

Maximum independent sets of the 120-cell and other regular polytopes

S. Debroni, E. Delisle, W. Myrvold, A. Sethi, J. Whitney, J. Woodcock [†]
*Department of Computer Science, University of Victoria, P.O. Box 3055, Stn CSC, Victoria, BC
Canada V8W 3P6*

P.W. Fowler, B. de La Vaissière [‡]
Department of Chemistry, University of Sheffield, Sheffield S3 7HF, UK

M. Deza
CNRS and LIGA, École Normale Supérieure, 45 rue d'Ulm, 75230 Paris, France

Abstract

A d -code in a graph is a set of vertices such that all pairwise distances are at least d . As part of a study of d -codes of three- and four-dimensional regular polytopes, the maximum independent set order of the 120-cell is calculated. A linear program based on counting arguments leads to an upper bound of 221. An independent set of order 110 in the antipodal collapse of the 120-cell (also known as the hemi-120-cell) gives a lower bound of 220 for the 120-cell itself. The gap is closed by the computation described here, with the result that the maximum independent set order of the 120-cell is 220. All maximum d -code orders of the icosahedron, dodecahedron, 24-cell, 600-cell and 120-cell are reported.

Keywords: Graphs, independent sets, polyhedra, polytopes.

Math. Subj. Class.: 05C69, 05C85, 05C10

1 Introduction

An *independent set* in a graph is a set of pairwise non-adjacent vertices such that all pairs are at distance two or more. A *clique* (a subset of the vertices that are pairwise adjacent) in a graph corresponds to an independent set in the complement of the graph. Hence algorithms for maximum clique can be applied to find maximum independent sets.

The problem of finding a maximum independent set in a graph is NP-Hard [1]. The DIMACS Clique Challenge arose from the need to find practical algorithms for the maximum clique problem,

[†]WM supported by NSERC.

[‡]PWF supported by Royal Society/Wolfson Research Merit Award, 2004-2009.

E-mail addresses: wendym@cs.uvic.ca (S. Debroni, E. Delisle, W. Myrvold, A. Sethi, J. Whitney, J. Woodcock), P.W.Fowler@sheffield.ac.uk (P.W. Fowler, B. de La Vaissière), Michel.Deza@ens.fr (M. Deza)

and the proceedings volume is an excellent place to start looking for information about practical algorithms for clique finding [4]. The challenge also included a database of difficult clique problems. Bomze, Budinich, Pardalos, and Pelillo [2] provide a comprehensive survey on the maximum clique problem. Östergård focusses on solving maximum clique problems on graphs arising from various combinatorial problems. Both surveys cite the problem of finding a maximum clique of the Keller graph of dimension 7 as an open problem. This problem has subsequently been solved [5].

In general, a d -code is a set of vertices such that pairwise distances are all at least d . The concept of distance between vertices a and b may be cast in terms of graphs (number of edges in the shortest path between the vertices), coding theory (the Hamming distance between $(0, 1)$ sequences representing coordinates of vertices), or sets (the cardinality of the symmetric difference between the two subsets that represent the vertices). A d -code in a graph G corresponds to an independent set in the graph H which has the same vertex set as G and the property that two vertices u and v are adjacent in H if and only if the distance between them in G is at most $d - 1$. The interest for applications is usually to find *maximum* d -codes, and one standard problem in the theory of error-correcting codes [7, 8] is to find the largest d -code in the n -cube. Here we consider the problem of finding largest d -codes in the graphs corresponding to other regular polytopes.

For polycyclic and polyhedral graphs in two and three dimensions, the construction of d -codes has applications to chemistry [9, 10, 11, 12]. For example, sets of codes with increasing d may be seen as templates for addition to an underlying molecular framework by addends of increasing steric demand. Codes have been presented for chemically relevant regular and semi-regular polyhedra [9] and arguments based on d -codes, coupled with spectral information, give useful rationalisations of the extent and symmetry of addition in fullerene chemistry, for example [13, 14, 15]. Although not invoked in chemistry so far, d -anticodes, defined by the requirement that pairwise distances should not exceed d , would model the opposite regime of attachment to a framework where the added groups cluster under strong inter-addend attraction. Extension of the existing lists [9] to d -codes in the graphs of all regular polytopes is a finite task, as there are only the following convex regular polytopes [16]: in dimension n the n -simplex (α_n), the n -cross-polytope (β_n), the n -cube (γ_n), and additionally in dimension 2 the regular polygons, and in dimensions 3 and 4, five sporadic polytopes. In dimension 3, α_3 is the tetrahedron, β_3 the octahedron and γ_3 the cube, and there is a dual pair of sporadic polyhedra: the icosahedron and the dodecahedron. In dimension 4, the analogues of the icosahedron and dodecahedron are the 600-cell (all of its independent sets have been enumerated previously [6]) and the 120-cell (again a dual pair), and there is also the self-dual 24-cell, without analogue in higher or lower dimensions [16].

