Satellite Communications and Multigraph Edge-Coloring

Dwight R. Bean[1]

and

Gernot M. Engel[2]

Abstract

An overview of satellite communication scheduling and its relation to edge-coloring of multigraphs is given. Then a theorem about a restricted class of multigraphs is proved to obtain conditions for scheduling in a satellite communications network of practical interest. Some limitations of the multigraph model are then discussed.

1 Introduction

In a common paradigm for a satellite communications network, one designated modem serves as network controller. Other modems then request communication services of various rates from the controller and the controller assigns to these modems time slots in which to burst their communications. To avoid scheduling conflicts time is divided into frames which are further subdivided into addressable slots of time. Typically, communications are buffered and then burst at a much higher rate than the requested rate so that the frame can be utilized as efficiently as possible. Furthermore, multiple carrier frequencies are used so that different modems can channel their communications through the satellite at the same time.

The underlying problem of the controller is to figure out which time slots and frequencies to assign modems to maximize the throughput of the network. For example, in [FSCS 89] the controller was required to schedule up to 900 requests per hour of 90 second average duration with start-up delay less than 9 seconds. This performance specification was written in 1989. Dramatic increases in demand have occurred since and will likely continue. To insure that the network is responsive to change, efficient scheduling algorithms are very important ([Acampora 78, Inukai 79]). However, results which predict the existence of a schedule are also very important. Requests may have assigned priorities. It is useful to have so-called "***doorman***" conditions to allow just the right number of requests to be handed to the scheduler so that scheduling heuristics can be invoked without further regard to priority.

For example, one obvious doorman condition is that total time slots requested not exceed the number of time slots in a frame times the number of available carrier frequencies. We will refer to this constraint below as the "area constraint". Another condition is that no modem may request to transmit or receive for more time slots than there are in a frame, to avoid simultaneous transmission or reception on two different carrier frequencies by the same modem. We refer to this below as the "modem constraint".

[1]University of San Diego

[2]Titan-Linkabit, Inc., San Diego

As in many other scheduling problems, chromatic graph theory provides an excellent model for solving satellite scheduling problems. Modems are represented by vertices. Service requests between modems are represented by edges connecting the vertices. Multiple edges are used, one per time slot requested. Thus, the underlying model is a ***multigraph***--a graph with multiple edges (but no edge connecting a vertex to itself in this application). A schedule corresponds to an ***edge-coloring*** of the multigraph. Each color represents a time slot. Edges sharing a vertex must be colored differently so that a modem will not be required to simultaneously participate in two different requests. A schedule will exist if the number of colors used does not exceed the number of time slots in the frame and if the number of edges colored the same color never exceeds the number of carrier frequencies allowed.

In general, modems may be ***full-duplex*** or ***half-duplex***, which means they may or may not be able to transmit and receive at the same time. The model described above may be applied to full-duplex modems if two vertices are introduced per modem, a transmit and a receive vertex. A network of full-duplex modems then corresponds to a ***bipartite*** multigraph, all edges connecting transmit vertices to receive vertices. Although the general problem of characterizing multigraphs having specified edge-colorings has not been solved, the case of bipartite multigraphs is well-understood. In fact, the area constraint and modem constraint are (necessary and) sufficient doorman conditions for a schedule to exist in the bipartite case. Various authors [Dulmage 69, Acampora 78, Inukai 79] have independently derived these conditions. [Acampora 78] and [Inukai 79] have also provided algorithms to schedule requests in satellite communications networks with only full-duplex modems.

General doorman conditions have not been found, but there are important special cases amenable to solution. A modem may be significantly restricted in the number of requests it can participate in, depending on the number of *ports* of the modem. This type of restriction has been exploited below to obtain doorman conditions for scheduling requests in a satellite communications network of practical significance.

It turns out that the general modem-to-modem satellite communications scheduling problem is purely a problem of minimizing the number of colors used to edge-color a multigraph. As long as the area constraint has been met, there is no need to worry about restricting the number of occurrences of a given color to the maximum number of available carrier frequencies. The reason, apparently first observed independently by [de Werra 71] and [McDiarmid 72], is that any edge-coloring of a multigraph using n colors can be transformed into an n-coloring which is as balanced as possible. If, for example, color α occurs at least two more times than color β, consider the subgraph determined by the edges colored with α and β. This subgraph must be a disjoint collection of chains and even-length cycles of edges alternating in the colors α and β, and so there must be an odd length chain in which the color α appears one more time than β does. Reversing the colors on this chain will bring the coloring more into balance, and iterating this process until no color appears two or more times than another yields a balanced coloring.

