PARTIAL MATCHING IN DEGREE-RESTRICTED BIPARTITE GRAPHS

Douglas B. West[1] and Prithviraj Banerjee[2]
University of Illinois
Urbana, Illinois 61801

Abstract

Given a set of n "input" nodes, let $f(n,s,d)$ be the minimum number of "output" nodes in a bipartite graph such that every set of at most s inputs is joined to at least s outputs, subject to the restriction that every output node has degree at most d. Let $\theta(s,d)=\lim_{n\to\infty} f(n,s,d)$. A general construction yields $\theta(s,d)\leq s/(s+d-1)$, and this is optimal when $s\leq 2$ or $d\leq 2$. For larger values of s and d, better constructions are given. If $s\leq d^r$, then $\theta(s,d)\leq r\lceil s^{1/r}/r\rceil/d$. For d fixed and s large, the problem is closely related to the construction of bounded concentrators and expander graphs, used in VLSI applications, and results there can be applied. A linear programming formulation is given.

Keywords: Bipartite, matching, bounded concentrator, expander graph.

Running head: PARTIAL MATCHING

---

[1]Department of Mathematics and Coordinated Science Laboratory.

[2]Department of Electrical and Computer Engineering and Coordinated Science Laboratory.
Research supported in part by an IBM Graduate Fellowship and by the Joint Services Electronics Program under Contract N00014-79-C-0424.

CONGRESSUS NUMERANTIUM, 49 (1985), pp.259-266

1. Introduction

Imagine n users trying to communicate with a central facility that has m service ports. At any time, a set of at most s users may want to use the facility, and one user can be connected to each port at a given time. We assume there are direct communication links from users to service ports, with at most d links incident to a single service port. This constraint on degree may model a fan-in limitation of some circuitry at the service port. The problem is to satisfy the constraints with the minimum number of service ports.

More formally, given positive integers n, s, d, we seek the minimum m allowing a bipartite graph with n input nodes, m output nodes of degree at most d, and each set of s inputs matchable to s distinct outputs. By Hall's Theorem on matching in bipartite graphs, this is the same as having each set of $k \leq s$ users joined to at least k outputs. Call this the (s,d)-matching problem, and let $f(n,s,d)$ be the minimum value of m.

The efficiency of the construction is described by the ratio m/n, so we let $\theta(n,s,d) = f(n,s,d)/n$. It also makes sense to define $\theta(s,d) = \lim_{n \to \infty} \theta(n,s,d)$, because a solution for any given (n_0,s,d) can be replicated, with the leftover inputs connected straight across, to obtain $\theta(n,s,d) < \theta(n_0,s,d) + n_0/n$. This is useful in proving lower bounds, because for $\theta(s,d)$ we may assume the components of an optimal graph all have the same input-output ratio.

In Section 2 we present a simple construction for general (s,d) yielding $\theta(s,d) \leq s/(s+d-1)$, which is of interest only because it is optimal for $s \leq 2$ and $d \leq 2$. Here we also present a construction useful whenever s is not arbitrarily larger than d, i.e. $s \leq d^r$. This yields $\theta(d^r,s,d) \leq r\lceil s^{1/r}/r \rceil/d$. Compare this to the trivial lower bound $\theta(s,d) \geq 1/d$. If s is bounded in terms of d by $\log s \leq \log d \log\log d$, then putting $r = \lceil \log d \rceil$ yields $\theta(d^r,s,d) \leq \lceil \log d \rceil/d$. We have not chosen the base of the logarithm here; choosing the largest base that satisfies $\log s \leq \log d \log\log d$ will yield the best upper bound on $\theta(s,d)$ obtainable from this construction. One would like to choose d as a base to make the upper bound equal the lower bound, but this can only be done when $s = 1$.

We present one further construction useful when d is large. In particular, if $d = (s-1)\binom{m-1}{s-2}$ for some m, then $\theta(s,d) \leq (s-1)/d$. Call integers of the form $(s-1)\binom{m-1}{s-2}$ *special numbers*. Since $\theta(s,d)$ is decreasing in d, this yields $\theta(s,d) \leq (s-1)/q$, where q is the largest special number less than or equal to d. Finally, Section 3 returns to small cases to mention some numerical results from a linear programming formulation.

