Level-set Based Design of Wang Tiles for Modelling Complex Microstructures ☆,☆☆

Martin Doškář [a,c,*], Jan Zeman [a,b], Daniel Rypl [a], Jan Novák [c]

[a] *Department of Mechanics, Faculty of Civil Engineering, Czech Technical University in Prague, Thákurova 7, 166 29 Prague 6, Czech Republic*
[b] *The Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodárenskou věží 4, 182 08 Prague 8, Czech Republic*
[c] *Experimental centre, Faculty of Civil Engineering, Czech Technical University in Prague, Thákurova 7, 166 29 Prague 6, Czech Republic*

ARTICLE INFO

ABSTRACT

Microstructural geometry plays a critical role in the response of heterogeneous materials. Consequently, methods for generating microstructural samples are increasingly crucial to advanced numerical analyses. We extend Sonon et al.'s unified framework, developed originally for generating particulate and foam-like microstructural geometries of Periodic Unit Cells, to non-periodic microstructural representations based on the formalism of Wang tiles. This formalism has been recently proposed in order to generalize the Periodic Unit Cell approach, enabling a fast synthesis of arbitrarily large, stochastic microstructural samples from a handful of domains with predefined microstructural compatibility constraints. However, a robust procedure capable of designing complex, three-dimensional, foam-like and cellular morphologies of Wang tiles has not yet been proposed. This contribution fills the gap by significantly broadening the applicability of the tiling concept.

Since the original Sonon et al.'s framework builds on a random sequential addition of particles enhanced with an implicit representation of particle boundaries by the level-set field, we first devise an analysis based on a connectivity graph of a tile set, resolving the question where a particle should be copied when it intersects a tile boundary. Next, we introduce several modifications to the original algorithm that are necessary to ensure microstructural compatibility in the generalized periodicity setting of Wang tiles. Having established a universal procedure for generating tile morphologies we compare strictly aperiodic and stochastic sets with the same cardinality in terms of reducing the artificial periodicity in reconstructed microstructural samples. We demonstrate the superiority of the vertex-defined tile sets for two-dimensional problems and illustrate the capabilities of the algorithm with two- and three-dimensional examples.

© 2020 Elsevier Ltd. All rights reserved.

---

1. Introduction

The geometrical details of a material composition drive many macroscopic phenomena such as crack initiation and propagation [1] or meta-behaviour [2]. Advanced computational strategies thus tend to incorporate knowledge of material microstructures[1] in order to enhance their predictive power. [3]

When micro and macro scales are well separated, a microstructural response is typically up-scaled through homogenization methods. Albeit computationally intensive, numerical homogenization methods [3] now supersede analytical methods [4] because of their ability to handle complex microstructural geometries and non-linear problems. With increasing computational power, numerical models that fully resolve the microstructural geometry in the whole macroscopic domain, e.g. [5,6], or in

---

[1] We use the term "microstructure" in a broader sense, without referring to a specific scale length.

regions of interest [7,8], have emerged, addressing problems without clear scale separation.

For both strategies, however, accuracy critically depends on the representativeness of the provided microstructural geometry, accenting the crucial role of microstructure modelling in multi-scale approaches.

Compared to materials with regular microstructures, characterized entirely by Periodic Unit Cells (PUCs), modelling random heterogeneous materials is more intricate; any finite-size representation automatically implies information loss. The optimal microstructure representation should capture intrinsic randomness and fluctuations in a microstructure while remaining computationally tractable (for numerical homogenization) or inexpensive to construct (for fully resolved simulations).

1.1. State-of-the-art in modelling random microstructures

One of the widely adopted approaches for modelling heterogeneous materials rests on an extension of PUC generated such that its spatial statistics match that of a reference microstructure. This procedure appears in the literature under various names such as Statistically Optimal Representative Unit Cell [9], Repeating Unit Cell [10], Statistically Similar Representative Volume Element [11], or Statistically Equivalent Periodic Unit Cell (SEPUC) [12], among others. The spatial statistics involved range from Minkowski functionals [13] to multi-point probability functions [4], out of which the two-point probability [14–16], two-point cluster [17], and lineal path [12,18–20] functions are the most frequently used.

Following Povirk's seminal work [21], the majority of cell representations are generated using optimization procedures, minimizing the discrepancy between the statistical characterization of the reference microstructure and its compressed, PUC-like representation. The particular choice of optimization algorithm currently varies with several options including simulated annealing [12,19,22], genetic [9,10,23,24] and gradient [21,25], or phase-recovery [26] algorithms.

The second approach to microstructure generation utilizes reference samples of the microstructure. New realizations are then obtained with a Markovian process, taking individual voxels [27] or a patch of voxels [28] from the provided reference samples according to the proximity of the spatial statistics computed for their surroundings. Alternatively, searching for statistics proximity can be replaced with a classification tree-based supervised learning model [29].

The previous two approaches suffer from high computational costs related either to optimization or to training the learning models. The applicability of their outputs is also sensitive to the spatial statistics considered, attesting to the ill-conditioning of the microstructure reconstruction problem itself. Achieving a good match in selected statistics does not automatically guarantee similar overall behaviour; for instance, Biswal et al. [30] demonstrated that realizations with similar two-point probability functions could have significantly different percolation characteristics that govern overall transport properties.

Complementary to the statistics-informed methods, a third approach to generating microstructural realizations relies on meta-modelling the genesis of a microstructure. These methods range in complexity and include the Monte-Carlo Potts [31] and phase field models [32] of grain growth; sedimentation-and-compaction models [30]; and various particle packing algorithms, e.g. [33,34, and references therein], based on either Random Sequential Adsorption (RSA)[2] [36,37] or molecular dynamics [38–40].

The relevance of packing algorithms extends beyond simple particle-matrix microstructures because the resulting packings often serve as initial seeds for tessellation-based models applicable to polycrystals [41,42], foams [34], and cell tissues [43,44]. Due to its straightforward implementation, Voronoï tessellation is the most common choice; however, the resulting geometry is oversimplified for many materials. For instance, Voronoï-based models overestimate overall stiffness for high porosity foams [45]. The curvature of cell walls [41,46] and heterogeneity in cell size [34,42] and wall thickness [42] must be additionally introduced to obtain realistic geometries. Similar effects can be achieved by modifying the distance measure used during tessellation, e.g. models based on the Laguerre variant generate microstructures with multi-mode cell size distribution [33,34, 42,47]. Inspired by Laguerre tessellation, Chakraborty et al. [44] proposed Adaptive Quadratic Voronoï Tessellation, attributing a distinct anisotropic metric to each seed and thus allowing for additional control over the resulting geometry.

The original RSA method [35] suffers from $\mathcal{O}\left(N^2\right)$ complexity for N particles due to overlap checks and is impractical for generating large, densely packed systems. Consequently, several accelerations have been proposed. For Dart Throwing Algorithm [35], which is a simplified case of RSA with equisized circular/spherical particles, Dunbar and Humphreys [48] introduced a scalloped sector representation of non-overlap guaranteed regions. In the same year, Jones [49] proposed an alternative bookkeeping of the regions based on an adaptively updated Voronoï tessellation. Moreover, both approaches utilize a tree data structure and improve the algorithm complexity to $\mathcal{O}(N \log N)$. For general RSA, Yang et al. [10] proposed an acceleration based on a combination of a spline description of particle shapes and hierarchically refined bounding boxes of each particle. Recently, Sonon et al. [50] introduced a method building on an implicit, level-set based description of particle shapes, achieving $\mathcal{O}(N)$ complexity (see Section 3.3 in [50]). Moreover, Sonon et al.'s method readily facilitates generating complex microstructures using linear combinations of the nearest neighbour distance functions and dedicated morphing operations [50–53]. In a sense, this approach introduces the anisotropic pseudo-metrics[3] of [44] in a geometrically-motivated way by considering arbitrarily-shaped particles. As a result, Sonon et al.'s method enables refined control over generated microstructure unattainable with standard Voronoï or Laguerre tessellations; see Section 3.2.

Albeit significantly faster than RSA or optimization-based approaches, the latter method still starts anew every time an additional realization is required, imposing overhead on, e.g., investigations of the Representative Volume Element (RVE) size that require multiple microstructural samples to be generated, see [54–56]. Alternatively, larger microstructural realizations can be assembled from (SE)PUC; however, such construction introduces non-physical, long-range, periodic artefacts in a microstructural geometry and its local response.

1.2. Wang tiling in microstructure modelling

Inspired by applications in computer graphics [57], we have introduced the formalism of Wang tiles as a suitable generalization of SEPUC representation of heterogeneous materials [58]. The formalism provides a compromise between the SEPUC approach and the use of a microstructure generator for each new realization. The concept of Wang tiles decomposes microstructure generation into an offline phase, which can possibly be computationally intensive, and a nearly instantaneous online phase. In the

---

[2] In its simplest setting of packing circular or spherical particles, RSA corresponds to the Dart Throwing Algorithm [35].

[3] Level-set description based on the signed distance to the nth nearest particle boundary does not fulfil all metric criteria.
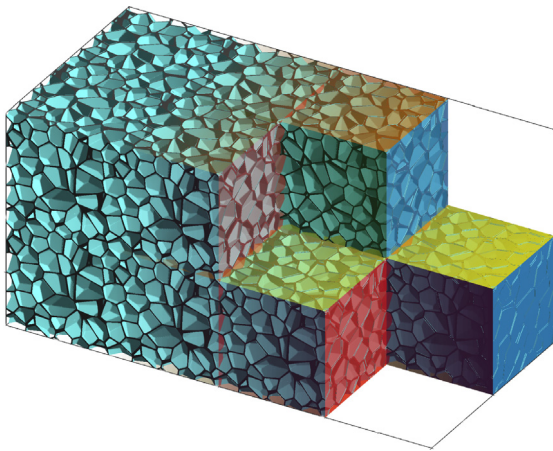
Fig. 1. Illustration of a synthesized sample of a closed-foam microstructure obtained as a 4 × 2 × 2 tiling assembled from a set of pre-generated Wang cubes. Individual face codes that encode the geometrical continuity of individual tiles and play the role of compatibility constraints during a tiling assembly are shown in semi-transparent colours.

offline phase, information regarding a heterogeneous microstructure is compressed into a set of smaller domains – Wang tiles – with predefined compatibility constraints on the compressed microstructural geometry. In the online phase, microstructural realizations are assembled from these domains with a fast, linear algorithm that produces stochastic realizations with suppressed periodicity. The merits of the tiling concept for RVE size analyses were demonstrated in [45,59].

Optimization-based approaches developed initially for the SEPUC design can be extended to incorporate generalized periodic boundary conditions and used to generate the morphology of tiles [58,60]. However, the extension amplifies the major weakness of optimization approaches – their computational cost – making them prohibitively expensive for complex three-dimensional models. As a remedy, we have proposed a method motivated by Cohen et al. [57] that combines a sample-based approach with quantitative spatial statistics [61]. While this method is by orders of magnitude faster than the optimization approach, it has difficulties handling complex, percolated microstructures such as foam, and produces corrupted ligaments in sample overlaps [45].

