

Venn Diagrams and Symmetric Chain Decompositions in the Boolean Lattice

Jerrold Griggs *

Department of Mathematics
University of South Carolina
Columbia, SC 29208
griggs@math.sc.edu

Charles E. Killian

Department of Computer Science, Box 8206
North Carolina State University
Raleigh, NC 27695
chip_killian@acm.org

Carla D. Savage †

Department of Computer Science, Box 8206
N. C. State University
Raleigh, NC 27695
savage@csc.ncsu.edu

Submitted: Aug 30, 2002; Accepted: Dec 11, 2003; Published: Jan 2, 2004
MR Subject Classifications: 03E99, 05610, 06A07, 06E10

Abstract

We show that symmetric Venn diagrams for n sets exist for every prime n , settling an open question. Until this time, $n = 11$ was the largest prime for which the existence of such diagrams had been proven, a result of Peter Hamburger. We show that the problem can be reduced to finding a symmetric chain decomposition, satisfying a certain cover property, in a subposet of the Boolean lattice \mathcal{B}_n , and prove that such decompositions exist for all prime n . A consequence of the approach is a constructive proof that the quotient poset of \mathcal{B}_n , under the relation “equivalence under rotation”, has a symmetric chain decomposition whenever n is prime. We also show how symmetric chain decompositions can be used to construct, for all n , monotone Venn diagrams with the minimum number of vertices, giving a simpler existence proof.

*Research supported in part by NSF grant DMS-0072187.

†Research supported by NSA grant MDA 904-01-0-0083

1 Introduction

1.1 Venn Diagrams

Following Grünbaum [Grü75], an n -Venn diagram is a collection of n simple closed curves in the plane, $\{\Theta_1, \Theta_2, \dots, \Theta_n\}$, with the property that for each $S \subseteq \{1, 2, \dots, n\}$ the region

$$\bigcap_{i \in S} \text{int}(\Theta_i) \cap \bigcap_{i \notin S} \text{ext}(\Theta_i)$$

is nonempty and connected, where $\text{int}(\Theta_i)$ and $\text{ext}(\Theta_i)$ denote the interior and exterior, respectively, of Θ_i . For this paper, we require that any two of the curves Θ_i intersect in only a finite number of points. Figure 1 shows four 3-Venn diagrams ((b), (c) and (d) are from [CHP96]). A *region* of the Venn diagram is a maximal connected subset of $R^2 - \cup_{i=1}^n \Theta_i$, where R^2 denotes the set of all points in the plane. Thus, a Venn diagram partitions $R^2 - \cup_{i=1}^n \Theta_i$ into exactly 2^n regions, one for each subset of $\{1, 2, \dots, n\}$.

A Venn diagram is called *simple* if no three curves have a common point of intersection. In Figure 1, the Venn diagram (a) is simple, but the others are not.

It is known that n -Venn diagrams exist for all $n \geq 1$ and constructions of Venn [Ven80] and Edwards [Edw89] are illustrated in [Rus97]. Most of the results, conjectures, and problems we mention in this paper can be found in [Rus97], an excellent survey and expository article on Venn diagrams by Frank Ruskey.

1.2 Symmetric Venn Diagrams

A *symmetric* Venn diagram is one with rotational symmetry. That is, there is a point p in the plane such that each of the n rotations of Θ_1 about p by an angle of $2\pi i/n$, $0 \leq i \leq n-1$, coincides with one of the curves $\Theta_1, \Theta_2, \dots, \Theta_n$. For example, in Figure 1, the Venn diagrams (a) and (c) are symmetric, but (b) and (d) are not.

Symmetric Venn diagrams have been considered by several researchers including Henderson [Hen63], Grünbaum [Grü92, Grü99], Ruskey [Rus97], Edwards [Edw98], and Hamburger [Ham02]. Henderson [Hen63] proved that symmetric Venn diagrams are not possible when n is not prime, but symmetric Venn diagrams are known for the primes $n = 3, 5, 7$ and, most recently, for $n = 11$ [Ham02]. It has been an open question whether a symmetric n -Venn diagram exists for every prime number n [Grü75]. The main result of this paper is a constructive proof to show that the answer is yes - symmetric n -Venn diagrams exist for every prime n .

For a survey of results on symmetric Venn diagrams, as well as Hamburger's construction of the first known symmetric Venn diagram for $n = 11$ sets see [Ham02].

1.3 Monotone Venn Diagrams

Two distinct regions of a Venn diagram are *adjacent* if their boundaries intersect at a set of positive length. For example, in Figure 1(d), the region corresponding to $\{1, 2, 3\}$ is

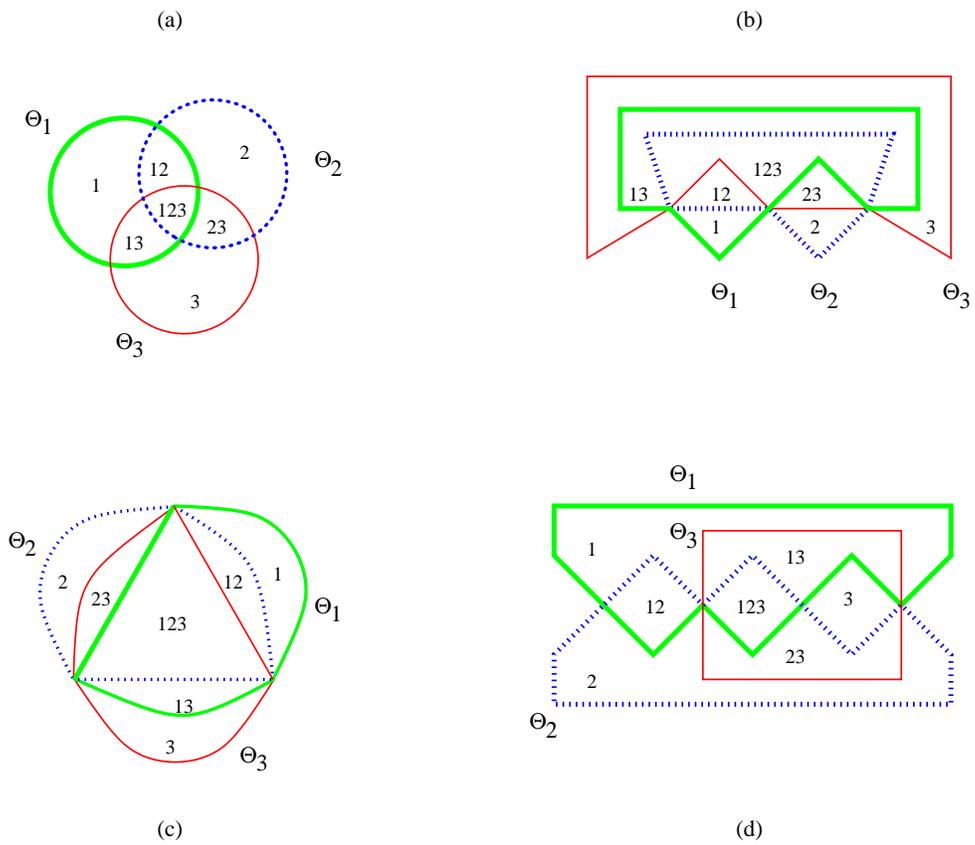


Figure 1: Four 3-Venn diagrams.

adjacent only to $\{1, 3\}$ and $\{2, 3\}$. A *monotone* Venn diagram is one in which every region corresponding to a subset of size k is adjacent to at least one region corresponding to a set of size $k - 1$ (if $k > 0$) and at least one region corresponding to a set of size $k + 1$ (if $k < n$). For example, in Figure 1, Venn diagrams (a), (b), and (c) are monotone, but (d) is not, since the region corresponding to $\{1, 2\}$ is not adjacent to the region $\{1, 2, 3\}$.

Monotone Venn diagrams are interesting because of their relationship to convex Venn diagrams. A Venn diagram is *convex* if the bounded region enclosed by each curve Θ_i is convex. If, in addition, the complement of the unbounded region is convex, the Venn diagram is *strongly convex*. Convex n -Venn diagrams for all n were shown to exist in [RRS51] and the existence of strongly convex n -Venn diagrams for all n was established by Grünbaum in [Grü75]. The Venn diagrams constructed in both of these papers were monotone as well. It was shown in [BGR98] that a Venn diagram is isomorphic to a convex Venn diagram if and only if it is monotone. In Figure 1, diagram (c) is a convex Venn diagram isomorphic to the (non-convex) monotone Venn diagram (b).

Define a *vertex* of a Venn diagram defined by the curves $\{\Theta_1, \Theta_2, \dots, \Theta_n\}$ to be a point in the plane where two or more of the curves Θ_i intersect. We can regard the Venn diagram as a plane graph P (that is, a planar embedding of a planar graph) whose vertices are these intersection points and where two vertices are joined by an edge in P if they are consecutive intersection points on one of the curves Θ_i . The number of vertices of the Venn diagram is the number of faces in any embedding of the dual graph of P .

In [BR98] Bultena and Ruskey ask for Venn diagrams with the minimum number of vertices. They show that a monotone Venn diagram always has at least $\binom{n}{\lfloor n/2 \rfloor}$ vertices and provide a construction which shows that monotone Venn diagrams which achieve this lower bound exist for all $n > 1$. The proof is a delicate inductive construction. It turns out that with the symmetric chain decomposition approach, we get a simpler proof of this result.

1.4 Symmetric Chain Decompositions in the Boolean Lattice

Consider a finite partially ordered set (poset) $\mathcal{A} = (A, \leq)$. For $x \neq y \in A$, y is said to *cover* x if $x \leq y$ and if there is no $z \in A$ such that $x < z < y$. The poset \mathcal{A} is *ranked* if one can define a function $r(x)$ on the elements $x \in A$ such that $r(x) = 0$ for all minimal elements x and $r(y) = r(x) + 1$ for all x, y such that y covers x . If \mathcal{A} is ranked, we say that $r(x)$ is the *rank* of x and the *rank of \mathcal{A}* is $\max_x r(x)$.

A *symmetric chain* in a ranked poset \mathcal{A} of rank n is a sequence of elements $x_1, x_2, \dots, x_t \in A$, such that x_{i+1} covers x_i for all i and $r(x_1) + r(x_t) = n$. A *symmetric chain decomposition* (SCD) of \mathcal{A} is a partition of the elements of A into symmetric chains.

The Boolean lattice $\mathcal{B}_n = (2^{[n]}, \subseteq)$ is the ranked poset consisting of all subsets of $[n] = \{1, 2, \dots, n\}$, ordered by inclusion. For $s \in 2^{[n]}$, $r(s)$ is the cardinality of s . It is well-known that \mathcal{B}_n has an SCD. Figure 2 illustrates (a) the Hasse diagram of \mathcal{B}_4 and (b) an SCD in \mathcal{B}_4 . In Section 3.1, we present the construction of Greene and Kleitman [GK76] for an SCD in \mathcal{B}_n . One of our main results in this paper is to show that this

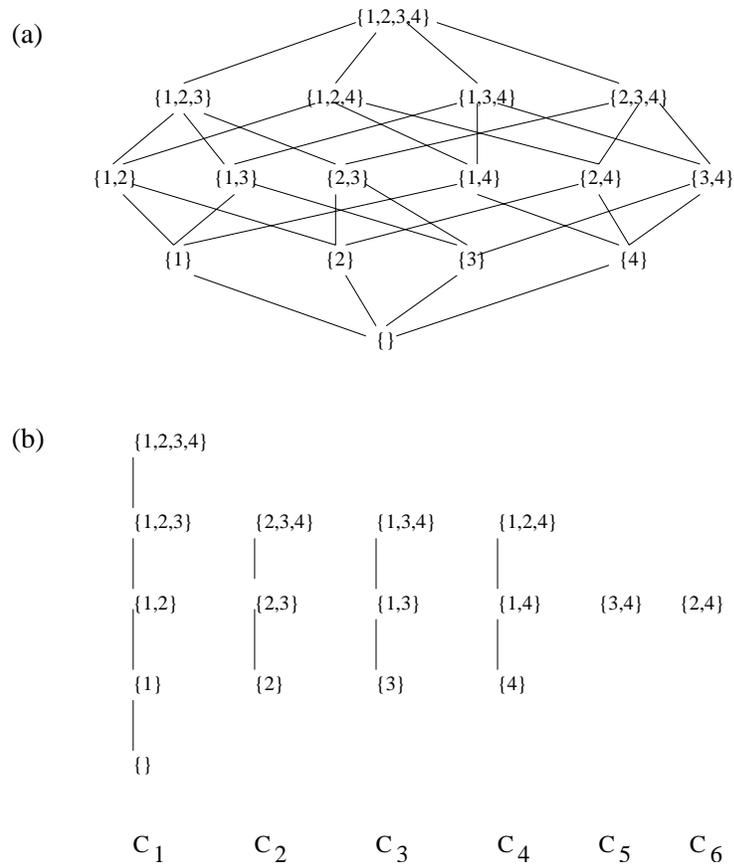


Figure 2: The Hasse diagram of \mathcal{B}_4 (a) with a symmetric chain decomposition (b).

