Designing the proof of Vizing's Theorem

by Josyula R. Rao & Edsger W. Dijkstra

We shall design the proof of the following theorem, due to V.G. Vizing.

Theorem For a finite undirected graph without auto-loops and without multiple edges, at any vertex of which fewer than N edges meet, N colours suffice for an edge colouring such that edges incident on the same vertex are of different colour.

* * *

The existence of such an edge colouring is demonstrated constructively. To this end it suffices to consider an algorithm that, given an acceptably coloured graph with one uncoloured edge, constructs an acceptable colouring of the whole graph. The graph being acceptably coloured is expressed by

(0) $(\underline{A}V:: acc.V)$,

where predicate acc on vertices is given by

acc.V ≡ (no two edges incident on vertex V have the same colour) .

In the rest of this note we confine ourselves to acceptably coloured graphs. In order not to interrupt the subsequent development of the algorithm too much, we first introduce a concept, the need of which will emerge, viz. the alternating path. An alternating path is a maximal path of at most two given colours. Because the colouring of the graph is acceptable, those two colours alternate along such a path, hence the name. One can show the following:

(i) an alternating path is either a cycle (i.e. without end points) or a simple path with two end points. (In order to avoid case analysis, we allow the end points to coincide, in which case the alternating path has length 0.)

(ii) a pair of colours and a vertex determine a unique alternating path of those colours and through that vertex. (We allow the two given colours to be equal, in which case the length of the alternating path is at most 1.)

(iii) swapping the colours of the edges of an alternating path maintains (0), i.e. leaves the colouring of the graph acceptable. (Here it is used that an alternating path with end points is maximal, i.e. cannot be extended at its end points.) The importance of swapping the colours of the edges of a noncyclic alternating path is that it changes the set of colours at the end points while leaving the colouring acceptable.

For each vertex V we define a colour c.V, in terms of which we define predicate f given by

$f.V \equiv$ (no edge incident on V has colour $c.V$) .

The fact that the number of colours is higher than the number of edges meeting at any vertex is exploited by choosing c in such a way that initially $(\underline{A}V{::}\ f.V)$ holds. In the following algorithm — which we explain afterwards — X, Y, and Z are variables of type vertex. Pre- and postcondition are given in full; intermediate assertions are only named and will be determined later.

{ $(\underline{A}V{::}\ acc.V) \wedge (\underline{A}V{::}\ f.V) \wedge$
  (XY is the only uncoloured edge)}
{P0: invariant}
<u>do</u> the $c.Y$-path ends in $Y \rightarrow$ {P1}
    determine Z so that edge
    XZ has colour $c.Y$ {P2}
  ; give edge XY colour $c.Y$ and
    uncolour edge XZ {P3}
  ; $Y := Z$ {P0}
<u>od</u> {P4}
; swap the colours along the $c.Y$-path {P5}
; give edge XY colour $c.Y$
{ $(\underline{A}V{::}\ acc.V) \wedge$ (all edges are coloured)}

The term

(0) $(\underline{A}V{::}\ acc.V)$.

is a conjunct of all intermediate assertions;

the term

(1) (XY is the only uncoloured edge)

is a conjunct of all, except P3. Note that, in view of the absence of autoloops, (1) implies that X and Y are two different vertices.

We now develop the algorithm, beginning with its last statement, viz. colouring XY with an acceptable colour. The colour is acceptable —i.e. the statement maintains (0)— provided it is incident on neither X nor Y. Somewhat asymmetrically, we decide to give edge XY colour c.Y, and take for P5

P5: $(0) \wedge (1) \wedge$ (c.Y not incident on X) $\wedge$ f.Y .

We now consider P5 as the postcondition to be established. Because the other conjuncts are implied by the program's precondition, we concentrate our attention on establishing

(c.Y not incident on X)

while maintaining the other three conjuncts of P5. If an edge incident on X has colour c.Y, its colour has to be replaced by a colour not incident on X. Since initially f.X, i.e. colour c.X not incident on X, we propose, in view of (iii), to swap under the initial validity of f.X the colours along the alternating path through X and with colours c.X and c.Y.

We call this "the c.Y-path"; for graphs with $(\underline{A}V :: acc.V) \wedge f.X$ we define —because we need the concept a number of times— for any colour p <u>the p-path</u> to be the alternating path through X with colours p and $c.X$. Because of $f.X$, colour $c.X$ is not incident on X , and hence

- no p-path is a cycle
- the p-path with $p = c.X$ is empty
- the non-empty p-path starts at X with an edge of colour p and ends at a vertex different from X .

In the above terminology, we proposed to establish ($c.Y$ not incident on X) by swapping the colours along the c.Y-path; this maintains $(0) \wedge (1)$, but maintains $f.Y$ only provided the c.Y-path does not end in Y . (Because of $f.Y$, a c.Y-path ending at Y does so with an edge of colour $c.X$; changing that colour into $c.Y$ would falsify $f.Y$.) Hence the following precondition suffices:

P4: $(0) \wedge (1) \wedge f.X \wedge f.Y \wedge$
(the c.Y-path does not end in Y) .

Viewed as postcondition, P4's last conjunct comes from the negation of the guard, the others have to come from the invariant. As we shall see shortly, all is available for the invariance of $(0) \wedge (1) \wedge f.X$; the invariance of $f.Y$, however, poses a problem. The invariance of $f.Y$ requires the conjunct $f.Z$ in $P3$, and the simplest way of justifying it there is by requiring $f.Z$ as conjunct in $P2$. (Besides this being the simplest way, our termination argument relies on the fact that no false f-value is truthified; see later.) For the justification of $f.Z$ in $P2$ we strengthen the invariant with a conjunct that restricts the occurrence of vertices V with $\neg f.V$ as follows

(2) for any colour p such that the p-path is not empty:
let the edge of colour p and incident on X be edge XV;
let the p-path end at vertex W;
then $c.W = p \Rightarrow f.V$.

