A Cantor-Bernstein theorem for paths in graphs

Reinhard Diestel & Carsten Thomassen

The Cantor-Bernstein theorem says that if for two infinite sets A and B there are injective functions $f\colon A \to B$ and $g\colon B \to A$ then there is a bijection $A \leftrightarrow B$. Perhaps the simplest and most intuitive proof considers the connected components of the bipartite graph whose vertex set is $A \cup B$ and whose edge set is

$$\big\{ \{ a, f(a) \} : a \in A \big\} \cup \big\{ \{ b, g(b) \} : b \in B \big\}.$$

As every vertex of this graph has one "outgoing" and at most one "incoming" edge, each of those components is a cycle or an infinite path. In each of these paths and cycles we now select every other edge to mark the desired bijection.

The Cantor-Bernstein problem, rephrased as above for graphs, has a natural generalization to paths. Let G be any graph, and let A and B be disjoint sets of vertices in G. Assume that we can find in G a set of disjoint paths from A to B that covers all of A (but not necessarily all of B), and a similar set of disjoint paths from all of B to A. Is there a set of disjoint A–B paths in G that covers both A and B?

Indeed there is. This was first shown in 1969 by Pym [3], and his proof is not short. Later [4], Pym also gave a short but indirect proof, which applies the *Rado Selection Principle* (an equivalent of the axiom of choice) to a suitably strengthened technical statement. Further interesting background, including a deduction of Pym's theorem from Tarski's fixed point theorem for lattices [5], can be found in Fleiner [2].

Our aim in this note is to give two short and direct proofs. Both are elementary, and they can be read independently. Our first proof is simpler, as long as readers are at ease with sequences indexed by ordinals and how to define such sequences inductively. The second proof avoids using the Axiom of Choice, which makes it a little more technical but perhaps also more illuminating.

A *path* in a graph G is a finite subgraph with distinct vertices $v_1, \dots, v_k$ and edges $v_1v_2, v_2v_3, \dots, v_{k-1}v_k$. We often refer to a path by the sequence of its vertices (in this or the reverse order); it then has a natural *first* and a natural *last* vertex. If these lie in sets A and B, and no other vertex of the path lies in $A \cup B$, we call it an *A–B path*. A set $\mathcal{P}$ of paths in G *covers* a set U of vertices if every vertex in U is the first or the last vertex of some path in $\mathcal{P}$. Any other notation we use can be found (online) in [1].

Theorem. *Let $G = (V, E)$ be a graph, and let $A, B \subseteq V$. Suppose that G contains a set $\mathcal{P}$ of disjoint A–B paths covering A, and a set $\mathcal{Q}$ of disjoint A–B paths covering B. Then G contains a set of disjoint A–B paths covering $A \cup B$.*

First proof. Our aim is to construct a sequence $(\mathcal{P}_\alpha)_{\alpha\leqslant\alpha^*}$ (indexed by ordinals) of sets of disjoint A–B paths each covering A so that $\mathcal{P}_{\alpha^*}$ also covers B. We shall define these $\mathcal{P}_\alpha$ "recursively", ie. for $\alpha = 0, 1, \dots$ in turn and so that the definition of $\mathcal{P}_\alpha$ may depend on that of $\mathcal{P}_\beta$ for any or all $\beta < \alpha$.

For each α, every path $P \in \mathcal{P}_\alpha$ will have the form $P = a \dots c \dots b$, where $c = c(P)$ is some specified vertex on P. The initial segment $a \dots c$ of P will always be an initial segment of some path in $\mathcal{P}$, and its final segment $c \dots b$ will be a final segment of some path in $\mathcal{Q}$. We write A_α for the set of all vertices on such initial segments $a \dots c$, ie. put $A_\alpha := \bigcup_{P\in\mathcal{P}_\alpha} V(Pc)$ where Pc denotes the initial segment $a \dots c$ of P. Note that $A \subseteq A_\alpha$, since by assumption $\mathcal{P}_\alpha$ covers A.

