ENUMERATION OF HOMOTOPY CLASSES OF FINITE T_0 TOPOLOGICAL SPACES

ALEX FIX AND STEPHEN PATRIAS

ABSTRACT. We present a method for enumerating all homotopy equivalence classes for finite spaces of a given size. This work expands on the methods of May, Barmak and Minian in classifying homotopy equivalence as an isomorphism of a certain subgraph of the Hasse diagram of the space called the minimal core. We thus convert the topological notion of homotopy equivalence into purely graph theoretical terms, allowing us to apply the methods of Brinkmann and McKay for enumerating all posets of a given size. An algorithm is developed, which we run to give exact figures for the number of homotopy classes for small spaces. Finally, we show that in fact, the number of homotopy classes of spaces is asymptotically equal to the number of all posets on n vertices.

CONTENTS

1. INTRODUCTION

The study of finite topological spaces provides a distinctive and largely unexplored view into the nature of spaces at large, particularly their homotopy theory. Indeed, the calculation of homotopy groups has proven to be quite difficult, especially the higher order homotopy groups. However, finite spaces are essentially combinatorial objects, as we shall show, and thus hopefully some of the difficult questions of Algebraic Topology can be translated into known questions about graphs and other finite objects.

The number of all finite topological spaces on a fixed set of n points is astounding for even small values of n. Of course, considering only the number of spaces up to homeomorphism decreases this number significantly (from 130023 spaces on 6

Date: DEADLINE AUGUST 22, 2008.

points to 318 spaces up to homeomorphism). This still leaves a large number of distinct spaces, and as a result, it is typical to examine another equivalence relation which narrows the field of study even further.

Definition 1.1. For two continuous maps $f, g : X \to Y$ between topological spaces, define f to be *homotopic* to g if there is a continuous map $h : X \times [0, 1] \to Y$ where $h(0) = f$ and $h(1) = g$; h is a homotopy between f and g and we write $f \simeq g$. Conceptually, this relation describes when one function can be continuously deformed into the other. We say that two topological spaces X and Y are *homotopy equivalent* when there exists two functions $f : X \to Y$ and $g : Y \to X$ such that $f \circ g \simeq \mathrm{Id}_Y$ and $g \circ f \simeq \mathrm{Id}_X$, and we write this relation as $X \simeq Y$.

Note that when studying finite spaces up to homotopy equivalence (as we do here), it suffices to consider only the T_0 spaces: that is, the spaces whose topologies distinguish points. To be precise, two points x and y are topologically indistinguishable in a space if for any open subset U, then either U contains both or neither of x and y. A space is T_0 if there are no two distinct topologically indistinguishable points.

In order to obtain a T_0 space from a non-T_0 space X, we note that inseperability forms an equivalence relation on the points of the space, and thus we can obtain X_0: the quotient space of X by this relation. By construction, this relation gives a quotient space with no inseparable points, so X_0 is a T_0 space. Furthermore, by a theorem of McCord[1] the quotient map $q_X : X \to X_0$ is itself a homotopy equivalence and thus when considering properties invariant under homotopy equivalence, it suffices to consider only the T_0 spaces, since every space is canonically homotopy equivalent to a T_0 space.

2. Finite Topological Spaces and Partial Orders

This section follows the approach of J. P. May[1] in describing a finite topological space in terms of a partial order. This is much closer to our goal of describing spaces in terms of combinatorial objects. Remember that a poset is a set with an associated partial order; that is, it has a relation which is reflexive, transitive and antisymmetric. Additionally, by the above argument, we assume henceforth that all finite topological spaces are T_0.

We develop the equivalence between finite topological spaces and posets as follows:

Definition 2.1. For each point x in a finite space X, we define the set U_x as the intersection of all open subsets of X which contain x. Then, we define a relation $\leq$ on X so that for any $x, y \in X$, $x \leq y$ whenever $x \in U_y$, or equivalently, whenever $U_x \subset U_y$

Note that the set U_x is always an open set in X, since there can be at most finitely many open sets which contain the point x, and the finite intersection of open sets is again open.

Furthermore, note that if the space X is T_0 then the order relation $\leq$ is a partial order. Specifically, this order is clearly reflexive and transitive, and we claim that T_0 implies antisymmetry. Indeed, let x, y be points of X such that $x \leq y$ and $y \leq x$. Then $U_x = U_y$, so any open subset of X either contains both U_x and U_y, or it contains neither, and specifically, it either contains both x and y or it contains neither. But since X has no indistinguishable points, $x = y$.

We also have the reverse construction: given any finite poset P, we construct a finite topological space as follows. Let the points of the space X be the points of the poset. For each point x, form the set $U_x = \{y \mid y \leq x\}$. The open sets of the topology are then the collection of finite unions of the U_x, so that the U_x form a basis for the topology.

