THEORY OF COMPUTING SYSTEMS, VOL. 33, PP. 295–312, 2000

On Multi-Dimensional Hilbert Indexings*

Jochen Alber Rolf Niedermeier†

Wilhelm-Schickard-Institut für Informatik, Universität Tübingen,

Sand 13, D-72076 Tübingen, Germany,

alber/niedermr@informatik.uni-tuebingen.de

Address for correspondence:

Jochen Alber
Universität Tübingen
Wilhelm-Schickard-Institut für Informatik
Theoretische Informatik/ Formale Sprachen
Sand 13
D-72076 Tübingen
Fed. Rep. of Germany

Email: alber@informatik.uni-tuebingen.de

Telephone: +49-7071-29-77568
Fax: +49-7071-29-5061

*An extended abstract of this work appeared in the proceedings of the Fourth Annual International Computing and Combinatorics Conference (COCOON'98), Springer, *LNCS* 1449, pages 329–338, held in Taipei, Taiwan, August 12–14, 1998.

†Work partially supported by a Feodor Lynen fellowship of the Alexander von Humboldt-Stiftung, Bonn, and the Center for Discrete Mathematics, Theoretical Computer Science, and Applications (DIMATIA), Prague.

Abstract

Indexing schemes for grids based on space-filling curves (e.g., Hilbert indexings) find applications in numerous fields, ranging from parallel processing over data structures to image processing. Because of an increasing interest in discrete multi-dimensional spaces, indexing schemes for them have won considerable interest. Hilbert curves are the most simple and popular space-filling indexing scheme. We extend the concept of curves with Hilbert property to arbitrary dimensions and present first results concerning their structural analysis that also simplify their applicability.

We define and analyze in a precise mathematical way r-dimensional Hilbert indexings for arbitrary $r \geq 2$. Moreover, we generalize and simplify previous work and clarify the concept of Hilbert curves for multi-dimensional grids. As we show, Hilbert indexings can be completely described and analyzed by "generating elements of order 1", thus, in comparison with previous work, reducing their structural complexity decisively. Whereas there is basically *one* Hilbert curve in the 2D world, our analysis shows that there are 1536 structurally different 3D Hilbert curves. Further results include generalizations of locality results for multi-dimensional indexings and an easy recursive computation scheme for multi-dimensional Hilbert indexings.

1 Introduction

Discrete multi-dimensional spaces are of increasing importance in computer science. They appear in various settings such as combinatorial optimization, parallel processing, image processing, geographic information systems, data base systems, and data structures. For many applications it is necessary to number the points of a discrete multi-dimensional space (which, equivalently, can be seen as a grid) by an indexing scheme mapping each point bijectively to a natural number in the range between 1 and the total number of points in the space. Often it is desirable that this indexing scheme preserves some kind of locality, that is, close-by points in the space are mapped to close-by numbers or vice versa. For this purpose, indexing schemes based on space-filling curves have shown to be of high value [2, 4, 5, 6, 8, 7, 9, 11, 12, 13, 14, 15, 16, 19].

In this paper, we study Hilbert indexings [10], perhaps the most popular space-filling indexing schemes. Properties of 2D and 3D Hilbert indexings have been extensively studied recently [5, 6, 7, 9, 12, 14, 15, 17]. However, most of the work so far has focused on empirical studies. Up to now, little attention has been paid to the theoretical study of structural properties of multi-dimensional Hilbert curves, the focus of this paper. Whereas with "modulo symmetry" there is only one 2D Hilbert curve, there are many possibilities to define Hilbert curves in the 3D setting [5, 15]. The advantage of Hilbert curves is their (compared to other curves) simple structure that may easily outweigh the asymptotically slightly better (concerning constant factors) locality properties of other space-filling curves. Also note that in defining indexing schemes for multi-dimensional grids, descriptional simplicity as provided by "pure" Hilbert indexing is a desirable property.

Our results can shortly be sketched as follows. We generalize the notion of Hilbert indexings to arbitrary dimensions. We clarify the concept of Hilbert curves in multi-dimensional spaces by providing a natural and simple mathematical formalism that allows combinatorial studies of multi-dimensional Hilbert indexings. For reasons of (geometrical) clearness, we base our formalism on permutations instead of e.g. matrices or other formalisms [3, 4, 5, 17]. So we obtain the following insight: Space-filling curves with Hilbert property can be completely described by simple generating elements and permutations operating on them. Structural questions for Hilbert curves in arbitrary dimensions can be decided by reducing them to basic generating elements. Putting it in catchy terms, one might say that for Hilbert indexings what holds "in the large" (i.e., for large side-length), can already be detected "in the small" (i.e., for side-length 2). In particular, this provides a basis for mechanized proofs of locality of curves with Hilbert property (cf. [15]). In addition, this observation allows the identification of seemingly different 3D Hilbert indexings [5], the generalization of a locality result of Gotsman and Lindenbaum [9] to a larger class of multi-dimensional indexing schemes, and the determination that there are exactly $6 \cdot 2^8 = 1536$ structurally different 3D

Hilbert curves. The latter clearly generalizes and answers Sagan's quest for describing 3D Hilbert curves [17]. Finally, we provide an easy recursive formula for computing Hilbert indexings in arbitrary dimensions and sketch a recipe for how to construct an r-dimensional Hilbert curve for arbitrary r in an easy way from two $(r-1)$-dimensional ones.

As a whole, our work lays foundations for future work dealing with combinatorial properties of multi-dimensional Hilbert curves and, in particular, a mechanized analysis of locality properties of multi-dimensional Hilbert curves. The main focus of this paper, however, is to provide a theoretical study of nice combinatorial properties of Hilbert curves in arbitrary dimensions and it is not to study e.g. locality properties in great depth, which may be the subject of future study.

The paper is organized as follows. Section 2 presents some basic facts on space-filling curves and grid indexings and, in particular, gives the construction scheme of 2D Hilbert curves. Section 3 contains our method to describe multi-dimensional Hilbert indexings by "generators" and permutations operating on a given corner-indexing of a cube. One of our main results shows that the structural analysis of multi-dimensional Hilbert curves can be completely reduced to the analysis of their (small) generating elements. In Section 4 we apply the methodology of Section 3 to derive several results concerning the structural analysis and computation of curves with Hilbert property. Finally, Section 5 draws some conclusions, outlines further generalizations, and gives some directions for future work.

2 Preliminaries

We focus our attention on cubic grids, where, in the r-dimensional case, we have n^r points arranged in an r-dimensional grid with side-length n. An r-dimensional *(discrete) curve* C is simply a bijective mapping $C: \{1, \ldots n^r\} \rightarrow \{1, \ldots, n\}^r$, thus providing a total ordering of the grid points. Note that, by definition, we do not claim the continuity of a curve. A curve C is called *continuous* if it forms a Hamilton path through the n^r grid points. An r-dimensional cubic grid is said to be of *order* k if it has side-length 2^k. Analogously, a curve C has order k if its range is a cubic grid of order k.

