A Convergence Proof

Jayadev Misra

June 6, 2016

The following problem was communicated to me by Edsger W. Dijkstra in 2000, and he had heard it from another scientist.

An undirected connected finite graph has a natural number v_i initially associated with each node i. There is a distinguished node, *anchor*. A non-anchor node may make a *move* by setting its value to 1 + the minimum value over all its neighbors. Using the notation $i \sim j$ to denote that i is a neighbor of j and $N_i = 1 + \min\{v_j \mid j \sim i\}$, the move is given by the action is $v_i := N_i$. The anchor node never makes a move. Show that the computation eventually converges so that no move changes any value.

I show two informal proofs in which the argument is over computation sequences. A third proof gives an inductive state-based argument that can be formalized in a logic such as UNITY.

1 First Proof (from August 30, 2000)

Lemma 1 The value of any node is bounded.
Proof: Let M be the maximum initial value of any node. We show that a node i at distance k from the anchor has a value at most $M + k$ at any moment. Proof is by induction on k.

- $k = 0$: Then i is the anchor and initial value is at most M. It never makes a move; hence its value is at most $M + 0$ at all times.

- $k + 1$: Initially $v_i \leq M$, which satisfies the bound. Node i has a neighbor at distance k, and, from the induction hypothesis, this neighbor's value is at most $M + k$. Therefore, $V_i \leq M + k + 1$, and any move of i assigns it a value of at most $M + k + 1$. □

A move that assigns value n to a node is called a n-move; a move that assigns a value at or below n is a $\leq n$-move; similarly $>n$-move. It follows from Lemma 1 that there is a value B such that every move is a $<B$-move so that $N_i \leq B$ for any i at all times.

Lemma 2 There is a finite number of $\lhd n$-moves, for any n.

Proof: Proof is by induction on n.

- $n = 0$: No move sets a node value to less than 0.
- $n+1,\ n \geq 0$: From the induction hypothesis, there is a finite number of $<\!n$-moves. Consider the point in the computation, p, at which all such moves have been made. We claim that any node performs at most one n-move beyond p. Hence, there is a finite number of $<\!(n+1)$-moves.

  Let the first n-move of y beyond p be at q. The next move of y beyond q, if there is one, is not an n-move, because consecutive moves of the same node assign it different values; nor is the move a $<\!n$-move because all such moves have been completed. Therefore, this move assigns y a value exceeding n. Hence, every neighbor of y has a value at least n at this point. Subsequently, since there is no $<\!n$-move, value of each neighbor remains at least n. Therefore, all subsequent moves of y assign it values exceeding n; i.e., they are not n-moves. □

From Lemma 1, each node value is bounded. That is, there is a value B such that node values are always below B. Hence each move is a $<B$-move. From Lemma 2, there is a finite number of $<\!B$-moves. Hence the computation is finite.

Note The point beyond which no more n-moves can be made can be written as a state formula, $(\forall i : i \neq \text{ anchor } : v_i = N_i \vee N_i > n)$. It can be shown that this predicate is stable, clearly any move in this state is a $>\!n$-move.

2 Second Proof

This proof replaces the Lemma 2 of the previous section by a more direct proof.

Lemma 3 The number of $\leq\!n$-moves for any node is finite.
Proof: We prove the result by induction on n.

- $n = 0$: No node is ever assigned value 0. So, the number of 0-moves for any node is 0.
- $n+1$: We claim that between any two $(n+1)$-moves of any node i there is some $\leq\!n$-move by a neighbor of i. Since i has a finite number of neighbors, using the inductive hypothesis, the number of $\leq\!n$-moves by all neighbors of i is finite. Hence the number of $(n+1)$-moves of i is finite.

  To see the claim, consider a point p and a subsequent point r where i makes $(n+1)$-moves. Then the value of i is different from $n+1$ at some intermediate point q, otherwise i would not make a move at r. We show that some neighbor of i makes a $\leq\!n$-move between q and r.

  Now i makes a move different from $n+1$ at q, so $N_i \neq n+1$ at q. If a neighbor of i then makes a $>\!n$-move, $N_i \neq n+1$ is preserved because the move either overwrites the minimum value of a neighbor so that $N_i > n+1$

or it does not alter the minimum value, leaving $N_i \neq n+1$. Therefore, if the neighbors make only $>n$-moves from q to r, $N_i \neq n+1$ at r. Since i makes a $(n+1)$-move at r, $N_i = n+1$ at r, contradiction. □

Theorem 1 The number of moves is finite.
Proof: Each move is a $<B$-move (Lemma 1), $<B$-moves for any node is finite (Lemma 3), and the graph is finite. □

3 Outline of a formal state-based proof

We show that every move decreases a function value that is well-founded.

Partition the non-anchor nodes into *bins* where node i is placed in bin N_i; since $1 \leq N_i \leq B$ there are B bins. Node i is *balanced* if $v_i = N_i$; it is *unbalanced* otherwise.