Finally, we have a correspondence between the morphisms of finite T_0 spaces and posets; namely, a function $f : X \to Y$ is continuous if and only if f is order-preserving relative to the associated partial orders of X and Y. However, for proof of this fact, and for more details on the technical nature of this correspondence we refer the reader to the lecture notes of May[1]. Instead, we sum up with a theorem giving the nature of this construction of partial orders from finite spaces.

Theorem 2.2. *The construction of partial orders from T_0 topologies gives a bijection between the set of all finite T_0 topologies and the set of all finite posets. Furthermore, two spaces X and Y are homeomorphic if and only iff there is an isomorphism of the associated posets P_X and P_Y.*

Proof. It is clear that the construction of posets from T_0 spaces, and the reverse construction back into spaces are in fact inverses of each other, so we do obtain a bijection.

Then, two spaces X and Y are homeomorphic when there is a continuous function $f : X \to Y$ such that $f^{-1} : Y \to X$ exists and is continuous. By the above theorem, this is true if and only if f and f^{-1} are order-preserving relative to the posets for X and Y. But a function which is both invertible and order-preserving in both directions is precisely an isomorphism of posets. □

3. Hasse Diagrams

A much more concrete way of seeing this correspondence between finite spaces and partial orders can be gained by looking at the graph which is associated with every partial order.

Definition 3.1. For every partial order P, we define its associated *Hasse diagram* H, a directed graph which captures all the relevant order information of P. Let the vertices of H be the points of P, and the edges of H are such that there is a directed edge from x to y whenever $y \leq x$ but there is no other vertex z such that $y \leq z \leq x$.

The edges of the graph H convey the notion of successor and predecessor. That is, there is an edge from x to y if x is a successor of y, *i.e.*, x is greater than y but there is no other point in the poset which is between x and y. We also see that the statement $x > y$ is equivalent to the existence of a path of directed edges from x to y. Indeed, if there is a path $x \to x_1 \to \cdots \to x_k \to y$ then $x > x_1 > x_2 > \ldots > x_k > y$ so $x > y$. Conversely, if $x > y$ and x is not a sucessor y, then we can find z so that $x > z > y$ and by doing this recursively (since the graph is finite), we can find $x > x_1 > \ldots > x_k > y$ so that each step is to a predecessor, and thus there is a path $x \to x_1 \to \cdots \to x_k \to y$ in H.

From this, we also see that the Hasse diagram is necessarily acyclic, that is, there are no directed cycles $x \to x_1 \to \cdots \to x_k \to x$ or else we would have $x > x$.

We can also go the other way, from a directed acyclic graph G back to a partial order P, by saying $x \geq y$ in P whenever there is a path (including trivial paths) from x to y in G. However, to do this uniquely, we need the following definitions.

Definition 3.2. We say that an edge $x \to y$ is a *shortcut* in a directed graph G if there is also a path $x \to x_1 \to \cdots \to x_k \to y$ with at least two edges between x and y. Furthermore, we say that a directed acyclic graph is a *partial order diagram* if it has no shortcuts.

Theorem 3.3. *The above construction of the Hasse diagram gives a bijection from partial orders to partial order diagrams. Furthermore, there is an isomorphism of partial orders between two posets P and Q if and only if there is a graph isomorphism between the associated diagrams H_P and H_Q.*

Proof. It is trivial to check that these two constructions are in fact inverses of each other, so that there is in fact have a bijection. Then, a bijection $\sigma : P \to Q$ is order-preserving if and only if it also preserves successors and predecessors (*i.e.*, it preserves edges in the graph), so that σ is an isomorphism of posets if and only if it is also a graph isomorphism of the associated Hasse diagrams. □

Corollary 3.4. *We have a bijection between finite T_0 topologies and Hasse diagrams, so that homeomorphism of spaces is equivalent to graph isomorphism of the two diagrams.*

It is also useful to have a convention for drawing these diagrams, as having an orderly presentation allows both a consistent visual understanding of their structure, and an additional handle for computation with these graphs.

Definition 3.5. For each vertex v in a Hasse diagram of a poset, we define the level of v to be the length of the longest chain $v_1 < v_2 < \cdots < v_k = v$ (so that in this example, $\text{level}(v) = k$).

We have the following important facts about levels:

(1) The level of a vertex v is also the length of the longest path beginning at v
(2) There is always an edge from a point v on level ℓ to some point v' on level $\ell - 1$.
(3) There can never be an edge from a point v in level ℓ to any point v' in level $\ell' \geq \ell$.
(4) Level 1 consists of precisely the minimal points of the graph.