Fig. 1 shows the smallest 2D continuous curve indexing a grid of size 4. This curve can be found in Hilbert's original work [10, 18] as a constructing unit for a whole family of curves. Fig. 2 shows the general construction principle for these so-called Hilbert curves: For any $k \geq 1$ four Hilbert indexings of size 4^k are combined into an indexing of size 4^{k+1} by rotating and reflecting them in such a way that concatenating the indexings yields a Hamilton path through the grid. Note that the left and the right side of the curve are symmetric to each other. Thus, as indicated in Fig. 2, we only need to keep track of the orientation of the

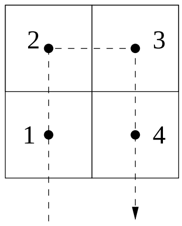

Figure 1: The generator Hil_1^2 and its canonical corner-indexing $\widetilde{\mathrm{Hil}}_1^2$.

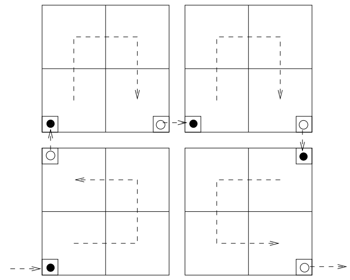

Figure 2: Construction scheme for the 2D Hilbert indexing.

edge which contains the start and end of the curve. As we will see later on, the above rule uniquely defines the 2D Hilbert indexing up to global rotation and reflection.

One of the main features of the Hilbert curve is its "self-similarity". Here "self-similar" shall simply mean that the curve can be generated by putting together identical (basic construction) units, only applying rotation and reflection to these units. In a sense, the Hilbert curve is the "simplest" self-similar, recursive, locality-preserving indexing scheme for square meshes of size $2^k \times 2^k$.

3 Formalizing Hilbert curves in r dimensions

In this section, we generalize the construction principle of 2D Hilbert curves to arbitrary dimensions in a rigorous, mathematically precise way. We restrict attention to indexing schemes of cubes with side-lengths 2^k for any natural number k, although generalizations are straightforward (see Section 5). We generate an r-dimensional curve filling a cubic grid with side-length 2^k with a sequence of 2^r subcurves filling grids with side-length 2^{k-1} each. For the generating subcurves we claim a certain similarity as given by the 2D Hilbert indexing. By "similar" we mean that the subcurves can be transformed by a symmetry mapping (reflection or rotation) into each other. We need a certain formalism to express these symmetry mappings. This, for example, can be done by means of permutations. Fixing a certain indexing of the corners in a multi-dimensional grid, such a symmetry transformation can be expressed by the action of a permutation on the given indexing. This is one of the most intuitive approaches to describe such automorphisms on the grid. Furthermore, there turns out to be a very simple relation between curves of the lowest order possible and such corner-indexings. We strongly believe that this at first sight maybe strange formalism used by us was the basis for deriving structural results on e.g. 3D Hilbert indexings as presented in Section 4. So we can hardly imagine a comparatively simple presentation of all structurally different 3D Hilbert curves as given in Table 1 (see Subsection 4.1) using other formalisms.

3.1 Classes of Self-Similar Curves and their generators

Let $V_r := \{x_1x_2\cdots x_{r-1}x_r \mid x_i \in \{0,1\}\ \}$ be the set of all 2^r corners of an r-dimensional cube coded in binary. Moreover, let $\mathcal{I} : V_r \longrightarrow \{1,\ldots,2^r\}$ denote an arbitrary indexing of these corners. To describe the orientation of subcurves inside a curve of higher order, we want to use symmetry mappings, which can be expressed via suitable permutations operating on such corner-indexings. Observe that any r-dimensional curve C_1 of order 1 naturally induces an indexing of these corners (see Fig. 1 and Fig. 3). We call the obtained corner-indexing the *canonical* one and denote it by $\widetilde{C_1} : V_r \longrightarrow \{1,\ldots,2^r\}$. Furthermore, let $W_{\mathcal{I}}$ $\big(\subset Sym(2^r)\big)$ denote the group of all permutations (operating on $\mathcal{I}$) that describe rotations and reflections of the r-dimensional cube. In other words, $W_{\mathcal{I}}$ is the set of all permutations that preserve the neighborhood-relations $n(i,j)$ of the corner indexing $\mathcal{I}$:

$$W_{\mathcal{I}} := \big\{\pi \in Sym(2^r)\ :\ n(i,j) = n(\pi(i),\pi(j))\ \ \forall i,j \in \{1,\ldots,2^r\}\big\}.$$

For a given permutation $\tau \in W_{\mathcal{I}}$, we sometimes write $(\tau : \mathcal{I})$ in order to emphasize that τ is operating on a cube with corner-indexing $\mathcal{I}$. The point here is that once we have fixed a certain corner-indexing $\mathcal{I}$, the set $W_{\mathcal{I}}$ will provide all necessary transformations to describe a construction principle of how to generate curves of higher order by piecing together a suitable curve of lower order. Obviously each permutation $(\tau : \mathcal{I})$ acting on a given corner-indexing $\mathcal{I}$ canonically induces a bijective mapping on a cubic grid of order k. In the following we do not distinguish between a permutation and the corresponding mapping on a grid.

We partition an r-dimensional cubic grid of order k into 2^r subcubes of order $k-1$. For each $x_1\cdots x_r \in V_r$ we therefore set

$$p^{(k)}_{(x_1\cdots x_r)} := (x_1\cdot 2^{k-1},\ldots,x_r\cdot 2^{k-1}) \in \{0,\ldots,2^k-1\}\times\ldots\times\{0,\ldots,2^k-1\}$$

to be the "lower-left corner" of such a subcube. Let C_{k-1} be an r-dimensional curve of order $k-1$ $(k\geq 2)$. Our goal is to define a "self-similar" curve C_k of order k by putting together 2^r pieces of type C_{k-1}. Let $\mathcal{I} : V_r \longrightarrow \{1,\ldots,2^r\}$ be a corner-indexing. We intend to arrange the 2^r subcurves of type C_{k-1} "along" $\mathcal{I}$. The position of the i'-th (where $i' \in \{1,\ldots,2^r\}$) subcurve inside C_k can formally be described with the help of the grid-points $p^{(k)}_{(x_1\cdots x_r)}$. Bearing in mind the classical construction principle for the 2D Hilbert indexing, the orientation of the constructing curve C_{k-1} inside C_k can be expressed by using symmetric transformations (that is reflections and rotations). For any sequence of permutations $\tau_1,\ldots,\tau_{2^r} \in W_{\mathcal{I}}$ we therefore define

$$C_k(i) := (\tau_{i'} : \mathcal{I}) \circ C_{k-1}\big(i \bmod (2^{k-1})^r\big) + p^{(k)}_{\mathcal{I}^{-1}(i')}, \tag{1}$$

where $i \in \{1,\ldots,(2^k)^r\}$ and $i' = (i-1)$ div $(2^{k-1})^r + 1$. The geometric intuition behind is that the curve C_k can be partitioned into 2^r components of the form C_{k-1} (reflected or rotated in a suitable way). These subcurves are arranged inside C_k "along" the given corner-indexing $\mathcal{I}$. The orientation of the i'-th subcurve inside C_k is described by the effect of $\tau_{i'}$ operating on $\mathcal{I}$.