SCD can be used to construct monotone n -Venn diagrams with the minimum number of vertices for every positive n . Note that any SCD in \mathcal{B}_n has exactly $\binom{n}{\lfloor n/2 \rfloor}$ chains, one for each element of rank $\lfloor n/2 \rfloor$ of \mathcal{B}_n .

1.5 Necklaces

It will be convenient at times to view the elements of \mathcal{B}_n as elements of $\{0, 1\}^n$, the set of all n -bit strings. With each $x = x_1x_2 \cdots x_n \in \{0, 1\}^n$, we associate a set, $S(x)$ defined by

$$S(x) = \{i \mid x_i = 1, 1 \leq i \leq n\}.$$

The Boolean lattice, then, is $\mathcal{B}_n = (\{0, 1\}^n, \leq)$, the poset whose base set is $\{0, 1\}^n$, and whose ordering, \leq is defined for $x, y \in \{0, 1\}^n$ by $x \leq y \iff S(x) \subseteq S(y)$. The Hasse diagram of the Boolean lattice \mathcal{B}_n is isomorphic to the n -cube.

For $x = x_1x_2 \cdots x_n \in \{0, 1\}^n$, let σ be the rotation of x defined by $\sigma(x) = x_2x_3 \cdots x_nx_1$. Let $\sigma^1 = \sigma$ and for $i > 1$, let $\sigma^i(x) = \sigma(\sigma^{i-1}(x))$. Define the relation “ \sim ” on $\{0, 1\}^n$ by $x \sim y$ iff $y = \sigma^i(x)$ for some integer $i \geq 0$. Then “ \sim ” is an equivalence relation on $\{0, 1\}^n$ and the equivalence classes are called *necklaces*.

Let N_n be the set of necklaces of $\{0, 1\}^n$ and define the *necklace poset*, \mathcal{N}_n , by $\mathcal{N}_n = (N_n, \preceq)$ with ordering \preceq defined for $\eta_1, \eta_2 \in N_n$ by $\eta_1 \preceq \eta_2$ if and only if some $x \in \eta_1$ differs from some $y \in \eta_2$ only in one bit i , where $x_i = 0$ and $y_i = 1$. As we discuss in Section 5, it can be shown from results in order theory that when n is prime, \mathcal{N}_n has a symmetric chain decomposition. It is well-known that when n is prime, each of the necklaces, other than the ones containing 0^n and 1^n , has exactly n elements.

We will require something stronger. One of our main results is to show that when n is prime, we can always select a set R_n , consisting of one representative string from each necklace N_n , so that the necklace-representative subposet (R_n, \leq) of \mathcal{B}_n induced by R_n has a symmetric chain decomposition with a certain *cover property*. Furthermore, from the SCD in this necklace-representative poset, we can construct a symmetric Venn diagram.

1.6 Overview of Main Results and Organization of Paper

Our main results in this paper are:

- For *all* $n \geq 1$, any symmetric chain decomposition in the Boolean lattice satisfying a certain “chain cover property” can be used to construct a monotone Venn diagram with the minimum number of vertices. A well-known symmetric chain decomposition of the Boolean lattice [Aig73, GK76] is shown to have this cover property.
- When n is *prime*, there is always a way to select a complete set R_n of necklace representatives so that the induced necklace-representative subposet (R_n, \leq) of the Boolean lattice has a symmetric chain decomposition which satisfies the required cover property.

- For all *prime* n , there is a symmetric Venn diagram for n sets which can be constructed from this symmetric chain decomposition in (R_n, \leq) .

The suggestion that symmetric chain decompositions might be useful in constructing symmetric Venn diagrams first appears in the paper of Hamburger [Ham02].

The remainder of this paper is organized as follows. Section 2 shows how to construct Venn diagrams from symmetric chain decompositions with the chain cover property: the monotone case for all n is covered in Section 2.1, and the symmetric case for n prime is covered in Section 2.2. Section 3 describes the Greene-Kleitman symmetric chain decomposition in \mathcal{B}_n and shows that it has the required cover property to construct monotone Venn diagrams for all n . Section 4 contains the key technical result of the paper: a proof of the existence, when n is prime, of a necklace-representative subposet of \mathcal{B}_n with a symmetric chain decomposition satisfying the required cover property. Section 5 suggests some related open questions.

2 Venn Diagrams from Symmetric Chain Decompositions

2.1 Monotone Venn Diagrams for all n

In this subsection we illustrate a connection between Venn diagrams and symmetric chain decompositions by using a symmetric chain decomposition of the Boolean lattice to give a simple construction for monotone n -Venn diagrams, with minimum number of vertices, for all n .

2.1.1 The Chain Cover Graph

Let \mathcal{C} be an SCD in a finite ranked poset $\mathcal{A} = (A, \leq)$ and for chain $C \in \mathcal{C}$, let *starter*(C) be the first element of C and let *terminator*(C) be the last element of C . Call the longest chains in \mathcal{C} the *root chains*. Say that \mathcal{C} has the *chain cover property* if whenever $C \in \mathcal{C}$ and C is not a root chain, then there exists a chain $\pi(C) \in \mathcal{C}$ such that

starter(C) covers an element $\pi_s(C)$ of $\pi(C)$ and

terminator(C) is covered by an element $\pi_t(C)$ of $\pi(C)$.

Call such a mapping π a *chain cover mapping*.

Note that the SCD of Figure 2 (b) has the chain cover property: define π by $\pi(C_2) = \pi(C_3) = \pi(C_4) = C_1$; $\pi(C_5) = C_3$; and $\pi(C_6) = C_2$. For example, for chain C_2 , $\pi_s(C_2) = \emptyset$ and $\pi_t(C_2) = \{1, 2, 3, 4\}$, since *starter*(C_2) = $\{2\}$ covers \emptyset from chain $\pi(C_2) = C_1$ and *terminator*(C_2) = $\{2, 3, 4\}$ is covered by $\{1, 2, 3, 4\}$ from chain $\pi(C_2) = C_1$.

We focus on the case where \mathcal{C} has a unique root chain (although this can be easily generalized). When the root chain is unique, π can be described by a rooted tree, $T(\mathcal{C}, \pi)$, called a *chain cover tree*, in which each node corresponds to a chain $C \in \mathcal{C}$ and the parent

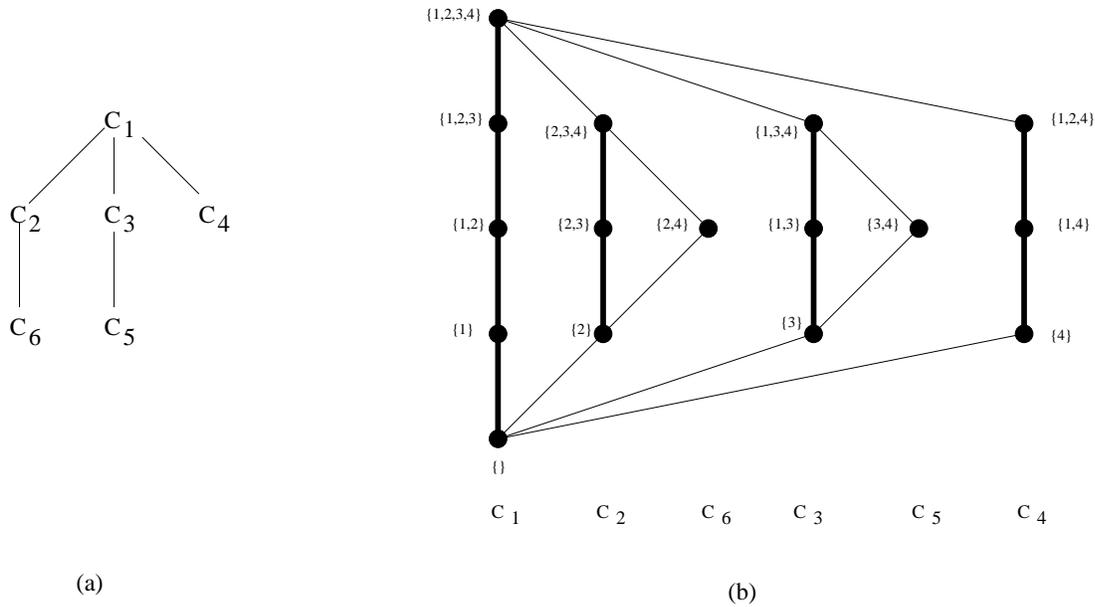


Figure 3: (a) A chain cover tree, $T(\mathcal{C}, \pi)$ for the SCD \mathcal{C} of Figure 2(b) and (b) a planar embedding, $P(\mathcal{C}, \pi)$, of $G(\mathcal{C}, \pi)$, with chain edges dark and cover edges light.

of node C is $\pi(C)$. Figure 3(a) shows the chain cover tree for the SCD of Figure 2(b) and the mapping π described above.

Let \mathcal{C} be an SCD for $\mathcal{A} = (A, \leq)$ with the chain cover property and let π be a chain cover mapping for \mathcal{C} . We consider the *chain cover graph*, $G(\mathcal{C}, \pi)$, whose vertices are the elements of A and whose edges consist of the covering edges in the chains in \mathcal{C} together with the *cover edges*, for each non-root chain $C \in \mathcal{C}$, from:

- *starter*(C) to $\pi_s(C)$ and
- *terminator*(C) to $\pi_t(C)$.

Figure 3(b) shows the chain cover graph for the chain cover tree in Figure 3(a).

First we show that the chain cover graph always has a planar embedding.

Lemma 1 *Let \mathcal{C} be a symmetric chain decomposition with the chain cover property for poset $\mathcal{A} = (A, \leq)$, and let π be a chain cover mapping for \mathcal{C} . The chain cover graph $G(\mathcal{C}, \pi)$ has a planar embedding $P(\mathcal{C}, \pi)$.*

Proof. We describe a planar embedding of $G(\mathcal{C}, \pi)$ by giving the coordinates of each vertex and then specifying that each edge is drawn as a straight line between its endpoints. Let $T = T(\mathcal{C}, \pi)$ be the chain cover tree. Order the children of each node in T from shortest chain to longest chain and perform a *preorder labeling* of the nodes of T . A preorder labeling of an ordered tree is a labeling $\lambda(v)$ of the nodes of the tree by consecutive integers in such a way that at every node v , if C_1, C_2, \dots, C_k is the ordered list of children of v , then for $1 \leq i < k$, $\lambda(v) < \lambda(u) < \lambda(w)$ for any nodes u, w in the subtrees rooted at

C_i, C_{i+1} , respectively. For example, a preorder labeling of the chain cover tree in Figure 3(a), with children ordered as shown, is:

$$\lambda(C_1) = 1; \lambda(C_2) = 2; \lambda(C_6) = 3; \lambda(C_3) = 4; \lambda(C_5) = 5; \lambda(C_4) = 6.$$

Now, if vertex s of $G(\mathcal{C}, \pi)$ is on chain $C \in \mathcal{C}$, embed s at the point with coordinates $(\lambda(C), \text{rank}(s))$. Embed all edges of $G(\mathcal{C}, \pi)$ as straight lines. The resulting embedding for the chain cover graph given by the chain cover tree of Figure 3(a) is the one shown in Figure 3(b).

