GRAPHS AND INTERCONNECTION NETWORKS: DIAMETER AND VULNERABILITY

J.C. Bermond, J. Bond, M. Paoli & C. Peyrat
Université de Paris-Sud, Laboratoire de Recherche en Informatique, E.R.A. 452 du C.N.R.S., 91405 Orsay, France.

Contents

1 INTRODUCTION

It is well known that telecommunication networks or interconnection networks can be modelled by graphs. Recent advances in technology, especially the advent of very large scale integrated (VLSI) circuit technology have enabled very complex interconnection networks to be constructed. Thus it is of great interest to study the topologies of interconnection networks, and, in particular, their associated graphical properties. If there are point-to-point connections, the computer network is modelled by a graph in which the nodes or vertices correspond to the computer centres in the network and the edges correspond to the communication links. When the computers share a communication medium such as a bus, the network is modelled by a hypergraph, where the nodes correspond to the computer centres and the (hyper)edges to the buses. Note that there exists a second important class of networks, the "multistage networks", but we will not consider them. For a survey of interconnection networks, we refer the reader to Feng (1981).

In the design of these networks, several parameters are very important, for example message delay, message traffic density, reliability or fault tolerance, existence of efficient algorithms for routing messages, cost of the networks, ...

One important measure of the power of an interconnection network is the length of the longest path that the messages must travel from one node to another in the network, i.e. the distance between the nodes. It is advantageous to make these distances as small as possible, since this will reduce the message delay and also the message traffic density in the links. Worst case distance corresponds to the diameter of the associated graph or hypergraph. Similarly other network characteristics correspond to parameters of graphs or hypergraphs, e.g. mean distance, symmetry, connectivity, ...

A direct approach to network construction is to consider the graph model of possible links, with lengths and costs. The problem of designing a minimal or even near-minimal diameter subgraph with limited total cost was proved to be NP-complete by Plesník (1981). Also, Yannakakis (1978) proved that determining a connected subgraph with bounded maximal degree (the case of most network applications) of a given graph is NP-complete.

The object of this paper is to survey the results concerning diameter and connectivity in graphs and hypergraphs, in particular those of some importance for interconnection networks. For other results on the diameter we refer the reader to Bermond & Bollobás (1981), and on the connectivity to Mader (1979). Furthermore, we will only consider the deterministic aspect of these problems, though it is worth noting that the probabilistic aspect is considered in many papers and is of great importance for the reliability properties of networks.

Definitions and notation

We use standard terminology (see for example Berge (1973)). We give below some important definitions and notation which might be different from the usual terms.

The graph $G = (X,E)$ has vertex-set X and edge-set E. We denote by $n = |X|$ the number of vertices and by $m = |E|$ the number of edges. The <u>degree</u> $d(x)$ of a vertex x is the number of edges incident with x. We denote by δ the minimum degree and by Δ the maximum degree of the graph. A path with endpoints x and y is called an x-y path. The <u>distance</u> between two different vertices x and y is the length of a shortest x-y path. The <u>diameter</u> is the maximum distance over all the pairs of vertices. It will be denoted by D.

2 (Δ,D)-GRAPHS

In interconnection networks a practical problem is to interconnect the maximum number of nodes, while minimizing the diameter and knowing that, because of technical constraints, a node must not be incident to more than a fixed number of links. When a graph is associated with the network, this problem, first posed by Elspas (1964) and known as the (Δ,D)-graph problem, can be stated in graph-theoretic terms. How many vertices can a (Δ,D)-graph have, where a (Δ,D)-graph is a graph with maximum degree Δ and diameter at most D ? Let $n(\Delta,D)$ be the maximum number of vertices of a (Δ,D)-graph.

