Dijkstra's proof of Hall's theorem

Jayadev Misra

September 22, 2021

1 Hall's theorem on distinct representatives

Dijkstra gives a remarakably simple proof of Hall's theorem on distinct representatives [1] in EWD1269 (page 4). He uses almost no mathematical notation, instead proving the result by appealing to the matrix structure of the problem. I redo his proof eliminating the matrix terminology, and, I believe, that this is a simpler proof than Dijkstra's.

1.1 Distinct representatives

Given bag B whose members are sets S_i, $0 \leq i < n$, a system of distinct representatives (SDR) is a set of distinct elements t_i, $0 \leq i < n$, where $t_i \in S_i$. B is a *bag* of sets instead of a *set* of sets because some of the sets in B may be identical. Hall gives the necessary and sufficient condition for the existence of an SDR in bag B: for every subbag of B the number of elements in its constituent sets is at least the size (the number of members) of the subbag.

There are other ways to describe the problem. Bag B can be represented by a bipartite graph where one set of nodes denote the sets and the other set of nodes the elements. An edge is incident on set S and element x if and only if $x \in S$. An SDR, if it exists, is a set of edges such that every element has an incident edge, and no two edges are incident on the same node. Dijkstra employs a matrix to represent the bag: each row represents a set, a member of the bag, and a column an element. A matrix entry of 1 denotes that the element corresponding to the column belongs to the row corresponding to the set, the entry is 0 otherwise. A set of rows *cover* a particular column if and only if at least one row of the set contains a 1 in that column. He states Hall's theorem as follows: "$\cdots$ the columns can be ordered in such a way that condition H is met, viz. for each row index i, the i^{th} row has a 1 in the i^{th} column". Actually, he permutes both rows and columns in his proof.

1.2 Proof of Hall's theorem

Notation and terminology For a bag of sets B, $|B|$ is its size, i.e. the number of member sets in B, and $\#B$ the number of elements in all the constituent sets of B. To distinguish between bag and set operations use $+$ instead $\cup$ for

bag union, so $G + H$ is the union of bags G and H. □

Hall condition (HC): Bag S meets HC if and only if the number of elements in the constituent sets of S is at least its size, $\#S \geq |S|$.

Theorem 1 [Hall] A bag has an SDR if and only if all its subbags meet HC.

Proof: If a bag has an SDR, then any subbag S has at least $|S|$ representatives in the SDR, so $\#S \geq |S|$.

The main part of the proof is the converse, that if all subbags of B meet HC, then B has an SDR. The empty bag has a trivial SDR. I prove the result for non-empty bags by induction on bag size.

Dijkstra uses a very simple idea in EWD1269 to arrive at a constructive solution. Given that all subbags of the initial bag meet HC and an empty set violates HC, if we repeatedly remove an arbitrary element from an arbitrary set of the bag it will eventually have a subbag that violates HC. So, at some point removal of an element x from a set r in B causes the transition from bag B in which all bags meet HC to bag B' in which some subbag G' violates HC.

Any subbag of B that does not include r is unaffected by the removal of x. Since G' violates HC, it includes $r' = r - \{x\}$. Let G be the subbag of B corresponding to G' before removal of x, $G = (G' - \{r'\}) + \{r\}$. Observe that G is not empty since it includes r, but it may be the full bag B. Define $H = B - G$. Then $H \neq B$ because it excludes r, though it could be empty. I show disjoint SDRs for G and H which constitute an SDR for B.