Because children of a node are ordered shortest chain to longest, and all chains are symmetric, it is easy to see that no edges cross and therefore this embedding is planar. The proof is by induction.

If the root r of T has no children, the graph is a vertical chain. Otherwise, let C be the last (longest) child of r . Assume inductively that the embedding described above is planar on the subgraph corresponding to the subtree T_C of T rooted at C and on the subgraph corresponding to T' , the tree T with the subtree rooted at C removed. All chains in the subtree rooted at C have preorder labels greater than all chains in T' , so nodes in chains in C 's subtree are embedded to the right of those in chains in T' . Let $s = \text{starter}(C)$ and $t = \text{terminator}(C)$. No node with a chain in C 's tree has y coordinate larger than $r(t)$ or smaller than $r(s)$, since C was longest. Furthermore, no node in T' not in r 's chain has y coordinate larger than $\text{rank}(t)$ or smaller than $r(s)$. Thus, the cover edges from chain C to chain r , which are the only edges in G connecting vertices in the two subtrees, do not cross any other edges. (These edges are: from $(\lambda(C), r(s))$ to $(\lambda(r), r(s) - 1)$ and from $(\lambda(C), r(t))$ to $(\lambda(r), r(t) + 1)$.) \square

2.1.2 The Venn Diagram

We now show that when the poset \mathcal{A} of Lemma 1 is the Boolean lattice, the dual of *any* planar embedding of the chain cover graph $G(\mathcal{C}, \pi)$ is a Venn diagram, a fact which is a straightforward consequence of classical theorems in graph theory. A *plane* graph is a planar embedding of a planar graph.

Following Harary, [Har69], the *geometric dual* of a plane graph, P is the graph P^* constructed by placing a vertex f^* in each face f of P (including the unbounded face) and, whenever two faces f and g have a common boundary edge e in P , joining vertices f^* and g^* of P^* with an edge e^* crossing only e . It is well-known that the dual of a plane graph is planar, that each face of P^* contains exactly one vertex of P and that P is connected if and only if P is isomorphic to $(P^*)^*$. For \mathcal{B}_4 , an embedding of the dual of $P(\mathcal{C}, \pi)$ in Figure 3 is shown, superimposed, in Figure 4.

A classical result of graph theory is the correspondence between *bonds* in a planar graph G and *simple cycles* in the dual of any planar embedding of G . (A *bond* in a graph G is a minimal set of edges whose removal disconnects G .) The formulation below is adapted from West [Wes96].

Lemma 2 *Edges in a connected plane graph P form a bond in P if and only if the corresponding dual edges form a cycle in P^* .* \square

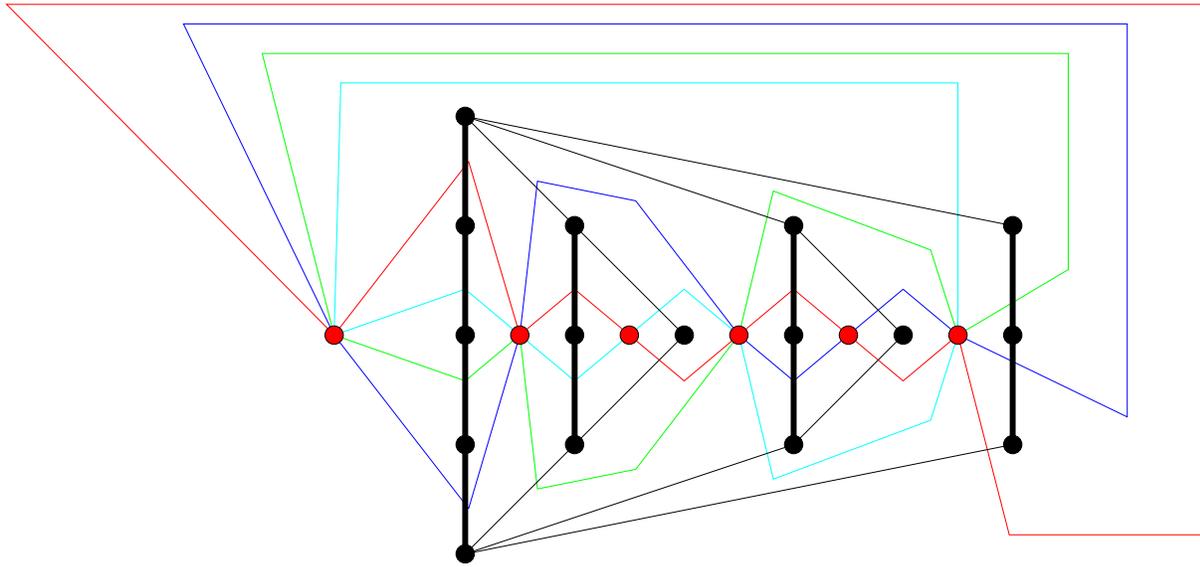


Figure 4: The geometric dual, P^* , of $P(\mathcal{C}, \pi)$ of Figure 3(b), indicated by the red vertices and the thin colored edges.

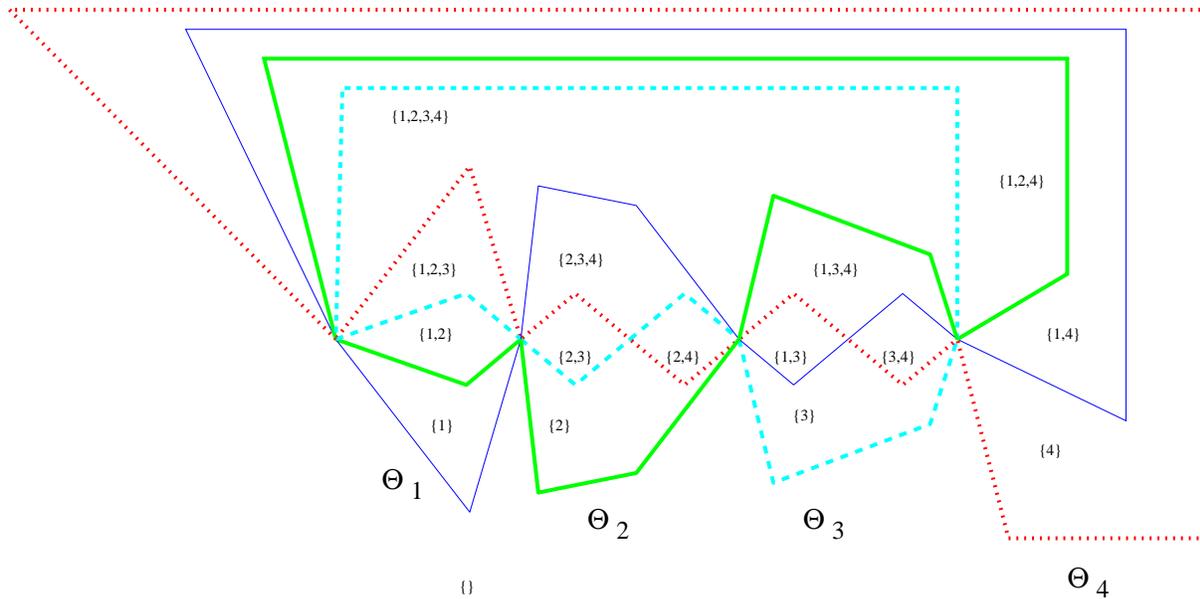


Figure 5: The 4-Venn diagram corresponding to Figure 4 with the curves $\{\Theta_1, \Theta_2, \Theta_3, \Theta_4\}$ highlighted.

Identifying the vertices of the n -cube with subsets of $[n]$, call a spanning subgraph G of the n -cube *monotone* if every subset of size k is adjacent to a subset of size $k - 1$ (if $k > 0$) and a subset of size $k + 1$ (if $k < n$). We then have the following.

Lemma 3 *If P is a plane, monotone spanning subgraph of the n -cube, the dual of P is a monotone Venn diagram.*

Proof. The dual of P , P^* , has exactly one (nonempty) face for each vertex of P , i.e., each subset of $[n]$, so P^* has the regions required of a Venn diagram. Without loss of generality, we may assume that \emptyset lies in the unbounded face of P^* . We need to show that the regions arise as the intersection of n simple closed curves in the plane.

Call an edge e in $P(\mathcal{C}, \pi)$ of the form $e = (S, S \cup \{x\})$ an x -edge of P and the corresponding edge e^* in P^* , an x -edge of P^* . We show that the x -edges of P^* form a simple cycle, Θ_x . It follows then that Θ_x is a simple closed curve. By Lemma 2, it suffices to show that the set of x -edges of P forms a *bond*. To see this, let V be the collection of vertices of P which are subsets of $[n]$ containing x and \bar{V} the collection of those which do not contain x . The only edges of P joining a vertex of V to a vertex of \bar{V} are x -edges, so removing the x -edges disconnects P . Further, the subgraphs of P , $P[V]$ and $P[\bar{V}]$, induced by V and \bar{V} , respectively, are connected: since P is monotone, for every $S \subseteq [n]$ there is a path in P from S to $[n] \in V$ consisting of a nested sequence of subsets and a path in P from $\emptyset \in \bar{V}$ to S consisting of a nested sequence of subsets. If $x \in S$, the nested path from S to $[n]$ contains no x edge; if $x \notin S$, the nested path from \emptyset to S contains no x edge. Thus no proper subset of the x -edges can disconnect P , and so the x -edges of P form a bond separating V and \bar{V} . It follows that Θ_x is a simple closed curve containing all faces corresponding to subsets containing x in its interior and those not containing x in its exterior. Every edge of P^* is an x -edge for exactly one $x \in [n]$ and therefore belongs to exactly one of the curves Θ_x . Thus, P^* is a Venn diagram. \square

Figure 5 illustrates the four curves $\Theta_1, \dots, \Theta_4$ which comprise the geometric dual P^* of $P(\mathcal{C}, \pi)$ from Figure 4.

Lemma 4 *Let \mathcal{C} be a symmetric chain decomposition with the chain cover property for \mathcal{B}_n , let π be a chain cover mapping for \mathcal{C} , and let P be a planar embedding of the chain cover graph, $G(\mathcal{C}, \pi)$. Then the dual, P^* , of P . is an n -Venn diagram. Moreover, it is a monotone Venn diagram with the minimum number of vertices.*

Proof. Clearly, $G(\mathcal{C}, \pi)$ is a spanning subgraph of the n -cube, planar by Lemma 1. To see that it is also monotone, note that in $G(\mathcal{C}, \pi)$, every vertex $S \neq [n]$ is adjacent to a set which covers it in \mathcal{B}_n (either its successor in its chain C or, if it is the last element of C , then $\pi_t(C)$). Similarly, in $G(\mathcal{C}, \pi)$, every vertex $S \neq \emptyset$ is adjacent to a set which it covers in \mathcal{B}_n (either its predecessor in its chain C or, if it is the first element of C , then $\pi_s(C)$). Thus, P is a plane, monotone spanning subgraph of the n -cube, and by Lemma 3, P^* is a monotone Venn diagram. The number of vertices of P^* is the number of faces of P , which is just the number of symmetric chains in \mathcal{C} : $\binom{n}{\lfloor n/2 \rfloor}$, that is, the

minimum possible number of vertices in a monotone Venn diagram, according to Bultena and Ruskey [BR98]. \square

In view of Lemma 4, we look closer in Section 3.1 at symmetric chain decompositions of \mathcal{B}_n . In Lemma 9 of Section 3.2 we show that we can always find one with the chain cover property and conclude in Theorem 2 of that section that we can always use this approach to construct monotone Venn diagrams.