The Moore bound and Moore graphs

A theoretical bound on $n(\Delta,D)$ was given by Moore:

$$n(2,D) \leqslant 2D+1$$

$$n(\Delta,D) \leqslant \frac{\Delta(\Delta-1)^D-2}{\Delta-2} \quad \text{if} \quad \Delta \geqslant 3 .$$

Let the theoretical Moore bound be denoted $n_o(\Delta,D)$. The (Δ,D)-graphs with $n_o(\Delta,D)$ vertices are called Moore graphs. This name was given by Hoffman & Singleton (1960), because E.F. Moore raised the question with Hoffman and thought that eigenvalue techniques could be used to solve the problem (personal communication of Hoffman).

We warn the reader of the fact that the term Moore graphs is also used for what are known as cages. A cage is a Δ-regular of girth g having the minimum possible number of vertices. It can be shown that if g is odd, $g = 2D+1$, the minimum number of vertices is also $n_o(\Delta,D)$ and the cages are of diameter D. Therefore in this case, cages are exactly Moore graphs. If g is even, $g = 2D$, the minimum number $n_1(\Delta,g)$ of vertices is

$$n_1(\Delta,g) = 2D \text{ if } \Delta = 2 ,$$

$$n_1(\Delta,g) = \frac{2(\Delta-1)^D-2}{\Delta-2} \quad \text{if} \quad \Delta \geqslant 3 .$$

We propose to call these graphs bipartite Moore graphs because they achieve the maximum number of vertices of a bipartite (Δ,D)-graph. For a survey on cages see Wong (1982).

It has been proved by various authors (see Biggs' book (1974)) that the Moore graphs exist only for $\Delta = 2$ (2D+1-cycles), or $D = 1$ (Δ+1-cliques), or $D = 2$ and $\Delta = 3$ (the Petersen graph), or $D = 2$

and $\Delta = 7$ (the Hoffman-Singleton graph) and possibly for $D = 2$ and $\Delta = 57$. Such a $(57,2)$-Moore graph has not been exhibited, however Aschbacher (1971) has shown that it cannot be distance transitive. Singleton (1966) proved that bipartite Moore graphs cannot exist unless $\Delta = 2$ (2Δ-cycles) or $D = 2$ (the complete bipartite graph $K_{\Delta,\Delta}$) or $D = 3$, 4 or 6. In the cases $D = 3$, 4 or 6, it has been shown, in particular by Benson (1966), that such cages (bipartite Moore graphs) exist if $\Delta = q+1$, q being a prime power. Therefore we know infinite families of bipartite Moore graphs.

Instead of undirected graphs, we can consider directed ones. If D is a digraph with indegree and outdegree at most Δ and diameter exactly D, the number of vertices is no greater than the directed Moore bound: $n_{Dir}(\Delta,D) = 1+\Delta+ \ldots + \Delta^D$. Bridges & Toueg (1980) have proved the impossibility of Moore digraphs (i.e. digraphs achieving the Moore bound) for $D > 1$ and $\Delta > 1$. This was also proved by Plesník & Znám (1974).

Large (Δ,D)-graphs

When Moore graphs do not exist, the determination of the exact value of $n(\Delta,D)$ appears to be a very difficult problem. So an interesting problem, especially for practical applications, is to find large lower bounds for $n(\Delta,D)$. For example, Erdös, Fajtlowicz & Hoffman (1980) ask the following question: given a non-negative number σ, is there a graph G with diameter 2, degree Δ and $\Delta^2+1-\sigma$ vertices? They show that it is false for $\sigma = 1$ except if $G = C_4$, and therefore $n(\Delta,2) \leqslant n_o(\Delta,2)-2$ if $\Delta \geqslant 4$ and $\Delta \neq 7$ or 57. In fact, this problem seems very difficult. A simpler one is to ask whether, for a fixed diameter D, there exist infinite families of (Δ,D)-graphs with about $C(D)\Delta^D$ vertices. That would mean $\liminf_{\Delta\to\infty} n(\Delta,D)\Delta^{-D} \geqslant C(D)$. We know by the Moore bound that $C(D) \leqslant 1$, and a question is: can we have $C(D) = 1$? Partial results on this problem will be given later on.

