

Empirical Analysis of Space-Filling Curves for Scientific Computing Applications

Daryl DeFord¹ Ananth Kalyanaraman²

¹Dartmouth College

Department of Mathematics

²Washington State University

School of Electrical Engineering and Computer Science

January 13, 2014

Outline

- ① Abstract
- ② Introduction
- ③ Definitions
- ④ Methodology
- ⑤ Results
- ⑥ References

Abstract

Abstract

Space Filling Curves are frequently used in parallel processing applications to order and distribute inputs while preserving proximity. Several different metrics have been proposed for analyzing and comparing the efficiency of different space filling curves, particularly in database settings. Here, we introduce a general new metric, called **Average Communicated Distance**, that models the average pairwise communication cost expected to be incurred by an algorithm that makes use of an arbitrary space filling curve. For the purpose of empirical evaluation of this metric, we modeled the communications structure of the Fast Multipole Method for n body problems.

Using this model, we empirically address a number of interesting questions pertaining to the effectiveness of space filling curves in reducing communication, under different combinations of network topology and input distribution settings. We consider these problems from the perspective of ordering the input data, as well as using space filling curves to assign ranks to the processors. Our results for these varied scenarios point towards a list of recommendations based on specific knowledge about the input data. In addition, we present some new empirical results, relating to proximity preservation under the average nearest neighbor stretch metric, that are application independent.

Overview

- History
- SFCs and Computational Efficiency
- Clustering Metrics
- Nearest Neighbor Metrics
- Average Communicated Distance

Overview

- History
- SFCs and Computational Efficiency
- Clustering Metrics
- Nearest Neighbor Metrics
- Average Communicated Distance

Overview

- History
- SFCs and Computational Efficiency
- Clustering Metrics
- Nearest Neighbor Metrics
- Average Communicated Distance

Overview


- History
- SFCs and Computational Efficiency
- Clustering Metrics
- Nearest Neighbor Metrics
- Average Communicated Distance

Overview


- History
- SFCs and Computational Efficiency
- Clustering Metrics
- Nearest Neighbor Metrics
- Average Communicated Distance

Definition 1 (SFC)

For our purposes, a Space-filling Curve (SFC) is a mapping from a multi-dimensional space to a linear ordering that allows for unique indexing of the points in that space.




(a) Hilbert Curve $\mathcal{H}_4$



(b) Z-Curve $\mathcal{Z}_4$



(c) Gray Order $\mathcal{G}_4$



(d)
Row/Column-Major

Figure : An example illustration of the Space-Filling Curves considered in our study.

Earlier Work

- **Linear Clustering of Objects with Multiple Attributes [10].**
- Analysis of the Hilbert Curve for Representing two-dimensional Space [11].
- Space-Filling Curves and their use in Designing Geometric Data Structures [2].
- Usefulness: Aluru, Hariharan, and Sevilgen
 - Parallel Domain Decomposition [1].
 - Compressed Octrees [8].
 - Bottom-Up Construction [6].
- Analysis of the Clustering Properties of Hilbert Space-Filling Curve [3].
- On the Optimality of Clustering Properties of Space-Filling Curves [9].

Earlier Work

- Linear Clustering of Objects with Multiple Attributes [10].
- Analysis of the Hilbert Curve for Representing two-dimensional Space [11].
- Space-Filling Curves and their use in Designing Geometric Data Structures [2].
- Usefulness: Aluru, Hariharan, and Sevilgen
 - Parallel Domain Decomposition [1].
 - Compressed Octrees [8].
 - Bottom-Up Construction [6].
- Analysis of the Clustering Properties of Hilbert Space-Filling Curve [3].
- On the Optimality of Clustering Properties of Space-Filling Curves [9].

Earlier Work

- Linear Clustering of Objects with Multiple Attributes [10].
- Analysis of the Hilbert Curve for Representing two-dimensional Space [11].
- Space-Filling Curves and their use in Designing Geometric Data Structures [2].
- Usefulness: Aluru, Hariharan, and Sevilgen
 - Parallel Domain Decomposition [1].
 - Compressed Octrees [8].
 - Bottom-Up Construction [6].
- Analysis of the Clustering Properties of Hilbert Space-Filling Curve [3].
- On the Optimality of Clustering Properties of Space-Filling Curves [9].