2.2 Symmetric Venn Diagrams from Symmetric Chain Decompositions

We would first like to emphasize that the basic idea to be applied here is not new. Virtually all approaches to constructing symmetric Venn diagrams which appear in the literature are based on the following scheme. (1) Construct a special planar graph G , whose vertices are n -bit necklace-representatives, (2) embed G in the plane in a pie slice of $(2\pi i)/n$ radians, (3) rotate the embedding about the center of the pie through $(2\pi i)/n$ for $1 \leq i < n$, and (4) embed the dual of the graph G , keeping the symmetry. The bottleneck to proving that symmetric Venn diagrams exist for all prime n has been finding the right graph G in Step (1). For a given value of n , this is usually done with an *ad hoc* approach. This scheme was used by Savage and Winkler in 1992 to construct one of the first simple symmetric 7-Venn diagrams, listed as **M1** in [Rus97], where a suitable G was found using a trial-and-error approach. In the mid-nineties, Ruskey developed the idea (described in the section, “Symmetric Diagrams and Necklaces” in [Rus97]) of constructing G by organizing necklaces into two opposing trees. Coupled with exhaustive search techniques, this allowed him to discover many new symmetric 7-Venn diagrams, listed in [Rus97]. Recently, in a tour de force, Hamburger [Ham02] extended the idea of two opposing trees to create the graph G in Step (1) for $n = 11$ and constructed for the first time a symmetric 11-Venn diagram. Hamburger refers to the dual of G as a *doodle*.

The main contribution in the current paper is to show that for every prime n there is a way to construct the graph G in Step (1) that guarantees we can carry out the remaining steps to get a symmetric Venn diagram. In this section, we show that when n is prime, we can modify the technique of Section 2.1 to find a Venn diagram with rotational symmetry *if* we can find a necklace-representative subposet of \mathcal{B}_n which has a symmetric chain decomposition with the chain cover property.

As in Section 1.5, it will be convenient at times to view the elements of \mathcal{B}_n as n -bit strings. Recall from Section 1.5 that the rotation σ^i of an n -bit string is defined by $\sigma^i(x_1x_2 \cdots x_n) = x_{i+1} \cdots x_nx_1 \cdots x_i$.

Lemma 5 *Let n be prime. If there exists a set R_n of necklace representatives for $\{0, 1\}^n$ such that the subposet $\mathcal{R}_n = (R_n, \leq)$ of \mathcal{B}_n has a symmetric chain decomposition with the chain cover property, then a symmetric n -Venn diagram can be constructed.*

Proof. Assuming that such a set R_n exists, let \mathcal{C} be an SCD in \mathcal{R}_n and let π be a chain cover mapping for \mathcal{C} . By Lemma 1, the chain cover graph $G(\mathcal{C}, \pi)$ has a planar embedding.

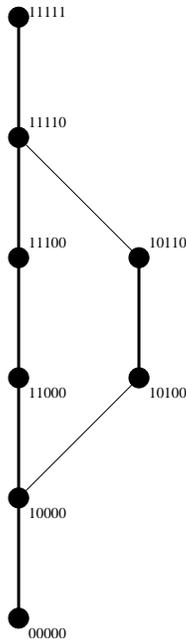


Figure 6: Embedding of the chain cover graph for a necklace-representative subsubset of \mathcal{B}_5 .

(See Figure 6 for an example when $n = 5$ and

$$R_n = \{00000, 1000, 11000, 11100, 11110, 11111, 10100, 10110\},$$

where the chains are indicated by bold edges and the chain cover edges are light.)

Fix a point p in the plane and partition the plane into n pie slices about p , each of $(2\pi)/n$ radians. We embed $G(\mathcal{C}, \pi)$ as a plane graph $P_0 = P_0(\mathcal{C}, \pi)$ in one of the pie slices, with the vertex $[n] = 11 \cdots 1$ at p and the vertex $\emptyset = 00 \cdots 0$ at the point at infinity. (If we view the embedding on the sphere, p is at the north pole and the point at infinity is the south pole.) Rotate the embedding of P (clockwise) through $(2\pi i)/n$ radians for each i , $1 \leq i \leq n - 1$. In the i -th rotation of P , relabel each vertex x of P by $\sigma^i(x)$ and let P_i be the resulting plane graph. Vertices $[n] = 11 \cdots 1$ and $\emptyset = 00 \cdots 0$ coincide in each P_i . (See Figure 7.)

Let P denote the union of the P_i , $0 \leq i < n$. Then P is a plane, monotone, spanning subgraph of the n -cube, so by Lemma 3, the dual of P is an n -Venn diagram. Furthermore, as in the proof of Lemma 3, the x -edges in P correspond to a simple cycle Θ_x in the dual of P , for each $x \in [n]$. However, we want the Venn diagram to have rotational symmetry. We follow the construction of the geometric dual, adjusting the topology to ensure symmetry. All faces of P incident with edges of P_0 are interior to the pie slice containing P_0 , except for the faces f_l and f_r , which are shared with P_{n-1} and P_1 , respectively. Embed a section of the dual, P^* , as follows: For each face $f \neq f_r$ of P that is incident with an edge of P_0 , vertex f^* of P^* is placed in the interior of the face f . Vertex f_r^* of P^* is placed in the

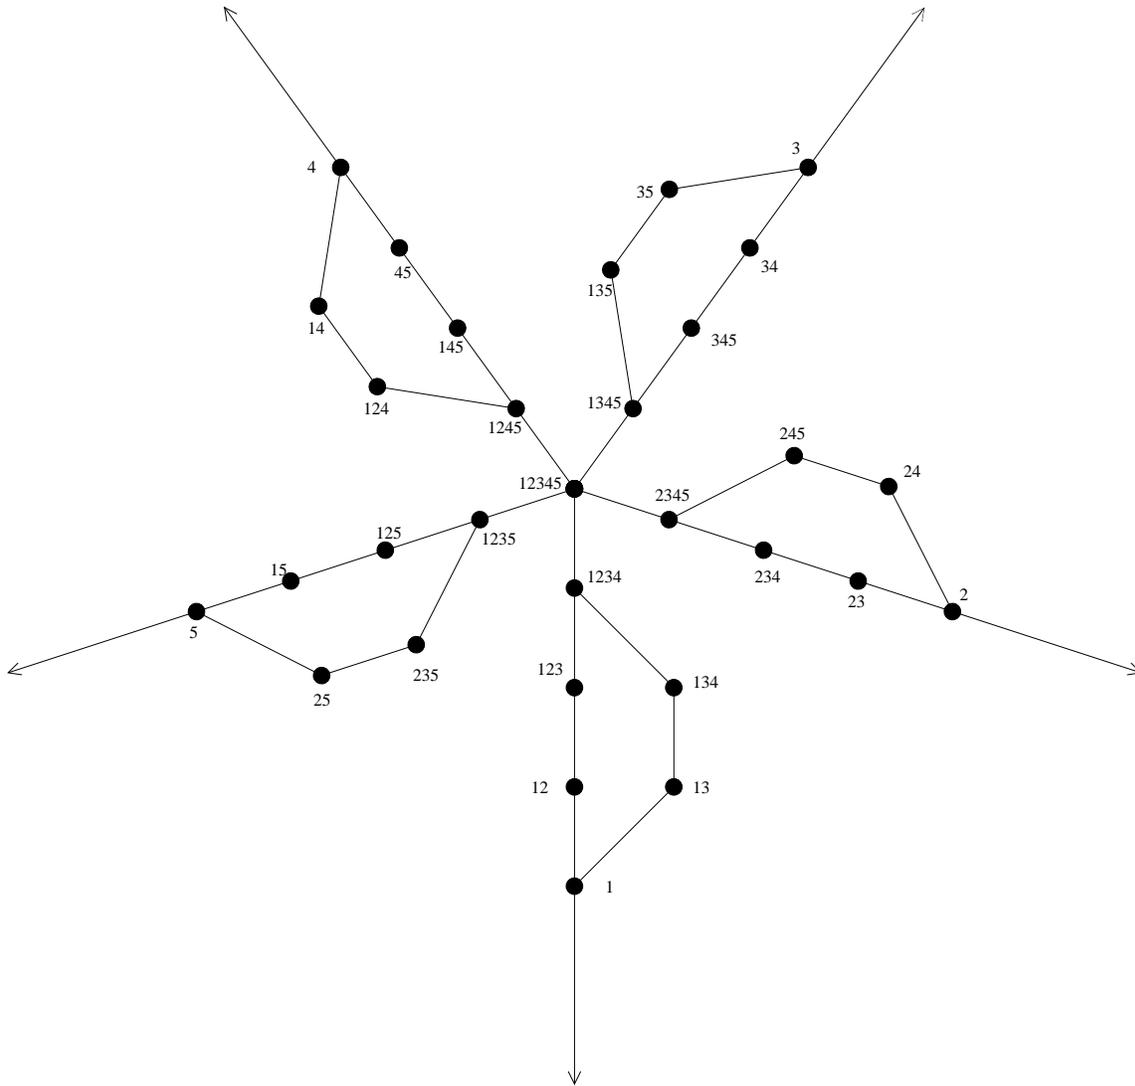


Figure 7: The planar graph P obtained by embedding the graph of Figure 6 in a pie slice and rotating it about the center by $(2\pi i)/n$ radians for each i , $0 \leq i \leq 4$. There is one more vertex at infinity, where the five edges shown by the arrows are incident.

interior of the face f_r of P at the point obtained by rotating point f_i^* by $(2\pi i)/n$ radians about p . Now a special point is chosen on each edge of P_0 . For each face f of P with an edge e of P_0 on its boundary, a “half-edge” of P^* is drawn from f^* to the special point on e , so that the half-edges incident at f^* are internally disjoint. Two half-edges of P^* meet at the special point of each edge e of P_0 to form the edge e^* of P^* . The complete embedding of P^* is now obtained by rotating the embedding of this section (clockwise) through $(2\pi i)/n$ radians for each i , $1 \leq i \leq n - 1$. (See Figure 8.)

The resulting embedding of P^* has rotational symmetry, but we want to show specifically that each of the n rotations of Θ_1 about p by an angle of $2\pi i/n$, $0 \leq i \leq n - 1$, coincides with one of the curves $\Theta_1, \Theta_2, \dots, \Theta_n$. Note that the j -edges in P_i , when rotated clockwise about p by an angle of $2\pi/n$, become the $(j - 1)$ -edges in $P_{i+1 \pmod n}$, if $j > 1$ or, if $j = 1$, the n -edges in $P_{i+1 \pmod n}$. Thus rotating Θ_j clockwise about p by an angle of $2\pi/n$ gives Θ_{j-1} if $j > 1$ or, if $j = 1$, Θ_n . (See Figure 9.) \square

Note 1. It can be checked that the resulting symmetric Venn diagram will also be monotone with the minimum number of vertices, owing to the SCD with the chain cover property of \mathcal{R}_n . This is not necessarily true in general: Grünbaum [Grü92] was the first to give examples of non-monotone symmetric Venn diagrams. Another 32 are reported by Ruskey in [Rus97].

3 Symmetric Chain Decompositions in the Boolean Lattice

3.1 The Greene-Kleitman Symmetric Chain Decomposition in \mathcal{B}_n

It is well-known that \mathcal{B}_n has a symmetric chain decomposition which can be constructed by the greedy lexicographic matching approach of Aigner [Aig73] or the parenthesis matching approach of Greene and Kleitman [GK76]. White and Williamson [WW77] and Griggs [Gri77b] have shown that both of these approaches give the same SCD, which, as shown in Greene and Kleitman [GK76], is the same as a natural recursive construction of deBruijn et al. [dBvETK51].

In this paper we make use of the formulation due to Greene and Kleitman. We go through the construction in detail since we will exploit several properties beyond those required to prove the SCD.

As a first step, in the string $x = x_1x_2 \cdots x_n$, regard the 0’s as left parentheses and the 1’s as right parentheses. Match parentheses in the usual way. That is, as string $x = x_1x_2 \cdots x_n \in \{0, 1\}^n$ is scanned from left to right, when a 0 is encountered, it becomes (temporarily) an unmatched 0. When a 1 is encountered, it is matched to the rightmost unmatched 0 to its left, if any, otherwise, it becomes an unmatched 1.