In this paper, we extend Sonon et al.'s method to Wang tiles in order to produce tile-based representations of microstructures intractable by the former methods [58,60,61], see Fig. 1. To this end, in Section 2, we review the fundamentals of the Wang tile concept and discuss, in detail, connectivity and mapping between vertex- and edge/face-based definitions of Wang tiles. Next, we outline Sonon et al.'s method [50,52] and describe modifications necessary to accommodate generalized periodicity, see Section 3. Finally, equipped with the adapted procedure, we illustrate sample outputs of the procedure in two and three dimensions. We also compare three tile sets – two stochastic and a strictly aperiodic one – with the same cardinalities in terms of periodicity artefacts, complementing our previous study [61] that dealt only with the distribution of tile types.

2. Wang tiles

Wang tiles constitute the building blocks of the abstract concept used in this work. Albeit the shape of the tile domains can be any parallelogram (or a parallelepiped in 3D), for simplicity's sake we assume only square (or cubic) domains in the sequel. All tiles from the *tile set* have compatibility codes attributed to

their edges (faces), illustrated using colours in Figs. 1 and 2. These codes play the role of constraints during an assembly of tile instances in a *tiling*—a portion of a plane (space) without any voids or overlaps; only tiles with the same codes on the abutting edges (faces) can be placed side by side. In addition, tiles can be neither rotated nor reflected during the assembly. Even though the last two requirements can be eliminated by modifying the definition of compatibility codes [62], we retain them for practical purposes since they preserve the orientation of the microstructure compressed within the tile set. The particular version of an assembly algorithm depends on the type of the tile set, discussed below.

Originally, the concept of Wang tiles was introduced in first-order predicate calculus as a visual surrogate to a decision problem of ∀∃∀ statements[4] [63,64]. The initial conjecture that a whole plane can be covered only if a periodic tiling exists [63] was disproved shortly after [65], triggering the pursuit of the smallest tile set allowing for strictly aperiodic tiling of the plane, see the classical monograph [66] and [67] for historical overviews. This quest for the smallest set appears to be over; supported by an extensive computer-aided search, Jeandel and Rao [67] announced a set of 11 tiles over 4 codes for each edge orientation, stating that no smaller set exists.

2.1. Stochastic tile sets

Similar to applications in biology [69], physics [70], and computer graphics [57,71], we use the concept only in its geometrical interpretation as a suitable formalism describing mutual compatibility of small domains. Except for the comparison study in Section 4.1, we also limit ourself to the stochastic tile sets introduced by Cohen et al. [57]. Besides the fact that deterministic tilings of the aperiodic sets often exhibit locally ordered patterns, the stochastic sets offer higher flexibility in terms of design, i.e. choosing the number of tiles and codes and their distribution within the set.[5]

The assembly algorithm for stochastic sets works sequentially: an initially empty grid is traversed in a scanline [73] way; at each grid node, possible candidates compatible with previously placed tiles are identified in the set and one candidate tile is randomly chosen and placed; see the illustration in Fig. 3. Thus, the only requirement for the set design is that there is at least one tile (but optimally two or more to preserve the stochastic nature of the assembly) for every possible combination of codes on the upper and left-hand edges. In principle, the random choice from the subset can be replaced with an informed selection preferring e.g. different phase volume fractions in different regions of a domain according to a pre-generated Gaussian random field. This modification allows for correlations in the microstructure at length scales larger than the tile size. For the sake of brevity and without loss of generality, however, we use only the standard stochastic algorithm in this work.

The tiling concept's ability to generate naturally looking patterns from a limited amount of samples – which proved highly appealing in computer graphics [57,71] – stems from the reduced periodicity in tiling assemblies. This feature complies well with our ambition to replace a SEPUC-based description of random heterogeneous materials with its generalization that would allow for fast synthesis of stochastic microstructural samples or microstructure geometries for entire macroscopic domains.

The design of a compressed microstructural representation consists of two steps. First, the cardinality of the tile set and a

---

[4] A statement containing one existential (∃) and two universal (∀) quantifiers.

[5] The limitations of aperiodic sets are especially critical in 3D, because only one aperiodic set of Wang cubes has been published to date [72].
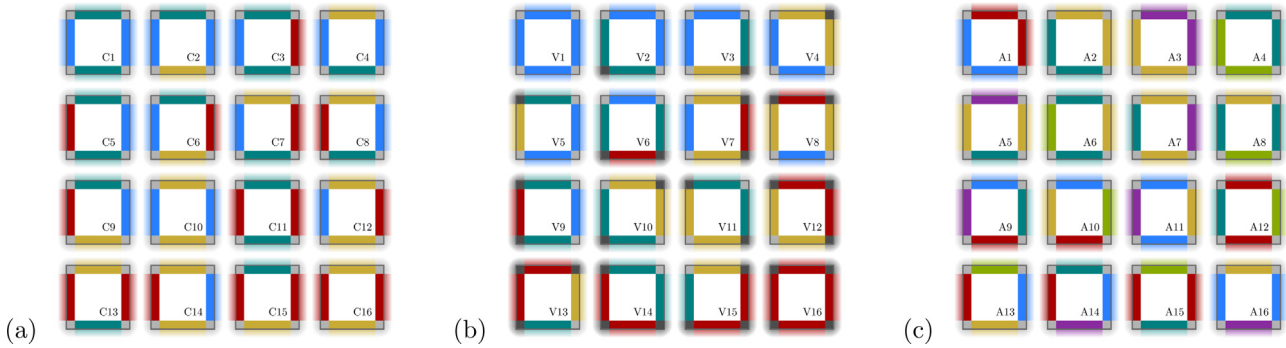
Fig. 2. Formal definition of three Wang sets with the equal cardinality of 16 tiles: (a) an edge-based set with two horizontal and vertical codes, (b) a vertex-based set over four horizontal and vertical codes, and (c) the aperiodic Ammann's set [68] over six horizontal and vertical codes. Similar to Fig. 1, the compatibility codes playing the role of constraints during the tiling assembly are illustrated with colours assigned to individual edges. The corresponding vertex codes are depicted in light and dark grey corner squares. While in the vertex-based set (b), edge codes were obtained by mapping two vertex codes uniquely to one edge code, the potentially different vertex codes for sets (a) and (b) were identified using the analysis from Section 2.3 of connectivity graphs shown in Fig. 5.
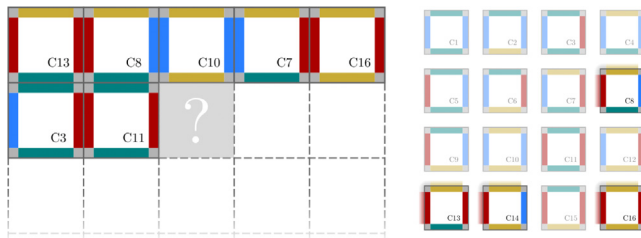

Fig. 3. Illustration of a step in the stochastic assembly algorithm [57]. A tiling of a given size is filled sequentially; for each position in a tiling, a final tile is chosen from candidate tiles that are filtered out from a given tile set based on the compatibility constraints posed by the previously placed tiles. The candidate tiles for the particular position denoted with the question mark are highlighted in the tile set shown on the right.

Fig. 4. Illustration of the generalized periodicity effect on the particle copy process. The three steps shown illustrate how a particle initially placed at the top left corner of the first tile (grey particle on the left-hand side of the figure) is sequentially copied to other loci following the edge codes it intersects. Dashed particle outlines loci where the particle will be copied in the next step.

particular distribution of edge (face) codes is chosen. This controls the frequency of tile occurrence in a tiling. Next, tile interiors are designed such that (i) the generated microstructure is continuous across the compatible edges and (ii) assembled tilings resemble the desired microstructure. Note that the latter requirement is not imposed on individual tiles. On the contrary, variability in the compressed representation is the main merit of the tiling concept. Tile interiors together with edge compatibility carry local microstructural characteristics, while the fluctuations over distances larger than the tile size are facilitated via the tile assembly algorithm.

2.2. Edge- vs. vertex-based tile definitions

Complementary to the edge-based specification introduced above, Wang tiles can be defined using vertex codes. Unlike the standard edge-based definition, the vertex-based definition allows for direct control of tile states across vertices, preventing pronounced repetitiveness when a visually distinctive microstructural feature falls into the vertex region [57]. To avoid this "corner problem", Cohen et al. [57] proposed marking the tile vertices with an additional set of codes, essentially overlaying two tile set definitions. Subsequently, Lagae and Dutré [73] retained only vertex codes, reporting superior spectral properties for assembled patterns compared to edge-based sets with the same cardinality.

Albeit the two definitions yield the same cardinality of tiles (C^4) in the complete set, i.e. the set that contains all combinations of C codes, they differ when it comes to minimal stochastic sets, i.e. sets that contain at least two tiles for each admissible code combination of already placed neighbouring tiles ($2C^2$ for
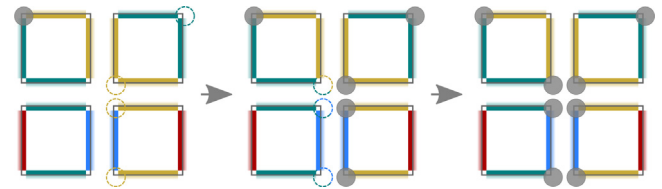
edge-based sets vs. $2C^3$ in the case of vertex-based ones). The difference is even more pronounced in three dimensions: a full face-defined set contains C^6 cubes compared to C^8 cubes in a vertex-defined set. Theoretically, Wang cubes can be also defined with edge codes in three dimensions; however, cardinality exceeds both the face- and vertex-based definitions without any known benefits.

Mapping vertex-defined tiles[6] to the original definition is straightforward: two vertex codes define one edge code. Hence all the above-mentioned design and assembly procedures directly apply to the vertex-based definition as well. In fact, the mapping provides an effortless way to produce a minimal stochastic edge-defined set with equal occurrence probability of each tile, the set characteristics sought for in [60]. On the other hand, because the mapping is injective only, it is generally not possible to map an edge-defined set to a vertex-based one.

2.3. Vertex analysis of a tile set

Here, we outline a tile set analysis capable of revealing the underlying vertex definition, if it exists. The motivation behind this analysis is the pragmatic question which arises when implementing generalized periodicity: If a particle intersects a tile boundary, what other tiles should it be copied to? While the answer is straightforward for a particle intersecting a tile edge, vertex overlap is more involved. A particle overlapping a vertex is carried to other tiles by both the horizontal and vertical edges. Because the particle overlapped a vertex in the original tile, its images overlap vertices of the other tiles as well. Consequently, these images are propagated further by the newly affected edge codes, see Fig. 4. Whether the particle is eventually copied to

---

[6] The same procedure holds also for Wang cubes, where four vertices define one face code.
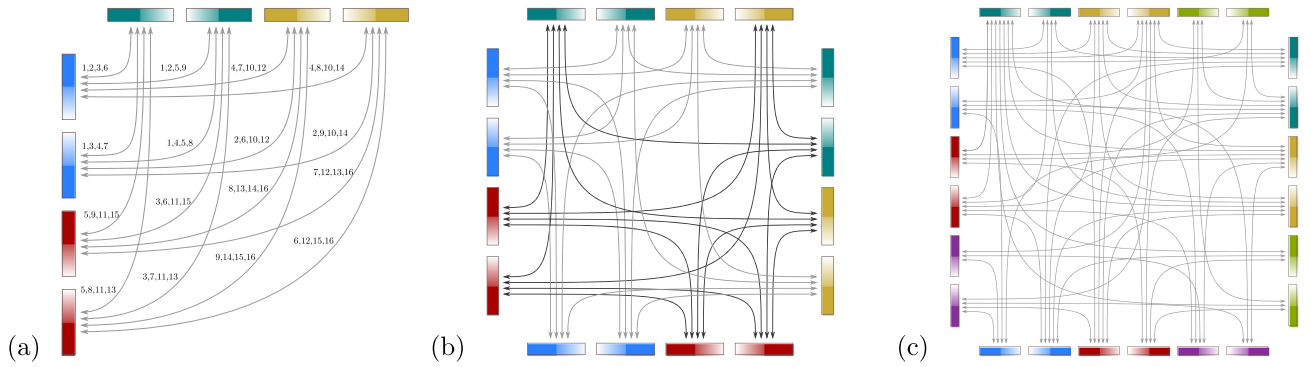
Fig. 5. Comparison of graphs pertinent to the tile sets from Fig. 2, identifying a potential vertex-based definition of the sets. Each graph node corresponds to a code on a half of a tile edge (either left/right for horizontal codes, or top/bottom for vertical codes), and graph arcs correspond to tile vertices. As there are multiple vertices – and hence parallel arcs – connecting the same nodes, the graphs should be drawn as a multigraph. However, multiple arcs are collapsed into one for the sake of brevity; only in (a) multiplicity is indicated by tile indices attached to each arc. Independent sub-graphs indicating the vertex character of the set are shown in distinct light and dark grey, which correspond to the vertex codes in Fig. 2.

all vertices or appears only in a selected subset depends on the allocation of codes to individual tiles.

