iT);
227     D2iT = D4i + (D3i * D1rT + D3r * D1iT);
228     valB = sqrt(D2rT * D2rT + D2iT * D2iT);
229     distance = -valA / (2. * valB);
230
231     } // inner dist calc
232
233     } // not card nor circ case
234
235     } // member case
236
237     // epilog: calculate result flag:
238
239     // case 1: point refugee && box refugee
240     if (distance > 0. && distance > sqrt(0.5) * boxSize) {result = 1.; }
241
242     // case 2: point refugee && box mixed
243     if (distance > 0. && distance <= sqrt(0.5) * boxSize) {result = 2.; }
244
245     // case 3: point member && box member
246     if (distance < 0. && -distance > sqrt(0.5) * boxSize) {result = 3.; }
247
248     // case 4: point member && box mixed
249     if (distance < 0. && -distance <= sqrt(0.5) * boxSize) {result = 4.; }
250
251     // case 5: chaotic point
252     if (!(distance > 0. && distance > sqrt(0.5) * boxSize) &&
253         !(distance > 0. && distance <= sqrt(0.5) * boxSize) &&
254         !(distance < 0. && -distance > sqrt(0.5) * boxSize) &&
255         !(distance < 0. && -distance <= sqrt(0.5) * boxSize) ) {
256         result = 5.; } // chaotic point
257
258     c[iGID] = result;
259     d[iGID] = (double) cycleCount;
260 }
261 }";

```

2. Definition of Mathematica Functions

The Mathematica functions are defined here, e.g. *decide* and *expand*. In addition, this Mathematica code initializes parameters.

```

1 (* apply opencl code *)
2 decide[{ list_ , startIter_ , maxDepth_ , boxSize_ , level_ , levelStep_ }]:= Block[{ tmp } ,
3   tmp=Join [ Transpose [ list ] , { startIter , maxDepth , boxSize , Length [ list ] } ] ;
4   tmp=Apply [ rawDecide , tmp ] ;
5   tmp=Transpose [ tmp ] ;
6   { tmp , startIter , maxDepth , boxSize , level , levelStep }
7 ] ;
8
9 (* quadtree generator *)
10 rawExpand=Compile [ { { list , _Real , 1 } , { boxSize , _Real } } ,
11   Module [ { delta = 0.25 * boxSize } ,
12     { { list [[ 1 ] - delta , list [[ 2 ] - delta , list [[ 3 ] , - 1 } ,
13       { list [[ 1 ] - delta , list [[ 2 ] + delta , list [[ 3 ] , - 1 } ,
14       { list [[ 1 ] + delta , list [[ 2 ] + delta , list [[ 3 ] , - 1 } ,
15       { list [[ 1 ] + delta , list [[ 2 ] - delta , list [[ 3 ] , - 1 } }
16   ] ,
17   CompilationTarget -> "C" ,
18   RuntimeAttributes -> { Listable } ,
19   Parallelization -> True ,
20   RuntimeOptions -> "Speed"
21 ] ;
22
23 (* apply quadtree generator *)
24 expand[{ list_ , startIter_ , maxDepth_ , boxSize_ , level_ , levelStep_ }]:= {
25   Flatten [ rawExpand [ list , boxSize ] , 1 ] ,
26   startIter ,
27   maxDepth ,
28   boxSize * 0.5 ,
29   level + 1 ,
30   levelStep
31 } ;
32
33 (* different compiled select functions for function saveDecide *)
34 rawSelector3=Compile [ { { list , _Real , 2 } } ,
35   Select [ list , #[ [ 3 ] ] == 3. & ] ,
36   CompilationTarget -> "C" , RuntimeOptions -> "Speed"
37 ] ;
38
39 rawSelector245=Compile [ { { list , _Real , 2 } } ,
40   Select [ list , #[ [ 3 ] ] == 2. || #[ [ 3 ] ] == 4. || #[ [ 3 ] ] == 5. & ] ,
41   CompilationTarget -> "C" , RuntimeOptions -> "Speed"
42 ] ;
43
44 rawSelector2=Compile [ { { list , _Real , 2 } } , Select [ list , #[ [ 3 ] ] == 2. & ] , CompilationTarget -> "C" , RuntimeOptions
45   -> "Speed" ] ;
46 rawSelector4=Compile [ { { list , _Real , 2 } } , Select [ list , #[ [ 3 ] ] == 4. & ] , CompilationTarget -> "C" , RuntimeOptions
47   -> "Speed" ] ;
48 rawSelector5=Compile [ { { list , _Real , 2 } } , Select [ list , #[ [ 3 ] ] == 5. & ] , CompilationTarget -> "C" , RuntimeOptions
49   -> "Speed" ] ;
50
51 rawSorter245 [ list_ ] := Join [ rawSelector5 [ list ] , rawSelector4 [ list ] , rawSelector2 [ list ] ] ;
52
53 (* create individual file name and directory *)
54 rawFilename [ list_ , level_ , status_ , workingDirectory_ ] := Module [ { hash , file } ,
55   hash = ToString [ Hash [ First [ list ] , "MD2" ] ] ;
56   file = FileNameJoin [ { workingDirectory , status <> "-" <> ToString [ level ] , hash <> ".mx" } ] ;
57   If [ ! ( DirectoryQ [ DirectoryName [ file ] ] ) , CreateDirectory [ DirectoryName [ file ] ] ] ;
58   file
59 ] ;
60
61 (* save individual todo and result "buzz" files *)
62 saveDecide [ { list_ , startIter_ , maxDepth_ , boxSize_ , level_ , levelStep_ } ] := Module [ { split , tally , hyper , data
63   , files , file , cpu , hdd } ,
64   (* timing *) split = AbsoluteTime [ ] ;
65   (* result "buzz" file handling *)

```