Earlier Work

- Linear Clustering of Objects with Multiple Attributes [10].
- Analysis of the Hilbert Curve for Representing two-dimensional Space [11].
- Space-Filling Curves and their use in Designing Geometric Data Structures [2].
- Usefulness: Aluru, Hariharan, and Sevilgen
 - Parallel Domain Decomposition [1].
 - Compressed Octrees [8].
 - Bottom-Up Construction [6].
- Analysis of the Clustering Properties of Hilbert Space-Filling Curve [3].
- On the Optimality of Clustering Properties of Space-Filling Curves [9].

Earlier Work

- Linear Clustering of Objects with Multiple Attributes [10].
- Analysis of the Hilbert Curve for Representing two-dimensional Space [11].
- Space-Filling Curves and their use in Designing Geometric Data Structures [2].
- Usefulness: Aluru, Hariharan, and Sevilgen
 - Parallel Domain Decomposition [1].
 - Compressed Octrees [8].
 - Bottom-Up Construction [6].
- Analysis of the Clustering Properties of Hilbert Space-Filling Curve [3].
- On the Optimality of Clustering Properties of Space-Filling Curves [9].


Earlier Work

- Linear Clustering of Objects with Multiple Attributes [10].
- Analysis of the Hilbert Curve for Representing two-dimensional Space [11].
- Space-Filling Curves and their use in Designing Geometric Data Structures [2].
- Usefulness: Aluru, Hariharan, and Sevilgen
 - Parallel Domain Decomposition [1].
 - Compressed Octrees [8].
 - Bottom-Up Construction [6].
- Analysis of the Clustering Properties of Hilbert Space-Filling Curve [3].
- On the Optimality of Clustering Properties of Space-Filling Curves [9].

Nearest Neighbor Stretch

In order to try to capture the efficiency of SFC's Xu and Tirthapura introduced a Nearest Neighbor Metric and proved some asymptotic results [8].

Figure : The Z -curve was shown to be within a constant factor of optimal for any SFC.



Defintion 2 (ACD)

Definition

Given a particular problem instance, the *Average Communicated Distance (ACD)* is defined as the average distance for every pairwise communication made over the course of the entire application. The communication distance between any two communicating processors is given by the length of the shortest path (measured in the number of hops) between the two processors along the network intraconnect.

Definition 3 (FMM)

Definition

The Fast Multipole Method (FMM) is an algorithm for computing the interactions in an n body problem [3]. We modeled the communications structure of this algorithm as a case study because it relies on computing the Near Field Interactions (NFI) and Far Field Interactions (FFI) separately. Each of these sets of computations has a different communications profile and requires distinct analysis under the ACD metric.

Research Questions

We addressed the following four research questions using our empirical models:

- Q1) *What is the nearest-neighborhood preservation efficacy achieved by different particle-order SFCs?*
- Q2) *What is the effect of different combinations of {particle-order, processor-order} SFCs on the Average Communicated Distance metric?*
- Q3) *What is the performance of each of the particle-order SFCs under the ACD metric, for a given network topology? Similarly, what is the performance of each of the network topologies under the ACD metric, for a given input distribution?*
- Q4) *How does the Average Communicated Distance vary as a function of processor size, input size and input distribution, for each SFC?*

Methodology

By varying the input parameters and running large scale successive experiments, we were able to generate meaningful results.

Some of the parameters included:

- Number of Points
- Number of Processors
- Underlying Distribution
- NFI Radius
- Space-Filling Curve for Point Separation
- Space-Filling Curve for Processor Ranking
- Spatial Resolution
- Network Topology

Probability Distributions




Figure : This figure shows examples of the two dimensional probability distributions considered in this paper.

Ordered Points




Figure : As an example of particle-ordering SFCs, this figure shows the linear order of the exponentially distributed particles displayed in Figure 2(c) by each of the SFCs respectively. It is interesting to observe the large “jumps” that occur in the orderings by the discontinuous curves, (b), (c), and (d), especially along the lines of symmetry [9].

FMM Abstraction (NFI)




Figure : ©MathWorks 2013

Nearest neighbor radius communications are a frequent model of parallel communication.