Definition 1. Whenever two r-dimensional curves C_{k-1} of order $k-1$ and C_k of order k satisfy equation (1) for a given sequence of permutations $\tau_1, \dots, \tau_{2^r} \in W_{\mathcal{I}}$ (operating on the corner-indexing $\mathcal{I} : V_r \longrightarrow \{1, \dots, 2^r\}$), we write

$$C_{k-1} \;\; {}_{(\tau_1,\dots,\tau_{2^r})}\!\!\overset{\mathcal{I}}{\ll} \;\; C_k$$

and call C_{k-1} the *constructor* of C_k.

Our final goal is to iterate this process starting with a curve C_1 of order 1. It's only natural and in our opinion "preserves the spirit of Hilbert" to fix the corner-indexing according to the structure of the defining curve C_1. Hence, in this situation we can specify our $\mathcal{I}$ to be the canonical corner-indexing $\widetilde{C_1}$. By successively repeating the construction principle in equation (1) k times, we obtain a curve of order k.

Definition 2. Let $\mathcal{C} = \{\, C_k \;\mid\; k \geq 1 \,\}$ be a family of r-dimensional curves of order k. We call $\mathcal{C}$ a *Class of Self-Similar Curves (CSSC)* if there exists a sequence of permutations $\tau_1, \dots, \tau_{2^r} \in W_{\widetilde{C_1}}$ (operating on the canonical corner-indexing $\widetilde{C_1}$) such that for each curve C_k it holds that

$$C_1 \;\; {}_{(\tau_1,\dots,\tau_{2^r})}\!\!\overset{\widetilde{C_1}}{\ll} \;\; C_2 \;\; {}_{(\tau_1,\dots,\tau_{2^r})}\!\!\overset{\widetilde{C_1}}{\ll} \;\; \cdots \;\; {}_{(\tau_1,\dots,\tau_{2^r})}\!\!\overset{\widetilde{C_1}}{\ll} \;\; C_{k-1} \;\; {}_{(\tau_1,\dots,\tau_{2^r})}\!\!\overset{\widetilde{C_1}}{\ll} \;\; C_k.$$

In this case, C_1 is called the *generator of the CSSC* $\mathcal{C}$ and we set

$$\mathcal{H}\big(\, C_1,\, (\tau_1, \dots, \tau_{2^r})\,\big) := \{\, C_k \;\mid\; k \geq 1 \,\}$$

as the CSSC generated by C_1 and $\tau_1, \dots, \tau_{2^r}$. A CSSC $\mathcal{C} = \{\, C_k \;\mid\; k \geq 1 \,\}$ is called *Class with Hilbert Property (CHP)* if all curves C_k are continuous.

Note that the CSSC $\mathcal{H}\big(\, C_1,\, (\tau_1, \dots, \tau_{2^r})\,\big)$ is well-defined, because any CSSC is uniquely determined by its generator C_1 and the choice of the permutations $\tau_1, \dots, \tau_{2^r} \in W_{\widetilde{C_1}}$. The nomenclature "Curve with Hilbert Property" is due to the fact that the constructing principle for a CHP grew out of the classical one for 2D Hilbert curves. Our concept for multi-dimensional CHPs only makes use of the very essential tools which can be found in Hilbert's context (cf. [10]) as rotation and reflection. We deliberately avoid more complicated structures (e.g., the use of different sequences of permutations in each inductive step, or the use of several generators for the constructing principle) in order to maintain conceptual simplicity and ease of construction and analysis. However, the theory which we develop in this paper doesn't necessarily restrict to the continuous case. That is the reason why all our definitions and theorems in Subsection 3.2 are held in the more general setting of non-continuous curves. In Subsection 3.2 we provide a necessary and sufficient condition on the generating elements of a CSSC (generator and sequence of permutations) such that the whole family consists of continuous curves only, i.e., is a CHP. We end this subsection with an example.

Example. One easily checks that the classical 2D Hilbert indexing can be described via

$$\mathcal{H}\big(\,\mathrm{Hil}_1^2,\ ((2\,4),\mathrm{id},\mathrm{id},(1\,3))\,\big) = \{\,\mathrm{Hil}_k^2 \ \mid\ k \geq 1\,\},$$

where the generator Hil_1^2 is given in Fig. 1.

As Theorem 4 will show, this is the only CHP of dimension 2 "modulo symmetry," which, once again, justifies the naming "Curve with Hilbert Property".

3.2 Disturbing the generator of a CSSC

In this subsection, we analyze the effects of disturbing the generator of a CSSC by a symmetric mapping. We will see that any disturbance of the generator will be hereditary to the whole CSSC in a very canonical way. And also the other way round: if two different CSSCs show a certain similarity in one of their members, this similarity can already be found in the structure of the corresponding generators. We illustrate this by the following diagram. Given two CSSCs $\mathcal{H}\big(\,C_1,\ (\tau_1,\ldots,\tau_{2^r})\,\big) = \{C_k \mid k \geq 1\}$ and $\mathcal{H}\big(\,D_1,\ (\tau_1,\ldots,\tau_{2^r})\,\big) = \{D_k \mid k \geq 1\}$, respectively.[1] Suppose there is a similarity at a certain stage of the construction, i.e., for some k_0 the curves C_{k_0} and D_{k_0} can be obtained from each other by a similarity transformation Φ. Can we conclude a vertical link between the curves of other orders? The investigations in this section will show that the inner structure of CSSCs are strong enough to yield the same behavior at the stage of any order. As a consequence, it will be sufficient to analyze the generating elements of a CSSC. Since all the information is encoded in the generator and the defining permutations, questions like continuity of a CSSC, structural similarity with other CSSCs can be answered by considering the generating elements only.

$$\begin{array}{ccccccc}
C_1 & \overset{\widetilde{C_1}}{\underset{(\tau_1,\ldots,\tau_{2^r})}{}} \ll & C_2 & \overset{\widetilde{C_1}}{\underset{(\tau_1,\ldots,\tau_{2^r})}{}} \ll & \cdots & \overset{\widetilde{C_1}}{\underset{(\tau_1,\ldots,\tau_{2^r})}{}} \ll & C_k \\
\Big\downarrow \Phi & \circlearrowright & \Big\downarrow \Phi & \circlearrowright & \Big\downarrow \Phi & \circlearrowright & \Big\downarrow \Phi \\
D_1 & \overset{\widetilde{D_1}}{\underset{(\tau_1,\ldots,\tau_{2^r})}{}} \ll & D_2 & \overset{\widetilde{D_1}}{\underset{(\tau_1,\ldots,\tau_{2^r})}{}} \ll & \cdots & \overset{\widetilde{D_1}}{\underset{(\tau_1,\ldots,\tau_{2^r})}{}} \ll & D_k
\end{array}$$

[1]Note that the τ's used in the definition of both CSSCs yield completely different automorphisms on the grid. Whereas in the first case they refer to the corner-indexing $\widetilde{C_1}$, in the second case they act on the corner-indexing $\widetilde{D_1}$, given by the generator D_1.