```

62 tally=list [[ All , 3]]; (* extract box type column from decided list *)
63 tally=Join[tally, {1., 2., 3., 4., 5.}]; (* housekeeping: get one with zero lengths *)
64 tally=Sort[Tally[tally]] [[ All , 2]] - 1; (* count member and refugee boxes *)
65 hyper=rawSelector3[list]; (* select members / hyperbolic components *)
66 hyper=Tally[hyper [[ All , 4]]]; (* extract and count cycle count column from decided list *)
67 data={startIter, maxDepth, boxSize, level, levelStep, tally, hyper, AbsoluteTime []}; (* collect data *)
68 file=rawFilename[list, level, "buzz", globalWD]; (* create unique filename *)
69 Export[file, data, "MX"]; (* save result "buzz" file *)
70 (* todo file handling *)
71 data=rawSelector245[list]; (* select !refugees, !members and unknowns *)
72 data=Partition[data, globalFileSize, globalFileSize, {1, 1}, {}]; (* divide large lists *)
73 data=Map[rawSorter245, data]; (* empiric: sort chaotic, member, refugee *)
74 files=Map[rawFilename[#, level+levelStep, "todo", globalWD]&, data]; (* create unique filename *)
75 data=Map[{#, startIter, maxDepth, boxSize, level, levelStep}&, data]; (* collect data *)
76 (* timing *) cpu=AbsoluteTime[] - split; split=AbsoluteTime[];
77 Map[Export[#[[1]], Compress[#[[2]], "MX"]&, Transpose[{files, data}]]; (* save compressed todo files
*)]
78 If[level > -1, file=rawFilename[list, level, "done", globalWD]]; (* housekeeping: init case *)
79 (* timing *) hdd=AbsoluteTime[] - split;
80 {cpu, hdd, Length[files]}
81 ];
82
83 (* get due todo file count *)
84 getTodoFiles[] := Block[{tmp},
85 tmp=FileNames["*", FileNameJoin[{globalWD, "todo" <> "-" <> "*" }], 1]; (* get all files *)
86 tmp=Map[{DirectoryName[#, #]&, tmp]; (* extract dirs *)
87 tmp=Split[Sort[tmp], #1[[1]] == #2[[1]]&][[1, All, 2]]; (* bin dirs and get first dir *)
88 tmp
89 ];
90
91 (* number formatter *)
92 floatForm[float_]:=StringTake[ToString[PaddedForm[Round[float, 0.001], {4, 3}, NumberPadding->{"", "0"}]], 5];
93 integerForm[integer_]:=StringTake[ToString[PaddedForm[Round[integer, 1], 4, NumberPadding->{"0", "0"}]], 5];
94
95 (* procedure for single todo file *)
96 processTodoFile[file_]:=Module[{split, todo, done, levelStep, timeImport, timeExpand, timeDecide,
timeTally, timeExport, numberFiles, logline},
97 (* timing *) split=AbsoluteTime[];
98 todo=Uncompress[Import[file]]; (* read todo file *)
99 (* timing *) timeImport=AbsoluteTime[] - split; split=AbsoluteTime[];
100 levelStep=todo[[6]]; (* extract level step from todo file *)
101 todo=Nest[expand, todo, levelStep]; (* expand todo file via nested quadtree *)
102 (* timing *) timeExpand=AbsoluteTime[] - split; split=AbsoluteTime[];
103 done=decide[todo]; (* TADA: decide todo file *)
104 (* timing *) timeDecide=AbsoluteTime[] - split;
105 {timeTally, timeExport, numberFiles}=saveDecide[done]; (* post process todo files *)
106 (* logging *) logline=DateString[{"Day", "-", "Hour", ":", "Minute", ":", "Second"}] <>
107 " l:" <> integerForm[done[[5]]] <>
108 " k:" <> integerForm[$KernelID] <>
109 " i:" <> floatForm[timeImport] <>
110 " m:" <> floatForm[timeExpand] <>
111 " g:" <> floatForm[timeDecide] <>
112 " c:" <> floatForm[timeTally] <>
113 " e:" <> floatForm[timeExport] <>
114 " n:" <> integerForm[numberFiles];
115 (* logging *) PutAppend[TextString[logline], FileNameJoin[{globalWD, "logfile.txt"}]];
116 RenameFile[file, StringReplace[file, "todo" -> "done"]]; (* move done todo file *)
117 ];
118
119 (* init main kernel *)
120 Needs["OpenCLLink"]; (* main kernel: load opencl library *)
121 $HistoryLength = 0; (* main kernel: reduce memory consumption *)
122
123 (* global parameters *)
124 globalWD = "D:\\Mandel3_Data_V1"; (* working directory *)
125 globalFileSize = 500000; (* max file size, approx. size * 2^levelstep < 500 000 *)
126 gobalTodoSetMax = 1000; (* max cycle count of kernels before restart *)
127

```