FMM Abstraction (FFI)




Figure : Interaction Lists: Figure showing two partitioned spatial resolutions. In the coarse resolution image (a), the interaction list of node 0 is $\{2, 3, 6, 7, 8 - 16\}$, or every node that it not in its quadrant. However, the interaction list of node 6 is $\{0, 4, 8, 12, 13, 14, 15\}$. At the finer resolution, nodes in the interaction list of x are marked with y and nodes in the interaction list of a are marked with b .

Computing the ACD

To effectively characterize the communication efficacies of different SFCs on to the FMM model, we study and evaluate the two interaction types — near-field and far-field — separately. The initial operation of our method is the same for either case and can be described as follows: Given an initial distribution of n particles in a $2^k \times 2^k$ spatial resolution:

- 1 Order the particles linearly with the specified particle-order SFC;
- 2 Partition the particles into p consecutive chunks of size $\frac{n}{p}$ each;
- 3 Order the processors with the specified processor-order SFC (applies only to mesh and torus topologies);
- 4 Distribute chunk i to processor i , for $1 \leq i \leq n$.

For NFI, we compute the neighborhood of each particle and determine the distance between each communication that occurs. For FFI, we use a log-tree in each quadrant to contact each processor that contains at least one particle in the quadrant.

Computing ACD (NFI)

For the near-field interactions:

- ⑤ For each particle x , construct a list of all neighbors y , of x , such that $d(x, y) \leq r$.
- ⑥ For each (x, y) pair, determine the communicated distance as the shortest path distance along the network (possibly zero) between the processor that contains x and the processor that contains y . Note that this manner of calculating the distance renders our model contention-unaware.
- ⑦ Output the sum of these communication distances for all (x, y) as the ACD value corresponding to all near-field interactions.

Computing ACD (FFI)


For the far-field interactions:

- ⑤ For each quadrant containing at least one particle, compute an ordered list of all of the processors that contain at least one particle in that quadrant.
- ⑥ Construct a log-tree (quadtrees in 2D) connecting the processors in each quadrant.
- ⑦ To capture the parent-child communication that happens during interpolation and antinterpolation, we compute the shortest path distance along the network between the two corresponding processors.
- ⑧ Construct the interaction list for each processor at each level of resolution.
- ⑨ For each processor, compute the distance along the network between that processor and each other processor in its interaction list.
- ⑩ Output the sum over all the communication distances — Interpolation, Antinterpolation, and Interaction List — as the ACD value corresponding to all far-field interactions.


Generality of the ACD

Although we have modeled the FMM algorithm in order to demonstrate the efficacy of the ACD metric, any communication bound parallel application can be evaluated with this metric. By abstracting different primitives of communications models, the ACD for most common types of parallel communication such as all-to-all and broadcast can be computed in advance for particular applications to allow algorithm designers to select the appropriate SFCs for data separation and processor ranking.

A1) ANNS Results



(a) Standard ANNS



(b) Large Radius ANNS

Figure : This figure shows the ANNS values [8] of the SFCs under consideration as the spatial resolution varies. Expanding the radius (b) does not affect the relative ordering of the SFCs. This confirms the theoretical calculations of Xu and Tirthapura on the Z and Row Major curves, and suggests that proximity preservation is not the best measure of SFC effectiveness for scientific computing [11].

A2) Main Results (NFI)

Table : A comparison of different particle/processor-order SFC combinations for NFI under various distributions. The lowest ACD value within each row is displayed in **boldface blue**, while the lowest ACD value within each column is displayed in *red italics*. The best option for each distribution is displayed in **bold green italics**.

	Particle Order			
Processor Order ↓	Hilbert Curve	Z-Curve	Gray Code	Row Major
Hilbert Curve	4.008	<i>4.308</i>	<i>4.939</i>	<i>13.117</i>
Z-Curve	5.486	5.758	6.573	18.127
Gray Code	5.802	6.010	6.970	19.220
Row Major	9.126	9.763	11.713	70.353

Table : Uniform Distribution

A2) Main Results (NFI) Continued

	Particle Order			
Processor Order ↓	Hilbert Curve	Z-Curve	Gray Code	Row Major
Hilbert Curve	8.561	<i>9.297</i>	<i>10.123</i>	<i>20.340</i>
Z-Curve	11.003	11.551	12.984	26.842
Gray Code	11.881	12.595	13.249	28.188
Row Major	20.143	22.221	24.053	66.719