Codes for the polytopes common to all dimensions (α_n , β_n and γ_n) are well studied. The 1-skeleton of α_n is the complete graph K_n , of diameter 1, and the 1-skeleton of β_n is the complete multipartite graph (the Cocktail-Party graph) $K_{n(2)} \equiv C_p(n) \equiv K_{2,2,\dots,2}$, of diameter 2, so coding problems are trivial for both. Codes for γ_n are the subject of classical coding theory [7]. As a bipartite graph with equal equal partite sets, γ_n has independence number 2^{n-1} . The coding problem is also trivial in two dimensions, where the order of the d -code is $\lfloor n/d \rfloor$ for the cycle of length n . It only remains to study the five exceptional regular polytopes in dimensions 3 and 4. Most of these problems are easy (see the summary in Table 3, §3). The difficult case is that of the independence number of the 600-vertex 120-cell, which does not appear to be computable in a reasonable amount of time by use of standard algorithms. The solution to this problem is described in the following.

2 Maximum Independent Sets of the 120-cell

The 120-cell is the largest regular polytope in 4 dimensions. Its properties are described in Coxeter’s book on Regular Polytopes [16] and Stillwell’s survey paper [17], for example. The 1-skeleton of the 120-cell is a 4-regular graph with 600 vertices, 720 pentagonal faces and 120 dodecahedral cells. Models exhibiting three-dimensional projections of the complete object have been constructed; photographs of Donchian’s models are shown in [16]. A partial model of the 120-cell is given as an example of a construction using Zome Models [18, Ch. 21]; this 330-vertex subgraph has 45 of the 120 dodecahedra, arranged in concentric shells of 1, 12 and 32 face-sharing cells.

In the following subsections, the steps leading to the solution of the problem of finding the maximum independent set order of the 120-cell are described. First, a description is given of how the vertices of the graph are numbered and how its automorphism group is computed (§2.1). Then (§2.2) a lower bound of 220 for the maximum independent set order is derived from an independent set of order 110 in the antipodal collapse of the 120-cell. An upper bound of 221 is established by use of a linear program (§2.3), and the information from the solution to the integer program is then exploited (§2.4) to infer structural information about a putative independent set of order 221. This information is subsequently used in the computational search described in the remaining subsections, which establishes that the maximum independent set order of the 120-cell is 220.

2.1 Numbering the Graph and its Automorphism Group

A special breadth-first search was used for numbering the 120-cell and finding a permutation representation of its automorphism group on the 600 vertices. The search in question is performed as follows:

Clockwise BFS Labelling Algorithm:

Input: an adjacency list for the 120-cell or its collapse.

Output: a canonical labelling for the graph and its automorphisms expressed in terms of permutations of the vertices of the canonically labelled graph.

To obtain the initial canonical labelling, select one vertex to be labelled as vertex 0 and then choose one way to label its four neighbours as 1, 2, 3 and 4. The remaining vertices are labelled using a breadth-first search starting at vertex 0, and visiting its neighbours 1, 2, 3, and 4 in order. In order to make the breadth-first search labelling deterministic, the neighbours of a vertex v are visited in an order which is decided as follows. Each unlabelled neighbour u of vertex v is in one pentagon with vertex v and the breadth-first search parent p of vertex v . Let the other two vertices of the pentagon be x and y so that the vertices of this pentagon in cyclic order are u, v, p, x, y . Since x is a neighbour of p (and p is the breadth-first search parent of v), x has already been labelled. The order of the neighbours of vertex v is selected so that the labels of the vertices indicated by x in their pentagons are sorted in increasing order. Figure 1 shows a portion of the 120-cell labelled this way.

The 120-cell has an automorphism group of order 14,400. Obtaining the permutations of the automorphism group is easy using the clockwise BFS as described above, as they correspond to choosing a start vertex for the BFS in each of the 600 possible ways, and for each start, considering each of the $4!$ permutations of its neighbours ($14,400 = 600 \times 4!$).

2.2 A Lower Bound from the Antipodal Collapse

Given a vertex v of a graph, its *antipodal* vertices are those at maximum distance from v . The 120-cell has the property that each of its vertices has a unique antipodal vertex. The *antipodal collapse* of

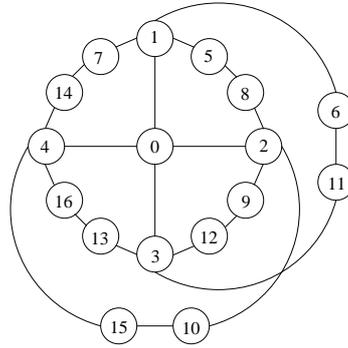


Figure 1: A subgraph of the 120-cell, as labelled by the Clockwise BFS Labelling Algorithm.

the 120-cell is obtained by identifying each vertex of the 120-cell with its antipodal vertex. If $\{u, u'\}$ and $\{v, v'\}$ are two sets of antipodal pairs of vertices of a graph G , then in the collapse, there is one edge between $\{u, u'\}$ and $\{v, v'\}$ corresponding to each edge of the form (u, v) , (u, v') , (u', v) , or (u', v') of the original graph G . Since multiple edges are inconsequential for the independent-set problem, each multiple edge is replaced by a single edge. The result is a 4-regular graph on 300 vertices which has the same local structure as the 120-cell. This graph is the 1-skeleton of the hemi-120-cell, one of the projective regular polytopes of rank 4 in projective 3-space [19, Section 6C].