These constructions for fixed s and d have the ratio m/n going to 1 as s grows. However, related to this problem is an important and well-studied problem in which $\theta(n,s,d)$ is of interest for a fixed value of $\alpha = s/n$ rather than fixed s. Work in this area began with graphs called *concentrators*, first studied by Pinsker [8] and Margulis [7]. More clearly related to our problem is the subsequent (n,θ,k,α,c)-*bounded strong concentrator* defined by Gabber and Galil [6] to be a bipartite graph with n inputs, θn outputs, and at most kn edges, such that any set S of at most αn inputs has at least $c|S|$ neighbors. (Our graphs are $(n,m/n,dm/n,s/n,1)$-bounded strong concentrators.) These graphs can be used in a recursive construction of *superconcentrators*, originally defined by Valiant [12], which are networks with n inputs and n outputs that have vertex-disjoint paths from any set of inputs to any set of outputs of the same size. One wishes to construct these graphs, which have many applications, with a small number of edges. One can construct an n-superconcentrator from an $n\theta k\alpha$ 1-bounded strong concentrator and a θn-superconcentrator [9]; hence the desire to minimize θ. Increasing α allows saving some edges in the recursive construction.

We have no new constructions to present for this situation; we merely cite results from the theory of concentrators that apply here. The simplest and most interesting case is $\alpha = s/n = \frac{1}{2}$, treated by Pippenger [9]. Chung [4] later obtained a more general result. It must be noted that these are nonconstructive results obtained from counting arguments.

LEMMA [9]: For every m there exists a bipartite graph with $6m$ inputs and $4m$ outputs in which every input has degree at most 6, every output has degree at most 9, and every set of at most $3m$ outputs is matchable.

In other words, $\theta(6m,3m,9) \leq 2/3$. The surprise here is that for $d=9$, $\theta(n,\alpha n,d)$ does not go to 1 as $n \rightarrow \infty$, whereas we have proved that for $d=2$ it goes to 1. Chung's result, which is not restricted to a single value of d, shows that, for sufficiently small α and sufficiently large n, $\theta(n,\alpha n,d) \leq 5/6$ when $d=6$. We ask whether $\theta(n,\alpha n,d)$ can be bounded away from 1 when $d=3,4,5$.

LEMMA [4]: Let m,t,a,b be integers with $a \geq b \geq 2$, and choose a real number $0 < \alpha < 1$. If m is sufficiently large and the conditions below hold, then there exists a bipartite graph with $n = am$ inputs, bm outputs, input degree bt, output degree $d = at$, and every set of at most $s = \alpha n$ inputs matchable. Let $H(z)$ be the entropy function $-z\lg z-(1-z)\lg(1-z)$. The conditions are

$$t > \frac{H(\alpha)+(b/a)H(a\alpha/b)}{bH(\alpha)-\alpha a H(b/a)} \qquad \alpha < \frac{b(b-2)}{ab-a-b} \qquad bt > 4$$

Variants of this problem include the design of connectors [10], nonblocking networks [3,4], generalized connectors [10], and expanders [1,2,6,11].

2. Small values of s

In this section, we focus on explicit constructions, which are useful when s is not too large. We begin with a construction that is not very good in general, but is valid for all (s,d) and is optimal when $s \leq 2$ or $d \leq 2$.

THEOREM 1. $\theta(s,d) \leq s/(s+d-1)$, with equality when $s \leq 2$ or $d \leq 2$.

Proof. The construction. By the remark in the introduction, it suffices to provide a component with $s+d-1$ inputs and s outputs. Construct such a component by joining output y_j to inputs $x_j, \ldots, x_{j+d-1}$. This component has the (s,d)-matching property, because the matching $M=\{x_i y_i\colon 1 \leq i \leq d\}$ can be altered iteratively to satisfy any set S of at most s inputs as follows: For $i=1,\ldots,d-1$, if $x_i \notin S$ and x_i is the $t+1$th input missed by S, alter the current M by replacing $\{x_j y_{j-t}\colon i \leq j \leq d+t\}$ by $\{x_{j+1} y_{j-t}\colon i \leq j \leq d+t\}$.

Optimality for $s \leq 2$, i.e. $\theta(n,2,d) \geq 2/(d+1)$ (obvious for $s=1$). Consider a component G of an optimal graph on n inputs. G has k inputs and $m=\theta k$ outputs. Let t be the number of inputs in G having degree 1; call them *leaves*. By the pigeonhole principle, $m \geq t$, else there is an unsatisfied pair of leaves. This means showing $t \geq 2k/(d+1)$ will complete the proof. If $t < 2k/(d+1)$, we get a contradiction by counting edges of G. Counting them by inputs yields at least $t+2(k-t)=2k-t>2kd/(d+1)$. Counting them by outputs yields at most $dm=d\theta k$. Together these imply $\theta \geq 2/(d+1)$.