The first significant results on edge-coloring multigraphs were made by [Shannon 49] and [Vizing 65]. Shannon's result is particularly easy to implement as a doorman condition. Let the **degree** $\rho(\nu)$ of a vertex ν in a multigraph be the number of edges containing ν. Let $\lfloor x \rfloor$ denote the greatest integer not greater than x and $\lceil x \rceil$ denote the least integer not less than x. Shannon's result states that if the maximum degree of any vertex is ρ then the multigraph can be edge-colored using at most $\lfloor 3\rho/2 \rfloor$ colors. Thus, sufficient doorman conditions to schedule any modem-to-modem network requests in a frame with n time slots and c carrier frequencies are as follows:

- area constraint: the total time slots requested does not exceed cn
- Shannon's constraint: no full-duplex transmitter, full-duplex receiver or half-duplex modem is allowed more than $\lfloor (2n+1)/3 \rfloor$ time slots

The example of three half-duplex modems, each communicating with the other two for one-third of the frame, shows that Shannon's condition is sharp. On the other hand, this restriction would rule out the quite feasible case of a pair of modems just communicating with each other for the whole frame.

Necessary and sufficient doorman conditions in the general case would require a characterization of multigraphs admitting n-colorings of edges, which does not presently exist. [Goldberg 84] made an interesting conjecture which, if true, would provide this characterization for all practical purposes. Let e(H) denote the number of edges in a subgraph H of multigraph G, and let v(H) denote the number of vertices. Since at most $\lfloor v(H)/2 \rfloor$ disjoint edges can be chosen from H, at least $\lceil e(H) / \lfloor v(H)/2 \rfloor \rceil$ colors are required. Denote by w(G) the maximum value of this expression for any subgraph H of G and let ρ be the maximum vertex degree. Goldberg conjectured that if the actual number of colors, n, required to edge-color G exceeds $w(G) + 1$ then $n = \rho$ and if n exceeds $\rho + 1$ then $n = w(G)$.

Thus (practically) necessary and sufficient conditions to schedule in n time slots using c carriers would be

- area constraint: total times slots $\leq$ cn
- modem constraint: total time slots requests for full-duplex transmitter or receiver or half-duplex modem $\leq$ n
- Goldberg's constraint: For any group of k modems in the network, if t is the total number of time slots of inter-requests among the k members, then $\lceil t / \lfloor k/2 \rfloor \rceil < n - 1$

provided that Goldberg's conjecture is true.

Of course, computing w(G) might take too long to be utilized by a responsive doorman.

A particular satellite communications network of practical interest is the ***star network***. The controller in such a network is a sophisticated full-duplex modem with a very large number of ports. The other members of the network are half-duplex modems, which, apart from their communications with the controller, have only one port for access to other members of the network. The star topology arises when weight, power, and cost constraints limit the capability of network members to supporting only one service at a time using a radio that is only capable of half-duplex operation. Manpack users who must carry their own batteries, antenna, and radio will choose to limit functionality in exchange for reduced weight and size. Manufacturers of mobile radio equipment will select reduced functionality in order to minimize costs to ensure there is a large market for their product ([Mobile 89]).

We will use the term ***star network*** also to refer to the associated multigraph. In the following section we characterize n-colorings of a star network, thus obtaining necessary and sufficient doorman conditions for scheduling a star network.

2 Star Network Theorem

Two vertices are ***neighbors*** if there is at least one edge connecting them. A ***subgraph*** of a given multigraph consists of a subset of the vertex set together with all edges of the original multigraph connecting pairs of vertices in the subset. The ***underlying graph*** of a given multigraph is obtained by replacing each non-empty set of multiple edges by a single edge. The diagrams of multigraphs shown below will actually show just the corresponding underlying graphs. A ***star network*** is a multigraph having two designated non-neighboring vertices a and b. Every other vertex is restricted to have at most three neighbors, and at most one neighbor other than a and b, referred to as its ***mate***. (a corresponds to the network controller transmitter and b corresponds to its receiver.) A vertex without a mate will be referred to as ***mateless***. Figure 1 shows (the underlying graph of) an example of a star network.