The automorphism group order of the collapse is $300 \times 4! = 7200$, and the Clockwise BFS Labelling Algorithm from Section 2.1 is first used to label the vertices and find the automorphisms. An independent set of order k in the antipodal collapse can be lifted to one of order $2k$ in the 120-cell (if a vertex is in the independent set in the collapse, then include the two corresponding vertices of the 120-cell).

A non-exhaustive computer search indicated that the antipodal collapse has at least 60 independent sets of order 110 up to isomorphism. This shows that the 120-cell has an independent set of order 220. The most symmetrical of the sets that we found in the antipodal collapse has stabiliser group of order 8. This set is lifted (Table 1) to give an independent set of order 220 in the 120-cell, with 16 automorphisms.

2.3 A Linear Programming Upper Bound

An upper bound of 221 is not difficult to prove by solving a linear programming problem which sets up necessary constraints for a maximum independent set of the 120-cell. The nine variables for this linear program are as follows:

R = the number of red (independent set) vertices

B_i for $i = 0, 1, 2, 3, 4$ = the number of blue (not in the independent set) vertices having i red neighbours.

P_i for $i = 0, 1, 2$ = the number of pentagons with i red vertices. Each of the nine variables is constrained to be non-negative. The LP has six further constraints (five equalities and one inequality). These are introduced after proving some theorems required to justify the sixth constraint. The other constraints are all trivial conditions.

A *blue pentagon* is a pentagon whose vertices are all blue (i.e., none of them are in the indepen-

dent set). An *isolated blue pentagon* is defined to be a blue pentagon such that all of its ten incident vertices (i.e., the ten vertices that are adjacent to a vertex of the pentagon but are not themselves in the pentagon) are red. A blue pentagon with at least one incident blue vertex is called a *non-isolated blue pentagon*. A blue vertex with one red and three blue neighbours is called a *key*. A blue vertex with four red neighbours is called an *isolated blue vertex*.

Remark 1. The independent set of order 220 listed in Table 1 has no isolated blue vertices and no keys.

Theorem 2. For any maximum independent set, the number of non-isolated blue pentagons is at most B_1 (the number of keys).

Proof. Note that for a maximum independent set, it is not possible to have a blue vertex v with four blue neighbours since otherwise the independent set order could be increased by colouring v red. Therefore, if there is a blue pentagon which is a non-isolated blue pentagon, there must be at least one blue vertex on that pentagon which has one red and three blue neighbours (a key). The number of vertices like this is B_1 . This does not complete the proof however because a key can be on 0, 1, 2, or 3 blue pentagons.

To finish the proof, start by assigning a weight to each vertex v of the graph which is a key: assign a weight of one if v is contained in at least one non-isolated blue pentagon and zero otherwise. The sum of the weights of the keys is at most B_1 .

Next, assign fractional weights to the non-isolated blue pentagons- if a key v is on r blue pentagons, this key contributes a weight of $1/r$ to each of its blue pentagons. The sum of the weights of the non-isolated blue pentagons is equal to the sum of the weights of the keys.

To finish the proof, we argue that for each of the non-isolated blue pentagons, the sum of the contributions from its keys is at least one, meaning that the number of non-isolated blue pentagons is at most B_1 . This argument is broken down into three cases according to the types of keys on each non-isolated blue pentagon P .

0	5	6	9	10	13	14	21	22	23	24	29	32	37	39
42	46	47	55	58	60	61	64	68	69	71	74	76	78	81
83	85	89	90	91	93	95	98	100	102	105	108	109	113	114
116	119	122	129	132	133	136	138	142	148	150	154	155	162	167
171	172	173	178	182	185	186	190	193	194	195	196	197	199	202
210	211	216	217	220	222	227	228	229	230	232	236	242	243	248
249	253	259	260	263	265	267	274	277	280	281	282	283	284	286
289	290	292	293	297	300	304	309	311	312	316	317	318	319	322
324	326	328	329	333	334	335	336	343	346	348	350	357	366	370
373	374	375	379	380	381	385	390	391	398	400	406	410	411	414
417	419	421	423	426	427	428	431	433	435	436	437	440	441	442
443	455	458	462	465	466	470	475	476	480	482	489	493	495	497
500	505	507	508	509	510	511	514	517	518	520	524	525	528	529
533	535	537	540	542	544	545	551	556	557	560	563	566	570	571
574	579	581	584	586	588	589	592	594	599					

Table 1: An independent set of order 220 in the 120-cell generated from a set of order 110 in the antipodal collapse.

Case 1: Pentagon P contains a key v which is only in one non-isolated blue pentagon. In this case, the weight that v contributes to P is one and so P has weight at least one.