We split the proof of the main theorem of this section into several steps, since each of these already contains some nice structural behavior of CSSCs. As a first step, we make a simple observation concerning the behavior of the construction principle of Definition 1 under the "symmetric disturbance" of a constructor:

Lemma 1. *Let C_{k-1} and C_k be curves of order $k-1$ and k, respectively. Suppose C_{k-1} is the constructor of C_k, i.e., $C_{k-1} \;{}_{(\tau_1,\dots,\tau_{2^r})}\!\overset{\mathcal{I}}{\ll}\; C_k$, for any sequence of permutations $\tau_1,\dots,\tau_{2^r} \in W_{\mathcal{I}}$ (acting on a given corner-indexing $\mathcal{I}$). Then for arbitrary $\phi \in W_{\mathcal{I}}$ we have*

$$(\phi:\mathcal{I}) \circ C_{k-1} \;{}_{(\tau_1\circ\phi^{-1},\dots,\tau_{2^r}\circ\phi^{-1})}\!\overset{\mathcal{I}}{\ll}\; C_k.$$

Proof. Since C_{k-1} is the constructor of C_k, by Definition 1 we have:

$$\begin{aligned}
C_k(i) &= (\tau_{i'}:\mathcal{I}) \circ C_{k-1}\left(i \bmod (2^{k-1})^r\right) + p^{(k)}_{\mathcal{I}^{-1}(i')}\\
&= (\tau_{i'}:\mathcal{I}) \circ (\phi^{-1}:\mathcal{I}) \circ (\phi:\mathcal{I}) \circ C_{k-1}\left(i \bmod (2^{k-1})^r\right) + p^{(k)}_{\mathcal{I}^{-1}(i')}\\
&= (\tau_{i'}\circ\phi^{-1}:\mathcal{I}) \circ \Big((\phi:\mathcal{I}) \circ C_{k-1}\Big)\left(i \bmod (2^{k-1})^r\right) + p^{(k)}_{\mathcal{I}^{-1}(i')}\,,
\end{aligned}$$

where $i \in \{1,\dots,(2^k)^r\}$ and $i' = (i-1) \operatorname{div} (2^{k-1})^r + 1$, proving the claim by Definition 1. □

Whereas, by Lemma 1, we investigated the influence of disturbing the constructor, we now, in a second step, analyze how transforming the underlying corner-indexing influences the construction principle. We will need such a result, since two different CSSCs (by definition) come up with two different corner-indexings, each of which given by the underlying generator.

Lemma 2. *Given the assumptions of Lemma 1 (that is: $C_{k-1} \;{}_{(\tau_1,\dots,\tau_{2^r})}\!\overset{\mathcal{I}}{\ll}\; C_k$ for two curves C_{k-1} and C_k of successive order), then for arbitrary $\phi \in W_{\mathcal{I}}$ and the modified corner-indexing $\mathcal{K} := \phi^{-1}\circ\mathcal{I}$ with $\Phi = (\phi:\mathcal{I}) = (\phi:\mathcal{K})$ we have*[2]

$$C_{k-1} \;{}_{(\tau_1\circ\phi,\dots,\tau_{2^r}\circ\phi)}\!\overset{\mathcal{K}}{\ll}\; \Phi \circ C_k\,.$$

Proof. First we deduce a simple transformation-rule for permutations out of our given relation $\mathcal{K} = \phi^{-1}\circ\mathcal{I}$. The effect of a given permutation $\pi \in W_{\mathcal{I}}$ acting on $\mathcal{I}$ is equivalent to the effect of the transformed permutation $\phi^{-1}\circ\pi\circ\phi$ operating on the transformed corner-indexing $\mathcal{K}$, i.e. $(\pi:\mathcal{I}) = (\phi^{-1}\circ\pi\circ\phi:\mathcal{K})$. Setting $\pi = \phi$, this particularly shows $(\phi:\mathcal{I}) = (\phi:\mathcal{K}) = \Phi$.

By assumption, C_{k-1} is the constructor of C_k which for all $i \in \{1,\dots,(2^k)^r\}$ and $i' = (i-1) \operatorname{div} (2^{k-1})^r + 1$ yields

$$\begin{aligned}
C_k(i) &= (\tau_{i'}:\mathcal{I}) \circ C_{k-1}\left(i \bmod (2^{k-1})^r\right) + p^{(k)}_{\mathcal{I}^{-1}(i')}\\
\Rightarrow\quad \Phi\circ C_k(i) &= \Phi\circ\Big((\tau_{i'}:\mathcal{I}) \circ C_{k-1}\left(i \bmod (2^{k-1})^r\right) + p^{(k)}_{\mathcal{I}^{-1}(i')}\Big)\\
&= \Phi\circ(\tau_{i'}:\mathcal{I}) \circ C_{k-1}\left(i \bmod (2^{k-1})^r\right) + p^{(k)}_{\mathcal{I}^{-1}(\phi(i'))}
\end{aligned}$$

[2]The fact that the corner-indexing is disturbed by ϕ^{-1} instead of ϕ is due to technical reasons only.

where the last equation is true, because the effect of the symmetry mapping Φ on a CSSC-curve C_k of order k can be split into its effect on the 2^r subcurves of order $k-1$ and the effect on the arrangement of these subcurves inside C_k. Whereas the i'-th subcurve of C_k lies next to the corner $\mathcal{I}^{-1}(i')$, the position of the i'-th subcurve of $\Phi \circ C_k$ is transformed according to ϕ. Therefore its new position is given by the corner $\mathcal{I}^{-1}(\phi(i'))$. Thus,

$$\begin{array}{rclcccc} \Phi \circ C_k(i) & = & (\phi \circ \tau_{i'} : \mathcal{I}) & \circ & C_{k-1}\big(i \bmod (2^{k-1})^r\big) & + & p^{(k)}_{(\mathcal{I}^{-1}\circ\phi)(i')} \\ & = & (\tau_{i'} \circ \phi : \mathcal{K}) & \circ & C_{k-1}\big(i \bmod (2^{k-1})^r\big) & + & p^{(k)}_{\mathcal{K}^{-1}(i')}, \end{array}$$

by applying the transformation-rule treated at the beginning with $\pi = \phi \circ \tau_{i'}$. By Definition 1 the last equation proves our claim. $\square$

Lemma 1 and 2 now allow the proof of the main result of this section. For its illustration we refer to the diagram at the beginning of this section. Do also recall the point made in the footnote there.

