

On a class of graphs with modest diameter

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 graphs, in which there is at most 1 edge between any two nodes. The degree of a node is defined as the number of edges meeting at that node; it equals the number of "direct neighbours" of that node. Note that summing the degrees of all nodes yields twice the number of edges.

The complete graph has a diameter equal to 1; the diameter of a star equals 2. Their maximum degree, however, is almost equal to the number of nodes. We are interested in graphs with a much smaller degree and yet a modest diameter. A well-known example of such a graph is the n -dimensional "cube": the degree of each of its 2^n nodes equals n , and so does its diameter. (Label the nodes with n -bit binary numbers and connect any two nodes whose labels differ in 1 bit. Then, firstly, each node has exactly n direct neighbours, and, hence, has degree n , and, secondly, the length of a shortest path between two nodes equals the number of bit positions in which their labels differ, and, hence, the diameter equals n .)

In [0] it was shown how the 3-cube could be modified so as to reduce the diameter to 2 without increase - actually: without change - of any of the degrees; without proof it was mentioned that, without changing the degrees, the n -cube could be transformed into a graph with a diameter equal to $(n \text{ div } 2 + 1)$. Looking for a constructive proof of that statement, I found a class of graphs for which the diameter is reduced by a factor of almost 3 and in which the number of edges is reduced by almost $1/3$; their maximum degree, however, is $n+1$. (See, however, the discussion of alternative constructions at the end of this note.) More precisely:

for n of the form $3 \cdot k + 2$ we construct a graph with 2^n nodes such that

- the diameter is $k+3$
- half the nodes have degree $n+1$
- half the nodes have degree $k+2$

The construction


We identify our nodes by strings of $k+1$ digits, more precisely; by 1 tetral (i.e. 4-valued) digit followed by k octal (i.e. 8-valued) digits. From now on we shall call our nodes "strings".

We introduce edges only between strings that differ in precisely one digit. A path thus consists of a sequence of strings, each next

string obtainable from its predecessor in the sequence by a one-digit modification.

In contrast to the n -cube, however, not every pair of strings that differ in one digit will be connected by an edge: this would yield a diameter of $k+1$ but, since the tetral digits can be modified into 3 other values and each octal digit into 7 other values, such freedom would imply for each string a degree of $3+7 \cdot k$ and would thus defeat our purpose of keeping the degrees small. We shall now detail which strings we connect by an edge. Strings differing in the tetral digit and strings differing in an octal digit will be dealt with separately.

The tetral digit Whether two strings differing only in their tetral digits are connected by an edge depends solely on these two tetral values. We choose some way of assigning the 4 tetral values to the corners of the following "tetral connectivity diagram"



and introduce an edge between two strings that start with tetral values connected in the above and are otherwise equal.

The letters at the corners serve a later pur-

pose of classifying strings according to their tetral digits: a string is an A-string, a B-string or a C-string according to the letter of the corner to which its tetral value is assigned.

The important property of the tetral connectivity diagram is that each corner can be reached from each - not necessarily different - corner via a path of length ≤ 3 that contains A, B, and at least one C corner. (The specific diagram given is primarily an existence proof for a connectivity diagram with the property just stated.)

The octal digits. Whether two strings differing only in one octal digit are connected by an edge depends on the combination of the pair of differing octal digits and the common tetral digit.

For the purpose of this description we write our octal digits in binary (i.e. ranging from 000 through 111). A pair of differing octal digits is an A-pair, a B-pair, or a C-pair according to the bit-wise sum (ranging from 001 through 111) of (the binary representations of) the two octal digits of the pair: for 3 values of the bitwise sum it is an A-pair, for 3 other values it is a B-pair, and for the last value it is a C-pair.

The specific way in which the 7 values of the bit-wise sum have been partitioned in $3+3+1$ is irrelevant. The important property is that keeping one octal digit fixed and letting the other one range over the 7 remaining values yields 3 A-pairs, 3 B-pairs and 1 C-pair. (The specific construction in terms of their bit-wise sum is primarily an existence proof of such a classification of pairs.)

Two strings differing in only one octal digit have the same tetral digit, and are thus both X-strings (with $X = A, B$ or C); we connect them by an edge if and only if the two differing octal digits form an X-pair.

And this completes the description of our construction of the graph, of which we shall now demonstrate that it enjoys the stated properties.

Proofs of the graph properties

The diameter We show how to construct a sequence of strings, starting with a given source string, ending with a given target string, and such that each pair of successive strings in the sequence is connected by an edge of the graph.

To this end we first compare in the k positions

the octal digits from the source string and the target string. In each of the k positions, we encounter 1 out of 4 situations: the two octal digits are equal - in which case the octal digit at that position remains constant - or the two digits form an A-pair, a B-pair, or a C-pair - in which case the octal digit in that position will change once (i.e. from initial value to final value).

So the path contains at most k edges whose traversal corresponds to the modification of an octal digit. But if source and target digit of such a modification form an X-pair, the corresponding edge is available if and only if the transition is between X-strings. Fortunately, the tetral digit can always be transformed from source value to target value in at most 3 steps so that all the string types needed for the octal modifications indeed occur in the sequence. So we never need to traverse more than $k+3$ edges; for $k \geq 2$, that number may actually be needed.

The degrees We determine the degree of a string by counting its direct neighbours.

Half the strings are A- or B-strings. From the tetral connectivity diagram we read that we can reach 3 neighbours by modifying the tetral


digit. Since in an A- or B-string each octal digit can be changed into 3 other values, we can reach $3 \cdot k$ neighbours by modifying an octal digit. Since all these neighbours are different, A- and B-strings have degree $3 \cdot k + 3 = n + 1$. A similar analysis shows that the C-strings have degree $k + 2$.

Discussion of alternatives

A slight modification brings the maximum degree back from $n + 1$ to n . We maintain the definition of X-strings, but modify the classification of the X-pairs: for 2 values of the bit-sum it is an A-pair, for 2 other values it is a B-pair, and for the last 3 values it is a C-pair. This does not affect the diameter; it yields for A- and B-strings a degree $2 \cdot k + 3$ and for C-strings a degree $3 \cdot k + 2 = n$. The price to be paid for this reduction is an increase in the total number of edges (though for $k \geq 2$ the number is less than in the n -cube).

Other variations are possible. For a graph with $5 \cdot 16^k$ nodes, we can characterize each node by a pental digit, followed by k hexadecimal ones. This is nice because $16 - 1$ is divisible. It is thus possible to achieve a diameter of $k + 5$ and for all nodes a degree of $3 \cdot k + 2$, and since $k \approx n/4$, this is asymptotically better. The above

can be achieved with the pental connectivity diagram



Paths ABCDEA, ABCDE, and ABCDED demonstrate that the pental digit need not be subjected to more than 5 modifications. Exploration of other variations - e.g. others than powers of 2 for the base of the k th power - is left to the reader

Acknowledgements I am indebted to J.L.A. van de Snepscheut for drawing my attention to the problem and to the members of the ATAC (= Austin Tuesday Afternoon Club) - in particular to Hamilton Richards Jr. and Amir Pnueli - for very helpful comments on an earlier effort of explaining my construction.

[0] J.L.A. van de Snepscheut, Private Communication.

Austin, 22 October 1986

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