On the Complexity of Isometric Embedding in the Hypercube

David Avis

School of Computer Science

McGill University

3480 University

Montreal, Quebec

H3A 2A7

ABSTRACT

A finite metric is $h-$embeddable if it can be embedded isometrically in the N-cube (hypercube) for some N. It is known that the problem of testing whether a metric is $h-$embeddable is NP-Complete, even if the distances are restricted to the set $\{2, 4, 6\}$. Here we study the problem where the distances are restricted to the set $\{1, 2, 3\}$ and give a polynomial time algorithm and forbidden submetric characterisation. In fact, we show these metrics are $h-$embeddable if and only if they are $11-$*gonal* and the sum of the distances arround any triangle is even. The so-called truncated metric case, where the distances are chosen from $\{1, 2\}$ is particularly simple, the only embeddable metrics arise from the graphs $K_{1,n-1}$, $K_{2,2}$, and $2K_n$ (K_n with all distances 2).

1. Introduction

For an integer n and a finite set $\mathbf{X} = \{x_1, \ldots, x_n\}$, let $(\mathbf{X}, d)$ be a metric space. In other words,

$$d(x_i, x_j) = d(x_j, x_i), \quad d(x_i, x_j) = 0, \text{ if and only if } i = j \text{ and}$$

$$d(x_i, x_j) \le d(x_i, x_k) + d(x_k, x_j) \quad \text{for distinct } i, j, k \in \{1, \ldots, n\}.$$

The (Hamming) $N-$*cube* Q_N is the set of all binary vectors of length N in R^N. The distance between two vertices of Q_N is the Hamming distance between the two binary vectors. We say that d is $h-$embeddable if there is a mapping

$$f: X \rightarrow Q_N$$

that preserves distances. The results in this paper in fact apply to more general *semi – metric spaces* where distinct points may have distance zero. It can easily be verified that such semi-metrics are h – embeddable if and only if the metric formed by contracting points at distance zero to single points is h – embeddable. For simplicity, we deal with metric spaces.

If d is rational, we say that $(\mathbf{X}, d)$ is L^1 – embeddable if and only if $(\mathbf{X}, kd)$ embeds isometrically into Q_N for some N and some $k > 0$. Problems concerning h – embeddable and L^1 – embeddable metrics have a long history, for a survey see Avis and Deza[1]. In this paper we will address some complexity issues related to testing whether a metric is h – embeddable. First we mention a related problem.

We can consider Q_N to be a graph on 2^N nodes by joining two nodes if the corresponding binary vectors differ in exactly one position. The problem of embedding graphs (not necessarily isometrically) into Q_N also has a long history, see for example Garey and Graham[7] for a discussion of forbidden subgraphs. Recently, this problem has again been raised in connection with Hypercube multiprocessors. It is known that the problem is NP-complete for graphs (Krumme et al.[9]), and even for embedding trees (Wagner and Corneil[10]). On the other hand, for the case of isometrically embedding a graph, the problem turns out to have a good characterisation and polynomial time algorithm (Djokovic[6]). When the graph structure is missing, the problem turns out to be NP-complete. This was shown by Chvátal[2] in the setting of intersection patterns, which we now review.

An *intersection pattern* is specified by an integer matrix $C = (c_{ij})$ of dimension $n \times n$. C is an intersection pattern if there exist finite sets A_i, $i = 1, \ldots, n$ such that

$$|A_i \cap A_j| = c_{ij} \qquad 1 \le i < j \le n.$$

The sequence $A_i, i = 1, \ldots, n$ is called a *realization*. Define $N = |\bigcup_{i=1}^{n} A_i|$ to be the size of the realization. The relationship between intersection patterns and hypercube embedding, pointed out by Deza[3], is given by the following.

Proposition 1.1

(a) If $C = (c_{ij})$ is an $n \times n$ intersection pattern with given N, then d is a metric on $n+1$ points embeddable in a Q_N, where

$$d(x_i, x_j) = c_{ii} + c_{jj} - 2c_{ij}.$$

(b) If d is a metric on $n+1$ points that is embeddable in Q_N, then $C = (c_{ij})$ is an $n \times n$ intersection pattern, with a realization of size N, where

$$c_{ij} = \frac{1}{2} (d(x_i, x_{n+1}) + d(x_i, x_{n+1}) - d(x_i, x_j)). \square$$

Results about the complexity of hypercube embedding follow from the following result on intersection patterns.