• SDR for G:

$$\begin{array}{ll}
 & \mathit{true} \\
\Rightarrow & \{G \text{ is a subbag of } B\text{; so it meets HC}\} \\
 & |G| \leq \#G \\
\Rightarrow & \{G \text{ and } G' \text{ differ in } r \text{ and } r'\text{, only in } x\text{. So, } \#G \leq \#G' + 1\} \\
 & |G| \leq \#G \leq \#G' + 1 \\
\Rightarrow & \{G' \text{ violates HC; so } \#G' + 1 \leq |G'|\} \\
 & |G| \leq \#G \leq \#G' + 1 \leq |G'| \\
\Rightarrow & \{|G| = |G'|\} \\
 & |G| \leq \#G \leq \#G' + 1 \leq |G| \\
\Rightarrow & \{\text{arithmetic}\} \\
 & |G| = \#G = \#G' + 1
\end{array} \qquad \text{(L)}$$

Bags G and G' differ only in r and r'. From (L) $\#G = \#G' + 1$. So, $G - \{r\}$ does not include x. Subbag $G - \{r\}$, is a strict subbag of B; so it and all its subbags meet HC. Inductively, $G - \{r\}$ has an SDR which does not include x. Adding x to this SDR yields an SDR for G in which x represents set r.

• SDR for H: To prove that H has an SDR that is disjoint from the SDR of G, call the elements of G *black* and the remaining elements *white*. So all white elements are in H and the SDR of G consists of only black elements. I show that any subbabg S of H meets HC counting only its white elements.

$$
\begin{array}{ll}
 & \text{number of white elements of } S \\
= & \{\text{all elements of } G \text{ are black}\} \\
 & \#(S+G) - \#G \\
\geq & \{\#(S+G) \geq |S+G|, \text{ from HC. } \#G = |G|, \text{ from (L).}\} \\
 & |S+G| - |G| \\
= & \{|S+G| = |S| + |G|\} \\
 & |S|
\end{array}
$$

Therefore, every subbabg of H meets HC counting only its white elements. Bag H is strictly smaller than B because it does not include set r, so, inductively, it has an SDR of white elements. □

A critique of EWD1269 Dijkstra's proof of Hall's theorem is actually very simple, justifying the title of EWD1269, "A simple proof of Hall's theorem". As I have shown, once a member set r and element x in it are identified, it is straightforward to define subbags G and H, and prove by induction that they have disjoint SDRs. My only criticism is that using matrices in the proof is unnecessary, and, in fact, it complicates the proof.

The subbags G and H are shown as submatrices in Dijkstra's proof, for which he has to permute the rows and columns of B. Representation by bags and sets, instead, avoids any mention of rearrangement because these structures are oblivious to permutations. A matrix representation can deal with only non-empty bags, so his induction can not start with empty subbags. His proof has a moderate amount of case analysis (I share part of the blame, see the last paragraph of the EWD), which could have been completely avoided. Finally, no important property of a matrix is used in the proof except that it is an orthogonal arrangement of rows and columns.

Dijkstra says in EWD1269:

```
This note presents Hall's Theorem and its proof in an
almost visual terminology, whose only defence is that
it enabled me to design this proof without pen and paper,
while still in bed on an early Sunday morning.
```

I remonstrated with him at the time that his is not a streamlined argument. I have no doubt, though, that he could have constructed a proof as simple as the one given here, or something even simpler, if he had walked from the bed to his desk.

A simple proof of Hall's Theorem

This note presents Hall's Theorem and its proof in an almost visual terminology, whose only defence is that it enabled me to design this proof without pen and paper, while still in bed on an early Sunday morning.

Take a matrix of 0s and 1s with N rows. We say that a set of rows cover a certain column iff at least one row of the set contains a 1 in that column. A matrix is called happy iff, for any n ($n \leq N$), any n-tuple of its rows covers at least n distinct columns. Hence (i) for $N=0$ the matrix is happy, and a happy matrix (ii) has no row of 0s only, (iii) has at least N columns, and remains happy if (iv) a row is taken away, (v) an all-0 column is taken away, or (vi) a 0 is turned into a 1.

Hall's Theorem states that in a happy matrix the columns can be ordered in such a way that condition H is met, viz. for each row index i, the ith row has a 1 in the ith column.