Theorem 3. *Let C_1 be the generator of the CSSC $\mathcal{H}\big(C_1, (\tau_1, \ldots, \tau_{2^r})\big) = \{C_k \mid k \geq 1\}$ and D_1 the generator of the CSSC $\mathcal{H}\big(D_1, (\tau_1, \ldots, \tau_{2^r})\big) = \{D_k \mid k \geq 1\}$. For an arbitrary permutation $\phi \in W_{\widetilde{C_1}}$ and the corresponding symmetric mapping $\Phi = (\phi : \widetilde{C_1}) = (\phi : \widetilde{D_1})$, the following statements are equivalent:*

(i) $\Phi \circ C_{k_0} = D_{k_0}$ *for some* $k_0 \geq 1$.

(ii) $\Phi \circ C_k = D_k$ *for all* $k \geq 1$.

Proof. (ii) $\Rightarrow$ (i) is trivial. For (i) $\Rightarrow$ (ii) we first show that statement (ii) is true for the generators C_1 and D_1: If $k_0 > 1$ we can divide the cubic grid of order k_0 into 2^r subgrids of order $k_0 - 1$. By the construction principle for CSSCs, the curves C_{k_0} and D_{k_0} traverse these subgrids "along" the canonical corner-indexings $\widetilde{C_1}$ resp. $\widetilde{D_1}$. Since, by assumption, $\Phi \circ C_{k_0} = D_{k_0}$, the corresponding relation also holds true for the corner-indexings $\widetilde{C_1}$ and $\widetilde{D_1}$, which finally yields the validity of the equation $\Phi \circ C_1 = D_1$, because of the isomorphisms $C_1 \simeq \widetilde{C_1}$ resp. $D_1 \simeq \widetilde{D_1}$.

We proceed proving (ii) by induction on k. Assuming that $D_k = \Phi \circ C_k$ we show this relation for $k+1$ by applying Lemma 1 and Lemma 2. Since $\{C_k \mid k \geq 1\}$ is a CSSC, we get

$$\begin{array}{ll} & C_k \;\; {}_{(\tau_1,\ldots,\tau_{2^r})}^{\quad\widetilde{C_1}}\!\ll \;\; C_{k+1} \\ \overset{\text{Lemma 1}}{\Longrightarrow} & \underbrace{\Phi \circ C_k}_{=D_k} \;\; {}_{(\tau_1\circ\phi^{-1},\ldots,\tau_{2^r}\circ\phi^{-1})}^{\qquad\qquad\widetilde{C_1}}\!\ll \;\; C_{k+1} \\ \overset{\text{Lemma 2}}{\Longrightarrow} & D_k \;\; {}_{(\tau_1,\ldots,\tau_{2^r})}^{\quad\widetilde{D_1}}\!\ll \;\; \Phi \circ C_{k+1}\ , \end{array}$$

where the last relation makes use of $\widetilde{D_1} = \phi^{-1} \circ \widetilde{C_1}$, which we immediately obtain from the given equation $D_1 = \Phi \circ C_1$.[3] This implies $D_{k+1} = \Phi \circ C_{k+1}$ because of the CSSC-property of $\{D_k \mid k \geq 1\}$. $\square$

In particular, the result of Theorem 3 implies that any questions concerning the structural similarity of two CSSCs can be reduced to the analysis of their generators. Any symmetric correspondence between two CSSCs in the large can be detected in the small, that is, in the structure of their generators. Thus, in order to give a classification of CSSCs where two families of curves that are equal modulo symmetry (rotation and reflection) are not distinguished, we need only distinguish between generators which differ modulo symmetry. We may therefore exclusively restrict our attention to the analysis of different types of generators and of suitable sequences of permutations. So, our result greatly simplifies the complete classification and the construction of all structurally different CSSCs. Moreover, it lays the foundations of a mechanized analysis of, for example, locality properties of multi-dimensional Hilbert indexings (cf. [15]).

4 Applications: computing and analyzing CHPs

First in this section, we attack a classification of all structurally different CHPs for higher dimensions. Whereas we can provide concrete combinatorial results for the 2D and 3D cases, the high-dimensional cases appear to be much more difficult. The basic tool for such an analysis, however, is given by Theorem 3. In the following subsections we sketch how to construct Hilbert indexings in higher dimensions and thus clarify the existence of such objects in arbitrary dimensions. Also in this section, we discuss computational aspects of Hilbert indexings and finally we conclude with locality properties of such curves. The general structural behavior of CHPs is sufficient to extend some results provided in previous work, such as Gotsman and Lindenbaum [9].

4.1 Classification Theorems for the two and three dimensional cases

Our first theorem investigates the two-dimensional setting. The result given below justifies the naming "class with Hilbert property" (CHP). Also, note that the subsequent proofs make decisive use of the geometric clearness provided by our formalism.

Theorem 4. *The classical 2D Hilbert indexing $\mathcal{H}\big(\, Hil_1^2,\ ((2\,4), id, id, (1\,3))\,\big)$ is the only CHP of dimension 2 modulo symmetry.*

[3] A disturbance by Φ implies a transformation of the corner-indexings by ϕ^{-1}, which can be checked.easily.

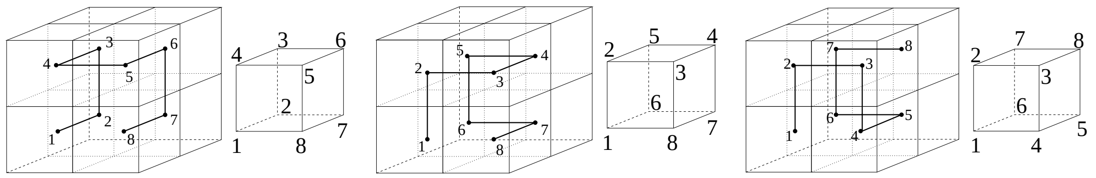

generator Hil^3_1.A generator Hil^3_1.B generator Hil^3_1.C

Figure 3: Continuous 3D generators Hil^3_1.x and their canonical corner-indexings $\widetilde{\mathrm{Hil}}^3_1$.x.