P1. Intersection Pattern.

Instance: Integer $n \times n$ matrix C.

Question: Is C an intersection pattern?

Complexity: NP-complete, even if each entry $c_{ij} \leq 3$. Solvable in polynomial time if each entry $c_{ij} \leq 2$, see Chvátal[2].

This result can be reformulated as a graph theory problem. Interpret C to be the adjacency matrix of a multigraph with n vertices and c_{ij} edges between vertices i and j. If C is an intersection pattern, then the edges of the multigraph can be partitioned into cliques. The sets A_i realizing the intersection pattern correspond to a list of the cliques containing the vertex i. The size N of the realization corresponds to the number of cliques in the partition. By the above result, the clique partition problem is still hard even if each vertex is contained in at most 3 cliques. If each vertex is to be contained in at most 2 cliques, then the problem reduces to line graph recognition, which is solvable in polynomial time. Using the equivalence between intersection patterns and hypercube embedding given in Proposition 1.1, and contracting points at zero distance, we obtain the following.

P2. Hypercube Embedding

Instance: Integer metric d on n points.

Question: Is d embeddable in a Hamming N-cube?

Complexity: NP-complete, even if each $d(x_i, x_j) \in \{2, 4, 6\}$. Solvable in polynomial time if each $d(x_i, x_j) \in \{2, 4\}$ and some point is at distance 2 from all points. Also solvable in polynomial time for metrics arising from bipartite graphs, Djokovic[6].

Using the graph theoretic interpretation given above, other complexity results for edge partitions of graphs can be restated in terms of hypercube embedding or intersection patterns. For example, Hoyler[8] showed that it is NP-complete to determine if the edges of a graph can be partitioned into complete subgraphs of fixed size t, for all $t \geq 3$. In terms of intersection patterns, this implies that it is NP-complete to decide if a matrix is an intersection pattern realizable by subsets $A_i \subseteq \{1, \ldots, N\}$ such that each element is contained in exactly t subsets.

In this paper, we extend the cases for which P2 is solvable in polynomial time. We first consider the case where the distances are chosen from $\{1, 2\}$. This is the so-called truncated graph metric: for any graph (connected or not) on n vertices the distance between adjacent vertices is one and between non-adjacent vertices it is two. We give a simple characterisation of $h-$embeddable metrics of this form. Our main result in an extension to metrics where the distances are chosen from $\{1, 2, 3\}$. This shows the vital importance of scaling, given the above mentioned NP-completeness for metrics with distances from $\{2, 4, 6\}$. In conclusion, we note that the complexity of determing if a metric is L^1-embeddable is unknown.

2. Preliminaries

Although no necessary and sufficient conditions are known for testing whether a metric is $h-$embeddable there are two useful necessary conditions, first discovered by Deza[5]: the even condition and the hypermetric condition.

Even Condition. If a metric d is $h-$embeddable then the sum of the distances arround any triangle must be even. □

Hypermetric Condition If a metric d is h – embeddable then

$$\sum_{1 \le i < j \le n} b_i b_j d(x_i, x_j) \le 0$$

for all integers $b_1, \ldots, b_n$ satisfying

$$\sum_{i=1}^{n} b_i = 1. \qquad \square$$

A metric d is $k - gonal$ if it satisfies the hypermetric condition for all integers b_i such that

$$\sum_{i=1}^{n} |b_i| \le k.$$

As an illustration, Figure 1 shows that $K_{2,3}$ does not satisfy the hypermetric condition, and in fact is not 5-gonal. The labels on the vertices are the integers b_i. Observe that all metrics are 3-gonal.

The above conditions are often useful for showing that a metric is *not* h – embeddable. In order to show that a metric is h – embeddable it suffices to give the embedding. For this it is convenient to assign sets, rather than binary vectors, to points in the metric space. Suppose a metric d is embeddable in Q_N. We represent this by choosing a set of N *atoms* $\{a, b, c, \cdots\}$, where a corresponds to the first coordinate of the embedding, b corresponds to the second coordinate of the embedding, etc. Each point in the metric space has a corresponding binary vector of length N defining its vertex in Q_N. We replace this vector by the subset of atoms formed by choosing the coordinates of the vector containing a one, see Figure 2. Now if vertices x_i and x_j receive as labels, sets S_i and S_j, we have $d(x_i, x_j) = |S_i \,\Delta\, S_j|$, where Δ represents symmetric difference.