(a) Normal Distribution

	Particle Order			
Processor Order ↓	Hilbert Curve	Z-Curve	Gray Code	Row Major
Hilbert Curve	5.238	<i>5.654</i>	<i>6.271</i>	<i>14.943</i>
Z-Curve	6.943	7.070	8.235	20.851
Gray Code	7.276	7.663	8.760	22.269
Row Major	12.483	13.017	15.289	61.227

(b) Exponential Distribution

A2) Main Results (FFI)

Table : A comparison of different particle/processor-order SFC combinations for FFI under various distributions. The lowest ACD value within each row is displayed in **blue boldface**, while the lowest ACD value within each column is displayed in *red italics*. The best option for each distribution is displayed in **bold green italics**.

	Particle Order			
Processor Order ↓	Hilbert Curve	Z-Curve	Gray Code	Row Major
Hilbert Curve	19.494	<i>20.841</i>	<i>22.572</i>	<i>31.124</i>
Z-Curve	24.217	24.793	27.787	37.709
Gray Code	24.622	25.446	27.997	39.282
Row Major	44.513	48.762	50.118	57.880

Table : Uniform Distribution

A2) Main Results (NFI) Continued


	Particle Order			
Processor Order ↓	Hilbert Curve	Z-Curve	Gray Code	Row Major
Hilbert Curve	26.336	<i>26.824</i>	31.963	<i>32.542</i>
Z-Curve	29.160	28.036	34.241	36.663
Gray Code	29.449	27.981	<i>31.909</i>	37.291
Row Major	43.639	44.636	49.133	45.475

(a) Normal Distribution


	Particle Order			
Processor Order ↓	Hilbert Curve	Z-Curve	Gray Code	Row Major
Hilbert Curve	18.960	<i>19.841</i>	<i>23.007</i>	<i>31.368</i>
Z-Curve	24.672	23.316	26.315	37.576
Gray Code	23.762	24.076	27.973	37.863
Row Major	42.447	44.067	46.872	50.963

(b) Exponential Distribution

A3) Topology Comparison




(a) Near-Field Interactions




(b) Far-Field Interactions

Figure : The charts show the results of comparing different network topologies for a) NFI and b) FFI, respectively. All experiments were performed using 1,000,000 uniformly distributed particles on a 4096×4096 spatial resolution. This plot is representative of all the experiments we performed to evaluate the topologies. It is important to note that quadtree structures have disproportionately large issues with contention in high volume communications.

A4) ACD Scaling



(a) NFI



(b) FFI

Figure : These plots show ACD values for a) NFI, and b) FFI, as a function of the number of processors and the SFC used. The input used was fixed at 1,000,000 uniformly distributed particles. This demonstrates the effect scale on processor ranking SFCs. Some of the row-major data has been excluded from these plots because for this SFC, the ACD values at larger processor numbers were significantly higher than the other data-points.

Analysis

Our results point towards a set of recommendations for designers of parallel algorithms for scientific computing. When the scientist has full control over both the data distribution and processor ranking, using the Hilbert Curve at both stages gives the lowest ACD values. Unfortunately, such control is not always feasible or desirable, in which case we present the following recommendation for SFC selection based on the ACD values:

$$\{Hilbert \approx Z\} < Gray \ll Row-major.$$

Future Work and Extensions

We intend to further extend our results by considering the following extensions:

- Adding a weighting function to evaluate data intensive applications
- Extending our metric to consider contention based communications models
- Extending our evaluation to real world implementations and applications other than FMM.
- Providing a closed, asymptotic expression for the ANNS of more complex curves.
- One of the interesting notions encountered in this work is the mapping of points from a multi-dimensional space to a 2D torus or mesh. This is unlike the traditional SFC problem, and does not appear to have been explored yet in theory. In this paper, we used SFCs to move from 2D to a linear ordering back to 2D, but certainly there appears to be no restriction on a direct mapping into the processor space. This raises theoretical questions for further study.
- Finally, while we expect the conclusions of most of the studies conducted in this paper to extend to 3D, further experimentation is needed to corroborate such trends.

Conclusions

Our results empirically validate previously published theoretical results. In addition, based on our results, we provided a list of recommendations that could serve as benchmarks for effective use of SFCs in FMM-type applications. Our findings suggest both theoretical avenues of inquiry for future research and practical applications of particular SFCs, both for distributing the input data among parallel processors, and for canonical labeling of processors on a particular network topology, with an overall goal of minimizing communication network usage. In particular, the ACD metric presented here represents an important contribution to the study of SFCs for scientific computing.