In order to give a lower bound for $n(\Delta,D)$, we have to show the existence of (Δ,D)-graphs with many vertices. For this purpose, there are basically two methods: use random graphs or design such graphs. For applications in interconnection networks, the second method is certainly more interesting than the first one since it really gives good graphs, whereas the first one, though perhaps more powerful, guarantees only their existence. However, we will note later on a few results on diameter in random graphs; for more details we refer the reader to

Klee & Larman (1981), Bollobás (1981) and Bollobas & de la Vega (1983).

Construction of large (Δ,D)-graphs

We will now survey the different types of method which have been used for the construction of large (Δ,D)-graphs.

Geometric methods: for instance, the existence of projective planes of order q, when q is a prime power, permits us to construct graphs with degree $\Delta = q+1$, diameter $D = 2$ and q^2+q+1 vertices. For these methods, see Delorme (1983a).

Direct methods: a simple example is given by the generalized de Bruijn graphs (de Bruijn (1946)). Let $q > 0$, $m \geqslant 2$ and $Q = \{0,1,\ldots,q-1\}$. The de Bruijn graph $G = (Q^m,E)$ is the graph with vertex-set Q^m and edge-set E where $(\alpha,\beta) \in E$, if and only if $\alpha = a\gamma$ and $\beta = \gamma b$ for some $a,b \in Q$ and $\gamma \in Q^{m-1}$. It has maximum degree $\Delta = 2q$, diameter $D = m$ and q^m vertices. Since $q^m = \left(\frac{\Delta}{2}\right)^D$, we have: $\lim_{\Delta\to\infty} N(\Delta,D)\Delta^{-D} \geqslant 2^{-D}$. These graphs have been generalized by Lam & van Lint (1978), Memmi & Raillard (1982), Delorme & Farhi (1983) and Quisquater (1983a).

Recently, we received a copy of Fiol, Alegre & Yebra (1983). They construct regular digraphs of out- and indegree Δ, diameter D, having $\Delta^D+\Delta^{D-1}$ vertices. In the case $D = 2$, these graphs have the maximum possible number of vertices $(\Delta^2+\Delta)$ because Moore digraphs (with $\Delta^2+\Delta+1$ vertices) do not exist. If we consider the associated undirected graphs, we have, for Δ even, $n(\Delta,D) \geqslant \left(\frac{\Delta}{2}\right)^D+\left(\frac{\Delta}{2}\right)^{D-1}$. The authors also show that their digraphs can be viewed as iterated line digraphs of complete symmetric digraphs. Still more recently, Y. Raillard informed us that Kautz (1969) used these same graphs, which he referred to as state diagrams of shift registers, for the same purpose.

Addition of vertices: in certain graphs, vertices can be added without changing the degree or the diameter of the initial graph. See for example Delorme (1983a, 1983b). Recently, one of us (Bond) has shown that vertices can be added to the de Bruijn graphs. For there is a loop in these graphs for each vertex $(a,a,\ldots,a)$ and a double-edge between two vertices of the form $(a,b,a,b,\ldots)$ and $(b,a,b,a,\ldots)$. Eliminating these loops and double-edges, there remain q points of degree $2q-2$ and $q(q-1)$ of degree $2q-1$. Then we add q points x_i, $i = 0,1,\ldots q-1$, with q edges between x_i and vertices $(\alpha,i,\alpha,i,\ldots)$ for any $\alpha \in Q$, so that x_i has degree q. We add a vertex x_q (of degree q) connected to $(\alpha,\alpha,\ldots,\alpha)$ for any $\alpha \in Q$. So we can add q other vertices y_i $(i \in Q)$

such that they form a complete bipartite graph $K_{q+1,q}$ with the (q+1) vertices x_i, $i = 0,1,\ldots,q$. Now, we have a graph with q^m+2q+1 vertices, all of degree 2q except y_i. The resulting completed de Bruijn graph still has diameter m.