Case 2: Pentagon P contains a key v which is in two non-isolated blue pentagons. Let A and B be the two non-isolated blue pentagons containing v and let (v, x) be the edge common to A and B . Vertex x is also a key. If it is a key which is in exactly two non-isolated blue pentagons then the weight of P is at least one, since each of v and x contributes $1/2$ to the weight of P . If x is in three blue pentagons, then consider Case 3 instead of Case 2.

Case 3: Pentagon P contains a key v which is in three non-isolated blue pentagons. Let the three blue neighbours of v be x , y and z where x and y are the vertices which are on P . Since v is on three non-isolated blue pentagons, x and y are either on two or three non-isolated blue pentagons and hence they contribute at least $1/3$ to each pentagon they are on. Since P has contributions of at least $1/3$ from v , x , and y , the sum of its contributions is at least one, as required. \square

Corollary 3. *For any maximum independent set of the 120-cell, the number of isolated blue pentagons is at least $P_0 - B_1$.*

Theorem 4. *For a maximum independent set of the 120-cell, if I is the number of isolated blue pentagons, then $2I \leq P_1$.*

Proof. In a dodecahedron, an isolated blue pentagon P appears as a blue pentagon with five incident red vertices. This means that the only possibility for another blue pentagon in the dodecahedron is the pentagon Q antipodal to P (all other pentagons contain at least one of the five reds). But the vertices in the dodecahedron incident to Q cannot be red (they have neighbours which are red) and therefore, a dodecahedron contains at most one isolated blue pentagon.

Any edge (u, v) of the pentagon Q antipodal to P with both endpoints blue is in a pentagon P' with exactly one red vertex in the dodecahedron which contains P and Q . Since the pentagon Q has at least one edge with both endpoints blue, there is at least one pentagon P' with exactly one red vertex in the dodecahedron with P and Q .

Each pentagon of the 120-cell falls into exactly two dodecahedra. To finish the proof, we argue that a pentagon P' with exactly one red vertex occurs as one which must be present as described above because of an isolated blue pentagon in at most one of its two dodecahedra. Suppose that P' corresponds to isolated blue pentagons in both of its two dodecahedra. Then the picture must be as in Figure 2 where the isolated blue pentagons are A and B , and P' is the pentagon with the bold edges. Vertex x is incident to A and vertex y is incident to B so both x and y must be red. This is a contradiction since x and y are adjacent to each other in the 120-cell.

The conclusion is that each isolated blue pentagon maps to at least one pentagon with exactly one red in each of its two dodecahedra. Further, such pentagons with one red correspond to at most one isolated blue pentagon of the graph. This implies that $2I \leq P_1$. \square

We now have the necessary theory to justify an integer programming problem which provides necessary constraints on a maximum independent set of the 120-cell. The conditions for the integer programming problem are:

1. $B_1 + B_2 + B_3 + B_4 + R = 600$
2. $P_0 + P_1 + P_2 = 720$
3. $4R = 1B_1 + 2B_2 + 3B_3 + 4B_4$
4. $6R = 0P_0 + 1P_1 + 2P_2$

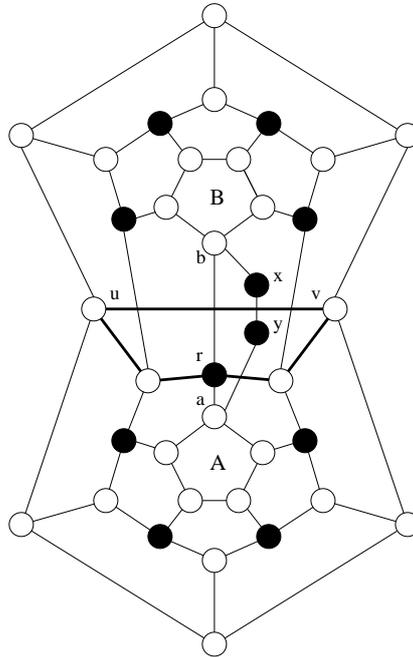


Figure 2: Two isolated blue pentagons A and B sharing a pentagon P' that has one red vertex (P' outlined in bold).

$$5. 5P_0 + 2P_1 = 3B_1 + 1B_2$$

$$6. P_1 \geq 2(P_0 - B_1)$$

The justifications for these constraints are:

1. The 120-cell has 600 vertices and for a maximum independent set $B_0 = 0$ as noted earlier.
2. The 120-cell has 720 pentagons.
3. Each red vertex is incident to four blue vertices. Hence, four times the number of red vertices is equal to the number of times a blue vertex is adjacent to a red one.
4. Each vertex is in six pentagons. Hence, six times the number of red vertices is equal to the number of times a red vertex occurs in a pentagon.

5. A *blue 2-path* is a path on 3 vertices (and hence two edges) whose vertices are all blue. Since each 2-path of the graph is in a unique pentagon, the number of blue 2-paths is $5P_0 + 2P_1$. The number of blue 2-paths is also equal to $3B_1 + 1B_2$ since a blue vertex with three blue neighbours is the centre of three blue 2-paths, a blue vertex with two blue neighbours is the centre of one blue 2-path, and a blue vertex with zero or one blue neighbours is not the central vertex of any blue 2-path.