Convention 3.6. When drawing the Hasse diagram of a poset, we always draw level 1 at the bottom, and each subsequent level ℓ immediately above its predecessor, level $\ell - 1$. Thus, all edges in the graph point downwards in the graph, allowing us to omit specifying the directions of edges.

4. Homotopy Equivalence of Finite Spaces

In order to systematically find all possible homotopy equivalences of finite spaces, we proceed by identifying a certain type of point which can be removed from the space without affecting the homotopy type. The following definitions and theorems are all from May [1].

Definition 4.1. Let X be a finite space. We say that a point $x \in X$ is an *upbeat* point if there exists a $y > x$ such that for all $z > x$ then $z \geq y$. Similarly, we say that x is a *downbeat* point if there exists $y < x$ such that for all $z < x$, then $z \leq y$.

Definition 4.2. A finite space X is *minimal* if it is T_0 and contains no upbeat or downbeat points. A *core* of a T_0 space X is a minimal subspace X_0 which is also a deformation retract of X.

Theorem 4.3. *Any finite space X has a core.*

Proof. See [1] for details, but the core of the argument is that if x is a beat point of X, then $X \setminus \{x\}$ is a deformation retract of X. Since the space is finite, we can continue this process inductively until there are no more beat points, and the resulting space will be a deformation retract of the original space. □

Theorem 4.4. *If X and Y are minimal finite spaces and $f : X \to Y$ is a homotopy equivalence then f is a homeomorphism.*

The above two theorems allow the following immediate corollary:

Corollary 4.5. *In order to enumerate all the finite spaces with n points up to homotopy equivalence, it suffices to enumerate the minimal spaces with at most n points up to homeomorphism.*

Proof. Since any finite space X on n points has a core, and this core is a deformation retract of the original space, X is homotopy equivalent to a minimal space on no more than n points. Thus, there is at least one minimal space in every homotopy equivalence class. Additionally, if there are two minimal spaces X and Y in the same class, then there is a homotopy equivalence $f : X \to Y$. But by the above theorem, f is a homeomorphism. So if we enumerate the minimal spaces up to homeomorphism, we pick exactly one representative from each homotopy equivalence class. □

5. Minimal Spaces as Graphs

We now begin the process of converting these topological notions into graph theory, from which actual computations can be made. Primarily, we wish to categorize a minimal space as a property of the associated Hasse diagram. We start first with a description of upbeat and downbeat points as they appear in the graph.

Theorem 5.1. *A point x in a finite space X is an upbeat point if and only if it has in-degree one in the associated Hasse diagram (that is, it has only one incoming edge). Similarly, x is downbeat if and only if it has out-degree one (it has only one outgoing edge).*

Proof. Assume that x is upbeat. Then there exists $y > x$ such that for all $z > x$, $z \geq y$. First, we have that y is a successor of x, since there cannot be any z with $y > z > x$. Thus, there is an edge $y \to x$ in the Hasse diagram. We claim that there is no other edge $y' \to x$ with $y' \neq y$. If there were, then $y' > x$ so since x is upbeat, $y' > y$. But since $>$ is equivalent to the existence of a path, we have that there exists a path $y' \to \cdots \to y$. Hence there is both a path $y' \to \cdots \to y \to x$ and an edge $y' \to x$ which violates the requirement that the Hasse diagram have no shortcuts. Thus, x has exactly one incoming edge.

Conversely, assume there is exactly one y such that $y \to x$. Then for any $z > x$ we have that there is a path $z \to \cdots \to x$. But since there is only one vertex y such that $y \to x$, this path must actually be $z \to \cdots \to y \to x$ so there is also a path from z to y so $z \geq y$. Thus x is upbeat.

The proof for the second claim is exactly symmetric. □

Corollary 5.2. *A space is minimal if and only if for every vertex in its associated Hasse diagram, the in-degree and out-degree are both not equal to one.*

Definition 5.3. Henceforth, we will refer to such a graph as a *minimal graph* for brevity.

We can derive several useful consequences from this classification. For starters, be can begin enumerating the minimal spaces by explicitly constructing graphs which satisfy the above condition (which we will do in the next section). However, we can also use this theorem to derive addition facts about the structure of minimal graphs which might otherwise be difficult to derive using only topological arguments.

Proposition 5.4. *Let G be a minimal graph with at least two vertices. Then each level of G contains at least two vertices.*

Proof. Assume first that level 1 has exactly one vertex v. Then, since G has at least two vertices, there is some vertex v' in level 2. But every vertex in level 2 has an edge to a vertex in level 1, so $v' \to v$ is an edge in the graph. However, since all edges in the graph go from level ℓ' to ℓ, where $\ell' > \ell$, all edges of v' must be to vertices in level 1. But there is only one such vertex, so v' has exactly one downwards edge, contradicting the minimality of G.