Products of graphs: in general, the classic products do not give good results, but Bermond, Delorme & Farhi (1982, 1983) have defined a new product which produces a lot of interesting (Δ,D)-graphs for small values of D and Δ.

Delorme (1983b) considers a product which, from two bipartite graphs, gives a bipartite graph. It consists of taking one of the connected components of the Kronecker product of two graphs; in what follows we will call it also a Kronecker product.

Let us consider two bipartite graphs $(X \cup X',E)$ and $(Y \cup Y',F)$ with diameters D and D' respectively, and degrees Δ_X, $\Delta_{X'}$, Δ_Y and $\Delta_{Y'}$ for the vertices of X,X',Y and Y' respectively. The Kronecker product of these two graphs has vertex set $(X \times Y) \cup (X' \times Y')$, and (x,y) is adjacent to (x',y') if and only if $(x,x') \in E$ and $(y,y') \in F$. The new graph is bipartite with maximum degree $\max(\Delta_X\Delta_Y,\Delta_{X'}\Delta_{Y'})$ and diameter $\max(D,D')$. Thus, if we take the Kronecker product of a bipartite graph $G = (X \cup X',E)$ and its opposite $(X' \cup X,E)$, the resulting graph has degree $\Delta_X\Delta_{X'}$ and diameter D(G).

Delorme uses for $(X \cup X', E)$ the bipartite Moore graph with diameter 4 and degree (q+1), where q is a prime power, and with $|X'| = |X| = (q+1)(q^2+1)$. He obtains thus a bipartite graph $P = (A \cup B,\mathcal{E})$ of diameter 4, degree $(q+1)^2$, and $2(q+1)^2(q^2+1)^2$ vertices. This graph P has a polarity, i.e. an automorphism f such that $f(A) = B$ and $f(B) = A$ and $f^2 = id$; here $f((x,x')) = (x',x)$. The quotient of this graph P by the polarity is the graph with vertex set A where two distinct vertices a and a' are adjacent if and only if a and f(a') are adjacent in P. The quotient Q has degree $(q+1)^2$, $(q+1)^2(q^2+1)^2$ vertices and diameter D equal to 3. Thus, we obtain a family of graphs of degree Δ, diameter 3 and with about Δ^3 vertices. Therefore: $\liminf_{\Delta\to\infty} n(\Delta,D)\Delta^{-3} = 1$. (We recall that because of the Moore bound $\liminf_{\Delta\to\infty} N(\Delta,D)\Delta^{-D} \leq 1$). In the same way, using the bipartite Moore graphs of degree q+1 and diameter 6, Delorme obtained by the same methods a family of graphs with degree Δ, diameter 5 and about Δ^5 vertices. For $D = 1,2,\ldots,10$ he gave improved lower bounds u_D for the asymptotic values $\liminf_{\Delta\to\infty} N(\Delta,D)\Delta^{-D}$.

D	1	2	3	4	5	6	7	8	9	10
u_D	1	1	1	$3^3.2^{-7}$	1	$2.5^6.6^{-6}$	$6^6.7^{-7}$	3.2^{-8}	14.3^{-8}	5.2^{-10}

Compounding methods: they have been developed by Bermond, Delorme & Quisquater (1982b), and also by Uhr (1983). In general, they consist of replacing the vertices in a large graph by copies of the same small graph or by different graphs. Some of these constructions were also considered by Wegner (1976); he proved, for example, that $n(\Delta,D) \geq (\Delta+1)n(\Delta-1,D-2)$.

