

Graphs of modest diameter and degrees

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

The diameter of a graph is defined as the minimum path length, measured in number of edges, that suffices to connect any two nodes; we confine ourselves to undirected connected graphs with at most 1 edge between any two nodes. Nodes connected by an edge are called each other's neighbours. The degree of a node is defined as the number of edges incident on that node; it equals the number of neighbours of that node. Note that summing the degrees of all nodes yields twice the number of edges.

At the one extreme we have the complete graph: it has the minimum diameter, but the maximum number of edges. At the other extreme we have the string: it has the maximum diameter and the minimum number of edges. We are interested in graphs with modest diameter and modest maximum degree.

A well-known example of such a graph is the n -dimensional "cube": its diameter equals n and so does the degree of each of its 2^n nodes. To see this, we identify the nodes with the 2^n n -digit binary labels and connect any two labels that differ in exactly 1 digit. Thus each node has n neighbours and, hence, degree n ; the length of a shortest path between two such labels equals the number of bit positions in which these labels differ and, hence, the diameter equals n . We shall now construct graphs whose diameter is smaller than the 2^{\log_2} logarithm of the number of nodes while the maximum degree remains at that bound or is decreased as well.

To this end we generalize the solution presented by the hypercube in three different ways.

Firstly, we replace our binary labels by mixed-radix labels, in which one digit, which we denote as "the leading digit", plays a special role. For the purpose of this exposition, we can confine ourselves to labels consisting of a leading

x -ary digit followed by k y -ary digits; the graph has then $x * y^k$ nodes.

Secondly, the edges correspond to a subset of the pairs of labels that differ in exactly 1 digit. (It is this subsetting that allows the departure from binary digits while keeping the degrees down.) Note that, as in the case of the hypercube, a path consists of a sequence of labels in which each next label differs in exactly 1 digit from its predecessor.

Thirdly, the restriction that on the paths considered each digit position is subjected to at most 1 transition --viz. from source value to target value-- is waived for the leading digit.


Our construction has the following characteristics:

- whether two labels differing only in one of the y -ary digits are connected or not depends on the pair of y -ary digits in that position and on the (common) leading digit; the dependence is such that there exists at least one value for the leading digit such that the two labels would be connected; we refer to such a leading digit value as "enabling" the y -ary transition in question;
- whether two labels differing only in their leading digits are connected or not depends only on the pair of leading digit values; among the x possible values of the leading digit these connections form a connected graph, to which we shall refer as "the x -ary connectivity".


To determine a path from a given source label to a given target label, we determine on the x -ary connectivity a shortest path P from source leading digit to target leading digit, such that each of the at most k y -ary transitions required is enabled by at least one x -ary value on P . We can now effectuate the stepwise transformation from source label to target label by performing (in order) the x -ary transitions of P , interleaved with the required y -ary transitions in such a way that each of the latter is enabled by the leading digit of the intermediate label. As a result, the diameter is at most $k +$ the maximum length of P . (The proviso "at most" has been included because for very few required y -ary transitions the above procedure sometimes does not lead to the shortest path between source and target label; only for very small k , the proviso may not be vacuous.) In the following, L will denote the maximum length of P .

The purpose of the remainder of this note is to show how we can choose the parameters --primarily x , y , k , and the enabling rules-- so as to gain on the hypercube.

(i) We begin by exploring the case $y = 4$ by partitioning the 6 possible y -ary transitions in classes A, B, and C as follows:




The property of this classification is that each y -ary value can be subjected to exactly 1 transition from each class. The fact that we have 3 classes suggests that we consider the case $x = 3$, with as x -ary connectivity



here each x -ary value has been marked with the class of y -ary transitions it enables. For sufficiently large k , L is 3, and hence the diameter is $3 + k$. The degree is for each node $2 + k$. Compared with the hypercube with 2^5 nodes, the graph with $k = 2$ has the same diameter ($= 5$), but more nodes (48 instead of 32) and a lower degree (4 instead of 5). With increasing size, degree and diameter grow for the hypercube twice as fast as for this type of graph. (End of (i).)


(ii) We now consider $y = 4$ (with the same classes A, B, and C) and $x = 2$. The number of classes being one larger than the number of x -ary values, we are led to the x -ary connectivity



For $k = 1$, L equals 2, the diameter equals 3, for 4 nodes the degree is 2 and for the 4 other nodes the degree is 3. Compared with the cube we have saved 2 edges. For even k , we can construct graphs in which each

node has the same degree $1 + 1.5*k$ by interchanging for half of the y -ary digits the enabling roles of the two x -ary values. For $k = 2$, this graph has 2^5 nodes, but diameter and degree both equal to 4 (instead of 5 for the hypercube). Compared with the hypercube, the diameter grows half as fast, the degree with $3/4$ of the speed. This example has been included because it illustrates a method for distributing the edges more evenly among the nodes; it does not reduce the total number of edges, nor the diameter, but brings minimum and maximum degree closer to the average degree. (End of (ii).)