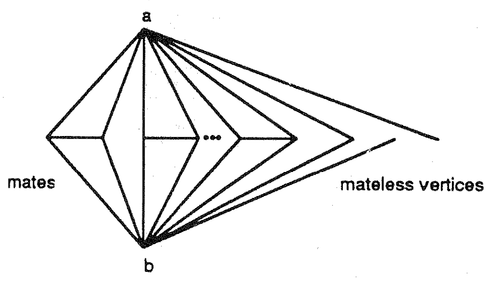

Figure 1. Star Network Example (Underlying Graph)

Remark: An important special case of a star network in which modems only communicate via the network controller is easy to analyze since the corresponding multigraph is bipartite: all edges join a and b to the other vertices of the multigraph.

Let G be a star network. A ***triangle*** is a subgraph of G with vertices a, c_1 and c_2 or b, c_1 and c_2, where each of the three vertices is a neighbor of the other two. See figure 2. A ***diamond*** is a subgraph of G with vertices a and b and two other vertices c_1 and c_2 which are mates, such that at least edges $\{a,c_1\}$ and $\{b,c_2\}$

or $\{a,c_2\}$ and $\{b,c_1\}$ are also present. See figure 3. A ***special*** subgraph of G consists of k diamonds, for some k>0, an additional vertex (referred to as the ***odd vertex***) which is a neighbor of a and b, and the edges connecting a and b to the odd vertex. See figure 4.

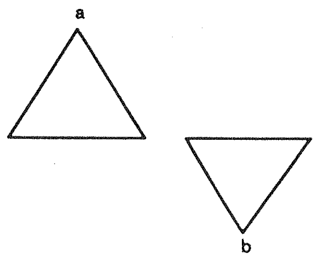

Figure 2. Triangles

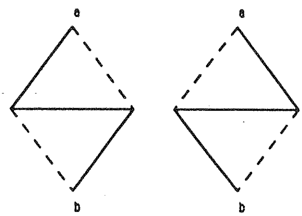

Figure 3. Diamonds

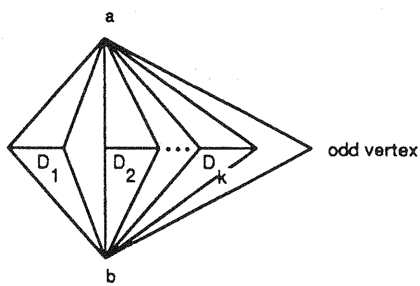

Figure 4. Special Subgraph with k Diamonds

Theorem. A star network G can be edge-colored with n colors iff $\rho(v) \leq n$ for every vertex v of G, $e(T) \leq n$ for every triangle T of G and $e(S) \leq (k+1)n$ for every special subgraph S of G with k diamonds.

Proof: The vertex and triangle conditions are obviously necessary. The conditions on a special graph with k diamonds are necessary since it contains exactly 2k+3 vertices, so at most k+1 edges can be colored with the same color.

A vertex v is ***unavoidable*** if $\rho(v) = n$. A triangle T is ***unavoidable*** if $e(T)=n$. A special subgraph S with k diamonds is ***unavoidable*** if $e(S) = (k+1)n$.

To establish sufficiency we will show:

(*) There is a set E of disjoint edges of G such that for every unavoidable vertex v there is an edge e of E such that v is in e, and E contains an edge of each diamond and of each unavoidable triangle, and E contains k+1 edges of each unavoidable special subgraph with k diamonds.

This suffices because one color can be assigned to this disjoint set of edges, and the remaining graph will satisfy the conditions of the theorem with n replaced by n-1. (For the base case, when n = 1, a graph can obviously be edge-colored with one color if $\rho(v) \leq 1$ for each vertex v).

Note that if v is unavoidable and has more than one neighbor w, then if w is unavoidable it also has another neighbor. This fact will be used a number of times below.

Since $\rho(a)$, $\rho(b) \leq$ n there can be no more than two unavoidable mateless vertices in G. Therefore, we need consider only three cases, depending upon how many unavoidable mateless vertices are present. We will assume a and b are unavoidable since it is always possible to make them unavoidable by introducing new neighbors for a and b without violating the conditions on G. We will also assume G is connected without loss of generality.

<u>case 0: There are no unavoidable mateless vertices.</u>

Here we consider separately the cases where there is or is not an unavoidable special subgraph.