References



S. ALURU AND F. SEVILGEN: *Parallel Domain Decomposition and Load Balancing Using Space-Filling Curves*, Fourth International Conference on High Performance Computing, (1997) 230–235.



T. ASAON, D. RANJAN, T. ROOS, E. WELZL, AND P. WIDMAYER: *Space Filling Curves and Their Use in the Design of Geometric Data Structures*, Journal of Theoretical Computer Science, **181 (1)** (1997), 3–15.



R. BEATSON AND L. GREENGARD: *A Short Course on Fast Multipole Methods*, Wavelets, Multilevel Methods and Elliptic PDEs, Oxford University Press, Oxford, (1997) 1–37.



P. CALLAHAN AND S. RAO KOSARAJU: *Algorithms for Dynamic Closest Pair and n –Body Potential Fields*, Proc. Symposium Of Discrete Algorithms, (1995), 263 – 272.



F. GRAY: *Pulse Code Communication*, U.S. Patent 2,632,058, (1953).



L. GREERGARD AND V. ROKHLIN: *A Fast Algorithm for Particle Simulations*, Journal of Computational Physics, **135** (1997), 280–292.



B. HARIHARAN AND S. ALURU: *Efficient Parallel Algorithms and Software for Compressed Octrees with Applications to Hierarchical Methods*, Parallel Computing, **31** (2005), 311–331.



B. HARIHARAN, S. ALURU, AND B. SHANKER: *A Scalable Parallel Fast Multipole Method for Analysis of Scattering from Perfect Electrically Conducting Surfaces*, Proceedings of the 2002 ACM/IEEE conference on Supercomputing, (2002), 1–17.



D. HILBERT: *Ueber die stetige Abbildung einer Linie auf ein Flchenstck*, Mathematische Annalen, **38** (1891), 459–460.



H. JAGADISH: *Linear Clustering of Objects with Multiple Attributes*, Proceedings ACM SIGMOD (1990), 332–342.



H. JAGADISH: *Analysis of the Hilbert Curve for Representing Two-Dimensional Space*, Information Processing Letters **62** (1997), 17–22.

References



I. LASHUK, A. CHANDRAMOWLISHWARAN, H. LANGSTON, T. NGUYEN, R. SAMPATH, A. SHRINGARPURE, R. VUDUC, L. YING, D. ZORN, AND G. BIROS: *A Massively Parallel Adaptive Fast-Multipole Method on Heterogeneous Architectures*, Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, (2009), 1–12.



T. MATTSON, R. VAN DER WIJNGAART, AND M. FRUMKIN: *Programming the Intel 80-core network-on-a-chip terascale processor*, Proceedings of the 2008 ACM/IEEE conference on Supercomputing (SC '08). IEEE Press, Piscataway, NJ, USA, Article 38, 11 pages.



B. MOON, H. JAGADISH, C. FALOUTSOS, AND J. SALTZ: *Analysis of the Clustering Properties of Hilbert Space-Filling Curve*, IEEE Transactions on Knowledge and Data Engineering, **13** (1) (2001), 124–141.



G. MORTON: *A Computer Oriented Geodetic Data Base and a New Technique in File Sequencing*, IBM, New York, (1996).



G. PEANO: *Sur une courbe, qui remplit toute une aire plane*, Mathematische Annalen **36** (1890), 157–160.



H. SUNDAR, R. SAMPATH, AND G. BIROS: *Bottom-Up Constructions and 2:1 Balance Refinement of Linear Octrees in Parallel*, SIAM Journal on Scientific Computing, **30** (5) (2008), 2675–2708.



TILERA CORPORATION: *Tilera*, <http://tilera.com>, accessed September 29, 2012.



P. XU AND S. TIRTHAPURA: *A Lower Bound on Proximity Preservation by Space Filling Curves*, Proc. International Parallel and Distributed Processing Symposium 2012, (2012), 1295 – 1305.



P. XU AND S. TIRTHAPURA: *On the Optimality of Clustering Properties of Space Filling Curves*, Proc. Principles of Database Systems 2012, (2012), 215–224.