6. This constraint comes from combining Corollary 3 with Theorem 4.

To get an upper bound for the maximum independent set order, the objective function is to maximize the value of R. Solving the linear programming relaxation gives an upper bound of $2880/13 = 221.538\dots$ which gives an upper bound of 221 on the integer programming problem. (The optimum solution is attained for the vector $P_0 = 360/13, P_1 = 720/13, P_2 = 8280/13, B_1 = 0, B_2 = 3240/13, B_3 = 1680/13, B_4 = 0$, and the polytope thus defined is three-dimensional and has nine-vertices.) Applying the same tactic to the antipodal collapse gives an upper bound of

P_0	P_1	P_2	B_1	B_2	B_3	B_4
25	64	631	0	253	126	0
26	62	632	0	254	124	1
27	60	633	0	255	122	2
28	58	634	0	256	120	3
26	62	632	1	251	127	0
27	60	633	1	252	125	1
28	58	634	1	253	123	2
29	56	635	1	254	121	3
27	60	633	2	249	128	0
28	58	634	2	250	126	1
29	56	635	2	251	124	2
28	58	634	3	247	129	0
29	56	635	3	248	127	1
30	54	636	3	249	125	2
29	56	635	4	245	130	0
30	54	636	4	246	128	1
30	54	636	5	243	131	0
31	52	637	5	244	129	1
31	52	637	6	241	132	0
32	50	638	7	239	133	0

Table 2: Solutions to the Linear Program that would correspond to an independent set of order 221.

110 for the collapse, implying that the independent set of order 110 found in §2.2 is a maximum independent set of the antipodal collapse.

2.4 Exploiting the Integer Program Information

The example in §2.2 gives a lower bound of 220 for the order of a maximum independent set of the 120-cell. On the other hand, §2.3 proves an upper bound of 221. This implies that if the independent set from §2.2 is not optimal, then there is an independent set of the 120-cell of order 221. The next step is to examine the solutions of the integer programming problem from §2.3 which have the number R of red vertices equal to 221 to gain structural information as to what a solution of order 221 must look like.

Table 2 shows all the integer solutions that could result in an independent set of order 221. Correctness of the LP code is not an issue since it is not hard to check the final solutions by hand.

Scanning the table of solutions, we observe that $P_0 - B_1$ is always at least 25. From Corollary 3, the implication is that any independent set of order 221 has at least 25 isolated blue pentagons. Observe also that all cases satisfy the constraint that the number B_1 of keys plus two times the number B_4 of isolated blues is at most seven.

The existence of an independent set of order 221 requires that there is some way to add at least 25 isolated blue pentagons to the 120-cell without creating too many keys or isolated blue vertices ($B_1 + 2B_4 \leq 7$). The next two sections explain how we first tried planting a smaller number of isolated blue pentagons in part of the graph in all ways up to isomorphism and give an account of how the search for the 25 isolated blue pentagons was completed.

2.5 Planting Blue Pentagons

A typical approach to trying to plant 25 isolated blue pentagons into the 120-cell that covers all possibilities up to isomorphism is to choose some smaller number of pentagons (for example, seven) that are placed in all ways up to isomorphism and then add the rest without concern for duplication since at some point, isomorphism screening is too costly for the amount of duplication prevented. This approach was taken first and it resulted in too many cases for a practical solution. The next strategy applied was to consider only a subgraph of the 120-cell for which it is possible to prove that at least some number k of isolated pentagons must be present in order to reach an independent set of order 221, and then to place the k pentagons in this region in all ways up to isomorphism.

It is assumed that vertices of the 120-cell are labelled by the Clockwise BFS Labelling Algorithm from Section 2.1 The *restricted region* for consideration is defined by first sorting the pentagons. Before sorting, a list of five integers is created for each pentagon which contains the labels of its vertices in reverse sorted order (which is not necessarily the same as a cyclic ordering). Then these lists are compared lexicographically while sorting the pentagons. This gives a sorted order of pentagons P_0, P_1, \dots, P_{719} . The first six pentagons are the ones pictured in Figure 1. The sequences used to sort them are:

- $P_0 : 8, 5, 2, 1, 0$
- $P_1 : 11, 6, 3, 1, 0$
- $P_2 : 12, 9, 3, 2, 0$
- $P_3 : 14, 7, 4, 1, 0$
- $P_4 : 15, 10, 4, 2, 0$ and
- $P_5 : 16, 13, 4, 3, 0.$

The last two pentagons (illustrating how lexicographic order is used to break ties) are:

- $P_{718} : 599, 598, 596, 592, 591$ and
- $P_{719} : 599, 598, 597, 594, 593.$

This (slightly unnatural) ordering was selected so that the maximum vertex number occurring in the pentagons numbered P_0, P_1, \dots, P_k is minimized given a chosen value of k . Intuitively, this helps to compress the first k pentagons into a small subgraph of the 120-cell.