Jerrum & Skyum (1983) have compounded de Bruijn graphs with a graph G', which have good properties of average path length, to obtain for Δ fixed, graphs with a diameter equal to $\lambda_\Delta \log_2 n+O(1)$. Their values of λ_Δ, which depend on Δ only, improved those of the de Bruijn graphs and are now the best obtained by constructive means. For example, they obtained cubic graphs with diameter $1.472169 \log_2 n+O(1)$; a previous construction giving a diameter of $1.5 \log_2 n+O(1)$ was given by Leland & Solomon (1982). However, the results on random graphs of Bollobas and de la Vega (1983) give better values. One of their main results is that, for $\Delta \geq 3$ and $\varepsilon > 0$ fixed, if $D = D(n)$ is the least integer satisfying $(\Delta-1)^{D-1} \geq (2+\varepsilon)\Delta n \log n$, then almost every Δ-regular graph of order n has diameter at most D. As a corollary of this result, for a fixed degree Δ, the minimum diameter of a graph of order n is bounded above by $\frac{\log n}{\log(\Delta-1)} + O(\log\log n)$. This is the best possible result because of the Moore bound: $D \geq \frac{\log n}{\log(\Delta-1)}$. It would be interesting to construct families of graphs which attain this bound.

Computer searches: these have been done by Quisquater (1983a) using Y-graphs and recently by Doty (1982) using chordal ring. Other algorithmic methods which need the help of a computer have been developed before by Toueg & Steiglitz (1979), Imase & Itoh (1981) and Arden & Lee (1982). Different tables of the largest known (Δ,D)-graphs, for $\Delta \leq 15$, $D \leq 10$, have been given by Storwick (1970), Leland, Finkel, Qiao, Solomon & Uhr (1981), Memmi & Raillard (1982) and Bermond, Delorme & Quisquater (1982a).

Minimum number of edges of a (Δ,D)-graph

Another problem close to the (Δ,D)-graph problem was proposed by Erdős & Rényi (1962). Suppose that there is a (Δ,D)-graph of order n. Let $e_D(n,\Delta)$ be the minimum number of edges of such a graph. An interconnection network whose associated graph achieves this bound is interesting because it has a minimum cost. Bollobás (1978, ch IV) presented, in the case $D \leq 3$, results of Erdős & Rényi (1962), Erdős, Rényi & Sós (1966) and Bollobás (1971). More recently, the case $D = 2$ has been settled almost completely by Pach & Surányi (1981). They defined a function $g(c)$ on $[0,1]$ and proved that $\lim_{n\to\infty} e_2(n,\lfloor cn \rfloor)/n = g(c)$ for every c, $0 < c < 1$, except for $c = c_1, c_2, \ldots$ where (c_k) is a sequence tending to 0.

Mean distance

We may also consider that the important parameter in an interconnection network is not the maximal time spent to travel from one node to another, but the average time. Therefore, we must study instead of the diameter of the associated graph, its average path length or mean distance. It has been defined by Doyle & Graver (1977):

$$\mu(G) = \frac{1}{\binom{n}{2}} \sum_{v,w \in X} d(v,w) .$$

Few things have been done on this parameter except for certain particular cases. The mean distance is easily calculated if the graph has great properties of symmetry, for example if it is distance degree regular, which means that each vertex has exactly D_j vertices at distance j from it. As this sort of graph is far from our topic, for other properties of them, we refer the reader to a survey of Bloom, Kennedy & Quintas (1981). Buckley & Superville (1981) have determined the mean distance for various classes of graphs, and Buckley (1981) for their line-graphs.

Cerf, Cowan, Mullin & Stanton (1975) introduced a problem similar to the (Δ,D)-graph problem: given the number n of vertices and the degree Δ of a graph, what is the lower bound of the mean distance? The graphs which achieve this bound are called generalized Moore graphs. The same authors (1974a, 1974b, 1976) described these graphs or established their non-existence up to $n = 34$, for $\Delta = 3$.

Other problems of distance have been investigated in the literature, for example problems of radius, eccentricity or of "locations in networks" (see the survey of Tansel, Francis & Lowe (1983) which contains many references).