Let $U_0(x)$, $U_1(x)$, and $M(x)$ represent, respectively, the sets of indices of the unmatched 0’s, the unmatched 1’s, and the matched pairs for the entire string x . For example, if

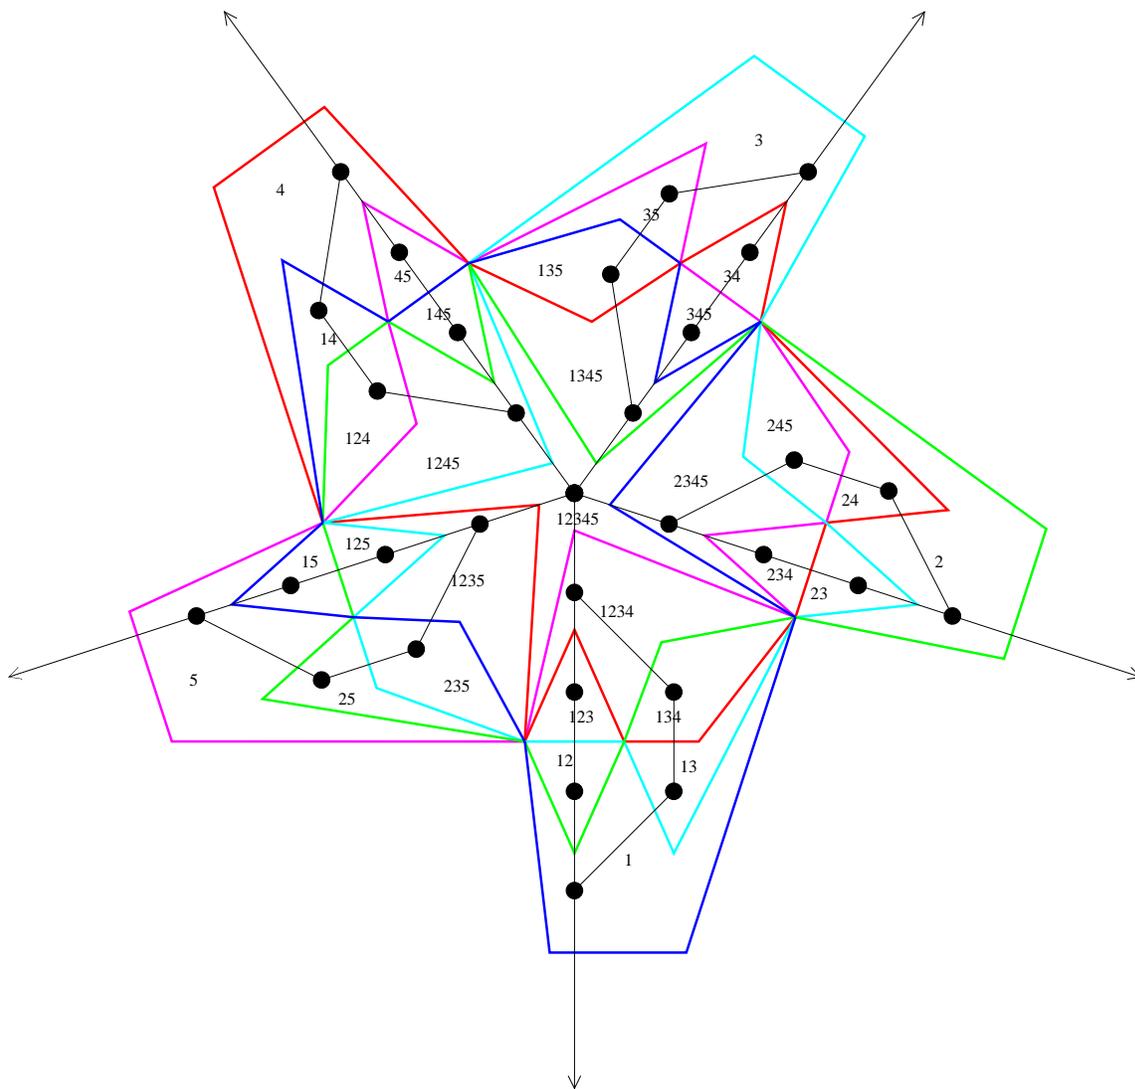


Figure 8: The geometric dual of the graph of Figure 7, embedded to have rotational symmetry, as described in the proof of Lemma 5.

$x = 01100101110010$, then

$$\begin{aligned} U_0(x) &= \{11, 14\}, \\ U_1(x) &= \{3, 10\}, \\ M(x) &= \{(1, 2), (5, 6), (7, 8), (4, 9), (12, 13)\}. \end{aligned}$$

Lemma 6 For $x \in \{0, 1\}^n$, if $U_1(x), U_0(x) \neq \emptyset$, then $\max(U_1(x)) < \min(U_0(x))$.

Proof. If $j \in U_1(x)$, x_j is an unmatched 1, so there is no unmatched 0 to the left of x_j . \square

The symmetric chain decomposition of \mathcal{B}_n is now defined by specifying, for each $x \in \{0, 1\}^n$, the successor of x , $\tau(x)$, on the chain containing x (we say $\tau(x) = nil$ if x is the last element of its chain):

$$\tau(x) = \begin{cases} nil & \text{if } U_0(x) = \emptyset, \text{ else} \\ x_1 \cdots x_{i-1} 1 x_{i+1} \cdots x_n & \text{where } i = \min(U_0(x)) \end{cases}$$

In the example above with $x = 01100101110010$, since $U_0(x) = \{11, 14\}$, we obtain $\tau(x) = 01100101111010$. Because of Lemma 6 we can make the following observations about τ .

Lemma 7 For $x \in \{0, 1\}^n$, let $U_1(x) = \{i_1, \dots, i_j\}$ and $U_0(x) = \{i_{j+1}, \dots, i_t\}$, where $i_1 < i_2 < \dots < i_t$. Then if $U_0(x) \neq \emptyset$,

$$\begin{aligned} S(\tau(x)) &= S(x) \cup \{i_{j+1}\} \\ U_1(\tau(x)) &= \{i_1, \dots, i_{j+1}\}, \\ U_0(\tau(x)) &= \{i_{j+2}, \dots, i_t\}, \\ M(\tau(x)) &= M(x). \end{aligned}$$

\square

Then chains in the Greene-Kleitman SCD are constructed starting from bitstrings with no unmatched 1 as follows.

Lemma 8 For $x \in \{0, 1\}^n$ with $U_1(x) = \emptyset$, let $k = |U_0(x)|$. Then

$$C_x = x, \tau(x), \tau^2(x), \tau^3(x), \dots, \tau^k(x),$$

is a symmetric chain ending at the string with 1 in position i if and only if $i \in S(x) \cup U_0(x)$.

Proof. By repeated application of Lemma 7, if the elements of $U_0(x)$, in increasing order, are i_1, i_2, \dots, i_k , then for $t \leq k$, $S(\tau^t(x)) = S(x) \cup \{i_1, i_2, \dots, i_t\}$ and $U_0(\tau^t(x)) = U_0(x) - \{i_1, i_2, \dots, i_t\}$. The set $U_0(\tau^t(x))$ becomes empty only when $t = k$. So, the chain starts at $S(x)$ and ends at $S(x) \cup U_0(x)$. Note that $|S(x)| = |U_1(x)| + |M(x)| = |M(x)|$, since $U_1(x) = \emptyset$ and $n - |S(x)| = |U_0(x)| + |M(x)|$, so

$$|S(x)| + |S(x) \cup U_0(x)| = |M(x)| + (n - |M(x)|) = n,$$

and, therefore, the chain is symmetric. \square

Theorem 1 (Greene and Kleitman [GK76]) *The following is a symmetric chain decomposition of \mathcal{B}_n :*

$$\mathcal{C} = \{C_x \mid x \in \{0, 1\}^n, U_1(x) = \emptyset\}. \quad (1)$$

Proof. To show that the chains are disjoint, define

$$\tau^{-1}(x) = \begin{cases} nil & \text{if } U_1(x) = \emptyset, \text{ else} \\ x_1 \cdots x_{j-1} 0 x_{j+1} \cdots x_n & \text{where } j = \max(U_1(x)) \end{cases} \quad (2)$$

It follows from Lemmas 6 and 7 that if $U_0(x) \neq \emptyset$ then $\tau^{-1}(\tau(x)) = x$ and if $U_1(x) \neq \emptyset$ then $\tau(\tau^{-1}(x)) = x$. Thus, the chains in \mathcal{C} are disjoint, and by the previous lemma they are symmetric. Furthermore, any element $x \in \{0, 1\}^n$ is in the chain C_y , where, by repeated application of τ^{-1} and Lemma 7, y satisfies $S(y) = S(x) - U_1(x)$ and $U_1(y) = \emptyset$. \square

Figure 2(b) shows the resulting SCD in \mathcal{B}_4 .

3.2 Monotone Venn Diagrams from the Greene-Kleitman SCD

We now extend Theorem 1 to show that this SCD has the chain cover property.

Lemma 9 *The Greene-Kleitman symmetric chain decomposition of \mathcal{B}_n has the chain cover property, with chain cover mapping π , defined for $C_x \in \mathcal{C}$ with $S(x) \neq \emptyset$ by*

$$\pi(C_x) = (C_y),$$

where $S(y) = S(x) - \{\max(S(x))\}$.

Proof. Since $C_x \in \mathcal{C}$, $U_1(x) = \emptyset$. Since we are given that $S(x) \neq \emptyset$, let $k = \max(S(x))$, so that $S(y) = S(x) - \{k\}$. We show that $C_y \in \mathcal{C}$ and the mapping $\pi(C_x) = (C_y)$ satisfies the required covering property.

Since in x , $x_k = 1$ was matched to a 0 in some position $m(k) < k$, this pair becomes unmatched in y , but no other matched pair is affected, so

$$U_0(y) = U_0(x) \cup \{k, m(k)\}; \quad U_1(y) = \emptyset; \quad M(y) = M(x) - \{(m(k), k)\}.$$

Since $U_1(y) = \emptyset$, y is a chain starter and $C_y \in \mathcal{C}$. Furthermore, by Lemma 8,

$$\text{terminator}(C_y) = S(y) \cup U_0(y) = (S(x) - \{k\}) \cup (U_0(x) \cup \{k, m(k)\}) = S(x) \cup U_0(x) \cup \{m(k)\},$$

which covers $\text{terminator}(C_x) = S(x) \cup U_0(x)$ in \mathcal{B}_n . Clearly x covers y in \mathcal{B}_n , so the required properties of π are satisfied. \square

Lemma 9 was the final piece required to prove the following.

Theorem 2 *For any $n \geq 1$, a monotone Venn diagram with minimum number of vertices can be obtained as the dual of any planar embedding of the planar graph $G(\mathcal{C}, \pi)$, where \mathcal{C} is the Greene-Kleitman SCD of \mathcal{B}_n , and π is the chain cover mapping for \mathcal{C} defined in Lemma 9.*

Proof. Combine Lemma 4 from the previous section with Lemma 9. □

Note 2. It follows from the proof of Lemma 9 that if $|C|$ denotes the length of a chain then $|C| = |\pi(C)| - 2$ for every non-root chain C in the Greene-Kleitman SCD. Thus, in the chain cover tree, $T(\mathcal{C}, \pi)$, all children of any node all have the same length. Thus *any* ordering of the children gives a planar embedding, by the technique described in Lemma 1. Indeed, any ordering of the children and their parent can be used! So, although some may be isomorphic, we can obtain many different monotone Venn diagrams by independently permuting the children at nodes in the chain cover tree.

3.3 Further Properties of the Greene-Kleitman SCD

We next identify some additional properties related to the Greene-Kleitman SCD which will be crucial for construction of symmetric Venn diagrams in Section 4.

Lemma 10 *For $x \in \{0, 1\}^n$, if $x_i = 0$ and either $x_{i+1} = 1$ or $x_{i-1} = 0$, then $i \neq \min(U_0(x))$.*

Proof. If $x_i = 0$, as x is scanned left-to-right, $x_i = 0$ is unmatched when it is first scanned. If $x_{i+1} = 1$, it will be matched with x_i , so $i \notin U_0(x)$. If $x_i = x_{i-1} = 0$, then both x_i and x_{i-1} are unmatched when x_i is scanned. As the scan continues beyond x_i , x_i has priority over x_{i-1} for being matched with a 1. Thus, if $x_i \in U_0(x)$, then so is x_{i-1} , so $\min(U_0(x)) \leq i - 1$. □

Similarly, we can prove the following.