After some experimentation, it was decided that planting seven pentagons in all ways up to isomorphism was the best compromise between the number of cases created and the difficulty for finishing the cases. The restricted region for planting these pentagons is shown to consist of the first 173 pentagons (P_0, P_1, \dots, P_{172}) in the following lemma.

Lemma 5. *If the 120-cell has an independent set of order 221 then it is possible to find an independent set of order 221 such that there are at least seven isolated blue pentagons in the first 173 pentagons (P_0, P_1, \dots, P_{172}).*

Proof. We already know from the results in §2.3 that the entire graph contains at least 25 isolated blue pentagons if there is an independent set of order 221. The idea for this proof is to count the number of isolated blue pentagons in the graph by considering the sets of pentagons numbered P_0, P_1, \dots, P_{172} for each of the automorphisms of the graph. If the average count over each of these sets P_0, P_1, \dots, P_{172} is greater than six, then we can conclude that there is at least one automorphism of the graph such that the count for P_0, P_1, \dots, P_{172} is at least seven.

Owing to the structure of the automorphism group of the graph, taking into consideration the sets of pentagons labelled P_0, P_1, \dots, P_{172} over all automorphisms accounts for each pentagon the same number of times; each is included $14,400 \times 173/720$ times. If the graph has 25 or more isolated blue pentagons, then the sum of the number of isolated blue pentagons over each choice for

P_0, P_1, \dots, P_{172} is equal to at least $25 \times 14400 \times 173/720$. Hence, the average term is equal to at least $25 \times 173/720$. But $25 \times 173/720 > 6$ and therefore, since the average is greater than six, at least one case must be greater than six. □

The total number of ways to plant seven isolated blue pentagons in the set P_0, P_1, \dots, P_{172} is equal to 8,211,380. It is a little more difficult than usual to define a canonical form for these, because some of the automorphisms of the graph can map a choice of seven pentagons selected from P_0, P_1, \dots, P_{172} to another choice of seven pentagons which is lexicographically smaller, but is no longer a selection from the pentagons P_0, P_1, \dots, P_{172} because the new set contains a pentagon numbered 173 or higher. To accommodate this difficulty, the canonical form is selected so that it is the lexicographically minimum set of seven pentagons with the additional property that the pentagon with the largest number in the set corresponds to some P_k for $k \leq 172$.

The algorithms used for this phase were very simple. A nested set of seven loops was used to plant all possible choices for seven isolated blue pentagons. For each isolated blue pentagon selected, the ten incident vertices are coloured red. Vertices adjacent to a red vertex are coloured blue. To determine if an additional choice for an isolated blue pentagon is compatible with a previously chosen set, it suffices to check if its ten incident vertices can all legally be coloured red (that is, they are either uncoloured or red already, but cannot be blue).

Then the 8,211,380 ways to place the isolated blue pentagons were run through a screen which kept only the canonical cases. For this step, the automorphism group of the graph was precomputed as described in §2.1. As a check on the computation, for each of the 1,379,646 cases retained, we determined the number of valid images it had (that is, the number of ways to renumber it with an automorphism such that the largest label on a pentagon is 172). The sum of these was equal to 8,211,380 (the number of cases possible without removing duplicates), a necessary condition for correctness. As an additional check of correctness, the number of cases to consider up to isomorphism matches that from a computation done earlier with a different approach.

2.6 Finishing the Search by a Backtrack

For each of the 1,379,646 non-isomorphic ways of planting seven blue pentagons in the pentagons P_0, P_1, \dots, P_{172} (described in §2.5), the next step is to determine if it is possible to extend the configuration so that it contains at least 25 isolated blue pentagons. The possibilities for an independent set of order 221 outlined in §2.4 indicate that a solution of order 221 does not have many keys or isolated blues, more specifically, that $B_1 + 2B_4 \leq 7$.

A backtracking routine was used to try to extend each of the cases with the seven isolated blue pentagons to 25 isolated blue pentagons without creating too many isolated blues or keys in the process. Some tricks were used to make this computation finish in a reasonable amount of time.

The backtracking algorithm at level k considers two cases: one where the pentagon P_k is not included as an isolated blue pentagon, and if feasible, a second case where the pentagon P_k is included as an isolated blue pentagon (which means that its ten incident vertices are coloured red).

The colour of a vertex is recorded as an integer which is 0 if a vertex is not coloured. The colour is decremented each time a vertex is coloured red, or incremented each time a vertex is coloured blue. This permits the algorithm to colour vertices then backtrack by uncolouring the vertices without using a data structure such as a stack to indicate vertices with a status change. Only blue or white vertices can legally be coloured blue. Only red or white vertices can legally be coloured red. If a vertex is coloured red, then its neighbours are immediately coloured blue. When the colour of a vertex returns to zero, it returns to the uncoloured status.

As the algorithm progresses, certain vertices can safely be coloured red. These are characterized in the next theorem.