Assume an undirected graph[7] where each node represents the code on a particular half of an edge (either top/bottom for vertical or left/right for horizontal codes, resulting in two occurrences of each edge code) and each arc corresponds to a tile vertex. The graph is by definition bipartite, because each vertex connects horizontal and vertical codes, and represents a multigraph since there are usually more vertices with the same adjacent edge codes in the tile set. To answer the aforementioned question, we identify connected components of the graph using the Depth-First Search algorithm [74, Section 18.2]. Each independent sub-graph then corresponds to a distinct vertex code and the arcs pertaining to the sub-graph determine the vertices to which a vertex-overlapping particle will be propagated. See Fig. 5 for a comparison of three graphs pertinent to the tile sets depicted in Fig. 2. The subgraphs (if present) are plotted in distinct grey colours, which correspond to the colours of vertex codes shown in Fig. 2.

In three dimensions, a set of Wang cubes can be analysed analogously with only a minor modification: a graph node represents one corner of a cube face. Thus, each face code appears four times in the graph. In addition, cube edges with the same direction vector must be also classified, addressing the situation when a particle intersects a cube edge. The classification follows the exact same procedure as vertex identification in two dimensions, neglecting the codes on faces perpendicular to the analysed edges.

3. Tile design using level-set functions

As mentioned in Section 1.1, Random Sequential Adsorption (RSA)[8] is one of the most frequent algorithms for generating particle packings and microstructural geometries. It follows the simple idea of throwing particles of an arbitrary shape into a given domain and keeping those that do not overlap with previously placed ones. In its original setting, however, the algorithm poses a critical drawback for higher volume fractions: the success rate for accepting the particle position rapidly decays in later stages because the majority of the randomly selected positions collide with the already placed particles. Optimally, the

remedy would be to sample the new particle position directly from a domain that is certified to result in a non-overlapping composition.

Several such remedies have been listed in Section 1.1. Here, we recall the approach of Sonon et al. because we have extended their methodology to Wang tiles. Sonon et al. [50–53] adopted the implicit, level-set-based description of a microstructural geometry and demonstrated that it enables – in addition to the desired sampling from a valid domain loci – generating complex microstructural geometries (such as open and closed foams) in a unified framework, using suitable morphing operations. Moreover, both features are independent: if a particle packing is the desired output, the morphing operations can be neglected. On the other hand, any distribution of points or particles obtained by different packing methods can serve as the input for the morphing operations.

In the the remainder of Section 3, we revise individual steps from the Sonon et al.'s framework and introduce necessary modifications facilitating the compatibility constraints arising in the concept of Wang tiles. For the sake of clarity, all procedures are presented in the two-dimensional setting with comments about three-dimensional problems when necessary. Section 4 then presents both two- and three-dimensional results.

3.1. Generating particle packings

The original Sonon et al.'s approach [50] rests on describing the geometry of particle $\mathcal{P}$ with a level-set function $\mathcal{L}^{\mathcal{P}}$, which gives the signed distance of point $\boldsymbol{x}$ from the nearest particle boundary $\Gamma^{\mathcal{P}}$ (with negative values inside the particle), such that

$$\mathcal{L}^{\mathcal{P}}(\boldsymbol{x}) = \min_{\boldsymbol{y} \in \Gamma^{\mathcal{P}}} d(\boldsymbol{x}, \boldsymbol{y}), \tag{1}$$

with the signed distance function $d(\boldsymbol{x}, \boldsymbol{y})$ given by

$$d(\boldsymbol{x}, \boldsymbol{y}) = \begin{cases} -\|\boldsymbol{x} - \boldsymbol{y}\| & \text{if } \boldsymbol{x} \in \mathcal{P}, \\ \|\boldsymbol{x} - \boldsymbol{y}\| & \text{otherwise}. \end{cases} \tag{2}$$

Consequently, the zero iso-line of $\mathcal{L}^{\mathcal{P}}$ represents the particle boundary $\Gamma^{\mathcal{P}}$, as depicted in Fig. 6. Note that unlike many applications of the level-set method where the implicit geometry description evolves in time according to the Hamilton–Jacobi equation, see [76, and references therein], $\mathcal{L}^{\mathcal{P}}$ is computed for the given geometry of a particle and remains fixed.

Assume for now that we have domain Ω that contains set $\mathcal{R}$ of already placed particles and their geometry is encoded in a single level-set field $\mathcal{L}^{\Omega}$,

$$\mathcal{L}^{\Omega}(\boldsymbol{x}) = \min_{\mathcal{P} \in \mathcal{R}} \mathcal{L}^{\mathcal{P}}(\boldsymbol{x}), \quad \forall \boldsymbol{x} \in \Omega; \tag{3}$$

---

[7] Instead of the standard vertex-edge nomenclature for a graph, we use node and arc terms in order to avoid confusion with similar geometrical notions related to Wang tiles.

[8] Sometimes, the term Random Sequential Addition, e.g. [75], is used interchangeably.
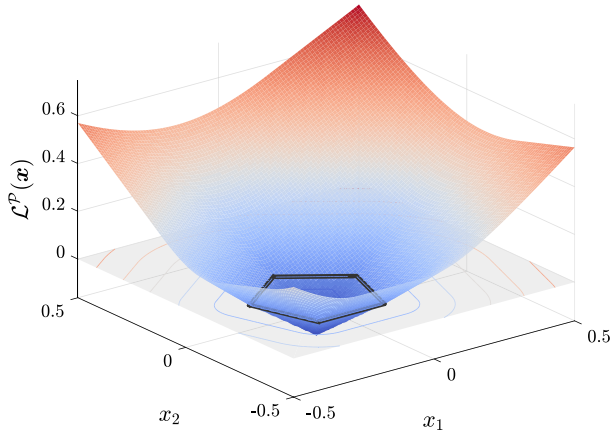
Fig. 6. Implicit representation of particle geometry with the level-set function $\mathcal{L}^{\mathcal{P}}$. The boundary of particle $\mathcal{P}$ is obtained as a zero-value contour of $\mathcal{L}^{\mathcal{P}}$ (plotted in a black line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
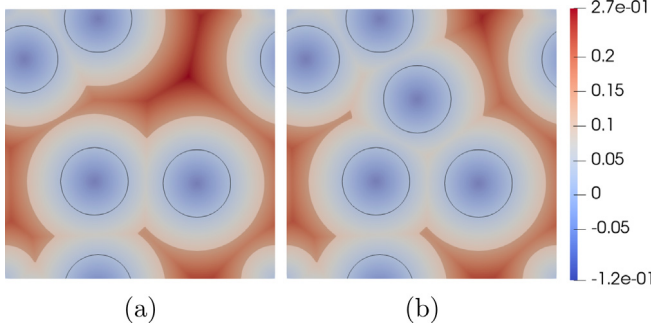

Fig. 7. Two consecutive steps of the level-set based RSA algorithm for a Periodic Unit Cell: (a) level-set field $\mathcal{L}^{\Omega}$ with previously placed particles (values of $\mathcal{L}^{\Omega}$ outside the admissible sampling domain $\mathcal{A}$ are displayed as translucent) at the nth step; (b) updated state of $\mathcal{L}^{\Omega}$ and $\mathcal{A}$ after placing a new particle at step $n+1$.

see Fig. 7 for an illustration. Knowing the radius r of the smallest circumscribed circle as the only characterization of the new particle, overlaps with the existing particles can be readily prevented by sampling its centre from the domain

$$\mathcal{A} = \left\{ \boldsymbol{x} \in \Omega \mid \mathcal{L}^{\Omega}(\boldsymbol{x}) \geq r \right\} . \tag{4}$$

After placing the new particle $\mathcal{P}$, the level-set field $\mathcal{L}^{\Omega}$ is updated with $\mathcal{L}^{\mathcal{P}}$,

$$\mathcal{L}^{\Omega}(\boldsymbol{x}) = \min \left(\mathcal{L}^{\Omega}(\boldsymbol{x}), \mathcal{L}^{\mathcal{P}}(\boldsymbol{x}) \right), \quad \forall \boldsymbol{x} \in \Omega . \tag{5}$$

This procedure repeats until either a pre-defined stopping criterion (e.g. a desired volume fraction) is met or $\mathcal{A} = \emptyset$.

To increase the volume fraction that can be achieved with this procedure, two additional restrictions can be posed on the admissible subdomain $\mathcal{A}$. First, a maximal distance ρ from existing particle boundaries can be added to Eq. (4), preventing too large gaps between particles. (If requested, a minimal distance κ between two particles can be further enforced as well.) Second, to prevent locally jammed states with large interparticle gaps, an additional field $\mathcal{L}_{II}^{\Omega}$ storing the shortest distance to the surface of the second nearest particle can be included in Eq. (4) and limited from above with σ, see [50] for additional details. With all these constraints in action, Eq. (4) takes the form

$$\mathcal{A} = \left\{ \boldsymbol{x} \in \Omega \mid r + \kappa \leq \mathcal{L}^{\Omega}(\boldsymbol{x}) \leq r + \rho, \ \mathcal{L}_{II}^{\Omega} \leq r + \sigma \right\} . \tag{6}$$
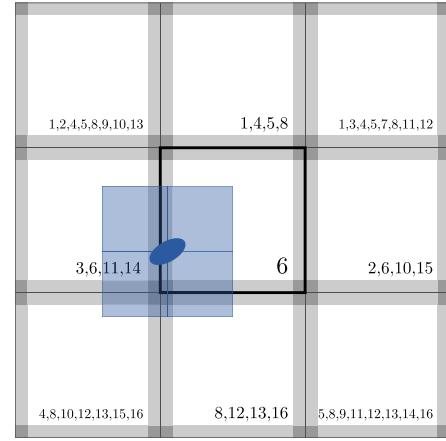

Fig. 8. Neighbourhood grid identified for each tile in a tile set (here shown for tile V6 from Fig. 2b) containing codes of tiles that can occur at given positions. Light and dark grey regions depict individual entities in the bookkeeping structure, which prevents calculations. (During the update procedure described in the text, the blue particle can be accounted for either in tile 6 at the centre, or from tiles 3, 6, 11, or 14 at the first column, second-row position. With the bookkeeping structure considered, only one update is performed.) The blue square illustrates the use of acceleration via a precomputed patch; values of $\mathcal{L}^{\mathcal{P}}$ are computed only once and used for all particle occurrences during an update. If requested, values outside the patch are computed in a regular way.