Optimality for $d=2$, i.e. $\theta(n,s,2) \geq s/(s+1)$ (obvious for $d=1$). Again, consider an optimal component G with k inputs and $m=\theta k$ outputs. We may assume $k > s$, else $\theta \geq 1$. Again we count edges. Since the component is connected, there must be at least $m+k-1$, but there are at most $2m$. Hence $m \geq k-1$, or $\theta \geq (k-1)/k \geq s/(s+1)$. □

Now we suppose d is somewhat larger, but we prevent s from growing arbitrarily in relation to d, i.e. $s \leq d^r$. For this case we provide only a construction. Ignoring the ceiling brackets may make it look essentially optimal, since the degree restriction places a trivial lower bound of $1/d$ on $\theta(s,d)$ and $s^{1/r}$ can be made arbitrarily small by letting r grow, but $\lceil s^{1/r}/r \rceil$ cannot be less than 1, and there is certainly no further improvement after reaching that point.

THEOREM 2. For $s \leq d^r$, $\theta(d^r,s,d) \leq r\lceil s^{1/r}/r \rceil /d$.

Proof. We construct a component using d^r inputs and $rd\lceil s^{1/r}/r \rceil$ outputs. Arrange the inputs in an r-dimensional array with side d. Assign an output to the d inputs in each 1-dimensional row of the array; i.e. fixing $r-1$ coordinates. We use $\lceil s^{1/r}/r \rceil$ copies of this set of rd^{r-1} outputs. The outputs all have degree d; we need only verify that every set of at most s inputs is satisfied.

Each input neighbors r outputs in each copy. Given a set of $k \leq s$ inputs, suppose its positions in the array meet $p = \Sigma p_i$ outputs in each copy, p_i in the ith direction. Then it has $p\lceil s^{1/r}/r \rceil$ output neighbors. For a given k, p is minimized when the positions are arranged in an r-dimensional cube of side $k^{(r-1)/r}$. In particular, $p \geq rk^{(r-1)/r}$. This implies $p\lceil s^{1/r}/r \rceil \geq k$, as desired. $\square$

Remarks. 1) We can do a bit better by not using identical copies of the outputs, but rather distributing them differently over the coordinates in each copy. For example, in the jth copy, let the input sets covered by an output be a point x and all translates of $x \bmod d$ by $(t_{1,j}, \ldots, t_{r,j})$, where t_{ij} is 0 or relatively prime to d. By appropriate choice of the t_{ij}, an s-set configured to hit the minimum number of outputs in one copy will hit more than the minimum in other copies. In other words, fewer copies will be needed. Similarly, we can shave a little off the outputs in each copy when $s^{1/r}/r$ is just a little more than an integer.

2) The fact that this construction is valid for an arbitrary number of dimensions suggests that it may be possible to improve on the *expander graphs* of [6,7] by generalizing their constructions to higher dimensions. Also, the expander graphs themselves may yield better constructions for the graphs studied here.

Finally, since this construction only yields $\theta(s,d) \leq \log d/d$ for large d and fixed s, we present another that yields about $(s-1)/d$. As noted in the introduction, we get $\theta(s,d) \leq (s-1)/q$, where q is the largest special number less than or equal to d.

THEOREM 3. If $d = (s-1)\binom{m-1}{s-2}$, then $\theta(s,d) \leq (s-1)/d$.

Proof. We construct a component with m outputs and $n=(s-1)\binom{m}{s-1}$ inputs. Partition the inputs into $\binom{m}{s-1}$ disjoint classes of $s-1$ inputs each. Put these classes in 1-1 correspondence with the subsets of size $s-1$ of the outputs. For each output set of size $s-1$, join each of its vertices to each of the vertices in the corresponding class of $s-1$ inputs.

Any set of inputs in a single class has size at most $s-1$ and has $s-1$ output neighbors. On the other hand, any set of inputs intersecting more than one class has at least two distinct sets of $s-1$ neighbors, hence at least s neighbors altogether. Furthermore, since each output belongs to $\binom{m-1}{s-2}$ sets of size $s-1$, each output has degree $(s-1)\binom{m-1}{s-2}$. Finally, $m/n=m/(s-1)\binom{m}{s-1}=m/m\binom{m-1}{s-2}=(s-1)/d$.

3. Integer programming formulation

The constraints and objective of the (s,d)-matching problem can be formulated as a 0,1-integer programming problem with linear constraints. It is a "multi-covering" problem; the dual is a weighted packing problem.