For each α, every $b \in B$ will have an 'index' $i_\alpha(b) \in \mathbb{N}$, defined as follows. Given $b \in B$, let Q be the path in $\mathcal{Q}$ ending at b. Let x be the last vertex of Q in A_α; this exists, because Q begins in $A \subseteq A_\alpha$. Then let $i_\alpha(b)$ be the length of xQ, the final segment of Q starting at x. (If $\mathcal{P}_\alpha$ happens to cover b, then xQ coincides with the final segment $c \dots b$ of the path in $\mathcal{P}_\alpha$ covering b.) We shall define the sets $\mathcal{P}_\alpha$ in such a way that for all $\beta < \alpha$ we have $i_\beta \leqslant i_\alpha$ (pointwise) and $i_\beta(b) < i_\alpha(b)$ for some $b \in B$. In particular, $i_\beta \neq i_\alpha$, giving $|\alpha^*| \leqslant \aleph_0^{|B|}$. Thus, our process of definition will terminate.

We start the definition of the $\mathcal{P}_\alpha$ with $\mathcal{P}_0 := \mathcal{P}$, putting $c := b$ for each path. For the recursion step at successor ordinals $\alpha+1$, let $\mathcal{P}_{\alpha+1}$ be obtained from $\mathcal{P}_\alpha$ as follows. If $\mathcal{P}_\alpha$ covers B, put $\alpha =: \alpha^*$ and stop the recursion. Suppose now that some $b' \in B$ does not lie on any path in $\mathcal{P}_\alpha$. Let Q' be the path in $\mathcal{Q}$ ending in b', and let x be the last vertex of Q' that lies on some path $P = a \dots c \dots b$ in $\mathcal{P}_\alpha$ (where $c = c(P)$). As $c \dots b$ is a final segment of some path $Q \neq Q'$ in $\mathcal{Q}$ but x does not lie on any other path in $\mathcal{Q}$, the vertex x precedes c on P (Figure 1). Let $\mathcal{P}_{\alpha+1}$ be obtained from $\mathcal{P}_\alpha$ by replacing P with $P' := aPxQ'b'$, and put $c' = c(P') := x$.

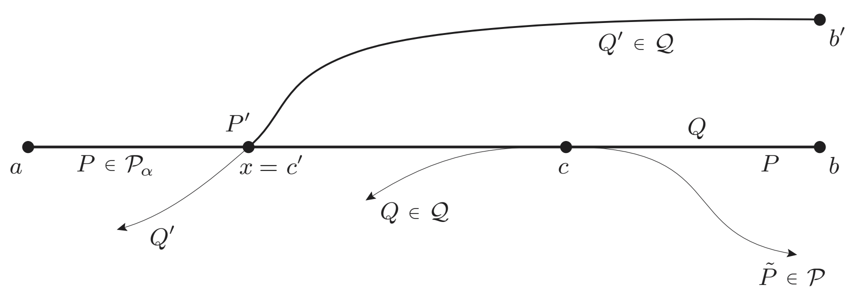

FIGURE 1. Modifying $P \in \mathcal{P}_\alpha$ into $P' \in \mathcal{P}_{\alpha+1}$

Clearly the new path P' again has the required form $a \dots c' \dots b'$, and $\mathcal{P}_{\alpha+1}$ covers A. Moreover, we have $i_{\alpha+1}(b) > i_\alpha(b)$. Indeed, $i_\alpha(b)$ is the length of the final segment cQ of the path $Q \in \mathcal{Q}$ ending in b. But cQ avoids $\mathcal{P}_{\alpha+1}$, so the final segment yQ of Q whose length is $i_{\alpha+1}(b)$ contains cQ properly. Finally, we have $i_{\alpha+1} \geqslant i_\alpha$ on all of B, because $A_{\alpha+1} \subseteq A_\alpha$.