Now, assume that some level $\ell > 1$ has exactly one vertex v. This vertex has a neighbor v' on level $\ell - 1$, so $v \to v'$ is an edge in the graph. Now, assume there is some other $w \neq v$ such that $w \to v'$ is also an edge in the graph. Since all edges proceed downwards, we have that w is on some level $k > \ell - 1$. Level ℓ has exactly one vertex and w is not it, so in fact, $k > \ell$. We claim that this implies that there is in fact a path $w \to \cdots \to v$ in the graph, so that the edge $w \to v'$ is a shortcut of the path $w \to \cdots \to v \to v'$, which is not allowed.

To prove this claim, we induct on k: for a vertex w on level $k = \ell + 1$, we have that w must have a neighbor on the next lowest level. But this is level ℓ, so $w \to v$ is an edge in the graph, and hence also a path. Then, for w on level $k > \ell + 1$, we again have that w has a neighbor on the next lowest level, so there is some w' on level $k - 1$ such that $w \to w'$ is an edge. But by induction, $w' \to \cdots \to v$ is a path in G, so $w \to w' \to \cdots \to v$ is also a path in G. □

6. Constructing Posets

Intuitively, we expect that as the number of points in the poset grows large, the number of neighbors of each point in the graph should grow large as well, and that cases where a point has exactly one neighbor should be very rare. We will examine this probabilistic reasoning rigorously in the final section, but for now, it seems a good heuristic that the large majority of graphs will be minimal once n grows large enough, and that non-minimal graphs will be the exception. Thus, it makes sense to try to count the number of minimal graphs by first enumerating all posets of a given size, and then checking to see whether each such generated graph is minimal.

As a reminder, by Corollary 4.5 we are interested in enumerating the minimal spaces up to homeomorphism, and by Corollary 3.4, homeomorphism of spaces is equivalent to graph isomorphism of the constructed Hasse diagrams.

Definition 6.1. Since an isomorphism between graphs is equivalent to relabeling the vertices in a consistent fashion, an equivalence class of graphs under graph isomorphism is called an *unlabeled graph*.

Since any relabeling of a minimal graph produces another minimal graph (as it does not change the in or out degree of any of the vertices), we can treat an unlabeled minimal graph as the equivalence class of a minimal graph under graph isomorphism. This represents the same object as the equivalence class of a minimal space under homeomorphism, so our task is to produce exactly one representative for each unlabeled minimal graph.

Fortunately, a very fast algorithm for producing exactly one representative of each unlabeled Hasse diagram has already been proposed by Brinkmann and McKay[3], and used to enumerate all unlabeled posets on up to 16 points. The remainder of this section will be a summary of these results.

The algorithm works by a method called the canonical construction path which, for every unlabeled poset P on n points, gives a canonical unlabeled poset Q on $n-1$ points such that Q can be obtained from P by deleting a point from the top level. This essentially turns the set of all unlabeled posets into a tree, whereby each poset on n points has a unique parent with $n-1$ points, turning the task of enumeration into a search on this tree.

In order for this construction to work, it is necessary to be able to reconstruct all children of a given poset, and to only construct exactly one example of each child graph (so that we do not produce two different labelings of the same graph, and consider them as different children). It is relatively straightforward to construct the set of all possible children for a graph, however, to reject possible isomorphisms between these candidates we require a device called a canonical choice function.

Definition 6.2. Let C be a set of candidates, each of which is a poset on n points, with vertex set $[n] = \{1, 2, \ldots, n\}$. Then a function $f : C \to 2^{[n]}$ (from candidates to subsets of $[n]$) is a *canonical choice function* if

(1) For each candidate G, the set $f(G)$ is an orbit under the automorphisms of G consisting of vertices on the highest level of G.
(2) For any two candidates G, G' , if $\sigma : G \to G'$ is an isomorphism of graphs, then σ maps $f(G)$ onto $f(G')$.

Definition 6.3. The *parent* of a graph G is the unlabeled graph formed by removing a point v in $f(G)$ from the graph.

Definition 6.4. Conversely, a graph G' is a *candidate child* of a graph G if we can add a point v to G to obtain G', and so that v is on the highest level of G'.

Since the point removed will be on the highest level, we will remove only downwards pointing edges from the graph, so we cannot create any shortcuts or cycles. Thus the parent of a Hasse diagram is again a Hasse diagram.

Also, the parent of a graph is uniquely defined, regardless of which point we remove from $f(G)$ to obtain it. Since $f(G)$ is an orbit of G, if v, w are both in $f(G)$ then there is an automorphism σ such that $\sigma(v) = w$. But then, the two parents, $G \setminus \{v\}$ and $G \setminus \{w\}$ are isomorphic by σ, so they are actually the same unlabeled graph.