An embedding, or labelling as we will call it, is certainly not unique. Clearly we can permute the set of atoms. Less trivially, let $S_1, \cdots, S_n$ be the label sets assigned to the n points, and let T be any subset of atoms. Then we get another valid set of labels $S_i \,\Delta\, T$, $i = 1, \ldots, n$. Call a metric *rigid* if all labellings are equivalent up to permutations of the labels and symmetric difference. In the next section we use the fact[4] that the metric induced by 5 points with all interpoint distances two, denoted $2K_5$, is rigid. The reader may discover that $2K_4$ is not rigid.

In order to illustrate the above concepts, let us consider the case where all distances are chosen from the set $\{1, 2\}$. Consider a metric d on n points with this property. Let $G = (V, E)$ be the graph defined on the n points, where two points are adjacent if and only if they have distance one. It is immediate from the even condition that G is a complete bipartite graph (possibly with no edges, ie. $K_{0,n}$). Let the vertices in the two parts be denoted S and T. We may assume that S is non-empty and $|S| \geq |T|$. Observe that all inter-point distances in S or T are two. If $|T| = 1$ then the graph is $K_{1,n-1}$ which is $\mathrm{h-embeddable}$: assign a unique atom to each point in S and the empty set to the point in T. If T is empty, the metric has all distances 2, which we denote $2K_n$. This is easily seen to be $\mathrm{h-embeddable}$: assign a unique atom to each vertex. If $|S| = |T| = 2$ then we have $K_{2,2}$ which is $\mathrm{h-embeddable}$: assign labels $\{a\}$ and $\{b\}$ to S and labels $\{a, b\}$ and $\emptyset$ to T. Finally if $|S| > |T| \geq 2$ then we have as a subgraph $K_{2,3}$, which does not satisfy the hypermetric condition so is not $\mathrm{h-embeddable}$. We can summarize the discussion in the following theorem.

Theorem 2.1 A metric with distances one or two is $\mathrm{h-embeddable}$ if and only if it is one of $K_{1,n-1}$, $2K_n$, or $K_{2,2}$. $\square$

3. Metrics with distances $\{1, 2, 3\}$.

Throughout this section we restrict ourselves to metrics d on n points in which all distances are from the set $\{1, 2, 3\}$, and to avoid some special cases we assume that $n \geq 9$. We will characterize which metrics are $\mathrm{h-embeddable}$ in terms of forbidden submetrics, and give an efficient embedding algorithm. We procede as in the discussion before Theorem 2.1 by defining a graph $G = (V, E)$, where V is the set of n points in a metric d and two points are adjacent if and only if they have distance one or three. We will also refer frequently to the subgraph $H = (V, \bar{E})$ of G consisting of just the edges corresponding to distance one in the metric. It is again immediate from the even condition that G is again a complete bipartite graph. Let the parts be $S = \{s_1, s_2, \ldots, s_{|S|}\}$ and $T = \{t_1, t_2, \ldots, t_{|T|}\}$. As before we may assume that S is non-empty and $|S| \geq |T|$. Since $n \geq 9$ we have $|S| \geq 5$. Observe that all inter-point distances in S or T are two.

As was remarked in the introduction, the submetric induced by S is rigid and so has a unique labelling (up to permutation and symmetric difference). We will use the atoms $\{a_1, \ldots, a_{|S|}\}$ assigning s_i the label $\{a_i\}$.

In the rest of this section, we concentrate on the subgraph H based on a metric d that is h – embeddable. Therefore adjacency should be interpreted with respect to H, and the notation $N(v)$ for $v \in V$ refers to the subset of vertices adjacent to v in H. We call a vertex $v \in T$ a *universal* vertex if $N(v) = S$. The following lemmas describe the structure of H.

Lemma 3.1 There is at most one universal vertex.

Proof: Suppose there are two universal vertices. Then d must contain the metric induced by $K_{2,3}$ as a submetric, since $|S| \geq 3$. As remarked in the previous section, this metric does not satisfy the hypermetric condition. □

If there is a universal vertex, it will receive the empty set as a label, which is seen to be consistent with the labelling of S. Let T_i denote the vertices of T which each have exactly i neighbours in H.