Theorem 6. *Suppose a 120-cell has an independent set of vertices coloured red, the neighbours of these are coloured blue, and the remaining vertices are uncoloured. If there is an uncoloured vertex v with three blue neighbours and one uncoloured neighbour w , then if there is a maximum independent set of the 120-cell which is consistent with the entire colouring, there is a maximum independent set with v coloured red.*

Proof. If v is red in the maximum independent set then there is nothing to prove. If v is not red, then w is red because if w is blue instead, v is a blue vertex with four blue neighbours, contradicting the maximality of the independent set (colouring v red increases the independent set order). An independent set of the same order can then be found by changing the colouring so that v is red and w is blue. □

The algorithm first inserts the initial seven isolated blue pentagons, waiting until they are all included before applying Theorem 6. The delay is needed because applying the theorem earlier can result in a colouring inconsistent with the initial pentagons. If there are two uncoloured vertices u and v which are adjacent to each other and also each is adjacent to three blues, there are two choices for how to apply Theorem 6, and it is possible that only one of these is consistent with the initial selection of the seven isolated blue pentagons.

During the course of the backtrack, each time an isolated blue pentagon is added to the current configuration, a queue is used to record vertices which evolve to being white with three blue neighbours. As the goal is to try to add 25 isolated blue pentagons, 25 queues suffice.

After addition of the ten incident reds of the isolated blue pentagon, the algorithm traverses the queue, and each vertex in the queue which is not blue is coloured red (as noted in the last paragraph, applying Theorem 6 at a vertex may prevent its subsequent use at another vertex). This process can trigger the addition of further vertices to the queue. When the isolated blue pentagon is removed (when the computation backtracks), the queue is first traversed in the reverse order to undo these changes.

New isolated blue vertices are recorded at the point when the fourth neighbour of the isolated blue is initially coloured red. The number is decremented when this fourth neighbour becomes uncoloured again. Keys arise either when an uncoloured vertex with three blue neighbours is coloured blue or when a third neighbour of a blue vertex is initially coloured blue. To facilitate the detection of isolated blue vertices and keys, respectively, the number of red neighbours and the number of blue neighbours of each vertex are maintained.

The algorithm takes exponential time to run, which is not surprising as the problem is hard. A careful selection of the data structures results in an approach such that the work it does to maintain the data structures is at most a constant times the number of times a vertex is coloured red.

The algorithm also used some precomputed upper bounds. We determined the maximum number of isolated blue pentagons possible if the pentagons are chosen from $P_k, P_{k+1}, \dots, P_{719}$ such that the *penalty* (equal to $B_1 + 2B_4$) is at most seven (as required for an independent set of order 221). There is no point in continuing this computation past the point where 18 isolated blue pentagons are possible: since we start with seven, only 18 more are required. Theorem 6 was not used for computing these upper bounds, owing to its interference with what we were trying to compute. At a given level of the backtrack for placing the 25 blue pentagons, if the number of isolated blue

Polytope	n	m	f	r	g	D	$ C_d $
Icosahedron	12	30	20	5	3	3	3, 2
Dodecahedron	20	30	12	3	5	5	8, 4, 2, 2
24-cell	24	96	96	8	3	3	8, 2
600-cell	120	720	1200	12	3	5	24, 8, 3, 2
120-cell	600	1200	720	4	5	15	220, 120, 48, 28, 24, 10, 8, 5, 5, 3, 2, 2, 2, 2

Table 3: Exceptional polytopes in dimensions three and four. The columns n , m , and f give the numbers of vertices, edges and two-dimensional faces of the polytope; r , g and D are the vertex degree, girth and diameter of the graph. The entries $|C_d|$ are the maximum orders of d -codes for $d = 2, 3, \dots, D - 1$.

pentagons included so far plus the bound for the level is less than 25, the algorithm backs up, since it is necessary to have at least 25 isolated blue pentagons for an independent set of order 221.

It is possible in the course of the algorithm that an isolated blue pentagon which has yet not been considered ends up with all ten of its incident vertices red. This however does not preclude the algorithm from adding it: the incident vertices just get coloured red more than once.

The algorithm for this last backtrack was coded independently twice to ensure correctness. The 1,379,646 cases were split into 64 slices, and run in parallel on a 64-processor array, with the run of the C program for a typical slice taking 16 – 18 hours. Both programs concluded that it is not possible to include 25 blue pentagons in the 120-cell with a penalty of seven or less after applications of Theorem 6 are taken into account. Because this must be possible for an independent set of order 221 to exist, the maximum independent set order of the 120-cell is 220.