Enforcing periodicity of the packing is also straightforward: upon placement of a new particle, the domain level-set field is also updated with the particle's periodic images (eight in 2D; recall the grid in Fig. 8, and 26 in 3D).

3.1.1. Extension to Wang tiles

Compared to generating PUCs, the additional compatibility constraints among individual Wang tiles necessitate several modifications in the original Sonot et al.'s algorithm to generate continuous microstructural morphology of a tile set. For the Poisson discs distribution, Lagae and Dutré [77] modified the Dart Throwing algorithm by partitioning tile domains into separate regions pertinent to vertices, edges, faces, and tile interiors, and filling the regions sequentially, starting with the vertex regions. Here we present a different approach better suited to packing particles of different sizes and shapes that processes all tiles simultaneously and does not require a priori partitioning of tile domains.

As $\mathcal{L}^{\Omega}$ we use a level-set field $\mathcal{L}^{\mathcal{S}} = \bigcup_{\mathcal{T} \in \mathcal{S}} \mathcal{L}^{\mathcal{T}}$ defined on a regular grid for each tile $\mathcal{T}$ in the tile set $\mathcal{S}$. For each tile in a pre-processing stage, we identify potential neighbouring tiles and store their indices in a grid (3×3 or $3 \times 3 \times 3$ in two or three dimensions, respectively), see Fig. 8, which will be used later in the updating phase.

For further description, we define a *copy-inducer* as a geometrical entity that is responsible for inducing particle copies to other tiles. Hence, a copy-inducer can be either a vertex, edge, face, or formally the tile itself and is related to a code identified in Section 2.3. The particular type of the copy-inducer stems from particle's intersections with a tile boundary (if a particle does not intersect a boundary, the tile is the copy-inducer, indicating that no copy is necessary. If a particles intersects one edge, the edge is the copy-inducer, and so on).

Modification 1: Update. The main structure of the modified algorithm follows the original one. After sampling a particle location from $\mathcal{A} \subseteq \mathcal{S}$, which yields tile $\mathcal{T}_i$ and centre $\boldsymbol{x}_c^{\mathcal{P}} \in \mathcal{T}_i$ where the new particle will be placed, the particle's intersections with the boundary of $\mathcal{T}_i$ are determined, and – based on the intersections

– a copy-inducer is identified. Images of the particle are subsequently copied to the relevant tiles following the matching codes on the corresponding copy-inducers. A level-set field of the whole set is then updated tile-wise as described below.

For each tile, we sequentially loop over all tile positions in the neighbours grid shown in Fig. 8 (including the centre position as well) and check each of the potential tiles at that position for a new particles added in the current step. If any, the central tile's level-set field is updated according to Eq. (5). Clearly, a new particle can appear several times in the same place. To avoid duplicated calculations, we set up a bookkeeping structure that records which copy-inducer in the neighbouring grid was already used during the update and run only the unperformed updates, see Fig. 8.

Two implementation approaches are possible at this stage: (i) $\mathcal{L}^{\mathcal{P}}$ is computed anew for each tile and position in the bookkeeping structure, or (ii) provided that $\mathcal{L}^{\mathcal{S}}$ is represented with discrete values on a regular grid, $\mathcal{L}^{\mathcal{P}}$ is precomputed for an auxiliary domain/patch aligned with the particle centre before any update and this field is then reused for all particle instances. Moreover, we combine both approaches with the pre-screening acceleration proposed in [50] that resorts to computing potentially demanding surface-distance functions only when the $\mathcal{L}$ value of a circumscribed circle/sphere – which is fast and easy to calculate – is less than the actual $\mathcal{L}^{\Omega}$ value for a given $\boldsymbol{x}$.

While the pre-screening acceleration is significant in all cases, the benefits of a precomputed patch depend mainly on the patch size relative to the particle circumscribed radius, the $\mathcal{L}^{\mathcal{S}}$ field resolution, and the cardinality of the tile set. Especially in the later stages of the algorithm, when only relatively small parts of $\mathcal{L}^{\mathcal{S}}$ are updated, pre-computing a patch for a particle that is not copied to other tiles might cause unnecessary overhead. Therefore, we use the patch approach only for particles with either vertex, edge, or face copy-inducers.

For the three-dimensional set of 16 cubes used in Section 4.3 with ellipsoidal particles of circumscribed radius $r = 0.15$, the patch pre-calculation delivered 10% saving in computational time on average. Note also that the multi-query problem of finding the shortest distance to a particle boundary from a set of points is trivially parallelizable, changing the trade-off between patch pre-computations and direct calculations with available threads.

When all tiles are updated, the algorithm proceeds with identifying $\mathcal{A}$ for the next particle.

Modification 2: Artificial level-set field. Definition of the admissible domain $\mathcal{A}$ must also be modified for Wang tiles; particle centres cannot be sampled directly from $\mathcal{A}$ obtained from $\mathcal{L}^{\mathcal{S}}$. The reason for this is that the near-boundary parts of each tile must be informed about the interiors of related tiles in order to prevent insertions similar to the one depicted in Fig. 9a, where a copied instance of a newly placed particle collides with interior particle(s) of other tiles. The problem arises only when a particle image is added to a copy-inducing entity with the same orientation as the actual inducer (left-hand side edges in Fig. 9a); particles copied to the opposite copy-inducers automatically meet the non-overlap requirement thanks to the first update modification described above.

To avoid this problem, we replace $\mathcal{L}^{\mathcal{S}}$ in Eq. (4) or Eq. (6)[9] with a modified field $\widetilde{\mathcal{L}}^{\mathcal{S}}$ which is equal to $\mathcal{L}^{\mathcal{S}}$ except for boundary regions, whose width is dictated by the radius r of the circle/sphere circumscribed to the current particle. Each boundary region is then constructed as a point-wise minimum over all regions related to the same copy-inducers with the same code and
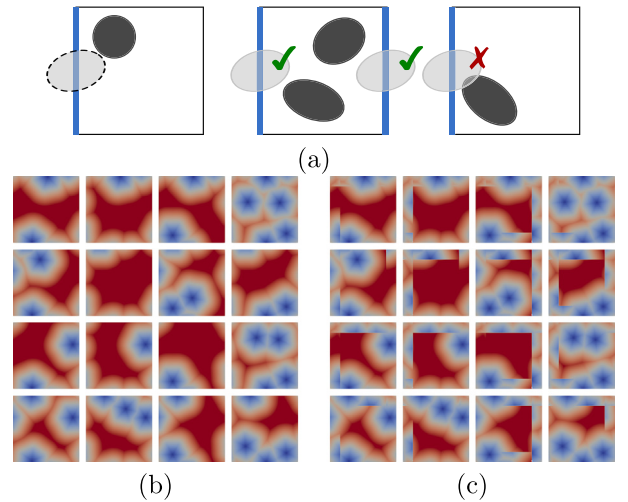

Fig. 9. (a) Illustrations of particle intersections that can happen due to copying particle images without considering artificially updated level-set field $\widetilde{\mathcal{L}}^{\mathcal{S}}$ in the definition of $\mathcal{A}$. The problematic overlaps occur only at inducer places (highlighted in blue) with the same orientation as the particle copy-inducer (dashed outlined grey particle on the left-hand side). Figures (b) and (c) depict the original $\mathcal{L}^{\mathcal{S}}$ and its modified version $\widetilde{\mathcal{L}}^{\mathcal{S}}$ that communicates necessary data across boundary regions to prevent intersections from (a). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the same orientation; see Fig. 9c. This construction propagates the close-boundary state across relevant tiles and thus prevents overlapping insertions similar to the one depicted in Fig. 9a. Note that a similar modification is not required for $\mathcal{L}_{\mathrm{II}}^{\Omega}$ because all information necessary to prevent the collisions is already contained in $\widetilde{\mathcal{L}}^{\mathcal{S}}$.

Modification 3: Breaking regular grid. The level-set field $\mathcal{L}^{\mathcal{S}}$ is typically implemented on a regular grid and the particle centres $\boldsymbol{x}_{c}^{\mathcal{P}}$ thus end up aligned with the grid. To break this artificial ordering, after sampling a new particle centre, we check whether the surrounding points belong to $\mathcal{A}$. The belonging points then define quadrants/octants in which the particle can be moved freely. We generate a random shift within these quadrants/octants and update the particle centre accordingly. However, this modification is possible only if the grid spacing is sufficiently small compared to the particle size; otherwise, the implicitly assumed linear approximation of $\mathcal{L}^{\mathcal{S}}$ is inaccurate and might result in particle intersections at intermediate locations.

The modifications described above are sufficient for extending the original Sonon et al.'s level-set-based packing algorithm to the concept of Wang tiles. Note that all control variables, e.g. particle radius r or minimal and maximal distances ρ and σ, can change while the algorithm operates. For instance, to achieve denser packings for multi-modal particle size distributions, it is preferential to start placing large particles first and sequentially proceed to the smaller ones.

Nonetheless, our extension inherits the weakness of the original algorithm when it comes to particles whose shapes are significantly different than their circumscribed circle/sphere, e.g. prolonged ellipsoids. In particular, the more the particle volume deviates from the volume of the circumscribed sphere, the less dense the final packing can be. While the geometry of the already placed particles is described exactly with $\mathcal{L}^{\mathcal{S}}$ (up to inaccuracies due to grid spacing in $\boldsymbol{x}$), the newly placed particle is represented only approximately during selection of a centre, yielding a pessimistic estimate of mutual intersections. The potential remedy

is to translate the particle centre right after sampling according to some heuristic rule aimed at producing denser packings; however, we leave this issue unaddressed in this paper.

3.2. Morphing operations

The acceleration of RSA by adopting the level-set approach is attractive on its own; however, the main appeal of the Sonon et al.'s level-set framework is the elegance with which complex microstructures can be generated. As the simplest example, adding a constant γ to $\mathcal{L}^{\mathcal{S}}$ enables (i) fine-tuning of particle volume fractions, (ii) smoothing of sharp corners of polyhedral particles, and (iii) particle coating, e.g. for defining the Interfacial Transition Zone.

Combinations of $\mathcal{L}^{\mathcal{S}}$ and $\mathcal{L}^{\mathcal{S}}_{\mathrm{II}}$ permit rendering interparticle bridges and controlled Voronoï-like tessellations; see [50], where these morphing operations were used to mimic microstructural geometry of clay/sand mixed soils and irregular masonry. Considering also the shortest distance to the boundary of the third nearest particle $\mathcal{L}^{\mathcal{S}}_{\mathrm{III}}$, Sonon et al. [52] demonstrated that their framework can produce highly adjustable models of foam-like microstructures, including features such as a smooth transition from open to closed foams, concavity of foam ligaments, their coatings and hollow interiors, and variable thickness of foam ligaments and cell walls.