Lemma 3.2 For $i = 3, 4, \ldots, |S| - 1$, $T_i = \emptyset$. At most one of T_0 and T_2 is non-empty.

Proof: The metrics (i) and (ii) of Figure 3 do not satisfy the hypermetric condition and so are not h – embeddable. Interpret the bottom row of 5 points as vertices from $|S|$. This shows that both T_3 and T_4 must be empty. For $5 \leq i \leq |S| - 1$, if T_i is non-empty, d must have (ii) as a submetric. Suppose now that both T_0 and T_2 are non-empty. Then we must have the non-hypermetric submetric shown in Figure 3(iv). □

Lemma 3.3 If $u, v \in T_2$ are distinct, then $|N(u) \cap N(v)| = 1$.

Proof: Choose u, v as in the statement of the Lemma. If $|N(u) \cap N(v)| = 0$ then we have the non-hypermetric submetric in Figure 3(v), where the bottom vertices are from $|S|$ and the top vertices are u, v. In the case $|N(u) \cap N(v)| = 2$ then we have the non-hypermetric submetric in Figure 3(vi), with vertices labelled similarly. □

We define a label for each $u \in T_2$ as follows: if $N(u) = \{s_i, s_j\}$ then label u with the set $\{a_i, a_j\}$. It follows from the Lemma 3.3 that the labels in T_2 are mutually consistent and it can be checked that they are consistent with the other vertices labelled so far.

Lemma 3.4 If $|T_1| \geq 2$ then either

(a) all $v \in T_1$ are adjacent to some common vertex $s_i \in S$, or

(b) all $v \in T_1$ have distinct neighbours in S.

Proof: If the lemma is false we must have three vertices $u, v, w \in T_1$ such that u and v have a common neighbour in S which is distinct from the neighbour in S of w. In this case we get the forbidden submetric shown in Figure 3(v), where bottom three vertices are u, v, w and the top two vertices are their neighbours in S. □

Depending on which case of Lemma 3.4 applies, we apply labels to T_1. In case (a) we give labels $\{a_i, b_0\}$ $\{a_i, b_1\}$ ... $\{a_i, b_{|T_1|-1}\}$. In case (b) we give labels $\{a_{i_1}, b_0\}$ $\{a_{i_2}, b_0\}$ ... $\{a_{i_{|T_1|}}, b_0\}$, where we assume the vertices in T_1 are adjacent to $s_{i_1}, \ldots, s_{i_{|T_1|}}$. In case $|T_1| = 1$, both labellings are identical. It is quite easy to check that the labelling of T_1 is consistent within T_1 and with S and the possible universal vertex. Its consistency with T_2 follows from the next lemma.

Lemma 3.5 If T_1 and T_2 are non-empty then for all $v \in T_1$ and all $w \in T_2$, $N(v) \cap N(w) \neq \emptyset$. If case (b) of Lemma 3.4 applies, we must have $|T_1| = 2$ and $|T_2| = 1$.

Proof: For the first statement, with v and w as given, if $N(v) \cap N(w) = \emptyset$ then we obtain again the forbidden subgraph in Figure 3(v). This time the top two vertices are w and v and the bottom vertices are from S. Now suppose case (b) of Lemma 3.4 applies. Then there can be at most two vertices in T_1 since every vertex in T_2 must be adjacent to each of their neighbours in S, which are distinct by assumption. Finally $|T_1| = 2$ and $|T_2| \geq 2$ is impossible because all vertices in T_2 would have to have the same two neighbours in S, namely the two distinct neighbours of the vertices in T_1. This is impossible by Lemma 3.3. □

Note that the labelling given above is valid for the metrics satisfying the second statement of Lemma 3.5. We also observe that case(a) of Lemma 3.4 does not cause any difficulty and the given labelling gives a valid embedding, as illustrated in Figure 4(i).

Finally we label vertices in T_0 with the labels $\{b_0, b_1\}$, $\{b_0, b_2\}$,..., $\{b_0, b_{|T_0|}\}$. It is quite easy to check that the labelling of T_0 is consistent within T_0 and with S and the possible universal vertex. Since both T_0 and T_2 cannot be simultaneously non-empty (Lemma 3.2) no inconsistency can arise between these two subsets. Its consistency with T_1 follows from the next lemma.