<u>subcase 0: There is no unavoidable special subgraph.</u> Choose edges connecting the mates of G. It is easy to check that all conditions (*) are met except that a and b are not members of any of the edges.

If there is any diamond D in G then replace the edge connecting the mates in D with a guaranteed pair of disjoint edges in D and (*) is satisfied.

Otherwise, since a and b are unavoidable, there must exist distinct vertices v_1 and u_1 and edges $\{a,v_1\}$ and $\{b,u_1\}$. v_1 and u_1 cannot be mates since there is no diamond, and neither can have an unavoidable mate since that would yield a triangle with more than n edges. Therefore, all conditions (*) are satisfied by augmenting the chosen edges with $\{a,v_1\}$ and $\{b,u_1\}$, replacing previous edges chosen including v_1 or u_1 (if they have mates).

<u>subcase 1: There is an unavoidable special subgraph.</u> Since there is no unavoidable mateless vertex, there must be a diamond D such that $e(D)>n$. Every unavoidable special subgraph must contain D as a

diamond or one of D's mates as an odd vertex. (Otherwise, there is an unavoidable, special subgraph edge-disjoint from D with k diamonds, and augmenting that subgraph with D yields a special subgraph S with k+1 diamonds such that e(S)>(k+2)n.) In either case, including an edge for each pair of mates not in D and including a guaranteed pair of disjoint edges from D will give a set of disjoint edges satisfying (*).

case 1. There is exactly one unavoidable mateless vertex v.

Note here that v must be the odd vertex of any unavoidable special subgraph, and furthermore, that for any diamond D, e(D) $\leq$ n, which in turn implies that the mates of D are not unavoidable and there is no unavoidable triangle in D. Without loss of generality assume a is a neighbor of v. Then choose an edge for each pair of mates in G and choose $\{a,v\}$ and all conditions of (*) will be satisfied except that b is not a member of any chosen edge. Since b is unavoidable it neighbors some vertex $v_1 \neq v$. If v_1 is mateless then add $\{b,v_1\}$ and (*) is satisfied. Otherwise, if v_2 is the mate of v_1, replace $\{v_1,v_2\}$ by $\{b,v_1\}$ and (*) is satisfied. (v_2 cannot be unavoidable, nor can $\{a,v_1,v_2\}$ determine an unavoidable triangle: If $\{a,v_1,v_2\}$ formed a triangle v_1 and v_2 would be in a diamond. If not then an unavoidable v_2 means edge $\{b,v_2\}$ must exist, and then the fact that the sum of the edges of triangle $\{b,v_1,v_2\}$ $\leq$ n is violated.)

case 2. There are exactly two unavoidable mateless vertices u and v.

Assume, without loss of generality, that a and u are neighbors. Then it is easy to check that b and v are neighbors. Choose $\{a,u\}$ and $\{b,v\}$ to satisfy (*). (The only other possible edges in the underlying graph of G are $\{a,v\}$ and $\{b,u\}$ since $\rho(a) = \rho(b) = $ n.)

Since there cannot be more than two unavoidable mateless vertices, the proof is complete.

Figure 5 shows a star network (multiple edges labeled by numbers to indicate their multiplicity) which is not edge-colorable with n colors even though $\rho(v) \leq$ n for every vertex v, e(T) $\leq$ n for every triangle T and e(S) $\leq$ (m+1)n for every special subgraph S with m diamonds, for m ***strictly less*** than k, the number of diamonds in G. This example indicates that the technical conditions of the theorem are probably warranted.

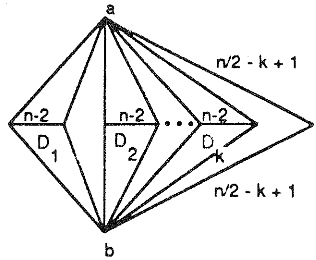

Figure 5. A Star Network Not Edge-Colorable Using n Colors

3 Doorman and Scheduler Complexity

Let m be the number of modems in the network and s be the number of service requests. Recall that there are n time slots to schedule and c carrier frequencies. Then a doorman exists with complexity $O(\max\{n,m,s\})$ and a scheduler exists with complexity $O(cn)$.