3.1 Preliminary Results

Henceforth, α_j denotes the action $v_j := N_j$.

Proposition 1 For $i \neq j$, execution of α_j does not affect v_i.
For $i \not\sim j$, execution of α_j does not affect N_i.

Proposition 2 A balanced node that stays in its own bin after execution of an action remains balanced.

Proof: Suppose for node i, $v_i = N_i = m$ before execution of α_j, and it remains in bin m after the execution. For $i = j$ given that i is balanced, execution of α_i has no effect, thus keeping i balanced. For $i \neq j$, v_i is unchanged, so $v_i = m$ and i stays in bin m so $N_i = m$; hence $v_i = N_i$.

Proposition 3 On execution of α_j, j in bin n, every node in bin m, $m \leq n$, stays at or above m. And a node in a bin above n ($m > n$) stays above n.

Proof: The proposition follows from the following result. For nodes i and j, not necessarily distinct, $\{N_i, N_j = m, n\}\ \alpha_j ::\ v_j := N_j\ \{N_i \geq \min(n+1, m)\}$.

For $i \not\sim j$, α_j does not affect N_i, so $N_i = m \geq \min(n+1, m)$ is preserved.

For $i \sim j$, $N_i = \min(1 + v_j, 1 + \min\{v_k \mid k \sim i, k \neq j\})$. Using the rule of assignment, we need to show

$$N_i, N_j = m, n \ \Rightarrow\ \min(1 + N_j, 1 + \min\{v_k \mid k \sim i, k \neq j\}) \geq \min(n+1, m).$$

$$\begin{array}{ll} & N_i, N_j = m, n \\ \Rightarrow & \{N_i = \min(1 + v_j, 1 + \min\{v_k \mid k \sim i, k \neq j\}) \text{ and } N_i = m\} \\ & 1 + \min\{v_k \mid k \sim i, k \neq j\} \geq m\ \wedge\ N_j = n \\ \Rightarrow & \{N_j = n \Rightarrow 1 + N_j \geq n+1\} \end{array}$$

$$
\begin{array}{ll}
& 1 + \min\{v_k \mid k \sim i, k \neq j\} \geq m \ \wedge \ 1 + N_j \geq n + 1 \\
\Rightarrow & \{\text{arithmetic}\} \\
& \min(1 + N_j, 1 + \min\{v_k \mid k \sim i, k \neq j\}) \geq \min(n + 1, m)
\end{array}
$$

3.2 Main Proof

Consider execution of action α_j where j is in bin n. Let u_i, b_i be the number of unbalanced and balanced nodes, respectively, in bin i before the move and u'_i, b'_i the corresponding values after the move. We show that there is a bin m such that for all i, $1 \leq i < m$, $(u'_i, b'_i) = (u_i, b_i)$ and $(u'_m, b'_m) \prec (u_m, b_m)$, where $\prec$ is the lexicographic order. Therefore, the tuple $\langle (u_1, b_1), \cdots (u_i, b_i), \cdots (u_B, b_B) \rangle$ decreases lexicographically with each move. Since each of u_i and b_i is bounded from below, the number of moves is finite.

Proposition 4 For any m at or below n if $(u_i, b_i) = (u'_i, b'_i)$ for all i, $1 \leq i < m$, then $(u'_m, b'_m) \preceq (u_m, b_m)$.

Proof: From Proposition 3, no node moves from a bin above m to any bin at or below m. Given $(u_i, b_i) = (u'_i, b'_i)$ for all i, $1 \leq i < m$, it follows by induction on m that no node moves out of any bin below m. Thus, no node moves from above or below into bin m. From Proposition 2, any balanced node in bin m that stays in bin m stays balanced. So, $u'_m \leq u_m$, and $(u'_m, b'_m) \preceq (u_m, b_m)$. □

Theorem 2 There is a bin m such that for all i, $1 \leq i < m$, $(u'_i, b'_i) = (u_i, b_i)$ and $(u'_m, b'_m) \prec (u_m, b_m)$.

Proof: Consider two cases.

- Case 1) There is a bin m, $m < n$, such that $(u'_m, b'_m) \neq (u_m, b_m)$:

Let m be the lowest such bin. Then $(u'_i, b'_i) = (u_i, b_i)$ for all i, $1 \leq i < m$. From Proposition 4, $(u'_m, b'_m) \preceq (u_m, b_m)$. Given that $(u'_m, b'_m) \neq (u_m, b_m)$, $(u'_m, b'_m) \prec (u_m, b_m)$.

- Case 2) For all bins m, $m < n$, $(u'_m, b'_m) = (u_m, b_m)$:

From Proposition 4, $(u'_n, b'_n) \preceq (u_n, b_n)$. Node j goes from being unbalanced to balanced while staying in bin n, so $u'_n < u_n$. Therefore, $(u'_n, b'_n) \prec (u_n, b_n)$. □