Now the time has come to list and justify assertions $P0$ through $P3$.

$P0$: $(0) \wedge (1) \wedge (2) \wedge f.X \wedge f.Y \wedge X \neq Y$.

Assertion $P0$ is implied by the precondition: its first two conjuncts occur in the precondition, the next three follow from $(\underline{A}V:: f.V)$, and $X \neq Y$ follows from (1) and the absence of autoloops. $P0$'s reestablishment at the end of the repeatable statement will be discussed after $P3$.

P1: P0 $\wedge$ (the c.Y-path ends in Y) $\wedge$
(the c.Y-path is at least 2 edges long) $\wedge$
(the c.Y-path starts at X with an edge of colour c.Y) .

The first two conjuncts follow from the topology of the program; the next conjunct follows from the c.Y-path being a coloured connection between X and Y, and hence different from the unique edge XY, which is uncoloured; the last conjunct is a property of nonempty p-paths.

P2: P1 $\wedge$ (XZ has colour c.Y) $\wedge$ $X \neq Z$ $\wedge$ $Y \neq Z$ $\wedge$
f.Z $\wedge$ $c.Z \neq c.Y$.

The first conjunct P1, which does not refer to Z, is maintained and implies that Z is properly defined; XZ is the first edge on the c.Y-path and hence has colour c.Y; XZ being an edge, $X \neq Z$; the path length being ≥ 2, $Y \neq Z$. The conjunct f.Z follows from (2) with the instantiation $p, V, W := c.Y, Z, Y$. From f.Z and the fact that XZ has colour c.Y, we conclude $c.Z \neq c.Y$ (from which $Z \neq Y$ could have been concluded).

P3: (0) $\wedge$ (XZ is the only uncoloured edge) $\wedge$
(XY has colour c.Y) $\wedge$ f.X $\wedge$ $\neg$f.Y $\wedge$ f.Z $\wedge$ (2).

We now justify the conjuncts of P3 in order.

For (0) we consider X, Y, Z since for all other vertices V, acc.V remains unchanged; acc.X remains unchanged because the bag of colours incident on X is unchanged; acc.Y is not falsified on account of f.Y in P2 ; acc.Z is not falsified since the bag of colours incident on Z is decreased.

The next two conjuncts follow from (1), the statement, and the fact that XY and XZ are different edges.

The term f.X is maintained since the bag of colours incident on X is left unchanged; f.Y is falsified because colour c.Y is given to an edge incident on Y; f.Z is maintained because the bag of colours incident on Z is decreased.

For the invariance of (2), we distinguish two cases.

In the case $p \neq c.Y$ we observe that the colouring of XY and the uncolouring of XZ leaves the p-path unchanged; in particular, if the p-path is not empty, its V and W remain the same and, because of the absence of multiple edges, $V \neq Y$. Since c is a constant function and Y is the only vertex whose f-value changes, the invariance of

$$c.W = p \Rightarrow f.V$$

for $p \neq c.Y$ follows.

In the case $p = c.Y$ we observe that the c.Y-path XZ...Y is replaced by the c.Y-path

$XY\ldots Z$, i.e. we have to demonstrate

$$c.W = p \Rightarrow f.V$$

for the instantiation $p, V, W := c.Y, Y, Z$. We observe

$$
\begin{array}{ll}
 & c.W = p \\
= & \quad \{ p, W := c.Y, Z \} \\
 & c.Z = c.Y \\
= & \quad \{ P2 \} \\
 & false \\
\Rightarrow & \quad \{ \text{predicate calculus} \} \\
 & f.V
\end{array}
$$

The verification of $[P3 \Rightarrow wp.\text{“}Y := Z\text{”}.P0]$ is left to the reader.

For the termination argument we observe that, the function c being constant, f-values can only be changed by changing edge colours. The (un)colour statement in the repetition decreases $(\underline{N}V :: f.V)$ by 1 ; hence the repetition terminates.

And this concludes our proof of Vizing's Theorem.

Acknowledgements We thank R.E. Tarjan for having drawn our attention to the ugliness of the existing proofs; for constructive criticism we thank the ATAC –in particular K.L. Calvert and J. Misra– and the ETAC –in particular W.H.J.

Feijen — ; for active assistance we thank A.J.M. van Gasteren and David Gries.

In retrospect

The above is our second effort and the result is much better than our first presentation (which was probably still too much influenced by the published proofs).

The algorithm is more subtle than it reveals at first sight. It is clear that the repeatable statement has to change the uncoloured edge, and once it has been decided to uncolour XZ, the text of the repeatable statement follows. The more subtle point is that at the beginning of the repeatable statement, the situation seems symmetric in X and Y: the alternating path with colours c.X and c.Y through X ends at Y, but the alternating path of those colours through Y ends at X! The symmetry is destroyed by invariant (2) via the definition of the p-path, and while Y is a variable , X is a constant, and the set of p-paths is a linear collection.

The crux of the argument is, of course, conjunct (2) of the invariant. Again, (2) is in its formulation more subtle than might be appreciated at first sight. Its mathematical contents is easily established: strong enough to allow the inference of $f.Z$ in P2, weak enough to be maintained: $(\underline{A}V:: f.V)$ would be too strong. It was the formulation of (2) that required careful design. The one-point rule makes it possible to simplify the formulation by eliminating the dummy p, but this simplification of (2) makes the case analysis in its proof of invariance much harder to capture.

Finally we observe with satisfaction that we needed only very few variables (viz. X, Y, and Z, of which X is constant and Z in essence a variable local to the repeatable statement).

Nuenen, 13 August 1990

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