Definition 6.5. If G' is a candidate child of G, formed by adding a point v, we say that f *accepts* G' if and only if v is in $f(G')$, where f is a canonical choice function. If we have fixed some f beforehand, we say that G' is an (actual) *child* of G if f accepts G'.

This definition allows us to actually use the canonical choice function to distinguish between the children of a graph to accept only one representative from the unlabeled children of a graph.

Lemma 6.6. *If H and H' are distinct children of a graph G,* i.e.*, both are accepted by some canonical choice function f, then H and H' are not isomorphic.*

The only remaining task is to ensure that we actually construct all possible candidate children of a graph, and accept at least one from each isomorphism class. To do this, we must consider all ways in which we can add a point to G such that the new point is now on the highest level.

First, note that if G has ℓ levels, then the new point must have an edge to some point on level $\ell - 1$ or level ℓ, or else the new point would not be on the highest level of G'.

Second, the new edges we add between our new point and its neighbors cannot create any shortcuts, since G' must be a Hasse diagram. So, if x and y are both neighbors of our new point, we cannot have $x > y$ or $y > x$. Thus, the neighbors of our new point must be pairwise incomparable. In graph theory, we call such a set an antichain. Each antichain with a point on the highest or next-highest level gives a valid set of neighbors for a new point on the top level, so these antichains describe all ways of connecting a new point to a graph to get a point on at the highest level.

Finally, if we pick two antichains A and A' such that there is a graph automorphism σ that sends A to A', then the resulting graphs formed by connecting a new point to each of A and A' will be isomorphic by the same permutation σ (extended to send the new vertex to itself). Thus, it suffices to consider only one representative from each orbit of the antichains under group automorphism.

From the above considerations, we have the following algorithm:

Algorithm 6.7. To construct all children of an unlabled poset P with ℓ levels:

(1) Find a representative from each orbit of antichains that contains a point on level ℓ or $\ell - 1$.
(2) Connect a new point v to each antichain computed in step 1 in turn.
(3) Compute the canonical choice function for each candidate constructed in step 2. A candidate is a child of P if and only if the new point v is in $f(P)$.

To actually enumerate all unlabeled posets with at most n points, begin with the graphs consisting of no more than n points all on the first row, and then performing a depth-first search on the children of each graph that we find.

The proof of the correctness of this algorithm is due to Brinkmann and McKay [3], but for now, the assertion that it does generate exactly one example of each unlabeled poset should suffice to justify our modifications to count minimal graphs.

7. Constructing Minimal Graphs

Since we are not in fact trying to count all posets, but only a subset of them, we really only need to generate graphs which are minimal, or some of whose children will eventually be minimal. If we can determine that a given graph will never have minimal descendants, then we can prune that node from our search, and not have to waste computation on branches which will never bear fruit. We can do this most easily by considering a slightly larger collection than the set of all minimal graphs.

Definition 7.1. We say that a graph is *non-downbeat* if there are no points with out-degree equal to 1. This is equivalent to the statement that the underlying topology has no downbeat points.

All minimal graphs are of course non-downbeat, so if we can construct all non-downbeat graphs and then check whether each one is non-upbeat as well, we will have accomplished our task of counting all minimal graphs.

The categorization of graphs as non-downbeat is useful primarily because it is a hereditary property:

Lemma 7.2. *If a graph G' is non-downbeat, then its parent G is non-downbeat as well.*

Proof. Let v be the vertex that we remove from G' to obtain G. Remember that v is on the top level, so there cannot be any edges $w \to v$, or else w would be on a higher level, thus in removing v from G', we do not change the out-degree of any point $w \neq v$. Thus since no points in $G' \setminus \{v\}$ have out-degree equal to 1, no points in G have out-degree 1 either. Thus G is non-downbeat. □

We can also categorize which children of a non-downbeat graph will also be non-downbeat (allowing us to not construct the other children in the first place).

Lemma 7.3. *If G is non-downbeat, and G' is obtained from G by adding a point v on the highest level, then G is non-downbeat if and only if v has two or more neighbors.*

Proof. Again, by adding a point at the top level, we do not change the out-degree of any of the points in G, so G' is non-downbeat if and only if v is not a downbeat point. Then, it is clear that v will not be a downbeat point if and only if it has two or more neighbors. □

Finally, we can identify a special case of child which will never produce any minimal descendants, even though the child itself is non-downbeat.

Lemma 7.4. *If G has exactly one point on the top level ℓ, and G' is obtained from G by adding a point to a new level $\ell + 1$, then no descendant of G' will ever be minimal.*

Proof. We claim that all descendants of G' will have exactly one point on level ℓ, but have a highest level $\ell' > \ell$. By Proposition 5.4, such graphs cannot be minimal.