As in Section 3.1, we present only the essentials of generating foam-like microstructures and introduce necessary modifications. The reader is referred to [50, Section 4] and [52, Sections 3 and 4] for details regarding the above-mentioned local adjustments, since they remain unchanged.

Assume we have at our disposal all three fields $\mathcal{L}^{\mathcal{S}}$, $\mathcal{L}^{\mathcal{S}}_{\mathrm{II}}$, and $\mathcal{L}^{\mathcal{S}}_{\mathrm{III}}$, containing the shortest distance to the boundary of three nearest particles (either computed a posteriori for a given particle assembly or already tracked during the packing algorithm).

In a manner similar to classical Voronoï tessellation, where a domain is partitioned by boundaries that have the same distance to the two closest seeds, thresholding the modified difference

$$\mathcal{F}_{\mathrm{c}} = \left(\mathcal{L}^{\mathcal{S}}_{\mathrm{II}} - \mathcal{L}^{\mathcal{S}}\right) + t_{\mathrm{c}} \tag{7}$$

yields a closed-cell, foam-like geometry with t_{c} indirectly controlling the thickness of foam cell walls. In the same spirit, the centrelines of the open-foam ligaments can be viewed as loci with the same distances to the boundary of three nearest particles; hence thresholding

$$\mathcal{F}_{\mathrm{o}} = \left(\tfrac{1}{2}(\mathcal{L}^{\mathcal{S}}_{\mathrm{III}} + \mathcal{L}^{\mathcal{S}}_{\mathrm{II}}) - \mathcal{L}^{\mathcal{S}}\right) + t_{\mathrm{o}}, \tag{8}$$

produces open-foam ligamentswith t_{o} governing their thickness. Finally, combining Eqs. (7) and (8)

$$\mathcal{F}(\boldsymbol{x}) = \min\left(\mathcal{F}_{\mathrm{c}}(\boldsymbol{x}), \mathcal{F}_{\mathrm{o}}(\boldsymbol{x})\right), \quad \forall \boldsymbol{x} \in \mathcal{S}, \tag{9}$$

allows for more realistic geometries with material concentration at wall intersections, see Fig. 10.

It is critical when extending the original framework to Wang tiles, with respect to the morphing procedures, to ensure that all three fields are continuous across the relevant tile edges/faces. As illustrated in Fig. 11, copying particle images introduced in the previous subsection is insufficient in this regard. To guarantee the required continuity, we define a wider domain margin as a subdomain related to individual copy-inducers. Instead of computing intersections with a tile boundary, we compute them for a virtual inset boundary (inset by r is usually enough) and copy the particles according to the virtual copy-inducers. Consequently, wider portions of tile domains are restrained near boundaries, resulting in restored continuity. This modification is also reflected in the construction of auxiliary field $\widetilde{\mathcal{L}}^{\mathcal{S}}$.
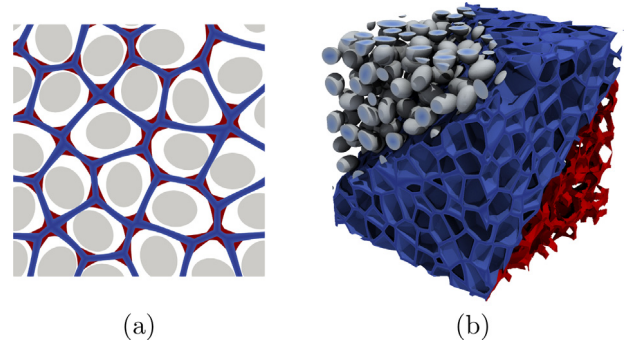

Fig. 10. Influence of morphing operations on resulting (a) two- and (b) three-dimensional geometry. From the original particle distribution (depicted with light grey), the closed-cell, foam-like geometry (shown in blue) is obtained by Eq. (7) and the open-like features (plotted in red) follow from Eq. (8). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
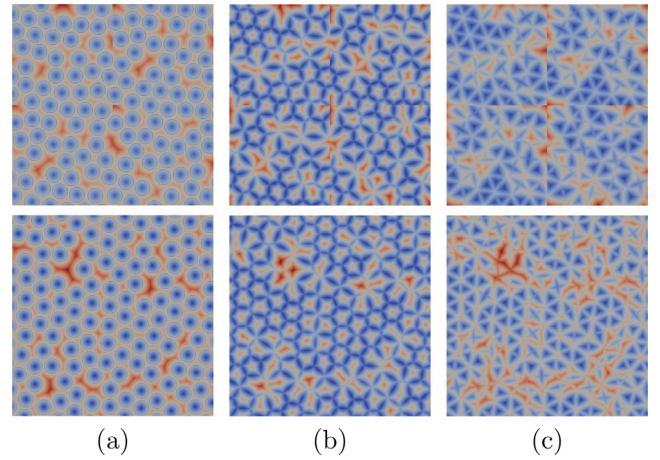

Fig. 11. Assembled level-set fields composed of 2×2 tiles without (top row) and with (bottom row) a virtual boundary inset taken into account during particle placement. Note that even the nearest-neighbour distance $\mathcal{L}$ (a) is not continuous across vertices when the virtual inset is not considered, albeit the particle boundaries, i.e. zero-value contours outlined in grey, are. Severe discontinuities then appear for the second-nearest $\mathcal{L}_{\mathrm{II}}$ (b) and the third-nearest $\mathcal{L}_{\mathrm{III}}$ (c) neighbour distance without the inset boundaries. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4. Results

4.1. Comparison of 16-tile sets

First, we compare the periodicity reduction in three planar sets depicted in Fig. 2, which have the same cardinality but different tile definitions. We supplement our earlier observations [61] regarding suppressed artificial periodicity compared to PUC-based reconstructions and better fitness of stochastic tile sets over their aperiodic counterparts. Namely, we consider: (C16) an edge-based stochastic set over four colours,[10] (V16) a vertex-based stochastic set over eight colours, and (A16) the aperiodic Ammann's set [68] over 12 colours.[11]

Assuming ergodicity of a microstructure, we quantify the artificial periodicity in the tiling assemblies by means of the secondary peaks in the two-point probability function $S_2(\boldsymbol{x})$ [4].

---

[10] This set corresponds to the set reported in [57].

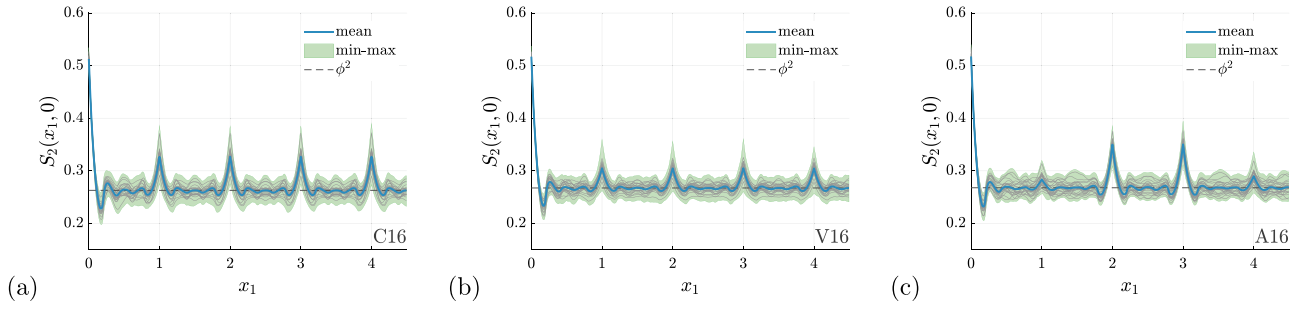[11] For the aperiodic assembly, we used the substitution rule from [78].

Fig. 12. S_2 cross-sections along an x_1 coordinate for 10×10 tilings for (a) the edge-defined tile set C16, (b) the vertex-defined set V16, and (c) the Amman set A16, considered in Section 2.3. Thin grey lines depict the average over 100 realizations of each tile set morphology, solid blue line shows the overall average (i.e. over realizations and morphologies) and the green area captures the range between minimal and maximal values. Dashed grey line marks the asymptotic value ϕ^2, where the average volume fraction ϕ is considered. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$S_2(\boldsymbol{x})$ states the chance of finding two points separated by $\boldsymbol{x}$ in the same phase. Hence, $S_2(\boldsymbol{x})$ attains its maximum at $\boldsymbol{x} = \boldsymbol{0}$, where it equals the volume fraction ϕ of a chosen phase. For microstructures without any internal ordering, $S_2(\boldsymbol{x}) \to \phi^2$ with $\|\boldsymbol{x}\| \to \infty$ because two sufficiently distant points are uncorrelated. On the other hand, if a microstructure is composed of a repeating (SE)PUC, $S_2(\boldsymbol{x})$ exhibits secondary peaks having the same magnitude as $S_2(\boldsymbol{0})$ at nodes of a regular grid whose spacing corresponds to the (SE)PUC size. Thus, in the intermediate case of Wang tiling, the magnitude of secondary peaks indicates the remaining artificial ordering in the reconstructed microstructure, see [58,61] for details. For computing $S_2(\boldsymbol{x})$, we used the Fast Fourier Transform (FFT) instead of random sampling, because the bias in the statistics caused by the implicit periodicity introduced with FFT is negligible for reasonably large microstructural realizations [79]. Recall also that the two-point probability function S_2, widely used for statistical quantification of material microstructures, e.g. [4], is the inverse Fourier Transform image of the power spectral density (or its estimate via periodogram) traditionally used in the Computer Graphics community, e.g., [57,71,77].

For each of the three tile sets we ran the presented algorithm 25 times and thus generated 25 different tile set morphologies. Aiming at capturing the artificial periodicity related to the limited number of edge codes and tiles, we resorted to the simplest set-up with circular inclusion of a fixed radius $r = 0.1$ in order to minimize the influence of particle shapes. In addition, we prevented the particles from overlapping tile vertices in order not to favour the vertex-based tile set a priori.[12] The level-set fields of each unit-size tile ($\Omega = [-0.5, 0.5]^2$) were discretized using a regular grid comprising 201×201 points. Finally, we posed two constraints, $\kappa = 0.01$ and $\rho = 0.03$, on $\mathcal{A}$, recall Eq. (6).

We assembled 100 microstructural realizations, i.e. tilings, containing 10×10 tiles for each tile set and tile morphology and computed S_2 statistics for the inclusion phase. Cross-sections of S_2 along the x_1 coordinate are plotted in Fig. 12 for the tile sets considered. As expected, the secondary peaks appear at loci with integer coordinates. Excluding the primary peak, we picked the second highest extreme $\hat{S}_2$ normalized against the averaged asymptotic value ϕ^2 for each realization and plotted the obtained data with box-and-whisker diagrams in Fig. 13.

While the Ammann's tile set A16 seems promising locally – note the suppressed secondary peaks at $x_1 = 1.0$ and $x_1 = 4.0$ in Fig. 12c – its deterministic aperiodic structure leads to pronounced secondary extremes in other regions, e.g. $x_1 = 3.0$. This observation corroborates our previous conclusions, based solely on the distribution of individual tiles within tilings, that
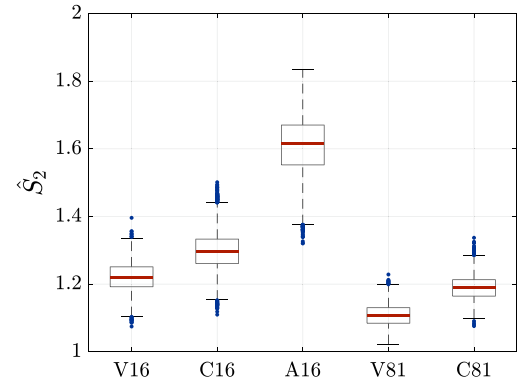
---

[12] From Fig. 5 it follows that a particle will be copied to all tiles if it overlaps any vertex for the A16 and C16 tile sets.