3 Other Results

Table 3 lists the orders of the maximum d -codes for all five exceptional polytopes. Apart from the 2-code of the 120-cell, the only case requiring special tactics is the 4-code of the 120-cell, which was solved as described in [20]. All five polytopes have antipodal pairs as their d -codes for $d = D$, the diameter of the graph. When the codes are considered in terms of their ‘contact polytopes’ [9], some interesting ‘Russian Doll’-like interconnections are seen. In the sense used in previous work [9], the contact polytope of a d -code has the same vertices as the independent set, and two vertices of the contact polytope are joined by an edge if the independent-set vertices are at distance d in the parent graph. Simplices of dimensions two, three and four appear: the triangle (α_2) is the contact polygon of the 3-code of the icosahedron, the 4-code of the 600-cell and the 11-code of the 120-cell; the tetrahedron (α_3) is the contact polyhedron of the 3-code of the dodecahedron; the four-dimensional simplex (α_4) is the contact polytope of the 9-code of the 120-cell. The cube appears (γ_3) appears as the contact polytope of the 2-code of the dodecahedron. The hyperoctahedron (β_4) appears as the contact polytope of 2-code of the 24-cell, 3-code of the 600-cell and the 8-code of the 120-cell. The 24-cell itself is the contact polytope of the 2-code of the 600-cell. The 3-code of the 120-cell is a *perfect* code [14] in the sense that each vertex of the code is at the centre of a ball of radius 1 containing one vertex of the 120-cell and its four nearest neighbours; the 120-cell is then partitioned exactly into 120 such balls, with centres on the vertex set of a 600-cell whose edges are paths of length 3 in the 120-cell. These observations are closely related to the fact, pointed out by Coxeter [16], that the vertices of the 120-cell embedded as an equilateral object in four-dimensional space include the vertices of all fifteen of the other regular polytopes in four dimensions, a property that has no analogy in three dimensions, where the dodecahedron contains the vertices of the cube and

tetrahedron, but not those of the octahedron or icosahedron.

References

- [1] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Co., San Francisco, 1979.
- [2] I.M. Bomze, M. Budinich, P.M. Pardalos and M. Pelillo, *The maximum clique problem* In: Handbook of combinatorial optimization, Supplement Vol. A, 1–74, Kluwer Acad. Publ., Dordrecht, 1999.
- [3] P.R.J. Östergård, Constructing combinatorial objects via cliques. *Surveys in combinatorics 2005. London Math. Soc. Lecture Note Ser.*, **327** (2005), 57–82, Cambridge Univ. Press, Cambridge.
- [4] D.S. Johnson and M.A. Trick, *Cliques, Coloring, and Satisfiability, Second DIMACS Implementation Challenge, Oct. 11–13, 1993* DIMACS Series in Discrete Mathematics and Theoretical Computer Science, (26), AMS, Providence, RI, 1996.
- [5] J. Debroni, J.D. Eblen, M.A. Langston, W. Myrvold, P. Shor and D. Weerapurage, A complete resolution of the Keller maximum clique problem, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 11)*, (2011), 129–135.
- [6] M. Dutour Sikirić and W. Myrvold. The special cuts of the 600-cell. *Beiträge Algebra Geom.* **49** (2008), 269–275.
- [7] F.J. MacWilliams and N.J.A. Sloane, *Theory of Error-Correcting Codes*, North-Holland, Amsterdam, 1978.
- [8] J.H. Conway and N. J. A. Sloane, *Sphere packings, lattices and groups*, Springer-Verlag, New York, 1988.
- [9] B. de La Vaissière, P.W. Fowler and M. Deza, Codes in Platonic, Archimedean, Catalan and related polyhedra: a model for maximum addition patterns in chemical cages, *J. Chem. Inf. Comp. Sci.* **41** (2001), 376–386.
- [10] J.D. Crane, Maximal Non-Adjacent Addition to fullerene-70: Computation of All the Closed Shell Isomers of $C_{70}X_{26}$, *Fullerene Sci. Technol.* **2** (1994), 427–435.
- [11] S. Fajtlowicz and C. E. Larson, Graph-theoretic independence as a predictor of fullerene stability, *Chem. Phys. Lett.* **377** (2003), 485–490.
- [12] P.W. Fowler, S. Daugherty and W. Myrvold, Independence number and fullerene stability, *Chem. Phys. Lett.*, **448** (2007), 75–82.
- [13] P.W. Fowler, P. Hansen, K.M. Rogers and S. Fajtlowicz, $C_{60}Br_{24}$: a chemical illustration of graph theoretical independence, *J. Chem. Soc., Perkin Trans. 2* (1998), 1531–1533.
- [14] P.W. Fowler, B. de La Vaissière and M. Deza, Addition patterns, codes and contact graphs for fullerene derivatives, *J. Mol. Graphics Mod.*, **19** (2001), 199–204.
- [15] P.W. Fowler, K.M. Rogers, K.R. Somers and A. Troisi, Independent sets and the prediction of addition patterns for higher fullerenes, *J. Chem. Soc., Perkin Trans. 2*, (1999), 2023–2027.
- [16] H.S.M. Coxeter, *Regular Polytopes*, 3rd Edition, Dover, New York, 1973.
- [17] J. Stillwell, The story of the 120-cell, *Notices of the AMS*, **48** (2001), 17–24.
- [18] G. W. Hart and H. Picciotto, *Zome Geometry: Hands-on Learning with Zome Models*, Key Curriculum Press, Emeryville, CA, 2000.
- [19] P. McMullen and E. Schulte, *Abstract Regular Polytopes*, Cambridge University Press, Cambridge, 2002.
- [20] W. Myrvold and P.W. Fowler, Fast enumeration of all independent sets up to isomorphism, Preprint.