Proof. Due to Theorem 3 it suffices to show that Hil^2_1 is the only continuous 2D generator, which is obvious. In addition, we have to check whether there is another sequence of permutations such that 4 generators Hil^2_1 can be arranged in a grid of order 2 along the canonical corner-indexing $\widetilde{\mathrm{Hil}}^2_1$ in a continuous way. A simple combinatorial consideration shows that no other sequence of permutations yields a continuous curve of order 2 whose starting- and endpoints are located at corners of the grid. However, any constructor for a continuous curve of higher order must have the property that both starting- and endpoint are corner-points of the grid. □

What about the 3D case? Are there any differences concerning the amount of possible CHPs? The analysis of the “Simple Indexing Schemes” (which are related to our CHPs) in Chochia and Cole [5] already shows that the number of CHPs in the 3D case grows drastically compared to the 2D setting. Lots of “Simple Indexing Schemes” in [5] now, by our analysis, turn out to be identical modulo symmetry. Our goal is to specify all structurally different CHPs, that is, all CHPs that are not identical modulo symmetry (rotation and reflection). Since, by Theorem 3, we find any symmetric similarities of two CHPs in the structure of their generating elements, we may restrict our attention to the investigation of the generators and all suitable sequences of permutations. In addition, Lemma 1 and Lemma 2 can be seen as helpful tools to describe symmetrically disturbed CHPs in a very constructive way. They at least provide formulas of how to calculate the sequence of permutations for a disturbed CHP out of the given sequence of the original CHP. The following theorem also generalizes and answers work of Sagan [17].

Theorem 5. *For the 3D case there are* $6 \cdot 2^8 = 1536$ *structurally different (that is: not identical modulo reflection and rotation) CHPs. These types are listed in Table 1.*

Proof. Theorem 3 says that we can restrict our attention to checking any continuous curves of order 1 which are different modulo symmetry. Given such a continuous generator C, the total amount of CHPs which can be constructed by

generator	version	τ_1	τ_2	τ_3	τ_4
Hil_1^3.A	(a)	(2 8)(3 5) / (2 4 8)(3 5 7)	(3 7)(4 8) / (2 8 4)(3 7 5)	(3 7)(4 8) / (2 8 4)(3 7 5)	(1 3)(6 8) / (1 3)(2 4)(5 7)(6 8)
	(b)	(2 8)(3 5) / (2 4 8)(3 5 7)	(3 7)(4 8) / (2 8 4)(3 7 5)	id / (2 4)(5 7)	(1 7 3)(4 6 8) / (1 7 5 3)(2 8 6 4)
	(c)	(2 8)(3 5) / (2 4 8)(3 5 7)	(3 7)(4 8) / (2 8 4)(3 7 5)	id / (2 4)(5 7)	(1 7)(4 6) / (1 7 5)(2 6 4)
	(d)	(2 8)(3 5) / (2 4 8)(3 5 7)	(3 7)(4 8) / (2 8 4)(3 7 5)	(3 7)(4 8) / (2 8 4)(3 7 5)	(1 3)(6 8) / (1 3)(2 4)(5 7)(6 8)
Hil_1^3.B	(a)	(2 8)(5 7) / (2 6 8)(3 5 7)	id / (2 6)(3 7)	(3 5)(6 8) / (2 8 6)(3 7 5)	(2 8)(5 7) / (2 6 8)(3 5 7)
	(b)	(2 8)(5 7) / (2 6 8)(3 5 7)	id / (2 6)(3 7)	(3 5)(6 8) / (2 8 6)(3 7 5)	(3 5)(6 8) / (2 8 6)(3 7 5)
generator	**version**	τ_5	τ_6	τ_7	τ_8
Hil_1^3.A	(a)	(1 3)(6 8) / (1 3)(2 4)(5 7)(6 8)	(1 5)(2 6) / (1 5 7)(2 4 6)	(1 5)(2 6) / (1 5 7)(2 4 6)	(1 7)(4 6) / (1 7 5)(2 6 4)
	(b)	(1 3 5)(2 6 8) / (1 3 5 7)(2 4 6 8)	id / (2 4)(5 7)	(1 5)(2 6) / (1 5 7)(2 4 6)	(1 7)(4 6) / (1 7 5)(2 6 4)
	(c)	(2 8)(3 5) / (2 4 8)(3 5 7)	id / (2 4)(5 7)	(1 5)(2 6) / (1 5 7)(2 4 6)	(1 7)(4 6) / (1 7 5)(2 6 4)
	(d)	(1 3 5)(2 6 8) / (1 3 5 7)(2 4 6 8)	id / (2 4)(5 7)	(1 5)(2 6) / (1 5 7)(2 4 6)	(1 7)(4 6) / (1 7 5)(2 6 4)
Hil_1^3.B	(a)	(1 3)(4 6) / (1 3 7)(2 4 6)	(1 3)(4 6) / (1 3 7)(2 4 6)	id / (2 6)(3 7)	(1 7)(2 4) / (1 7 3)(2 6 4)
	(b)	(1 7)(2 4) / (1 7 3)(2 6 4)	(1 3)(4 6) / (1 3 7)(2 4 6)	id / (2 6)(3 7)	(1 7)(2 4) / (1 7 3)(2 6 4)

Table 1: Description of all 3-dimensional CHPs.

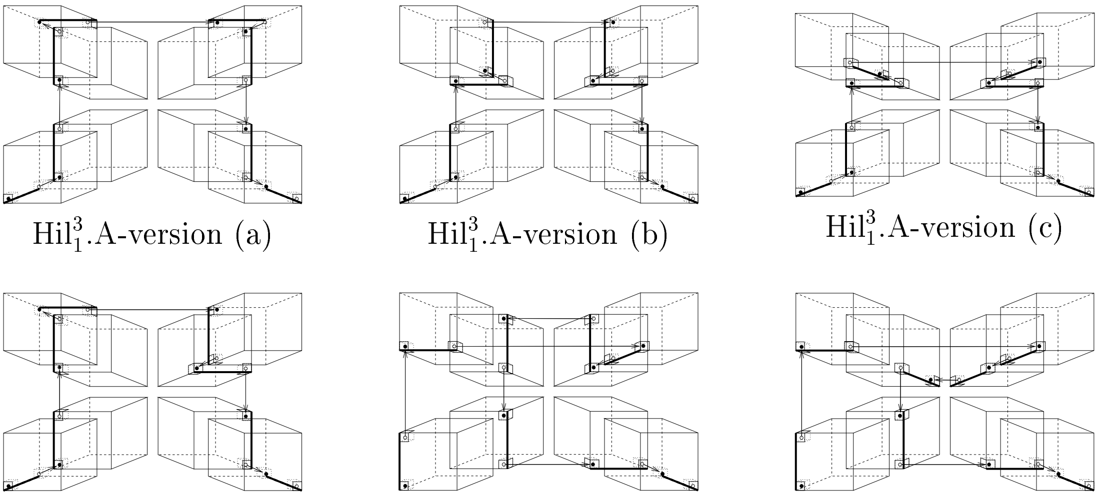

Hil_1^3.A-version (a) Hil_1^3.A-version (b) Hil_1^3.A-version (c)

Hil_1^3.A-version (d) Hil_1^3.B-version (a) Hil_1^3.B-version (b)

Figure 4: Construction principles for CHPs with generators Hil_1^3.A and Hil_1^3.B.