Fig. 13. Box-and-whisker plot of the normalized maximal secondary peaks $\hat{S}_2$ over 100 realizations and 25 tile set morphologies for each tile set. The central lines mark the median; edges of each box denote the first and third quartiles; and whisker lines contain all data not considered as outliers (marked as dark blue dots).

strictly aperiodic sets such as the Ammann's or Culik's set considered in [61] lead to higher secondary peaks than stochastic sets [57] with the same cardinality, which feature more uniform secondary peaks. Figs. 12 and 13 clearly show that, out of the two stochastic sets, the vertex-based definition performs better as expected, even despite the prevented vertex overlaps. This superiority of vertex-based tile sets in two dimensions is further supported by additional results comparing vertex- and edge-based sets comprising 81 tiles, see Fig. 13.

4.2. 2D example

Second, we present an example of a two-dimensional microstructure based on polygonal particles. The geometry of each particle was derived from an originally regular, randomly-oriented polygon with the number of vertices being sampled from a normal distribution with the mean value 0.6 and the standard deviation 0.5, $\mathcal{N}\left(6.0, 0.5^2\right)$, and rounded to the nearest integer. The angle between two rays connecting the particle centre and neighbouring vertices was perturbed with a value randomly chosen from $\mathcal{N}\left(0.0, 0.5^2\right)$. Finally, each vertex was placed on its corresponding ray at a distance sampled from $\mathcal{N}\left(0.95, 0.05^2\right)$ (and capped by 1.0) relative to the particle's circumscribed radius.

Following the outcomes of Section 4.1, we picked the vertex-based set with 16 tiles depicted in Fig. 2b. Similarly to the previous set-up, the size of each tile was $\Omega = [-0.5, 0.5]^2$ and the corresponding level-set fields were discretized using a regular grid comprising 201×201 points. The initial value of the
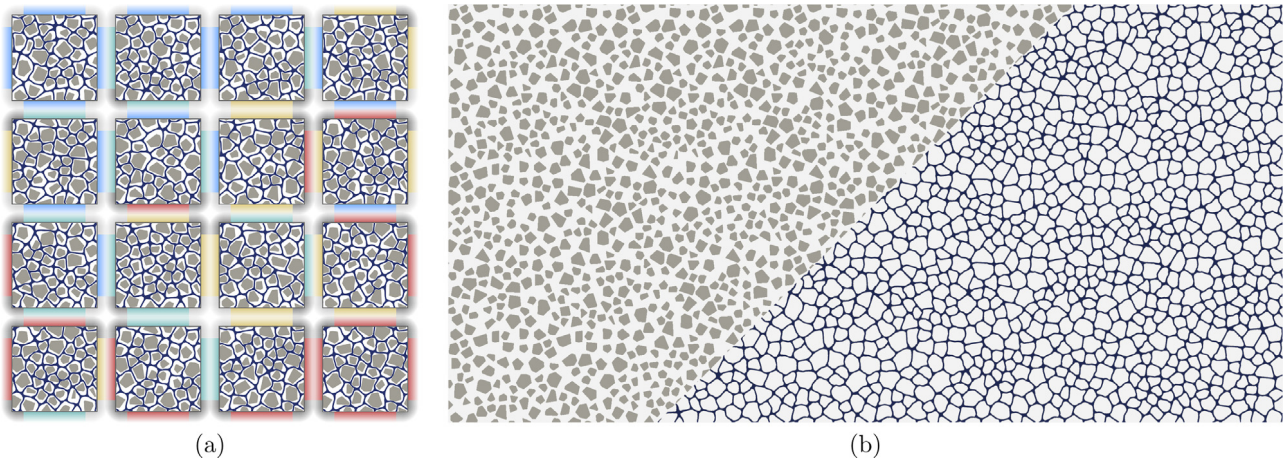
Fig. 14. Results of the algorithm set-up described in Section 4.2: (a) A composed view of the tile set with highlighted edge and vertex codes. The particulate assembly is drawn in grey; the corresponding foam-like microstructure is shown in blue (similar to Fig. 10); (b) An example of a 10 × 5 tiling. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

circumscribed radius was set to $r = 0.08$. After 40 algorithm steps, the radius was increased to 0.1, and eventually we reduced it to 0.06 after next 60 steps. The number of algorithm steps and related radii were chosen randomly to illustrate the algorithm's control options. Only the first two constraints on $\mathcal{A}$ from Eq. (6) were posed. We fixed the minimal particle distance at $\kappa = 0.01$; the maximal distance ρ was set to 0.05 in the first 50 steps and decreased to 0.02 later. The width of the virtual boundary inset was kept constant at 0.1.

Fig. 14 depicts a composed view of the resulting particle assembly and the related closed-cell, foam-like microstructural geometry obtained by considering $t_c = 0.015$ and $t_o = 0.020$ in Eq. (9).

4.3. 3D example

Finally, we demonstrate a three-dimensional output of the algorithm for a set of 16 Wang cubes with two codes for each face orientation uniformly distributed in the set.

Again, we used centred tile domains of a unit size discretized with a grid of $101 \times 101 \times 101$ points. For the packing part of the algorithm, we considered ellipsoidal particles with a random orientation and a fixed circumscribed radius $r = 0.10$. The ratio between the middle and the major semi-axes lengths was sampled from a uniform distribution $\mathcal{U}(0.7, 0.9)$; the minor to the major semi-axes length ratio was randomly generated from $\mathcal{U}(0.6, 0.7)$. The admissible domain at each algorithm step was dictated both by $\mathcal{L}^S$ and $\mathcal{L}^S_{II}$ with constraints $\kappa = 0.02$, $\rho = 0.05$, and $\sigma = 0.05$ in Eq. (6). Since we aimed at modelling a foam-like microstructure, the width of a virtual inset boundaries was set to 0.1.

The final microstructure was obtained by performing the morphing operations according to Eq. (9) with $t_c = 0.02$ and $t_o = 0.03$. For the sake of brevity, we do not show individual Wang cubes, and we plot only a microstructural sample comprising $5 \times 5 \times 3$ tiles in Fig. 15.

4.4. Computational cost

Our extension inherits the $\mathcal{O}(N)$ complexity of the Sonon et al.'s method [50,52], as shown in Fig. 16. However, the actual time needed to generate particle packings or complex morphologies depends on several factors. While the impact of the grid resolution is straightforward (computational time scales linearly with the number of grid points), the effect of other factors such
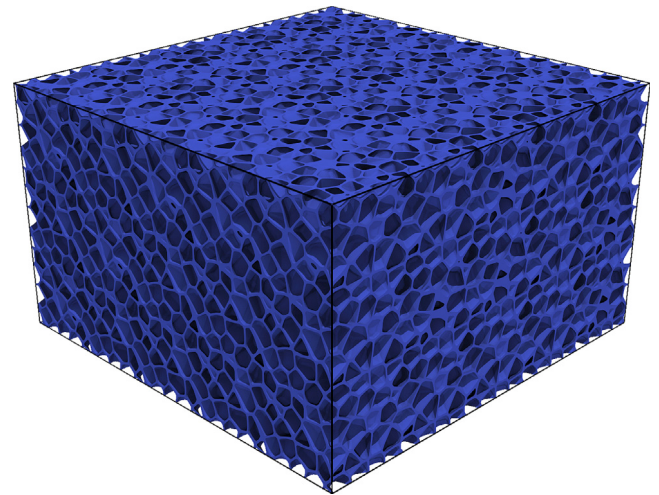

Fig. 15. A microstructural realization assembled as a $5 \times 5 \times 3$ tiling from the Wang tiles generated in Section 4.3.

as particle shape and compatibility constraints due to different definitions of a tile set is less obvious.

To illustrate the influence of these factors, wall-clock time for several settings is plotted in Fig. 16. Tile sets considered for the analysis are the two-dimensional ones used in Section 4.1 and the three-dimensional sets from Section 4.3; in addition, a PUC was included in the comparison as a trivial case of a tile set. For all settings, tile domains of unit size were discretized with 401×401 grid points in two and $101 \times 101 \times 101$ grid points in three dimensions. The circumscribed radius of particles was set to $r = 0.02$ for two-dimensional and $r = 0.05$ for three-dimensional problems. The algorithm was terminated once the desired number of placed particles was achieved. Computational time[13] for each setting was measured 10 times.

Interestingly, albeit of the same cardinality, individual tile sets from Section 4.1 exhibit significantly different computational times. This is due to the distribution of edge codes within a set,

---

[13] All computational times are reported for a workstation equipped with an Intel® Xeon® E31280 3.50 GHz processor and 16 GB RAM running Windows 10 version 1903. The algorithm was implemented in the latest standard of C++ language (C++17), relying on parallel algorithms from the C++'s Standard Template Library.
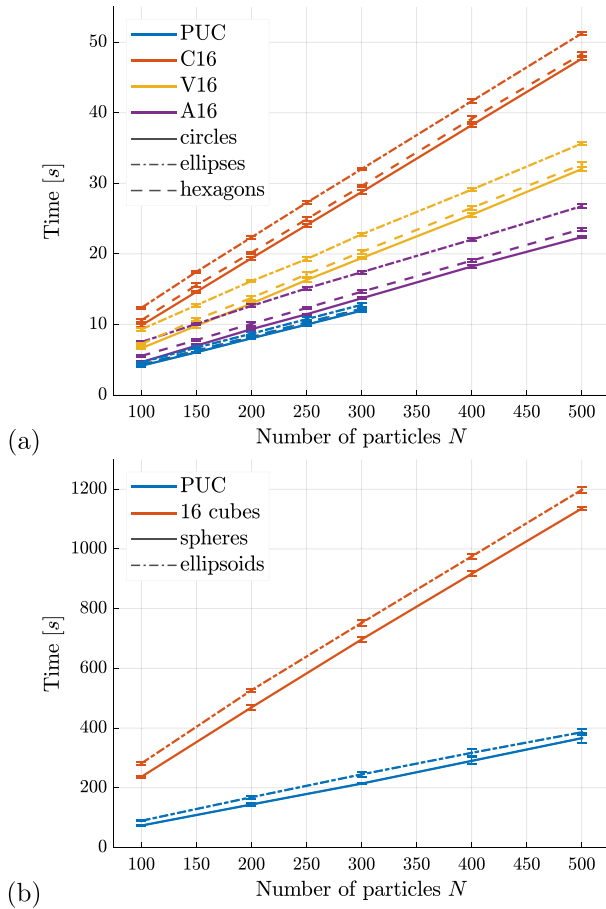
Fig. 16. Wall-clock time needed to pack *N* particles into (a) two-dimensional and (b) three-dimensional tile sets. Results for different particle shapes are drawn in distinct line styles; considered tile sets (including the trivial case of PUC) are distinguished with distinct colours. Error bars show the standard deviation for each setting. PUC results in two dimensions are limited by the number of particles that can be packed in a unit domain.

because a larger part of the tile set domain must be updated with each particle copy in sets with fewer codes, recall Section 3.1.1. Note that Ammann's set with six codes for both horizontal and vertical edges behaves nearly the same as PUC.