Remark Above statement of the theorem refers to "the ith row", which is ugly be-

cause overspecific: the order of the rows is irrelevant in the sense that subjecting the rows and the equal number of left-most columns to the <u>same</u> permutation transforms a matrix meeting condition H into a matrix that still satisfies H. Peccavi. (End of Remark.)

We prove Hall's Theorem (which is an existence theorem) by displaying an algorithm that, given a happy matrix of N rows, rearranges its rows and columns in such a way that condition H is met. Because for $N=0$ condition H is met, we can confine our attention to happy matrices with $N>0$.

Since such a matrix contains at least one 1, we can select a 1 and <u>remove</u> it – i.e. replace it by a (pink) 0 – if the resulting matrix is still happy. We do so repeatedly until –and that moment comes!– we have selected a 1 whose removal would transform a happy matrix into an unhappy one. More specifically this means that the selected element x occurs in some k-tuple of rows such that with $x=1$, the k-tuple covers at

least k columns, while with $x=0$, that k-tuple covers less than k columns. Since the transition from $x=1$ to $x=0$ reduces coverages by at most one, we conclude

(0) the number of columns covered by the k-tuple equals k if $x=1$, and equals $k-1$ if $x=0$.

We now distinguish two cases.

$\underline{k=N}$ By permuting rows and columns, we move x to the top-left corner:

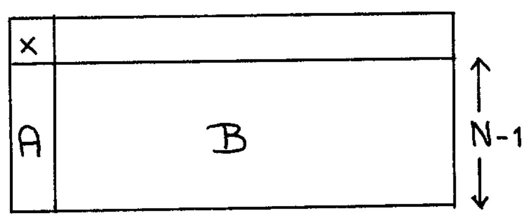

Because removal of the 1 at x would reduce the coverage of the whole matrix, the truncated column A consists of 0s only. Therefore, the set of columns covered by an n-tuple of rows from B remains the same when the rows are extended with their element from A; consequently the happiness of the whole matrix implies that B is happy. The

algorithm is applied recursively to matrix B , which has only N-1 rows.

<u>$k < N$</u> By permuting rows and colums, we move the rows of k-tuple to the top and the k columns they cover to the left:

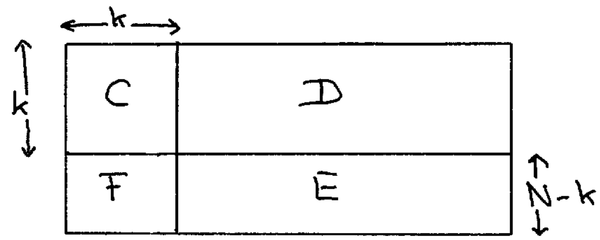

Element x occurs in C . Since (0) states that the k top rows cover the k left columns and nothing more, submatrix D consists of 0s only. We can now conclude the happiness of C and E as follows.

Because the whole matrix is happy, so is the top part formed by C and D, but since D , consisting of 0s only, does not contribute to the coverage, C is happy all by itself.

To establish the happiness of E , we consider an n-tuple of rows from the lower part formed by F and E , and extend this n-tuple with the k rows of the top part. Because

the whole matrix is happy, these $n+k$ rows cover at least $n+k$ columns, at least n of which lie in the right-hand part formed by D and E. Since D consists of 0s only, these n or more columns are covered via 1s in the E-part of the n-tuple of rows, hence E is happy. The algorithm is now applied recursively to matrices C and E, both of which have fewer than N rows.

And this concludes my proof of Hall's Theorem.

* * *

I thank Jayadev Misra for drawing my attention to the theorem and for suggesting later that my first proof (in which the two above partitions of the matrix were combined) could be simplified.

Austin, 3 January 1998

prof. dr Edsger W. Dijkstra
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712-1188
USA

References

[1] Philip Hall. On representatives of subsets. *J. London Math. Soc.*, 10(1):2630, 1935.