It remains to consider the limit step in our recursion. Let α be a (non-zero) limit ordinal, and assume that $\mathcal{P}_\beta$ has been defined as required for all $\beta < \alpha$. Recall that when P is changed into P' in the successor step, its initial segment $a \dots c$ gets shorter. Thus, the path containing a given vertex $a \in A$ changes only finitely often as β approaches α. Hence for every a there is a path P starting in a that lies in $\mathcal{P}_\beta$ for every β greater than some $\beta_0 < \alpha$, and we take this path as the path in $\mathcal{P}_\alpha$ starting at a. To define the function $i_\alpha \colon B \to \mathbb{N}$, notice that for every $b \in B$ the value of $i_\beta(b)$ is bounded by the length of the path in $\mathcal{Q}$ ending in b, and so again the values of $i_\beta(b)$ agree for all β greater than some $\beta_0 < \alpha$ depending on b. We may thus take as i_α the pointwise limit of the functions i_β $(\beta < \alpha)$. $\square$

Formally, the path system constructed in the above proof depends on the choices of the uncovered vertex $b' \in B$ made at each step in the recursion. One can show, however, that these choices influence only the (transfinite) route by which the proof arrives at the final path system: that system itself is actually independent of the choices made in its construction.

The above observation suggests that it should be possible to rewrite the proof in a way that does not appeal to the axiom of choice. This is indeed possible. In the following proof we define the paths of the final system directly. This complicates the proof somewhat, because we now have to show that our 'locally' defined paths are disjoint and cover B.

Second proof of the Theorem (avoiding AC).

We shall consider various families $(P_a)_{a \in A}$ of disjoint A–B paths such that $a \in P_a$ for all a; let us call such a family an *A-family*. Every such path P_a will have a specified vertex $c = c(P_a)$ such that its initial segment $P_a c$ is contained in a path from $\mathcal{P}$ and its final segment cP_a is contained in a path from $\mathcal{Q}$. (For the paths $P \in \mathcal{P}$ we specify their last vertex as $c(P)$.) We shall write $\bar{P}_a := P\mathring{c}$ for the initial segment of P_a up to but not including c.

If a vertex x lies on $\bar{P}_a \cap Q$ for some $Q \in \mathcal{Q}$, and replacing P_a with the path $P'_a := P_a x Q$ results in another A-family (which is the case iff $xQ \cap P_{a'} = \emptyset$ for all $a' \neq a$), we say that this new family is obtained from the old by a *switch at x*, and specify $c(P'_a) := x$ (Figure 2). Note that since x lies on at most one path P_a and on at most one $Q \in \mathcal{Q}$, this switch (i.e., the new A-family) is well defined just by the vertex x.

Lemma. If $(P_a), (P'_a), (P''_a)$ are A-families such that (P'_a) and (P''_a) are each obtained from (P_a) by a finite sequence of switches, then an A-family (R_a) with $\bar{R}_a = \bar{P}'_a \cap \bar{P}''_a$ for all a can be obtained from (P_a) by a finite sequence of switches.

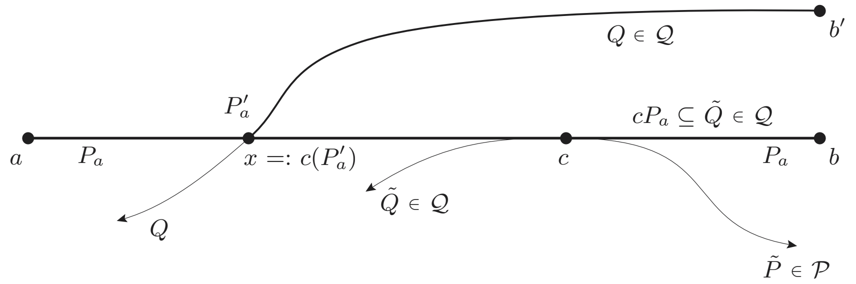

FIGURE 2. Changing P_a into P'_a by a switch at x

Proof of the lemma: Let (P''_a) have been obtained from (P_a) by switches at the vertices $x_1, \dots, x_n$ (in this order), with interim families (P^i_a) after switching at x_i. Now consider the family (P'_a), and apply switches at $x_1, \dots, x_n$ (in this order) whenever possible. More formally, we ask for $x_1, \dots, x_n$ in turn whether x_i defines a switch in the A-family (R^{i-1}_a) obtained from (P'_a) by switches at $x_1, \dots, x_{i-1}$ (whenever possible). If so, we perform this switch and call the resulting family (R^i_a); if not, we leave the current family unchanged, ie. put $R^i_a := R^{i-1}_a$ for all a. Induction on i shows that, for all i,