3 CONNECTIVITY

In the design and analysis of interconnection networks, one of the fundamental considerations is the reliability, in particular that the centres can communicate in case of link or node failure. Reliability has been defined in a number of different ways, either with deterministic or probabilistic criteria. Maybe the probabilistic criteria fit better with the real world, indeed vertices or links can fail randomly: one of the main problems is to compute the probability that there is an operating path between two nodes knowing the probability of vertex or link failure. We will not discuss this non-combinatorial aspect here; we refer the reader to the surveys of Frank & Frisch (1970, 1971), Wilkov (1972) and Hwang (1978) in the case of multistage networks. Note that Ball (1980) has recently shown that virtually all network reliability analysis problems of practical interest are NP-hard.

In the case of deterministic networks, the criteria indicate how difficult it is to disrupt the operation of the whole network (or some part of it). Therefore one aim is to maximize the number of nodes or links that must fail in order to disrupt the operation of the network, while taking into account that there are cost constraints. The main criteria are the classical connectivity or edge-connectivity and the local connectivity.

Classical results

Recall that an x-y separating set S is a non-empty subset of vertices of G whose deletion destroys all the paths between x and y. The local connectivity between x and y is defined as the minimum cardinality of an x-y separating set. The connectivity of a non-complete graph is the minimum value of the local connectivity over all pairs of vertices. The connectivity will de denoted by $\kappa(G)$. By convention we set $\kappa(K_n) = n-1$. Analogous definitions can be given for local edge-connectivity and edge-connectivity; the latter will be denoted by $\lambda(G)$.

For an excellent survey of connectivity see Mader (1979). We will mention here only a very small number of results, which are of interest for our purpose. Recall first that by Menger's theorem the local connectivity between x and y is equal to the maximum number of vertex-disjoint paths between x and y. A similar result holds with edges instead of vertices. Using this or flow techniques there exist good (i.e. polynomial) algorithms to determine the local or the global connectivity (or edge-connectivity) of a graph (see for example Even's book (1979) or Chvátal's book (1983)). An algorithm in $O(n^3)$ for finding maximum flows is given in Malhotra, Kumar & Maheshwari (1978).

Symmetric graphs

Another well-known result says that, in any graph G, $\kappa \leqslant \lambda \leqslant \delta$. When the graph has some symmetries the following results are interesting; in particular they ensure that the graph has the best connectivity possible without any computation. A graph is said to be vertex-transitive if for every pair of vertices x and y there exists an automorphism f of the graph such that $f(x) = y$. Edge-transitive graphs are defined analogously.

Watkins (1970) and Mader (1970) have shown that if G is a connected edge-transitive graph then $\kappa = \lambda = \delta$. In the case of connected vertex-transitive graphs they have proved that $\kappa \geqslant 2\left\lfloor\frac{\delta}{3}\right\rfloor + 2$. Furthermore Mader (1971) proved that in any connected vertex-transitive graph $\lambda = \delta$. Mader (1970) also showed that if G is a connected vertex-transitive graph without a complete graph on 4 vertices then $\kappa = \delta$.

This result of Mader (1970) can be used to show that each of the graphs defined by Memmi & Raillard (1982) (see §2) has its connectivity equal to its degree. This result has been proved directly by exhibiting the paths by Amar (1983).

Maximum connectivity

Note that as $\kappa \leqslant \lambda \leqslant \delta$ the maximum connectivity or edge-connectivity of a graph with n vertices and m edges is $\left\lfloor\frac{2m}{n}\right\rfloor$. Harary (1962) exhibited constructions of graphs having this maximum connectivity. In fact, that follows from the construction of graphs on n vertices, of connectivity κ, having the minimum number of edges (and corresponds to the problem of designing a minimum cost network, in which all the edges have the same cost with a uniform reliability). If the connectivity k is even, the best known of these graphs is $C_n^{\frac{1}{2}k}$ the $\frac{1}{2}k$ power of a cycle C_n (that is the graph in which two vertices i and j are joined if $|j-i| \leqslant \frac{1}{2}k\ [\text{mod } n]$). If the connectivity k is odd and the