We proceed by structural induction on the tree of descendants of G'. As a base case, this is trivially true of G'. Now, let H be a descendant of G' with exactly one point on level ℓ and with highest level $\ell' > \ell$. Then all children of H are formed by adding a point on level ℓ' or $\ell' + 1$, so all children of H still have exactly one point on level ℓ. □

These three Lemmas allow us to make the following changes to the above algorithm which will prune dead-ends. We call all children which are not known to be dead-ends by the above lemmas *useful children.*

Algorithm 7.5. To construct all useful children of a graph G with highest level ℓ:

(1) Find a representative from each orbit of antichains that contains a point on level $\ell - 1$. If G has more than one point on level ℓ, also find representatives from each orbit of antichains with a point on level ℓ.

(2) Connect a new point v to each antichain computed in step 1 whenever the antichain contains at least two vertices.
(3) Compute the canonical choice function for each candidate constructed in step 2. A candidate is a child of P if and only if the new point v is in $f(P)$.
(4) If the canonical choice function accepts, then verify that the graph is non-upbeat as well by checking that no point has in-degree 1. If the graph is non-upbeat, then increment our count of minimal graphs encountered. Even if the graph is contains upbeat points, it is still a useful child of G and could have minimal descendants, so we must recursively find its children as well.

By the above Lemmas, the children which we ignore are all such that they are not minimal, and will never have minimal descendants, so we can ignore those branches and still find representatives of all minimal graphs.

8. Computational Results

The above algorithm was actually implemented and run to obtain the exact counts of unlabeled minimal graphs with small numbers of points. Various optimizations described in [3] were implemented to expedite the computation of the canonical choice function, and in the construction of antichains. Canonical labeling of graphs (needed for the canonical choice function) was achieved by the using the graph isomorphism library `nauty` [6]. This is the same library used by Brinkmann and McKay in their original library [3].

Points	Minimal graphs	Homotopy classes	Unlabeled posets
1	1	1	1
2	1	2	2
3	1	3	5
4	2	5	16
5	4	9	63
6	11	20	318
7	36	56	2045
8	160	216	16999
9	954	1170	183231
10	7929	9099	2567284
11	92092	101191	46749427
12	1493102	1594293	1104891746

Table 1. Counts of minimal graphs and Homotopy classes

To ensure the correctness of these results, we used the C preprocessor to compile two different versions of the algorithm, one with our changes as described above, and one functionally identical to the original algorithm for enumerating all unlabeled posets. The unmodified algorithm successfully reproduced the counts for all unlabeled posets up to 11 points, but could not be run on higher inputs since it takes far longer to run than the modified version (This was the purpose of pruning branches in the first place). Since the code for the two versions is 99% identical, it is much more feasible for a human to check that the changes we implemented actually produce the desired result. Furthermore, at the beginning of researching this topic, one of the authors enumerated all minimal graphs up to 8 points by hand, and these counts were verified by the algorithm.

Table 8 gives the counts for the number of unlabeled minimal graphs with up to 12 points. Since the number of homotopy classes with n points is the number of unlabeled graphs with at most n points, their number is simply the sum of the counts of minimal graphs with at most n points. We also provide the number of unlabeled graphs (equal to the number of T_0 spaces up to homeomorphsim) from [3] for reference.

9. Asymptotic Enumeration

Kleitman and Rothschild's paper [4] has been used to describe the asymptotic behavior of posets as consisting of graphs with exactly three levels with 'roughly' $n/4, n/2$ and $n/4$ points on each of the three levels. However, the exact statement of the result will prove much more useful in describing the asymptotic behavior of minimal graphs.

Their paper describes a set of posets on a vertex set V of n points which formalizes this notion of three-leveled posets. The collection, $Q(V)$ consists of the posets P such that

(1) The vertices of P are the disjoint union, $S_1 \amalg S_2 \amalg S_3$ where points in S_i only have edges going to points in S_{i-1} or S_{i-2}
(2) The size of the partition is such that
  (a) $||S_i| - n/4| < (n-1)^{\frac{1}{2}} \log(n-1)$
  (b) $||S_2| - n/2| < log(n-1)$
(3) For every $u \in S_1 \cup S_3$, $||N(u) \cap S_2| - n/4| < (n-1)^{7/8}$, where $N(u)$ is the set of neighbors of u.
(4) For every $u \in S_2$, $||N(u) \cap S_i| - n/8| < (n-1)^{7/8}$ for $i = 1$ or $i = 3$

By a collection of logarithmic bounds given by their lemma, they find that the number of posets on n points, P_n, is asymptoticaly equivalent to the number of posets in $X(V)$, and that this is asymptoticaly equivalent to the number of posets in $Q(V)$. Specifically, if Q_n counts the number of posets in $Q(V)$ with n points, then $P_n = (1 + O(1/n))Q_n$.