- $\bar{R}^{i-1}_a = \bar{P}^{i-1}_a \cap \bar{P}'_a$ for all $a \in A$;
- x_i defines a switch in (R^{i-1}_a) whenever $x_i \in \bar{R}^{i-1}_a$ for some a;
- $\bar{R}^i_a = \bar{P}^i_a \cap \bar{P}'_a$ for all $a \in A$.

For $i = n$ this is yields the desired result with $R_a := R^n_a$ □

We start our second proof by rewriting $\mathcal{P}$ as an A-family $(P_a)_{a \in A}$. For each $d \in A$ separately, let x_d be the first vertex on P_d such that a suitable finite sequence of switches turns (P_a) into an A-family $(P^d_a)_{a \in A}$ with $c(P^d_d) = x_d$. We claim that $(P^a_a)_{a \in A}$ is an A-family covering B.

Every path $P = P^a_a$ is taken from some A-family, and hence has a vertex $c = c(P)$ such that Pc is an initial segment of a path in $\mathcal{P}$ and cP is a final segment of a path in $\mathcal{Q}$. To show that the paths P^a_a are disjoint for different a, let $a' \neq a$ and consider $P' := P^{a'}_{a'}$. It suffices to show that $cP \cap P'c' = \emptyset$, where $c' := c(P')$. The minimality of $x_{a'} = c'$ implies that $P'c' \subseteq P^a_{a'}$. But $P^a_{a'}$ lies in a common A-family with $P = P^a_a$, and hence avoids cP.

To show that the paths P^a_a cover B, consider any uncovered $b \in B$ and let Q be the path in $\mathcal{Q}$ containing b. Let x be the last vertex of Q that lies on $P^a_a =: P$ for some $a =: a_0$. Then $x \in \bar{P}$, since otherwise $P \supseteq xQ \ni b$. Consider the finite set $A' := \{ a \neq a_0 \mid xQ \cap P_a \neq \emptyset \}$. By our lemma, there is a finite sequence of switches that turns $\mathcal{P}$ into a family (P''_a) such that $\bar{P}''_{a'} = \bigcap_{d \in A'} \bar{P}^d_{a'} = \bar{P}^{a'}_{a'}$ for all $a' \in A'$. Since xQ avoids all these $P^{a'}_{a'}$ and $x \in \bar{P} \subseteq \bar{P}''_{a_0}$, it follows that x defines a switch in (P''_a). This switch produces an A-family containing the path PxQ with $c(PxQ) = x \in \bar{P}$, contradicting the minimality of $c(P) = x_{a_0}$ on P_{a_0}. □

References

1. R. Diestel, *Graph Theory* (2nd edition), Springer-Verlag 2000. http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html
2. T. Fleiner, A fixed-point approach to stable matchings and some applications, *Mathematics of Operations Research* **28** (2003), 103–126.
3. J.S. Pym, The linking of sets in graphs, *J. Lond. Math. Soc.* **44** (1969), 542–550.
4. J.S. Pym, A proof of the linkage theorem, *J. Math. Anal. Appl.* **27** (1969), 636–638.
5. A. Tarski, A lattice-theoretical fixpoint theorem and its applications, *Pac. J. Math.* **5** (1955), 285–309.
6. C. Thomassen, Infinite graphs, in (L.W. Beineke and R.J. Wilson, eds.): *Selected Topics in Graph Theory 2*, Academic Press 1983.

Version 29.9.2003

Reinhard Diestel
Mathematisches Seminar der
Universität Hamburg
Bundesstraße 55
D - 20146 Hamburg
Germany
`diestel@math.uni-hamburg.de`

Carsten Thomassen
Mathematical Institute
Technical University
DK - 2800 Lyngby
Denmark
`C.Thomassen@mat.dtu.dk`