Computational time related to the morphing operations depends on the operational mode of the algorithm. If the algorithm is used for both packing and morphing at the same time, the level-set fields used during the particle packing phase can be directly utilized for the morphing operations. On the other hand, when the morphing part runs separately, all the distance fields must be calculated again. Yet this recalculation is typically faster than the packing because all copies and compatibility constraints are already taken into account; for instance, generating the particle packing for the microstructure presented in Section 4.3 took 1280 s while the recalculation needed for the morphing operations finished in 420 s.

5. Summary

In this work, we have extended Sonon et al.'s level-set based framework [50] to microstructural models based on the formalism of Wang tiles. The extension enables generating compressed representations of complex microstructural geometries such as open and closed foams or cells, which have been nearly impossible [61] or very expensive [58] to generate for Wang tiles to date.

Advancing from the standard Periodic Unit Cell to the generalized periodicity of Wang tiles necessitated several modifications of the original algorithm that addressed difficulties originally not encountered in the case of PUC. Driven by the geometrically-motivated question of where a vertex-overlapping particle should be copied to, we have come up with a simple procedure based on graph analysis capable of revealing the underlying vertex definition of a tile set, if present. We have also demonstrated that a straightforward copying of boundary-intersecting particles according to the edge and vertex codes is insufficient for preventing spurious particle overlaps because the information about the interior of a tile is not automatically communicated across individual tiles. As a remedy, we have introduced an artificially updated level-set field that facilitates this necessary communication. The width of the updated region in this field is dictated by the required final geometry. If only a particle distribution is to be generated, the width equal to the particle radius *r* suffices. On the other hand, if foam-like microstructures are desired, the virtual boundary inset by *r* is necessary in order to preserve continuity of the level-set fields $\mathcal{L}_I^{\mathcal{S}}$ and $\mathcal{L}_{II}^{\mathcal{S}}$ across tile edges, resulting in the updated region of width $2r$.

Additionally, we have devised two minor modifications, one capable of breaking a regular underlying grid of possible particle centres by introducing their random shifts within an admissible region, the other accelerating the $\mathcal{L}^{\mathcal{S}}$ updates by pre-computing $\mathcal{L}^{\mathcal{P}}$ on a patch. The importance of the latter modification grows with higher $\mathcal{L}^{\mathcal{S}}$ resolutions and larger cardinality of a tile set.

Having a universal framework for generating Wang tile morphologies at our disposal, we have supplemented our previous comparison of strictly aperiodic and stochastic tile sets [61] and confirmed the superiority of vertex-defined tile sets in suppressing artificial periodicity in assembled microstructural samples. Given the same cardinality of the edge- and vertex-based tile sets in two dimensions, we recommend using the latter. On the other hand, the same comparison in three dimensions is more subtle and the choice should be always a compromise between available computational resources and required periodicity reduction, following e.g. the approximate formula proposed in [58].

Admittedly, the presented modification inherits certain limitations posed by the original framework; namely, it provides only indirect control of a resulting microstructural geometry and no spatial statistics are involved in the particle placing process. Albeit addressing these limitations remained out of the scope of this work, ideas conceptually similar to [10], i.e. optimizing the particle positions a posteriori to minimize the discrepancy between target and computed spatial statistics, could be potentially applied to steer the generated microstructure to the desired statistics.

Albeit our extension features linear complexity as in the original Sonon et al.'s framework, the generalized periodicity brings increased computational cost compared to PUC generation, see Section 4.4. Recall though that tile-based compression is intended primarily for generating numerous stochastic realizations with arbitrary sizes. Thus, more time can be spent on preparing the initial microstructure because the increased cost will be amortized later during repeated use of the same compression.

Our current focus is on developing a robust finite element discretization tool that would enable meshing both outputs of the framework, i.e. analytical geometries of particulate microstructures and complex geometries implicitly defined via level-set fields and processed with the marching-cube algorithm. In both cases, special attention will be paid to ensuring topological and geometrical compatibility of the resulting finite-element meshes across the corresponding faces/edges.

Acknowledgements

References

[1] Nguyen VP, Lloberas-Valls O, Stroeven M, Sluys LJ. Homogenization-based multiscale crack modelling: From micro-diffusive damage to macro-cracks. Comput Methods Appl Mech Engrg 2011;200(9–12):1220–36. http://dx.doi.org/10.1016/j.cma.2010.10.013.

[2] Barchiesi E, Spagnuolo M, Placidi L. Mechanical metamaterials: a state of the art. Math Mech Solids 2019;24(1):212–34. http://dx.doi.org/10.1177/1081286517735695.

[3] Matouš K, Geers MG, Kouznetsova VG, Gillman A. A review of predictive nonlinear theories for multiscale modeling of heterogeneous materials. J Comput Phys 2017;330:192–220. http://dx.doi.org/10.1016/j.jcp.2016.10.070.

[4] Torquato S. Random heterogeneous materials: microstructure and macroscopic properties. Interdisciplinary applied mathematics, vol. 16, New York: Springer; ISBN: 0-387-95167-9, 2002.

[5] Nguyen T, Yvonnet J, Bornert M, Chateau C, Bilteryst F, Steib E. Large-scale simulations of quasi-brittle microcracking in realistic highly heterogeneous microstructures obtained from micro CT imaging. Extreme Mech Lett 2017;17:50–5. http://dx.doi.org/10.1016/j.eml.2017.09.013.

[6] Lloberas-Valls O, Rixen D, Simone A, Sluys L. Domain decomposition techniques for the efficient modeling of brittle heterogeneous materials. Comput Methods Appl Mech Engrg 2011;200(13–16):1577–90. http://dx.doi.org/10.1016/j.cma.2011.01.008.

[7] Akbari Rahimabadi A, Kerfriden P, Bordas S. Scale selection in nonlinear fracture mechanics of heterogeneous materials. Phil Mag 2015;95(28–30):3328–47. http://dx.doi.org/10.1080/14786435.2015.1061716.

[8] Lloberas-Valls O, Rixen D, Simone A, Sluys L. Multiscale domain decomposition analysis of quasi-brittle heterogeneous materials. Internat J Numer Methods Engrg 2012;89(11):1337–66. http://dx.doi.org/10.1002/nme.3286.

[9] Lee H, Brandyberry M, Tudor A, Matouš K. Three-dimensional reconstruction of statistically optimal unit cells of polydisperse particulate composites from microtomography. Phys Rev E 2009;80(6):061301. http://dx.doi.org/10.1103/PhysRevE.80.061301.

[10] Yang M, Nagarajan A, Liang B, Soghrati S. New algorithms for virtual reconstruction of heterogeneous microstructures. Comput Methods Appl Mech Engrg 2018;338:275–98. http://dx.doi.org/10.1016/j.cma.2018.04.030.

[11] Balzani D, Scheunemann L, Brands D, Schröder J. Construction of two- and three-dimensional statistically similar RVEs for coupled micro-macro simulations. Comput Mech 2014;54(5):1269–84. http://dx.doi.org/10.1007/s00466-014-1057-6.

[12] Zeman J, Šejnoha M. From random microstructures to representative volume elements. Modelling Simulation Mater Sci Eng 2007;15(4):S325–35. http://dx.doi.org/10.1088/0965-0393/15/4/S01.

[13] Scheunemann L, Balzani D, Brands D, Schröder J. Design of 3D statistically similar representative volume elements based on Minkowski functionals. Mech Mater 2015;90:185–201. http://dx.doi.org/10.1016/j.mechmat.2015.03.005.

[14] Jiao Y, Stillinger F, Torquato S. Modeling heterogeneous materials via two-point correlation functions: Basic principles. Phys Rev E 2007;76(3):031110. http://dx.doi.org/10.1103/PhysRevE.76.031110.

[15] Zeman J, Šejnoha M. Numerical evaluation of effective elastic properties of graphite fiber tow impregnated by polymer matrix. J Mech Phys Solids 2001;49(1):69–90. http://dx.doi.org/10.1016/S0022-5096(00)00027-2.

[16] Rozman MG, Utz M. Efficient reconstruction of multiphase morphologies from correlation functions. Phys Rev E 2001;63(6):066701. http://dx.doi.org/10.1103/PhysRevE.63.066701.

[17] Jiao Y, Stillinger FH, Torquato S. A superior descriptor of random textures and its predictive capacity. Proc Natl Acad Sci 2009;106(42):17634–9. http://dx.doi.org/10.1073/pnas.0905919106.

[18] Lu B, Torquato S. Lineal-path function for random heterogeneous materials. Phys Rev A 1992;45(2):922–9. http://dx.doi.org/10.1103/PhysRevA.45.922.

[19] Kumar H, Briant C, Curtin W. Using microstructure reconstruction to model mechanical behavior in complex microstructures. Mech Mater 2006;38(8–10):818–32. http://dx.doi.org/10.1016/j.mechmat.2005.06.030.

[20] Havelka J, Kučerová A, Sýkora J. Compression and reconstruction of random microstructures using accelerated lineal path function. Comput Mater Sci 2016;122:102–17. http://dx.doi.org/10.1016/j.commatsci.2016.04.044.

[21] Povirk GL. Incorporation of microstructural information into models of two-phase materials. Acta Metall Mater 1995;43(8):3199–206. http://dx.doi.org/10.1016/0956-7151(94)00487-3.

[22] Yeong C, Torquato S. Reconstructing random media. Phys Rev E 1998;57(1):495–506. http://dx.doi.org/10.1103/PhysRevE.57.495.

[23] Basanta D, Miodownik MA, Holm EA, Bentley PJ. Using genetic algorithms to evolve three-dimensional microstructures from two-dimensional micrographs. Metall Mater Trans A 2005;36(7):1643–52. http://dx.doi.org/10.1007/s11661-005-0026-2.

[24] Kumar NC, Matouš K, Geubelle PH. Reconstruction of periodic unit cells of multimodal random particulate composites using genetic algorithms. Comput Mater Sci 2008;42(2):352–67. http://dx.doi.org/10.1016/j.commatsci.2007.07.043.

[25] Fullwood D, Kalidindi S, Niezgoda S, Fast A, Hampson N. Gradient-based microstructure reconstructions from distributions using fast Fourier transforms. Mater Sci Eng A 2008;494(1–2):68–72. http://dx.doi.org/10.1016/j.msea.2007.10.087.

[26] Fullwood DT, Niezgoda SR, Kalidindi SR. Microstructure reconstructions from 2-point statistics using phase-recovery algorithms. Acta Mater 2008;56(5):942–8. http://dx.doi.org/10.1016/j.actamat.2007.10.044.

[27] Liu X, Shapiro V. Random heterogeneous materials via texture synthesis. Comput Mater Sci 2015;99:177–89. http://dx.doi.org/10.1016/j.commatsci.2014.12.017.

[28] Tahmasebi P, Sahimi M. Cross-correlation function for accurate reconstruction of heterogeneous media. Phys Rev Lett 2013;110(7):078002. http://dx.doi.org/10.1103/PhysRevLett.110.078002.

[29] Bostanabad R, Bui AT, Xie W, Apley DW, Chen W. Stochastic microstructure characterization and reconstruction via supervised learning. Acta Mater 2016;103:89–102. http://dx.doi.org/10.1016/j.actamat.2015.09.044.