Lemma 11 *For $x \in \{0, 1\}^n$, if $x_i = 1$ and either $x_{i+1} = 1$ or $x_{i-1} = 0$, then $i \neq \max(U_1(x))$.* □

4 Symmetric Chain Decompositions in a Necklace-Representative Poset

We show for prime n in this section, how to identify a subset $R_n \subseteq \{0, 1\}^n$ of necklace representatives, so that $\mathcal{R}_n = (R_n, \leq)$, the subposet of \mathcal{B}_n induced by R_n , has an SCD \mathcal{J} with the chain cover property. This proof is constructive, so by Lemma 5, this proves that symmetric Venn diagrams exist for all prime n and provides a construction.

4.1 Block Codes for n -bit Strings

With each $x \in \{0, 1\}^n$, we associate a sequence $\beta(x)$ over the alphabet $\{2, \dots, n, \infty\}$ called the *block code* of x as follows.

If x has the form

$$x = 1^{a_1}0^{c_1}1^{a_2}0^{c_2} \dots 1^{a_k}0^{c_k}, \quad k > 0, \quad a_i > 0, \quad c_i > 0, \quad 1 \leq i \leq k, \quad (3)$$

then

$$\beta(x) = (a_1 + c_1, a_2 + c_2, \dots, a_k + c_k).$$

Otherwise, for convenience, we define $\beta(x) = (\infty)$.

We regard the block codes as ordered lexicographically, using $c < \infty$ for any integer c , so that the block code (∞) is strictly greater than $\beta(x)$ for any x of the form (3).

For example, considering the rotations of $x = 0011011$, we have

$$\begin{aligned} \beta(0011011) &= (\infty), & \beta(0110110) &= (\infty), & \beta(1101100) &= (3, 4), & \beta(1011001) &= (\infty), \\ \beta(0110011) &= (\infty), & \beta(1100110) &= (4, 3), & \beta(1001101) &= (\infty). \end{aligned} \quad (4)$$

If $b = (b_1, b_2, \dots, b_k)$ is a sequence of integers, we let $|b|$ denote the number of terms of b , $|b| = k$, and we let $\|b\|$ denote the sum of the integers comprising b , $\|b\| = b_1 + b_2 + \dots + b_k$. Analogous to the rotation σ of a string, define rotation σ on sequences by $\sigma^i(b) = (b_{i+1}, \dots, b_k, b_1, \dots, b_i)$ for $i < k$. Concatenation of sequences b and b' is denoted by bb' .

Lemma 12 *If $x \in \{0, 1\}^n$, y is a rotation of x , and both x and y start with '1' and end with '0', then $\beta(y)$ is a rotation of $\beta(x)$.*

Proof. Let $x = 1^{a_1}0^{c_1}1^{a_2}0^{c_2} \dots 1^{a_k}0^{c_k}$, $k > 0$, $a_i > 0$, $c_i > 0$, $1 \leq i \leq k$. Then for some t , $0 \leq t < k$, $y = \sigma^{a_1+c_1+\dots+a_t+c_t}(x)$, so $\beta(y) = \sigma^t(\beta(x))$. \square

The following result is well known (e.g. proof of Prop. 5.1.2 in [Lot83]).

Lemma 13 *If w is a sequence of length k and if $\sigma^i(w) = w$ where $1 \leq i \leq k - 1$, then there is a nonempty subsequence v of w such that $w = vv \dots v = v^t$ for some $t \geq 2$.* \square

Lemma 14 *If n is prime, no two strings of $\{0, 1\}^n$ in the same necklace have the same finite block code.*

Proof Assume $x \in \{0, 1\}^n$ with $(\infty) \neq \beta(x) = b = (b_1, b_2, \dots, b_k)$. Then $\|b\| = n$, which is prime. Suppose $y \neq x$ is a rotation of x and y starts with '1' and ends with '0'. Then by Lemma 12, $\beta(y) = \sigma^j(b)$ where $1 \leq j \leq k - 1$. If $\beta(x) = \beta(y)$, then $b = \sigma^j(b)$, so by Lemma 13, $b = v^t$ for some nonempty subsequence v of b and $t \geq 2$. Then $\|b\| = t\|v\|$. We must also have $\|v\| \geq 2$, since $b_i \geq 2$ for $1 \leq i \leq k$, by definition of block code. Thus, $\|b\|$ is not prime, a contradiction. \square

4.2 Choosing Necklace Representatives

We assume for the remainder of this section that n is prime. For $x \in \{0, 1\}^n$, let $\rho(x)$ denote the representative of the equivalence class of x under rotation. Choose necklace representatives as follows:

$\rho(x)$ is the rotation y of x for which $\beta(y)$ is minimum.

For example, if $x = 0011011$, then $\rho(x) = 1101100$, since it can be seen from (4) that $\beta(1101100) = (3, 4)$ which is the minimum over all rotations of 0011011.

By Lemma 14, if n is prime, each equivalence under rotation in $\{0, 1\}^n$ has a unique representative. We let R_n be the set of these representatives. Then

$$R_n = \{\rho(x) \mid x \in \{0, 1\}^n\} = \{x \in \{0, 1\}^n \mid \rho(x) = x\}.$$

Let $\mathcal{R}_n = (R_n, \leq)$ be the subposet of $\{0, 1\}^n$ induced by R_n . We call this poset the *necklace-representative poset*, and our goal is to show that it has a symmetric chain decomposition with the chain cover property.

4.3 A Symmetric Chain Decomposition for \mathcal{R}_n when n is Prime

We will make use of the Greene-Kleitman mapping τ from Section 3. Recall that if $S(x)$ is the set of positions of the 1's in x , if $U_0(x)$ is the set of positions of the unmatched 0's in x , and if $U_0(x) \neq \emptyset$, then $\tau(x)$ is defined by $S(\tau(x)) = S(x) \cup \{\min(U_0(x))\}$.

We first note that when τ is restricted to elements of $R_n^* = R_n - \{0^n, 1^n\}$ with at least two unmatched 0's, the block code is preserved by τ . Note that $x = 0^n$ and $x = 1^n$ are the only elements of R_n with $\beta(x) = (\infty)$.

Lemma 15 *If $x \in R_n^*$ and x has at least two unmatched 0's, then $\beta(x) = \beta(\tau(x))$.*

Proof. Since $x \in R_n^*$, $x_1 = 1$ and $x_n = 0$, so $n \in U_0(x)$ and $1 \notin U_0(x)$. We are given that $|U_0(x)| \geq 2$, so if $i = \min(U_0(x))$, we must have $2 \leq i \leq n - 1$ and $\tau(x) \neq nil$, so let $\tau(x) = y = y_1 \cdots y_n$. By definition of τ , the only difference between x and y is that $x_i = 0$ and $y_i = 1$. Furthermore, by Lemma 10 of Section 3, $x_{i-1}x_i x_{i+1} = 100$, and so $y_{i-1}y_i y_{i+1} = 110$. That is, $y = \tau(x)$ differs from x only in that one of the blocks $1^{a_j}0^{c_j}$ of x with $a_j \geq 1$ and $c_j \geq 2$ changes to $1^{a_j+1}0^{c_j-1}$. This does not change the block code. (It also implies that if x has the form (3), then so does y .) \square

Recall that if $U_1(x)$ is the set of unmatched 1's in x and if $U_1(x) \neq \emptyset$, then $\tau^{-1}(x)$ is defined by $S(\tau^{-1}(x)) = S(x) - \{\max(U_1(x))\}$. Similarly, when τ^{-1} is restricted to elements of $R_n^* = R_n - \{0^n, 1^n\}$ with at least two unmatched 1's, the block code is preserved.

Lemma 16 *If $x \in R_n^*$ and x has at least two unmatched 1's, then $\beta(x) = \beta(\tau^{-1}(x))$.*

Proof. Since $x \in R_n^*$, $x_1 = 1$ and $x_n = 0$, so $1 \in U_1(x)$ and $n \notin U_1(x)$. Since $|U_1(x)| \geq 2$, if $i = \max(U_1(x))$, we must have $2 \leq i \leq n - 1$ and $\tau^{-1}(x) \neq nil$, so let $\tau^{-1}(x) = y = y_1 \cdots y_n$. By definition of τ^{-1} , the only difference between x and y is that $x_i = 1$ and $y_i = 0$. Furthermore, by Lemma 11, $x_{i-1}x_i x_{i+1} = 110$, and so $y_{i-1}y_i y_{i+1} = 100$. That is, $y = \tau(x)$ differs from x only in that one of the blocks $1^{a_j}0^{c_j}$ of x with $a_j \geq 2$ and $c_j \geq 1$ changes to $1^{a_j-1}0^{c_j+1}$. This does not change the block code. (It also implies that if x has the form (3), then so does y .) \square

Now observe in the following corollary that when n is prime, τ maps elements of R_n^* with at least two unmatched 0's to elements of R_n^* and τ^{-1} maps elements of R_n^* with at least two unmatched 1's to elements of R_n^* .

Corollary 1 *If $x \in R_n^*$ and $|U_0(x)| \geq 2$, then $\tau(x) \in R_n^*$. Similarly, If $x \in R_n^*$ and $|U_1(x)| \geq 2$, then $\tau^{-1}(x) \in R_n^*$.*

Proof. It follows from Lemma 15 that $\beta(x) = \beta(\tau(x))$. Since $x \in R_n^*$, $\beta(x) \neq \infty$ and the sequence $\beta(x)$ is lexicographically smaller than any of its rotations. Thus, the same is true for $\beta(\tau(x))$, which, by definition of R_n , means that $\tau(x) \in R_n$. Furthermore, as noted at the end of the proof of Lemma 15, $\tau(x)$ starts with 1 and ends with 0, so $\tau(x) \in R_n^*$. A similar argument follows for $\tau^{-1}(x)$ from Lemma 16. \square

Theorem 3 *If n is prime, \mathcal{R}_n has a symmetric chain decomposition with the chain cover property.*

Proof. First consider $\mathcal{R}_n^* = (R_n^*, \leq)$, the poset \mathcal{R}_n with elements 0^n and 1^n removed. We show \mathcal{R}_n^* has an SCD.

Since for $x \in \{0, 1\}^n$, $U_1(x)$ is the set of unmatched 1's in x , for every $x \in R_n^*$, $1 \in U_1(x)$. Define, as chain starters, the set \mathcal{S}^* of strings in R_n^* for which the only unmatched 1 occurs in position 1:

$$\mathcal{S}^* = \{z \in R_n^* \mid U_1(z) = \{1\}\}.$$

By Lemma 7, $|U_0(x)| = |U_0(\tau(x))| + 1$. For $z \in \mathcal{S}^*$, let $k = k(z) = |U_0(z)| \geq 1$ (since $n \in U_0(z)$). If $k \geq 2$, $\tau^i(z)$ has at least 2 unmatched zeros for each i such that $0 \leq i \leq k - 2$. Thus by repeated application of Corollary 1, $\tau^i(z) \in R_n^*$ for $0 \leq i \leq k - 1$. Define the *chain of z* , J_z , by

$$J_z = z, \tau(z), \tau^2(z), \dots, \tau^{k-1}(z).$$

Clearly J_z is a chain in \mathcal{R}_n^* and, by Lemma 7, the terminator of J_z has its 1's in positions $S(\tau^{k-1}(z)) = S(z) \cup (U_0(z) - \{n\})$. To show J_z is symmetric we show that $|S(z)| + |S(\tau^{k-1}(z))| = n$. $S(z)$ is the set of ones in z , and since z has only one unmatched 1, z must have $|S(z)| - 1$ matched 0's and $|U_0(z)|$ unmatched 0's. So,

$$\begin{aligned} n &= 2|S(z)| - 1 + |U_0(z)| \\ &= |S(z)| + |S(z)| + |U_0(z)| - 1 \\ &= |S(z)| + |S(\tau^{k-1}(z))|. \end{aligned}$$

From Theorem 1, distinct elements of $\{0, 1\}^n$, each with at least one unmatched 0, cannot be mapped to the same element by τ , so this remains true when the domain of τ is restricted. Thus for distinct $z, z' \in \mathcal{S}^*$, the chains J_z and $J_{z'}$ have no common elements.