```

128 preIterations = 1024;
129 maxDepth = 9;
130 boxSize = 2.5;
131 levelStep = 1;
132
133 (* House keeping: start and init worker kernels *)
134
135 startKernels [] := Block[{},
136   LaunchKernels[];
137   DistributeDefinitions[rawDecide, cISourceDecide, globalWD];
138   ParallelEvaluate[Needs["OpenCLLink"]];
139   ParallelEvaluate[$HistoryLength = 0];
140   (* define opencl function from source *) ParallelEvaluate[
141     rawDecide = OpenCLFunctionLoad[
142       cISourceDecide,
143       "rawDecide",
144       {{_Real}, {_Real}, {_Real}, {_Real}, _Integer, _Integer, _Real, _Integer},
145       16,
146       "ShellOutputFunction" -> None,
147       "Platform" -> 1,
148       "Device" -> Mod[$KernelID, 4] + 1]
149   ];
150 ];

```

3. Initialisation of Working Files

This Mathematica code initializes the working files. This code is only called once.

```

1 (* save first todo file *)
2 todo = {{{-0.75, 1.25, 5., -1}}, preIterations, maxDepth, boxSize, 0, levelStep};
3 saveDecide[todo];

```

4. Mathematica Code for the Main Loop

This Mathematica code represents the main loop. This code can be restarted.

```

1 (* main loop *)
2 Do[
3   files = getTodoFiles[];
4   todoSets = Partition[files, gobalTodoSetMax, gobalTodoSetMax, {1, 1}, {}];
5   Print[DateString[], " ***** level ", i, " ***** Total files: ", Length[files]];
6   (* logging *)
7   logline =
8   DateString[{"Day", "=", "Hour", ":", "Minute", ":", "Second"}] <>
9   " ***** start new round, " <>
10  " counter:" <>
11  integerForm[i];
12  (* logging *)
13  PutAppend[TextString[logline],
14  FileNameJoin[{globalWD, "logfile.txt"}]];
15  Do[
16    startKernels[];
17    Print[DateString[], " * Kernel restart * todoset: ", todoSet, "/", Length[todoSets]];
18    (* logging *)
19    logline = DateString[{"Day", "=", "Hour", ":", "Minute", ":", "Second"}] <> " *****
20    restart kernels";
21    (* logging *)
22    PutAppend[TextString[logline], FileNameJoin[{globalWD, "logfile.txt"}]];
23    ParallelMap[processTodoFile[#] &, todoSets[[todoSet]], Method -> "FinestGrained"];
24    CloseKernels[];
25    Pause[5];
26    {todoSet, Length[todoSets]}

```

-
- [1] http://en.wikipedia.org/wiki/Mandelbrot_set
 - [2] <http://oeis.org/A098403>
 - [3] <http://www.mrob.com/pub/muency/pixelcounting.html>
 - [4] ...
 - [5] <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.53.2337>
 - [6] <http://en.wikipedia.org/wiki/OpenCL>
 - [7] <http://forums.wolfram.com/mathgroup/archive/2011/Jun/msg00615.html>
 - [8] https://en.wikipedia.org/wiki/Mandelbrot_set#Distance_estimates
 - [9] <http://mrob.com/pub/muency/distanceestimator.html>
 - [10] <http://www.moleculardensity.net/buddhabrot/appendix/2>
 - [11] https://en.wikipedia.org/wiki/Mandelbrot_set#Escape_time_algorithm
 - [12] https://en.wikipedia.org/wiki/Mandelbrot_set#Cardioid_2F_bulb_checking
 - [13] https://en.wikipedia.org/wiki/Koebe_quarter_theorem
 - [14] <http://www.mrob.com/pub/muency/pixelcounting.html>
 - [15] https://en.wikipedia.org/wiki/Fractal_dimension
 - [16] https://en.wikipedia.org/wiki/Minkowski?Bouligand_dimension
 - [17] https://en.wikipedia.org/wiki/Mandelbrot_set#Hyperbolic_components