In our enumeration we have been concerned with non-isomorphic, minimal, leveled digraphs (equivalently unlabeled, minimal Hasse diagrams) as these define the homotopy classes of T_0 spaces, yet Kleitman of Rothschilds result is using *labeled* Hasse Diagrams, which gives the number of all T_0 spaces. To make use of their result, we need to know the relation between the number of unlabeled graphs and labeled graphs. For this we make use of an exceedingly general result from Prömel [5], which states that in any large enough collection of labeled objects, the fraction of objects with non-trivial automorphism group goes to 0, and thus asymptotically, the ratio of labeled objects to unlabeled objects approaches $\frac{1}{n!}$.

Lemma 9.1. *Let $\mathscr{C}$ be a class of finite labeled structures (*i.e.*, a finite labeled set with a single binary relation) which is closed under substructures and isomorphisms. Let $C(n)$ count the number of such structures on sets with n points, and let $C^u(n)$ count the number of unlabeled structures on n points. If (C) satisfies the growth condition*

$$C(n) = cn^2 + dn + o(n)$$

where $c > 0$ and d is arbitrary, then

$$C^u(n) \sim \frac{C(n)}{n!}$$

Applied to the case of classes of posets, this lemma states that as long as our collection of labeled posets is large enough, we can directly derive asymptotic bounds on the growth of the collection of unlabeled posets. Since this condition is satisfied both by the set of all posets and by the set of posets in $Q(V)$ we have the immediate corollary:

Corollary 9.2. *The number of unlabeled posets in $Q(V)$, Q_n^u, is asymptoticaly equal to the number of unlabeled posets, P_n^u.*

Proof. We know, by Kleitman and Rothschild's result [4], that the number of all labeled poset P_n, is such that $\log(P_n) = \frac{n^2}{4} + \frac{3n}{2} + O(\log(n))$. So by the above lemma, $P_n^u \sim \frac{1}{n!} P_n$. Similarly, since $P_n \sim Q_n$, we have that $Q(V)$ satisfies the growth condition as well, so $Q_n^u \sim \frac{1}{n!} Q_n$. Also, $P_n \sim Q_n$ implies that $\frac{P_n}{n!} \sim \frac{Q_n}{n!}$ so

$$Q_n^u \sim \frac{Q_n}{n!} \sim \frac{P_n}{n!} \sim P_n^u$$

□

An asymptotic enumeration of the homotopy classes of finite T_0 spaces follow directly from this.

Corollary 9.3. *The number of homotopy classes of finite T_0 topological spaces is asymptotically equivalent to the number of all T_0 spaces up to homeomorphism.*

Proof. By definition, graphs in $Q(V)$ have the property that

(1) For every $u \in S_1 \cup S_3$, the number of neighbors of u in S_2 is greater than $n/4 - (n-1)^{7/8}$
(2) For every $u \in S_2$, the numbers of neighbors of u in S_1 and S_3 are each greater than $n/8 - (n-1)^{7/8}$

Thus, for n large enough, every point in the top row has out-degree at least 2, every point in the middle row has out-degree and in-degree at least 2, and every point in the bottom row has in-degree at least 2. Thus, every graph in $Q(V)$ with enough points is a minimal graph.

But then, every unlabeled graph in $Q(V)$ is an unlabeled minimal graph, so if we let M_n^u be the number of unlabeled minimal graphs with n points, then we have that $Q_n^u \leq M_n^u \leq P_n^u$. And $Q_n^u \sim P_n^u$ so by the squeeze theorem, $M_n^u \sim P_n^u$.

But remembering that M_n^u also counts the number of homotopy classes of finite spaces up to homotopy, and P_n^u counts the number of finite spaces up to homeomorphism, we have that almost every unlabeled graph on n vertices is minimal and therefore the number of homotopy classes of finite T_0 topological spaces is asymptotically equal to the number of all T_0 spaces. □

Before considering the implications of this, it is worth noting that the above method is not the only way to prove this result; instead, one only needs that almost every poset has three levels and that these levels monotonicaly increase in size as the poset grows.

Lemma 9.4. *Almost all graphs with 3 levels are minimal.*

Proof. Let $P = L_1 \amalg L_2 \amalg L_3$ be an unlabeled digraph with three levels, and let $|L_3| = j$, $|L_2| = k$, and $|L_1| = l$.

To determine the probability of this graph being minimal, consider that P is formed by taking the complete tri-partite graph on its levels, randomly deleting some number of edges, and possibly adding edges from L_3 to L_1.