Furthermore, any element $x \in R_n^*$ is in the chain J_y , where, by repeated application of Lemma 7 and τ^{-1} , y satisfies $S(y) = S(x) - (U_1(x) - \{1\})$ and $U_1(y) = \{1\}$. By repeated application of Corollary 1, y is also in R_n^* , so $y \in \mathcal{S}^*$.

This shows that every $x \in R_n^*$ is in a chain J_z for some $z \in \mathcal{S}^*$ and therefore the set of chains

$$\{J_z \mid z \in \mathcal{S}^*\}$$

is a symmetric chain decomposition of \mathcal{R}_n^* . We extend this to an SCD in \mathcal{R}_n by extending the chain $10^{n-1}, 110^{n-2}, \dots, 1^{n-2}00, 1^{n-1}0$ to the chain $0^n, 10^{n-1}, 110^{n-2}, \dots, 1^{n-2}00, 1^{n-1}0, 1^n$. Then for the SCD in \mathcal{R}_n , the chain starters are

$$\mathcal{S} = (\mathcal{S}^* - \{10^{n-1}\}) \cup \{0^n\},$$

and the chain terminator of the chain J_z starting at z is 1^n , if $z = 0^n$. Else it's the string y with $S(y) = S(z) \cup (U_0(z) - \{n\})$.

It remains to show that the SCD $\{J_z \mid z \in \mathcal{S}\}$ of the necklace-representative poset \mathcal{R}_n , has the chain cover property. Suppose $z \in \mathcal{S} - \{0^n\}$. Then $1 \in U_1(z)$ and (since $z \neq 10^{n-1}$) $|S(z)| \geq 2$. Let $\alpha(z)$ be the string obtained by changing the last 1 in z to a 0. Then z covers $\alpha(z)$ in \mathcal{B}_n . We show

- $\alpha(z) \in \mathcal{S}^*$, so either $\alpha(z) \in \mathcal{S} - \{0^n\}$ or $\alpha(z) = 10^{n-1}$, and
- the terminator of chain J_z is covered by the terminator of chain $J_{\alpha(z)}$, if $\alpha(z) \neq 10^{n-1}$, and by $1^{n-1}0$, otherwise.

Since $z_1 = 1, z_n = 0$, and $|S(z)| \geq 2$, the position l of the last 1 in z must satisfy $1 < l < n$. Thus,

$$\begin{aligned} z &= z_1 z_2 \cdots z_{l-1} 1 0 \cdots 0, \\ \alpha(z) &= z_1 z_2 \cdots z_{l-1} 0 0 \cdots 0. \end{aligned}$$

Also, since $z \in \mathcal{S}^*$, z_1 is the only unmatched 1 in z . So, since $l > 1$, z_l is a matched 1. Let $m < l$ be the position of the 0 matched to z_l . Then m and l become unmatched 0's in $\alpha(z)$, so $U_0(\alpha(z)) = U_0(z) \cup \{l, m\}$. However, no matched 1 in a position $j < l$ is affected and clearly $z_1 = 1$ cannot become matched. Thus $U_1(\alpha(z)) = U_1(z) = \{1\}$.

Since $U_1(\alpha(z)) = \{1\}$, to show $\alpha(z) \in \mathcal{S}^*$, it remains to show $\alpha(z) \in R_n^*$: if $z_{l-1} = 1$, then the last block of z has the form $1^{a_i} 0^{c_i}$ where $a_i \geq 2$ and $c_i \geq 1$, and so the last block of $\alpha(z)$ is $1^{a_i-1} 0^{c_i+1}$, so $\beta(\alpha(z)) = \beta(z)$; if $z_{l-1} = 0$ and if $\beta(z) = (b_1, b_2, \dots, b_k)$, then $\beta(\alpha(z)) = (b_1, b_2, \dots, b_{k-2}, b_{k-1} + b_k)$. In the first case, $\beta(\alpha(z))$ is the lexicographically smallest of all of its rotations, since $\beta(z)$ was, and therefore $\alpha(z) \in R_n^*$. We show this is also true in the second case. We know that since $z \in R_n^*$, for any $i < k$, $(b_1, \dots, b_k) < (b_{i+1}, \dots, b_k, b_1, \dots, b_i)$ and, thus, $(b_1, \dots, b_{k-i}) \leq (b_{i+1}, \dots, b_k)$. Then $(b_1, \dots, b_{k-i-2}) < (b_{i+1}, \dots, b_{k-2})$, or both $((b_1, \dots, b_{k-i-2}) = (b_{i+1}, \dots, b_{k-2})$ and $b_{k-i-1} \leq b_{k-1} < b_{k-1} + b_k$. Either way, $(b_1, \dots, b_{k-2}, b_{k-1} + b_k) < (b_{i+1}, \dots, b_{k-2}, b_{k-1} + b_k, b_1, \dots, b_i)$, so $\beta(\alpha(z))$ is the lexicographically smallest of all of its rotations. Therefore $\alpha(z) \in R_n^*$.

To show that the terminator t of J_z is properly covered, note first that if $\alpha(z) = 10^{n-1}$, then $|S(z)| = 2$, and, therefore, since J_z is symmetric, $|S(t)| = n - 2$. But also, $t \in R_n^*$, so $t_n = 0$. Thus, t is covered by $1^{n-1}0$ on the chain J_{0^n} .

Otherwise, if $\alpha(z) \neq 10^{n-1}$, then $J_{\alpha(z)}$ is a chain in the SCD of \mathcal{R}_n which terminates at the bitstring with ones exactly in the positions

$$\begin{aligned} S(\alpha(z)) \cup (U_0(\alpha(z)) - \{n\}) &= (S(z) - \{l\}) \cup ((U_0(z) \cup \{m, l\}) - \{n\}) \\ &= S(z) \cup (U_0(z) - \{n\}) \cup \{m\}, \end{aligned}$$

which, since $m \notin S(z) \cup U_0(z)$, covers the terminator t of J_z which has $S(t) = S(x) \cup (U_0(x) - \{n\})$.

Thus, for $z \in \mathcal{S} - \{0^n\}$, the mapping π defined by

$$\pi(J_z) = \begin{cases} J_{0^n} & \text{if } |S(z)| = 2, \text{ i.e. if } \alpha(z) = 10^{n-1} \\ J_{\alpha(z)} & \text{otherwise} \end{cases}$$

is a chain cover mapping for the SCD $\{J_z \mid z \in \mathcal{S}\}$, of \mathcal{R}_n , and, therefore, the necklace-representative poset has an SCD with the chain cover property. \square

Note 3. It is interesting that for each chain J_z in the SCD of \mathcal{R}_n , $C_{\tau^{-1}(z)}$ is a chain in the Greene-Kleitman SCD of \mathcal{B}_n (except for $z = 0^n$).

Note 4. It turns out that our SCD \mathcal{J} for \mathcal{R}_n and the chain cover mapping π for \mathcal{J} have the following property: if $\pi(J) = \pi(J')$ for chains $J, J' \in \mathcal{J}$, then J and J' have the same length. So, as in the case of Note 2, independently permuting the children of nodes in the chain cover tree $T(\mathcal{J}, \pi)$ gives rise to different (some possibly isomorphic) symmetric Venn diagrams.

From Lemma 5 and Theorem 3 we get our main result:

Theorem 4 *Venn diagrams with rotational symmetry exist for all prime n .* \square

The algorithm to compute and embed $G(\mathcal{J}, \pi)$ for arbitrary n has been programmed, and labeled versions of the output for $n = 11, 13$, and 17 are available as xfig files at: <http://www.csc.ncsu.edu/faculty/savage/>. We do show here the output for $n = 11$ and $n = 13$, with the vertex labels suppressed in consideration of the limited space. Figures 10 and 11, respectively, show planar embeddings of the chain cover graph $G(\mathcal{J}, \pi)$ resulting from the SCD \mathcal{J} and the chain cover mapping π in the necklace-representative poset \mathcal{R}_n when $n = 11$ and $n = 13$. These depict one “wedge” of the dual of the symmetric Venn diagram. Embedding this wedge in a pie slice of $(2\pi)/n$ radians and then rotating the embeddings through each of $(2\pi i)/n$ for $1 \leq i \leq n - 1$ gives the full symmetric dual of the Venn diagram. The resolution was too poor to show $n = 17$ on one page. Just for the record, we note that the symmetric 11-Venn diagram which results from Figure 10 is not isomorphic to the original one constructed by Hamburger in [Ham02], or to subsequent constructions for $n = 11$ [HS]. To see this, note that all wedges constructed by our method from symmetric chains in \mathcal{R}_n are necessarily symmetric about a line drawn between the middle levels, whereas the wedges of the constructions in [Ham02, HS] do not have that property.

5 Open Questions

Does there always exist a *simple* symmetric n -Venn diagram when n is prime? The answer is known to be yes for $n = 3, 5$ [Grü75], and 7 [Grü92], but is open beyond that.

Another interesting question to consider:

Does the necklace poset \mathcal{N}_n have a symmetric chain decomposition for all n (not just prime n)?

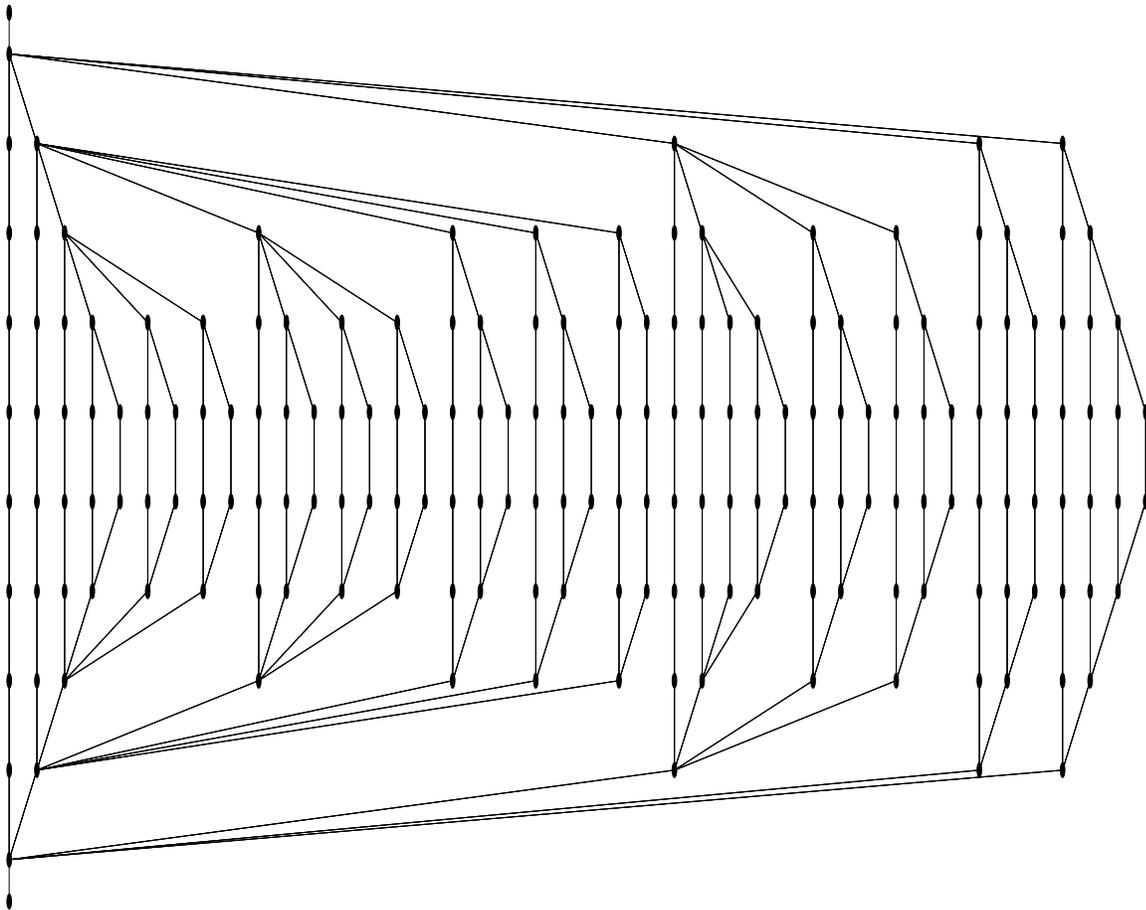


Figure 10: A planar embedding of the chain cover graph $G(\mathcal{J}, \pi)$ resulting from the SCD \mathcal{J} and the chain cover mapping π in the necklace-representative poset \mathcal{R}_{11} as described in Theorem 3.