To define the program, note that we may assume all the outputs have degree d, since extra edges thrown in cannot hurt. In running the programs, we have considered only the class of graphs where no output nodes share the same neighborhoods. We have not proved that the optimal configuration is of this type. We associate with each potential output one of the d-subsets of the n inputs. We want to choose the minimum number of these that will satisfy the constraints.

The constraint matrix is the incidence matrix between the d-subsets of the inputs (corresponding to possible output neighborhoods) and the subsets of size at most s of the inputs. The constraint for a set of size k is that it intersect at least k of the chosen output neighborhoods. Hence the constraints look like $\Sigma_i x_i a_{i,S} \geq |S|$, and the objective function is $\min \Sigma x_i$.

In the table below, we show the optimal number of outputs for some small values of n, with s and d varying from 3 to 5. These are the results of an integer linear program run using a Multi-Purpose Optimization System Package [5] running under the NOS operating system on a CDC CYBER-175 computer. Larger problems than the ones listed could not be run due to limited memory space. The columns are indexed by the number of inputs n, and the entries give the minimum m and the minimum output-input ratio m/n.

Optimum number of outputs and values of $\theta(n,s,d)$.

$n=$	5	6	7	8	9	10
$s=3$, $d=3$	3 .600	4 .677	4 .571	5 .625	5 .556	6 .600
$s=3$, $d=4$	3 .600	3 .500	4 .571	4 .500	5 .556	
$s=3$, $d=5$	3 .600	3 .500	3 .429	4 .500		
$s=4$, $d=3$	4 .800	4 .667	5 .714	5 .625		
$s=4$, $d=4$	4 .800	4 .667	4 .571	5 .625		
$s=4$, $d=5$	4 .800	4 .667	4 .571	4 .500		
$s=5$, $d=3$	5 1.000	5 .883	5 .714	6 .750		
$s=5$, $d=4$	5 1.000	5 .883	5 .714	5 .625		
$s=5$, $d=4$	5 1.000	5 .883	5 .714			

Acknowledgement

The authors would like to thank Mr. Prasoon Tiwari for valuable discussions of the problem.

References

[1] N. Alon, Eigenvalues, geometric expanders, sorting in rounds, and Ramsey theory, preprint. Also Expanders, sorting in rounds, and superconcentrators of limited depth, *Proc. 17th Ann. ACM Symp. Th. Comp.* 1985.

[2] N. Alon and V. D. Milman, Eigenvalues, expanders, and superconcentrators. *25th IEEE Conf. Foundations Comp. Sci.* (1984), pp. 320-322.

[3] L. Bassalygo and M. Pinsker, Complexity of an optimum nonblocking switching network without reconnections, *Problemy Peredachi Informatsii*, Vol. 9, no. 1, 1973, pp. 84-87, translated into English in *Problems in Information Transmission,"* Vol. 9, no. 1, 1974, pp. 64-66.

[4] F. R. K. Chung, On concentrators, superconcentrators, generalizers, and nonblocking networks, *The Bell System Technical Journal*, Vol. 58, No. 8, Oct. 1979, pp. 1765-1777.

[5] C. Cohen and J. Stein, *Multi-purpose optimization system, user's guide*, Vogelback Computing Center, Northwestern University, Evanston, IL, 1978.

[6] O. Gabber and Z. Galil, Explicit constructions of linear size superconcentrators, *Proc 20th IEEE Conf. Foundations of Comp. Sci.*, 1979, pp. 364-370, and *J. Comp. Sys. Sci.* 22(1981), 407-420.

[7] G. A. Margulis, Explicit construction of concentrators, *Problemy Peredachi Informatsii*, Vol. 9, no. 4, 1973, pp. 79-80, translated into English in *Problems of Information Transmission*, Vol. 9, no. 4, 1973, pp. 325-332.

[8] M. S. Pinsker, On the complexity of a concentrator, *Proc. 7th International Teletraffic Conference*, Stockholm, 1973, pp. 318/1-318/4.

[9] N. Pippenger, Superconcentrators, *SIAM Journal of Computing*, Vol. 6, no. 2, 1977, pp. 298-304.

[10] N. Pippenger, Geaeralized connectors, *SIAM Journal of Computing*, Vol. 7, no. 4, 1978, pp. 510-514.

[11] R. M. Tanner, Explicit concentrators from generalized N-gons, *SIAM Journal of Algebraic and Discrete Methods*, Vol. 5, no. 3, 1984, pp. 287-293.

[12] L. G. Valiant, On nonlinear lower bounds in computational complexity, *Proc. 7th Annual ACM Symposium on Theory of Computing*, Albuquerque, New Mexico, 1975, pp. 45-53.