C is given by all possibilities of piecing together 8 (rotated or reflected) versions of C ("subcurves") along its canonical corner-indexing $\widetilde{C}$. By exhaustive search, we get that there are 3 different (modulo symmetry) types of continuous generators, namely Hil_1^3.A, Hil_1^3.B and Hil_1^3.C (see Fig. 3). As described above, we now have to check whether there are continuous arrangements of these generators along their canonical corner-indexings. Beginning with type A, an exhaustive combinatorial search yields that there are 4 possible continuous formations of Hil_1^3.A along $\widetilde{\mathrm{Hil}_1^3}$.A. All possibilities are shown in Fig. 4, where the orientation of each subcube is given by the position of an edge (drawn in bold lines). For each subcube there are two symmetry mappings which yield possible arrangements for the generator within such a subgrid. The permutations expressing these mappings are listed in Table 1.

Analogously, we find out the possible arrangements for generator type B. Note that there are no more than 2 different continuous arrangements of this generator along its canonical corner-indexing. Finally we easily check that Hil_1^3.C cannot even be the constructor of a continuous curve of order 2. Table 1 thus yields that there are exactly $4 \cdot 2^8 + 2 \cdot 2^8 = 6 \cdot 2^8$ structurally different CHPs. □

A complete classification of the high-dimensional cases appears to be much more difficult. We end this section by sketching several further results based on our characterization of curves with Hilbert property.

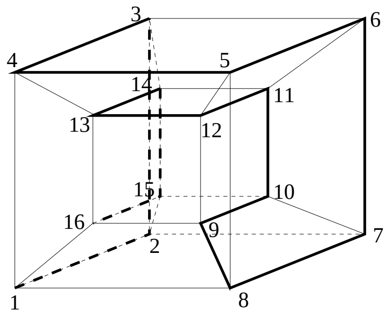

$$
\begin{aligned}
\tau_1 &= (2\ 16)(3\ 13)(6\ 12)(7\ 9)\\
\tau_2 &= (3\ 15)(4\ 16)(5\ 9)(6\ 10)\\
\tau_3 &= \tau_2\\
\tau_4 &= (1\ 3\ 13\ 11\ 9\ 7)(2\ 4\ 14\ 12\ 10\ 8)(5\ 15)(6\ 16)\\
\tau_5 &= \tau_4\\
\tau_6 &= (1\ 5\ 13\ 9)(2\ 6\ 14\ 10)(3\ 11\ 15\ 7)(4\ 12\ 16\ 8)\\
\tau_7 &= \tau_6\\
\tau_8 &= (1\ 7)(4\ 6)(10\ 16)(11\ 13)\\
\tau_9 &= \tau_8\\
\tau_{10} &= (1\ 9\ 13\ 5)(2\ 10\ 14\ 6)(3\ 7\ 15\ 11)(4\ 8\ 16\ 12)\\
\tau_{11} &= \tau_{10}\\
\tau_{12} &= (1\ 11)(2\ 12)(3\ 5\ 7\ 9\ 15\ 13)(4\ 6\ 8\ 10\ 16\ 14)\\
\tau_{13} &= \tau_{12}\\
\tau_{14} &= (1\ 13)(2\ 14)(7\ 11)(8\ 12)\\
\tau_{15} &= \tau_{14}\\
\tau_{16} &= (1\ 15)(4\ 14)(5\ 11)(8\ 10)
\end{aligned}
$$

Figure 5: Constructing elements for a 4-D CHP (generator Hil_1^4 and permutations).

4.2 Construction of an r-dimensional Hilbert curve

As already mentioned before, CHPs seem to outperform many other space-filling curves concerning their properties important for applications like data structures or parallel processing (e.g. computational effort, locality, etc.). Since such qualities might depend only weakly on the inside structure of a CHP, it, however, seems to be important to have at least one easily constructible CHP for each dimension. Without giving an explicit proof here, we just indicate how the construction of a high-dimensional CHP can be done inductively in an easy way: A continuous generator of dimension r can be derived inductively simply by "joining together" two continuous generators of dimension $r-1$. A similar consideration finally helps to specify the suitable permutations in order to obtain indexings of higher order. As an example we give a CHP of dimension 4, whose generator Hil_1^4 is constructed by joining together two generators Hil_1^3, version (a) (cf. Figure 3). The generator Hil_1^4 and a suitable sequence of permutations are shown in Fig. 5. Note that this construction principle can be extended to obtain Hilbert indexings in arbitrary dimensions in an expressive, easy, and constructive way: Following the construction principle of Hil_1^3, version (a), first pass through an $r-1$-dimensional structure, then in "two steps" do a change of dimension in the rth dimension, and finally again pass through an $r-1$-dimensional structure. This method applies to finding the generators as well as to finding the permutations. Thus the construction principle of Hil_1^3, version (a) in a sense is iterated $r-2$ times in order to generate an indexing of dimension r.

4.3 Recursive computation of CSSCs

Note that whenever a CSSC $\mathcal{C} = \{C_k \mid k \geq 1\}$ is explicitly given by its generator and the sequence of permutations, we may use the recursive formula (1) of Subsection 3.1 to compute the curves C_k. In other words, the defining formula (1) itself provides a computation-scheme for CSSC, which is parameterized by the generating elements (generator and sequence of permutations). This underlines the usefulness of the simple structure of CSSCs in particular with respect to aspects of computation.

4.4 Aspects of locality

The above mentioned parameterized formula might, for example, also be used to investigate locality properties of CSSCs by mechanical methods. The locality properties of Hilbert curves have already been studied in great detail. As an example for such investigations, we briefly note a result of Gotsman and Lindenbaum [9] for multidimensional Hilbert curves. In [9] they investigate a curve $C: \{1, \ldots, n^r\} \to \{1, \ldots, n\}^r$ with the help of their locality measure

$$L_2(C) := \max_{i,j \in \{1,\ldots,n^r\}} \frac{d_2(C(i), C(j))^r}{|i-j|}, \tag{2}$$

where d_2 denotes the Euclidean metric. In their Theorem 3 they claim the upper bound $L_2(H_k^r) \leq (r+3)^{\frac{r}{2}} 2^r$ for any r-dimensional Hilbert curve of order k, without precisely specifying what an r-dimensional Hilbert curve shall be. Since the proof of their result does not utilize the special Hilbert structure of the curve, this result can even be extended to arbitrary CSSCs.