[30] Biswal B, Manwart C, Hilfer R, Bakke S, Øren P. Quantitative analysis of experimental and synthetic microstructures for sedimentary rock. Physica A 1999;273(3–4):452–75. http://dx.doi.org/10.1016/S0378-4371(99)00248-4.

[31] Saito Y, Enomoto M. Monte Carlo simulation of grain growth. ISIJ Int 1992;32(3):267–74. http://dx.doi.org/10.2355/isijinternational.32.267.

[32] Krill III C, Chen L-Q. Computer simulation of 3-D grain growth using a phase-field model. Acta Mater 2002;50(12):3059–75. http://dx.doi.org/10.1016/S1359-6454(02)00084-8.

[33] Falco S, Jiang J, De Cola F, Petrinic N. Generation of 3D polycrystalline microstructures with a conditioned Laguerre-Voronoi tessellation technique. Comput Mater Sci 2017;136:20–8. http://dx.doi.org/10.1016/j.commatsci.2017.04.018.

[34] Alsayednoor J, Harrison P. Evaluating the performance of microstructure generation algorithms for 2-d foam-like representative volume elements. Mech Mater 2016;98:44–58. http://dx.doi.org/10.1016/j.mechmat.2016.04.001.

[35] Mitchell DP. Generating antialiased images at low sampling densities. ACM SIGGRAPH Comput Graph 1987;21(4):65–72. http://dx.doi.org/10.1145/37402.37410.

[36] Cooper DW. Random-sequential-packing simulations in three dimensions for spheres. Phys Rev A 1988;38(1):522–4. http://dx.doi.org/10.1103/PhysRevA.38.522.

[37] Segurado J, Llorca J. A numerical approximation to the elastic properties of sphere-reinforced composites. J Mech Phys Solids 2002;50(10):2107–21. http://dx.doi.org/10.1016/S0022-5096(02)00021-2.

[38] Lubachevsky BD, Stillinger FH. Geometric properties of random disk packings. J Stat Phys 1990;60(5–6):561–83. http://dx.doi.org/10.1007/BF01025983.

[39] Donev A, Torquato S, Stillinger FH. Neighbor list collision-driven molecular dynamics simulation for nonspherical hard particles. I. Algorithmic details. J Comput Phys 2005;202(2):737–64. http://dx.doi.org/10.1016/j.jcp.2004.08.014.

[40] Ghossein E, Lévesque M. Random generation of periodic hard ellipsoids based on molecular dynamics: A computationally-efficient algorithm. J Comput Phys 2013;253:471–90. http://dx.doi.org/10.1016/j.jcp.2013.07.004.

[41] De Giorgi M, Carofalo A, Dattoma V, Nobile R, Palano F. Aluminium foams structural modelling. Comput Struct 2010;88(1–2):25–35. http://dx.doi.org/10.1016/j.compstruc.2009.06.005.

[42] Chen Y, Das R, Battley M. Effects of cell size and cell wall thickness variations on the stiffness of closed-cell foams. Int J Solids Struct 2015;52:150–64. http://dx.doi.org/10.1016/j.ijsolstr.2014.09.022.

[43] Mebatsion H, Verboven P, Verlinden B, Ho Q, Nguyen T, Nicolaï B. Microscale modelling of fruit tissue using Voronoi tessellations. Comput Electron Agric 2006;52(1–2):36–48. http://dx.doi.org/10.1016/j.compag.2006.01.002.

[44] Chakraborty A, Perales MM, Reddy GV, Roy-Chowdhury AK. Adaptive geometric tessellation for 3D reconstruction of anisotropically developing cells in multilayer tissues from sparse volumetric microscopy images. In: Merks RM, editor. PLoS ONE 2013;8(8). e67202. http://dx.doi.org/10.1371/journal.pone.0067202.

[45] Doškář M, Novák J. A jigsaw puzzle framework for homogenization of high porosity foams. Comput Struct 2016;166:33–41. http://dx.doi.org/10.1016/j.compstruc.2016.01.003.

[46] Simone A, Gibson L. The effects of cell face curvature and corrugations on the stiffness and strength of metallic foams. Acta Mater 1998;46(11):3929–35. http://dx.doi.org/10.1016/S1359-6454(98)00072-X.

[47] Lyckegaard A, Lauridsen EM, Ludwig W, Fonda RW, Poulsen HF. On the use of Laguerre tessellations for representations of 3D grain structures. Adv Energy Mater 2011;13(3):165–70. http://dx.doi.org/10.1002/adem.201000258.

[48] Dunbar D, Humphreys G. A spatial data structure for fast Poisson-disk sample generation. ACM Trans Graph 2006;25(3):503. http://dx.doi.org/10.1145/1141911.1141915.

[49] Jones TR. Efficient generation of Poisson-disk sampling patterns. J Graph Tools 2006;11(2):27–36. http://dx.doi.org/10.1080/2151237X.2006.10129217.

[50] Sonon B, François B, Massart T. A unified level set based methodology for fast generation of complex microstructural multi-phase RVEs. Comput Methods Appl Mech Engrg 2012;223–224:103–22. http://dx.doi.org/10.1016/j.cma.2012.02.018.

[51] Sonon B, Massart T. A level-set based representative volume element generator and XFEM simulations for textile and 3D-reinforced composites. Materials 2013;6(12):5568–92. http://dx.doi.org/10.3390/ma6125568.

[52] Sonon B, François B, Massart TJ. An advanced approach for the generation of complex cellular material representative volume elements using distance fields and level sets. Comput Mech 2015;56(2):221–42. http://dx.doi.org/10.1007/s00466-015-1168-8.

[53] Massart TJ, Sonon B, Ehab Moustafa Kamel K, Poh LH, Sun G. Level set-based generation of representative volume elements for the damage analysis of irregular masonry. Meccanica 2018;53(7):1737–55. http://dx.doi.org/10.1007/s11012-017-0695-0.

[54] Kanit T, Forest S, Galliet I, Mounoury V, Jeulin D. Determination of the size of the representative volume element for random composites: statistical and numerical approach. Int J Solids Struct 2003;40(13–14):3647–79. http://dx.doi.org/10.1016/S0020-7683(03)00143-4.

[55] Gitman I, Askes H, Sluys L. Representative volume: Existence and size determination. Eng Fract Mech 2007;74(16):2518–34. http://dx.doi.org/10.1016/j.engfracmech.2006.12.021.

[56] Dirrenberger J, Forest S, Jeulin D. Towards gigantic RVE sizes for 3D stochastic fibrous networks. Int J Solids Struct 2014;51(2):359–76. http://dx.doi.org/10.1016/j.ijsolstr.2013.10.011.

[57] Cohen MF, Shade J, Hiller S, Deussen O. Wang tiles for image and texture generation. ACM Trans Graph 2003;22(3):287–94. http://dx.doi.org/10.1145/882262.882265.

[58] Novák J, Kučerová A, Zeman J. Compressing random microstructures via stochastic Wang tilings. Phys Rev E 2012;86(4):040104. http://dx.doi.org/10.1103/PhysRevE.86.040104.

[59] Doškář M, Zeman J, Jarušková D, Novák J. Wang tiling aided statistical determination of the representative volume element size of random heterogeneous materials. Eur J Mech A Solids 2018;70:280–95. http://dx.doi.org/10.1016/j.euromechsol.2017.12.002.

[60] Novák J, Kučerová A, Zeman J. Microstructural enrichment functions based on stochastic Wang tilings. Modelling Simulation Mater Sci Eng 2013;21(2):025014. http://dx.doi.org/10.1088/0965-0393/21/2/025014.

[61] Doškář M, Novák J, Zeman J. Aperiodic compression and reconstruction of real-world material systems based on Wang tiles. Phys Rev E 2014;90(6):062118. http://dx.doi.org/10.1103/PhysRevE.90.062118.

[62] Robinson RM. Undecidability and nonperiodicity for tilings of the plane. Invent Math 1971;12(3):177–209. http://dx.doi.org/10.1007/BF01418780.

[63] Wang H. Proving theorems by pattern recognition - II. Bell Syst Tech J 1961;40(1):1–41. http://dx.doi.org/10.1002/j.1538-7305.1961.tb03975.x.

[64] Wang H. Games, logic and computers. Sci Am 1965;213(5):98–106.

[65] Berger R. The undecidability of the domino problem. Providence, R.I.: American Mathematical Society; ISBN: 978-1-4704-0013-2, 1966, 1-4704-0013-8.

[66] Grünbaum B, Shephard GC. Tilings and patterns: Second edition. Mineola, New York: Dover Publications, Inc; ISBN: 978-0-486-46981-2, 2016.

[67] Jeandel E, Rao M. An aperiodic set of 11 Wang tiles. 2015, arXiv:1506.06492.

[68] Ammann R, Grünbaum B, Shephard GC. Aperiodic tiles. Discrete Comput Geom 1992;8(1):1–25. http://dx.doi.org/10.1007/BF02293033.

[69] Rothemund PWK, Papadakis N, Winfree E. Algorithmic self-assembly of DNA Sierpinski triangles. PLoS Biol 2004;2(12). e424. http://dx.doi.org/10.1371/journal.pbio.0020424.

[70] Leuzzi L, Parisi G. Thermodynamics of a tiling model. J Phys A: Math Gen 2000;33(23):4215–25. http://dx.doi.org/10.1088/0305-4470/33/23/301.

[71] Kopf J, Cohen-Or D, Deussen O, Lischinski D. Recursive Wang tiles for real-time blue noise. ACM Trans Graph 2006;25(3):509. http://dx.doi.org/10.1145/1141911.1141916.

[72] Kari J, Culik K. An aperiodic set of Wang cubes. JUCS - J Univ Comput Sci 1995;1(10):675–86. http://dx.doi.org/10.3217/jucs-001-10-0675.

[73] Lagae A, Dutré P. An alternative for Wang tiles: colored edges versus colored corners. ACM Trans Graph 2006;25(4):1442–59. http://dx.doi.org/10.1145/1183287.1183296.

[74] Sedgewick R. Algorithms in C++, part 5: Graph algorithms. 3rd ed.. Boston: Addison-Wesley; ISBN: 978-0-201-36118-6, 2002.

[75] Talbot J, Schaaf P, Tarjus G. Random sequential addition of hard spheres. Mol Phys 1991;72(6):1397–406. http://dx.doi.org/10.1080/00268979100100981.

[76] Osher S, Fedkiw RP. Level set methods and dynamic implicit surfaces. New York: Springer; ISBN: 978-0-387-22746-7, 2003.

[77] Lagae A, Dutré P. Poisson sphere distributions. In: Kobbelt L, Kuhlen T, Aach T, Westermann R, editors. Vision, modeling, and visualization 2006. Berlin: Akademische Verlagsgesellschaft Aka GmbH; ISBN: 3-89838-081-5, 2006, p. 373–9, 1-58603-688-2.

[78] Stam J. Aperiodic texture mapping. Tech. Rep. R046, 1997, p. 1–9, URL: www.autodeskresearch.com/publications/aperiodictexture.

[79] Gajdošík J, Zeman J, Šejnoha M. Qualitative analysis of fiber composite microstructure: Influence of boundary conditions. Probab Eng Mech 2006;21(4):317–29. http://dx.doi.org/10.1016/j.probengmech.2005.11.006.