(iii) We now turn to $y = 5$. A systematic way of partitioning the 10 possible y -ary transitions in classes A and B is as follows:



Having two classes, we choose $x = 2$ with the x -ary connectivity



In general, the diameter is $2 + k$, whereas the degree is $1 + 2*k$. Because of the latter factor of 2 --in a label each y -ary digit can be subjected to 2 transitions-- and 5 is only a little bit larger than 4, the gain in degree is not spectacular. (E.g. for $k = 4$, we have 1250 nodes, degree 9 and diameter 6, which is a uniform improvement over the hypercube with 1024 nodes.) The construction is of particular interest for the case $k = 1$: 10 nodes of degree 3, but with diameter only 2 because each disabled y -ary transition can be accomplished by two enabled ones. (This is the famous Peterson Graph.) (End of (iii).)

Before showing a few more graphs, we mention two theorems about possible partitionings of y -ary transitions.

Theorem 0 For $y = 2*n + 1$, the $y*(y-1)/2$ y -ary transitions can be partitioned into n classes such that each y -ary value can be subjected to 2 transitions from each class.

(This partitioning is a generalization of the one shown under (iii).)

Proof Distribute the $2*n + 1$ values evenly along the circumference of a circle. Pairs of values with the same distance belong to the same class; because $2*n + 1$ is odd, each value has the same distance to two other values. (End of Proof.)

Theorem 1 For $y = 2*n$, the $y*(y-1)/2$ y -ary transitions can be partitioned into $y-1$ classes such that each y -ary value can be subjected to 1 transition of each class.

(This partitioning is a generalization of the one shown under (i); it is sometimes referred to as The Tournament Problem.)

Proof Place one value at the centre of a circle and distribute the remaining $2*n - 1$ values evenly along the circle's circumference. Draw a line from the centre to one of the values and, in the orthogonal direction, the $n - 1$ chords that pair the remaining values. The connected pairs belong to the same class; the other classes are obtained by rotation of the line/chords pattern. Because each possible chord length occurs in the pattern, each y -ary value can be subjected to $y-1$ different transitions, each from a different class. (End of proof.)

(iv) With $y = 6$, we get, according to Theorem 1, 5 transition classes. Analogously to (ii) we can construct a graph of $2*6^2 = 72$ nodes of degree $1 + 2 + 3 = 6$ and of diameter $2 + 1 + 1 = 4$. (End of (iv).)


For larger values of y , $x = 2$ is too small if the degree should not be larger than in the comparable hypercube: with $x = 2$, each next 7-ary digit gives an average degree increase of $(7-1)/2 = 3$, but $7 < 2^3$.

(v) With $y = 7$, we get, according to Theorem 0, 3 transition classes. With the x -ary connectivity as in (i), we get graphs with $3*7^k$ nodes of degree $2 + 2*k$ and with diameter $3 + k$. (End of (v).)

(vi) For $x = 3$, $y = 10$ is interesting. We get --Theorem 1-- 9 transition classes, and graphs with $3*10^k$ nodes of degree $2 + 3*k$ and with diameter $3 + k$, e.g. for $k = 2$: 300, 8, and 5 respectively. (Here we have extensively

exploited the freedom, illustrated in (ii), of allowing x-ary digit values to enable y-ary transition from several classes as presented by our Theorems.) (End of (vi).)


For $x = 4$ and $x = 5$ the following x-ary connectivities are worth considering



where A, B, C, D, and E stand for the (groups of) classes into which the y-ary transitions have been partitioned. They have the property that $L = x$ can be achieved with degree 2. (Achieving $L = x$ for $x = 6$ requires nodes of degree 3.)

(vii) For $x = 5$, $y = 16$ we thus generate a class of graphs with $5 \cdot 16^k$ nodes of degree $2 + 3 \cdot k$ and with diameter $5 + k$. For larger values of k the gain on the hypercube --both in degree and in diameter-- becomes impressive. (So does the size of the graph.) (End of (vii).)

Finally, we mention one further freedom: different x-ary values may enable the same y-ary transition. The 4-ary connectivity



gives rise to $L = 3$. It may be considered as an alternative if three groups A, B, and C are appropriate, but a binary or ternary leading digit is not.

Acknowledgements I am indebted to J.L.A. van de Snepscheut for drawing my attention to the possibility of gaining on the hypercube and to the Austin Tuesday Afternoon Club for its constructive comments on three (very different) versions of the presentation of the above. (End of Acknowledgements.)

Austin, 27 February 1987