Moreover, apart from the given locality measure L_2 we can consider measures L_p (with $p = 1$ or $p = \infty$), replacing the Euclidean distance d_2 in definition (2) by the Manhattan metric d_1, and the Maximum metric d_∞, respectively.

When making use of the special CHP-property of a class of curves one even can get closer results. For the 2D case (see Theorem 4) Gotsman and Lindenbaum present a result (cf. [9, Theorem 4]) which can be improved to the following theorem. Its proof, which is based on a more detailed investigation than the one given in Gotsman and Lindenbaum's previous proof, can be found in [1].

Theorem 6. *For the 2D Hilbert indexing $\mathcal{H}^2 = \{Hil_k^2 \mid k \geq 1\}$ we have*

$$\begin{array}{ccccc} 6(1 - O(2^{-k})) & \leq & L_2(Hil_k^2) & \leq & 6\frac{1}{2} \\ 6(1 - O(2^{-k})) & \leq & L_\infty(Hil_k^2) & \leq & 6\frac{2}{5}. \end{array}$$

for all $Hil_k^2 \in \mathcal{H}^2$ with $ord(Hil_k^2) = k$.

This result is completed by a result for the Manhattan metric [6, 15], which in the above notation would be

$$L_1(\mathrm{Hil}_k^2) = 9 \quad \text{for all} \quad \mathrm{Hil}_k^2 \in \mathcal{H}^2.$$

5 Conclusion

There is no denying the fact that dealing with dimensions greater than 3 makes the study of multi-dimensional structures quite hard due to the loss of geometric intuition. In this paper we tried to provide a simple as possible mathematical mechanism to describe and analyze space-filling Hilbert curves in arbitrary dimensions. Using a formalism based on generating elements and permutations, which completely describe whole families of Hilbert curves, we were able to discover some nice combinatorial properties of Hilbert curves in arbitrary dimensions.

Our formalism still leaves a lot of freedom we have not made use of. So giving up the restriction to "pure" Hilbert curves, it would be fairly straightforward to also study generators with side-length b instead of 2 (cf. [1]). However, in this case the formalism would become a little more complicated because there is no longer such a simple isomorphism between corner indexings and generators. Note that, for example, Butz [3] studied locality in multidimensional curves with $b = 3$, paying less attention to a combinatorial study and structural issues of the curves as we did. From an application point of view, it may also be important to study non-cubic grids and the corresponding indexings. Here our formalism in principle also works, but one has to take care of the fact that in this case only a more restricted form of permutations applies. It would also be possible to make use of more than one generator as we do in the Hilbert case, thus also gaining curves with somewhat better locality properties than Hilbert ones (cf. [2, 5] for 2D and 3D cases). However, this probably would extremely complicate the combinatorial analysis while only obtaining a modest improvement in locality properties. Our paper lays the basis for several further research directions. So it could be tempting to determine the number of structurally different r-dimensional curves with Hilbert property for $r > 3$. Moreover, a (mechanized) analysis of locality properties of r-dimensional ($r > 3$) Hilbert curves is still to be done (cf. [15]). An analysis of the construction of more complicated curves using more generators or different permutations for different levels remains open.

References

[1] J. Alber. Locality properties of discrete space-filling curves: Results with relevance for computer science. Studienarbeit, Universität Tübingen, 1997. (in German). Available through http://www-fs.informatik.uni-tuebingen.de/ niedermr/publications/index.html.

[2] T. Asano, D. Ranjan, T. Roos, E. Welzl, and P. Widmayer. Space-filling curves and their use in the design of geometric data structures. *Theoretical Comput. Sci.*, 181:3–15, 1997.

[3] A. R. Butz. Space filling curves and mathematical programming. *Information and Control*, 12:314–330, 1968.

[4] A. R. Butz. Convergence with Hilbert's space filling curve. *Journal of Computer and System Sciences*, 3:128–146, 1969.

[5] G. Chochia and M. Cole. Recursive 3D mesh indexing with improved locality. Technical report, University of Edinburgh, 1997. Short version appeared in the Proceedings of *High-Performance Computing and Networking 1997* (HPCN'97), Lecture Notes in Computer Science 1225.

[6] G. Chochia, M. Cole, and T. Heywood. Implementing the Hierarchical PRAM on the 2D mesh: Analyses and experiments. In *IEEE Symp. on Parallel and Distributed Processing*, 1995.

[7] C. Faloutsos and S. Roseman. Fractals for secondary key retrieval. In *8th ACM Symposium on Principles of Database Systems*, pages 247–252, 1989.

[8] U. Feige and R. Krauthgamar. Stereoscopic families of permutations, and their applications. In *5th IEEE Israel Symposium on the Theory of Computing and Systems*, pages 85–95, 1997.

[9] C. Gotsman and M. Lindenbaum. On the metric properties of discrete space-filling curves. *IEEE Transactions on Image Processing*, 5(5):794–797, May 1996.

[10] D. Hilbert. Über die stetige Abbildung einer Linie auf ein Flächenstück. *Mathematische Annalen*, 38:459–460, 1891.

[11] P. Indyk, R. Motwani, P. Raghavan, and S. Vempala. Locality-preserving hashing in multidimensional spaces. In *Proceedings of the 29th ACM Symposium on Theory of Computing*, pages 618–625, 1997.

[12] H. V. Jagadish. Linear clustering of objects with multiple attributes. *ACM SIGMOD Record*, 19(2):332–342, June 1990.

[13] H. V. Jagadish. Analysis of the Hilbert curve for representing two-dimensional space. *Information Processing Letters*, 62:17–22, 1997.

[14] I. Kamel and C. Faloutsos. Hilbert R-tree: An improved R-tree using fractals. In *20th International Conference on Very Large Data Bases*, pages 500–509, 1994.

[15] R. Niedermeier, K. Reinhardt, and P. Sanders. Towards optimal locality in mesh-indexings. In B. S. Chlebus and L. Czaja, editors, *Proceedings of the 11th Conference on Fundamentals of Computation Theory*, number 1279 in Lecture Notes in Computer Science, pages 364–375, Kraków, Poland, Sept.

1997. Springer-Verlag. A long version is available as technical report IB 12/97, Universität Karlsruhe. Submitted for publication.

[16] L. K. Platzman and J. J. Bartholdi. Spacefilling curves and the planar travelling salesman problem. *Journal of the ACM*, 36(4):719–737, 1989.

[17] H. Sagan. A three-dimensional Hilbert curve. *Int. J. Math. Educ. Sci. Technol.*, 24(4):541–545, 1993.

[18] H. Sagan. *Space-Filling Curves*. Universitext. Springer-Verlag, 1994.

[19] P. Sanders and T. Hansch. Efficient massively parallel quicksort. In *4th International Symposium on Solving Irregularly Structured Problems in Parallel*, number 1253 in Lecture Notes in Computer Science, pages 13–24. Springer-Verlag, June 1997.