So $x \in L_3$ has between 1 and k edges leading to L_2, by definition of the levels of a graph; for $y \in L_2$ y has between 0 and j edges to it from L_3. A point in L_3 might have edges going to L_1 in addition to its edges going to L_2, so for any $x \in L_3$ prob(outdegree$(x) > 2) \geq 1 - \frac{1}{k}$. This bound is from the fact that there are k ways for x to have one edge, but also k ways for it to have any degree up to $k-1$ and so we get a very conservative bound by considering only one possiblility for each possible degree that x may have.

Each event (placing edges from a point in L_3 to points in L_2) is independent from the others, so

$$\text{prob}(\forall x \in L_3, \text{outdegree}(x) \geq 2) \geq (1-\frac{1}{k})^j = (\frac{k-1}{k})^j = \frac{k^j - jk^{j-1} + \cdots - (-1)^j k + (-1)^j}{k^j}$$

For a given j, $\lim_{k\to\infty}(\text{prob}(\forall x \in L_3, \text{outdegree}(x) \geq 2)) = 1$.

Then, we have that for any $y \in L_3$ prob(indegree$(x) \geq 2) > (1 - \frac{1}{j})^2$ and

$$\text{prob}(\forall y \in L_2, \text{outdegree}(y) \geq 2) > (1-\frac{1}{j})^{2k} = (\frac{j-1}{j})^{2k} = \frac{j^{2k} - 2kj^{2k-1} + \cdots - k + 1}{k^j}$$

And for a given k $\lim_{j\to\infty}(\text{prob}(\forall y \in L_2, \text{outdegree}(y) \geq 2)) = 1$.

Similarly $\lim_{l\to\infty}(\text{prob}(\forall y \in L_2, \text{outdegree}(y) \geq 2)) = 1$ and $\lim_{k\to\infty}(\text{prob}(\forall z \in L_1, \text{outdegree}(z) \geq 2)) = 1$.

These events are not probablisticaly independent, so we cannot just multiply the individual probabilities to obtain the probability of all 4 events happening simultaniously. However, we can take the union bound on the complement of these events, giving prob(P is not minimal) $\leq \epsilon_1 + \epsilon_2 + \epsilon_3 + \epsilon_4$ where $\epsilon_1 = \text{prob}(\exists x \in L_3, \text{outdegree}(x) < 2)$ $\epsilon_2 = \text{prob}(\exists y \in L_2, \text{indegree}(x) < 2)$ $\epsilon_3 = \text{prob}(\exists y \in L_2, \text{outdegree}(x) < 2)$ and $\epsilon_4 = \text{prob}(\exists z \in L_1, \text{outdegree}(z) < 2)$.

Then almost all such graphs P are minimal, provided that the size of each level increases as the graph itself grows, meaning graphs on n vertices $P = L_1 \amalg L_2 amalg L_3$ with $|L_3| = an$ $|L_2| = bn$ $|L_1| = cn$ st $a + b + c = 1$. □

Remark 9.5. The graphs in $Q(V)$ are of this form, but this proof is perhaps more intuitive.

Let us go back and consider this result. In some ways it is unsurprising to find this behavior; given a large space, the digraph representing it is large and thus has many more possible edges between vertices. In this way it makes sense that with enough edges on the graph, there is a good probability that every vertex has in-degree and out-degree at least 2. However, with respect to the topology, this result is startling; homotopy equivalence does not narrow down the classification of spaces any more than homeomorphism for large spaces. Nevertheless, when we look at the actual, numerical counts for number of spaces up to homotopy and homeomorphism, we see a large gap between the relative growth rates. For example, for spaces with 12 points, there are 1104891746 spaces up to homeomorphism, with only 1594293 distinct spaces up to homotopy equivalence (a factor of 70 difference). Thus, even though the asymptotic behavior of these two numbers is the same, the convergence for small values is very slow.

References

[1] May, J. P. Finite Topological Spaces. http://www.math.uchicago.edu/∼may/MISC/FiniteSpaces.pdf
[2] May, J. P. Finite Spaces and Simplical Complexes. http://www.math.uchicago.edu/∼may/MISC/SimpCxes.pdf
[3] Brinkmann, Gunnar and Brendan D. McKay. Posets on up to 16 Points. http://cs.anu.edu.au/ bdm/papers/posets.pdf
[4] D.J. Kleitman and B.L. Rothschild. Asymptotic Enumeration of Posets on a Finite Set. Transaction of The American Mathematical Society, Vol. 205 (April 1975), pp. 205-220.
[5] Prömel, Hans Jürgen. Counting Unlabeled Structures.
[6] Brendan D. McKay. The nauty page. http://cs.anu.edu.au/˜bdm/nauty/