Lemma 3.6 If T_0 and T_1 are non-empty and case (a) of Lemma 3.4 applies, then $|T_1| = 2$ and $|T_0| = 1$.

Proof: Assume case(a) of Lemma 3.4 applies and $|T_1| \geq 3$. Then we obtain the non-hypermetric submetric in Figure 3(vii). Here the bottom vertices are from S. If $|T_1| = 2$ and $|T_0| \geq 2$ then d has the non-hypermetric submetric shown in Figure 3(viii). □

Note that the labelling given above is valid for the metrics satisfying the statement of Lemma 3.6. We also observe that case(b) of Lemma 3.4 does not cause any difficulty and the given labelling gives a valid embedding, as illustrated in Figure 4(ii). Together, these lemmas can be used to prove the following theorem.

Theorem 3.1 Let d be a metric on $n \geq 9$ points with distances chosen from $\{1, 2, 3\}$, and which satisfies the even condition. Then the following are equivalent:

(i) d is h $-$ embeddable

(ii) d is L^1 $-$ embeddable

(iii) d is hypermetric

(iv) d is 11-gonal

(v) d does not contain any of the submetrics shown in Figure 3.

In $O(n^2)$ time a hypercube embedding for d can be found if one exists, or else one of the forbidden submetrics can be exhibited.

Proof: For all metrics d, it is known that

$$(i) \rightarrow (ii) \rightarrow (iii) \rightarrow (iv) \rightarrow (v)$$

Lemmas 3.1 - 3.6 show that $(v) \rightarrow (i)$.

For the complexity result, we first check that G is a complete bipartite graph, which can be done by a standard depth-first-search in $O(n^2)$ time. If G is bipartite, we take the smaller part as T and construct the sets T_i. We may now readily check the conditions of Lemmas 3.1 and 3.2. If these conditions are satisfied, we check whether condition (a) or (b) of Lemma 3.4 holds, in $O(n)$ time. If either condition holds, we can label the entire graph as specified above. We now check that the labelling gives the correct distances. Since each label contains at most two atoms, all distances can be calculated in $O(n^2)$ time, and checked against the input. In case the labelling fails, then one of the conditions of Lemmas 3.1 - 3.6 fails, and the appropriate forbidden submetric is produced. □

References

1. D. Avis and M. Deza, "The Cut Cone, L1-Embedability, Complexity and Multicommodity Flows," *Networks*, vol. 21, pp. 595-617, 1991. ftp://mutt.cs.mcgill.ca/pub/doc/AD91a.ps.gz

2. V. Chvátal, "Recognizing Intersection Patterns," *Annals of Discrete Maths.*, vol. 8, pp. 249-251, 1980.

3. M. Deza, "Matrices de formes quadratiques non negatives pour des arguments binaires," *C.R. Acad. Sc. Paris*, vol. 277, pp. 873-875, 1973.

4. M. Deza and N.M. Singhi, "On Rigid Pentagons in Hypercubes," *Graphs and Combinatorics*, vol. 4, pp. 31-42, 1988.

5. M. Deza(Tylkin), "On Hamming Geometry of Unit Cubes," *Doklady Akad. Nauk. SSR.*, vol. 134, pp. 1037-1040, 1960. English translation in Soviet Physics Dokl. 5 (1961), 940-943.

6. D.Z Djokovic, "Distance Preserving Subgraphs of Hypercubes," *J. Comb. Theory B*, vol. 14, pp. 263-267, 1973.

7. M.R. Garey and R.L. Graham, "On Cubical Graphs," *J. Comb.Theory B*, vol. 18, pp. 84-95, 1975.

8. I. Hoyler, "The NP-Completeness of Some Edge-Partition Problems," *SIAM J. Computing*, vol. 10, pp. 713-717, 1981.

9. D. W. Krumme, K. N. Venkataraman, and G. Cybenko, "Hypercube Embedding is NP-Complete," in *Hypercube Multiprocessors 1986*, pp. 148-157, SIAM, 1986.

10. A. Wagner and D.G. Corneil, "Embedding Trees in a Hypercube is NP-complete," in *TR 197/87*, University of Toronto, 1987.