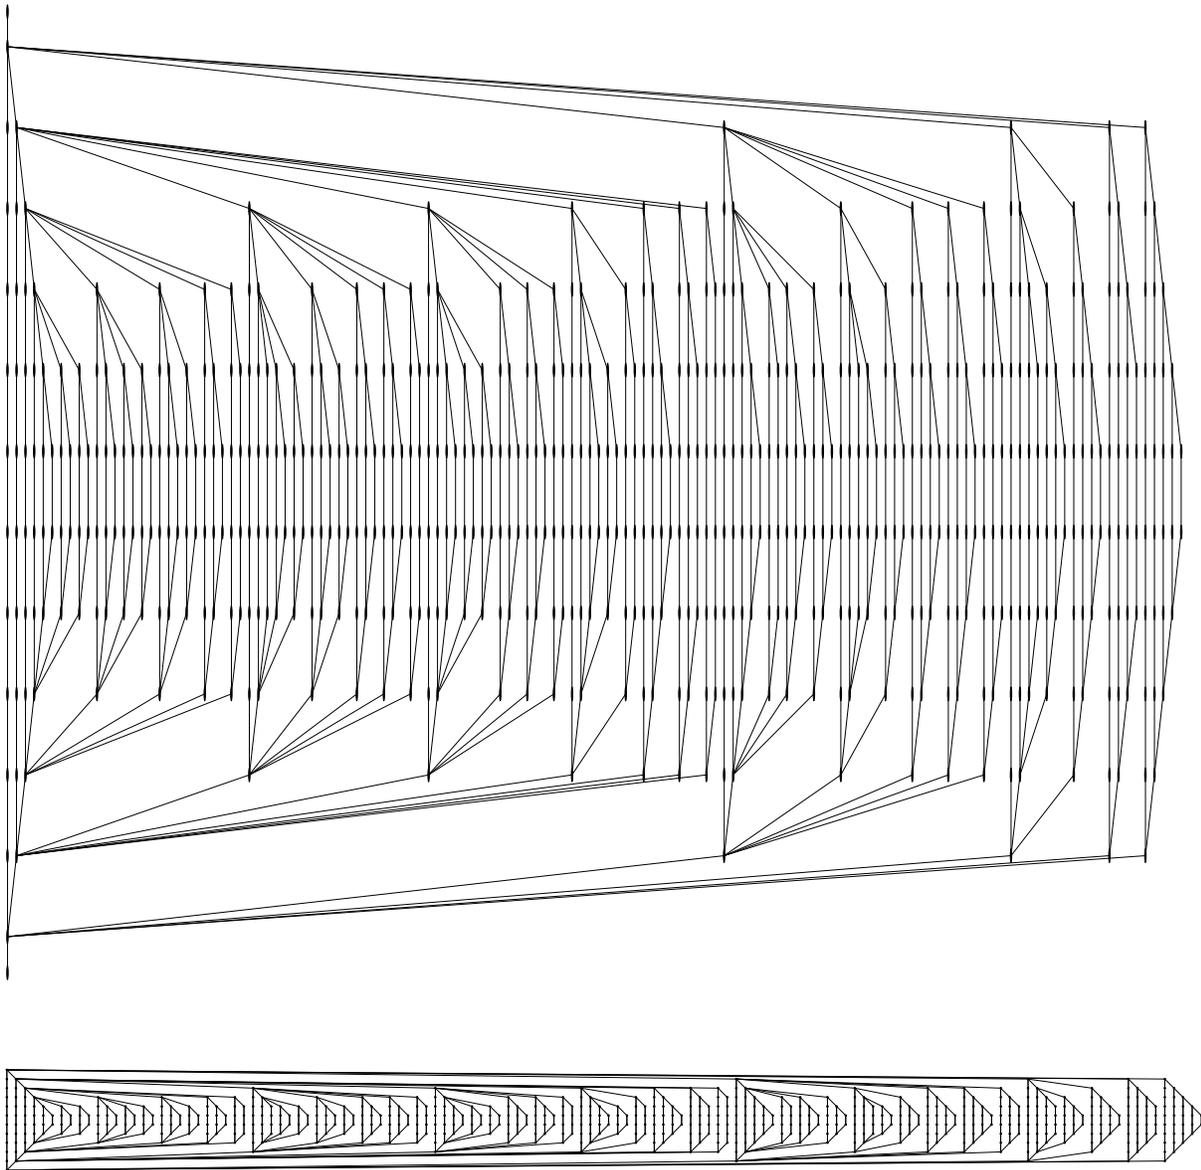


Figure 11: Two views of a planar embedding of the chain cover graph $G(\mathcal{J}, \pi)$ resulting from the SCD \mathcal{J} and the chain cover mapping π in the necklace-representative poset \mathcal{R}_{13} as described in Theorem 3.

Here is what we know so far for general n : A poset which has a symmetric chain decomposition is called a *symmetric chain order* (SCO). An SCO \mathcal{A} is rank-symmetric and rank-unimodal, meaning that if a_i denotes the number of elements of rank i ,

$$a_0 = a_n \leq a_1 = a_{n-1} \leq \cdots \leq a_{\lfloor n/2 \rfloor} = a_{\lceil n/2 \rceil},$$

where n is the rank of \mathcal{A} . It also holds that an SCO is strongly Sperner, which means that for all k , the union of k middle levels is a k -family (union of k antichains) of maximum size. A ranked poset with these properties of being rank-symmetric, rank-unimodal, and strongly Sperner is said to be a *Peck* poset. Thus, an SCO is Peck, but the converse is false in general.

The necklace poset \mathcal{N}_n is a *quotient* of the Boolean lattice \mathcal{B}_n , meaning that its elements are orbits of \mathcal{B}_n under the action of a group of automorphisms of \mathcal{B}_n , ordered in the natural way. For \mathcal{N}_n , the group is that induced on \mathcal{B}_n by the action of rotations of the ground set $\{1, \dots, n\}$, that is, the powers of the permutation $(12 \cdots n)$.

Stanley[Sta84] showed that any quotient of the Boolean lattice is Peck, hence this is true for the necklace poset, \mathcal{N}_n . For prime n , it is easy to show that \mathcal{N}_n satisfies an additional condition, the normalized matching condition, which is equivalent to what is called the LYM property. It turns out that the LYM Property, together with rank-symmetry and rank-unimodality (which are easily verified directly in the prime case without Stanley's result) implies the poset is a SCO[Gri77a]. In this paper we gave an explicit SCD of the necklace poset in the prime case. For composite n , however, the situation is still up in the air. But the fact that \mathcal{N}_n is Peck lends some support that it may also be a SCO for general n .

Acknowledgements. We would like to thank Frank Ruskey for many helpful discussions and for his advice on drawing the figures. For comments and suggestions on an earlier version of this paper we are grateful to several people including Branko Grünbaum, Frank Ruskey, Donald Knuth, and the referees.

References

- [Aig73] Martin Aigner. Lexicographic matching in Boolean algebras. *J. Combinatorial Theory Ser. B*, 14:187–194, 1973.
- [BGR98] Bette Bultena, Branko Grünbaum, and Frank Ruskey. Convex drawings of intersecting families of simple closed curves. 1998. preprint, submitted.
- [BR98] Bette Bultena and Frank Ruskey. Venn diagrams with few vertices. *Electron. J. Combin.*, 5(1):Research Paper 44, 21 pp. (electronic), 1998.
- [CHP96] Kiran B. Chilakamarri, Peter Hamburger, and Raymond E. Pippert. Venn diagrams and planar graphs. *Geom. Dedicata*, 62(1):73–91, 1996.

- [dBvETK51] N. G. de Bruijn, Ca. van Ebbenhorst Tengbergen, and D. Kruyswijk. On the set of divisors of a number. *Nieuw Arch. Wiskunde (2)*, 23:191–193, 1951.
- [Edw89] Anthony W. F. Edwards. Venn diagrams for many sets. *New Scientist*, 7:51–56, 1989.
- [Edw98] A. W. F. Edwards. Seven-set Venn diagrams with rotational and polar symmetry. *Combin. Probab. Comput.*, 7(2):149–152, 1998.
- [GK76] Curtis Greene and Daniel J. Kleitman. Strong versions of Sperner’s theorem. *J. Combinatorial Theory Ser. A*, 20(1):80–88, 1976.
- [Gri77a] Jerrold R. Griggs. Sufficient conditions for a symmetric chain order. *SIAM J. Appl. Math.*, 32(4):807–809, 1977.
- [Gri77b] Jerrold R. Griggs. Symmetric chain orders, Sperner theorems, and loop matchings. 1977. Ph. D. dissertation, M. I. T.
- [Grü75] Branko Grünbaum. Venn diagrams and independent families of sets. *Math. Mag.*, 48:12–23, 1975.
- [Grü92] Branko Grünbaum. Venn diagrams. II. *Geombinatorics*, 2(2):25–32, 1992.
- [Grü99] Branko Grünbaum. The search for symmetric Venn diagrams. *Geombinatorics*, 8(4):104–109, 1999.
- [Ham02] Peter Hamburger. Doodles and doilies, non-simple symmetric Venn diagrams. *Discrete Math.*, 257(2-3):423–439, 2002. Kleitman and combinatorics: a celebration (Cambridge, MA, 1999).
- [Har69] Frank Harary. *Graph theory*. Addison-Wesley Publishing Co., Reading, Mass.-Menlo Park, Calif.-London, 1969.
- [Hen63] D. W. Henderson. Venn diagrams for more than four classes. *American Mathematical Monthly*, 70:424–426, 1963.
- [HS] Peter Hamburger and Attila Sali. 11-doilies with vertex sets 275,276, ...,462. Preprint.
- [Lot83] M. Lothaire. *Combinatorics on words*, volume 17 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley Publishing Co., Reading, Mass., 1983. A collective work by Dominique Perrin, Jean Berstel, Christian Choffrut, Robert Cori, Dominique Foata, Jean Eric Pin, Guiseppe Pirillo, Christophe Reutenauer, Marcel-P. Schützenberger, Jacques Sakarovitch and Imre Simon, With a foreword by Roger Lyndon, Edited and with a preface by Perrin.

- [RRS51] A. Rényi, C. Rényi, and J. Surányi. Sur l'indépendance des domaines simples dans l'espace Euclidien à n dimensions. *Colloquium Math.*, 2:130–135, 1951.
- [Rus97] Frank Ruskey. A survey of Venn diagrams. *Electron. J. Combin.*, 4(1):Dynamic Survey 5 (electronic), 1997.
- [Sta84] Richard P. Stanley. Quotients of Peck posets. *Order*, 1(1):29–34, 1984.
- [Ven80] J. Venn. On the diagrammatic and mechanical representation of propositions and reasonings. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 9:1–18, 1880.
- [Wes96] Douglas B. West. *Introduction to graph theory*. Prentice Hall Inc., Upper Saddle River, NJ, 1996.
- [WW77] D. E. White and S. G. Williamson. Recursive matching algorithms and linear orders on the subset lattice. *J. Combinatorial Theory Ser. A*, 23(2):117–127, 1977.