The bottleneck to an efficient doorman is checking the special subgraph condition, but by reasoning as in case 0, subcase 1 of the proof in the previous section we can achieve the efficiency claimed. The idea is to only worry about the worst case special subgraph with a minimum number of diamonds, which involves precisely the diamonds with more than n edges and possibly a mateless vertex of maximum degree. We begin with the subgraph only including the edges including ***a*** or ***b***. We first partition the vertices other than ***a*** and ***b*** so that $V[i]=\{\text{mateless } v\text{: } \rho(v)=i\}$, $i \le n$. (All vertices start out being mateless so only communications with the controller are considered thus far). Then service requests are considered in priority order. Each new request causes the two vertices involved to leave the mateless partition V. If a diamond D is formed by the new request, and if $e(D) > n$, then the excess $e(D)-n$ is accumulated in E and r, the minimum $\rho(v)$, for v in such D's, is maintained. If R is the largest index such that V[R] is non-empty, it is easy to check that the special subgraph condition still holds with the new service included iff $E + R \le n$ and $E - (e(D)-n) + (e(D)-r) \le n$ (the latter condition corresponds to the case where the odd vertex is chosen to be the mate of the v with $\rho(v)=r$ in a D with $e(D)>n$, and simplifies to the condition $E \le r$).

The fast scheduler claimed to exist can be achieved by also partitioning the diamonds D so that $D[i]=\{\text{diamonds D: } e(D)=i\}$, $i \le 2n-1$. Then at any stage k from 0 to n-1 (corresponding to time slots in the frame or edge-colors of the graph) the construction given in the proof of the theorem can be used to produce a matching in $O(c)$ steps. The case to consider in the proof is given by the size of V[n-k] and in the non-trivial case 0, subcase 1, a disjoint pair of edges can be selected from a diamond in D[i] where i is largest such that D[i] is non-empty. Updates to the D and V arrays can be done in constant time for each of the at most c edges chosen in the matching.

4 Limitations of the Model

In making our claim that edge-colorings of multigraphs provide a good model for scheduling modem-to-modem satellite communications, we neglected a technical signal processing problem of locating the exact starting bit of a data stream. A common solution to this problem is to include a known preamble pattern at the beginning of each data burst for signal acquisition purposes. However, the overhead of such preambles can be considerable if communication between pairs of modems is broken up into multiple bursts. If we consider the correspondence of edge colors and consecutive time slot numbers, what we would really like to achieve is an ***interval*** edge-coloring, where multiple edges receive consecutive colors.

This extra constraint is sometimes impossible to meet. For example, it is easy to check that a star network with no controller communications and with c pairs of mates, each communicating for n-1 time slots, will require at least c-1 extra bursts (c-1 "***interrupts***") using c carriers and n time slots per carrier. [Schneider ar] has shown that there are examples of networks requiring on the order of c^2 interrupts on c carriers.

Remark: The problem of minimizing the number of interrupts in a schedule is NP-hard since it can be used to solve the bin packing problem: Given c bins of capacity n and k packages of sizes $s_1, s_2, \ldots, s_k$, define a star network with k pairs of mates, where the ith pair communicates for s_i time slots. Then the communication can be scheduled on c carriers of n time slots each with a minimum of zero interrupts if and only if the packages may be packed into the c bins.

Preambles are a particular nuisance in controller communications. One way around the preamble problem in this case is to allow the controller to ***broadcast*** to the other modems, i.e., put all transmissions into a single burst with one preamble and require all modems to simultaneously tune to the same carrier frequency to listen to this burst. General broadcasting would require a ***hypergraph*** model, which we won't discuss here. If we just limit broadcasting to the controller, we can divide the frame into the broadcast region (where only one carrier can be used) and the remaining region (in which there are no special subgraphs, so that scheduling is relatively easy.)

The penalty of only utilizing one carrier during the broadcast is severe if there are many frequencies available. A possible solution to this problem is to partition the modems into two or more groups and broadcast to each group separately. Then, during a broadcast to one group, communications among the modems not in the group could be scheduled on the other carriers.

References

[Acampora 78] Acampora, A. S. and Davis, B. R.
Efficient Utilization of Satellite Transponders via Time-Division Multibeam Scanning.
Bell System Technical Journal 57(8):2901-2914, October, 1978.

[de Werra 71] de Werra, D.
Equitable Colorations of Graphs.
Rev. Fran. Inf. Rech. Oper. 5:3-8, 1971.

[Dulmage 69] Dulmage, A. L. and Mendelsohn, N. S.
Some Graphical Properties of Matrices with Non-Negative Entries.
AEQ. MATH. 2:150-162, 1969.

[Fiorini 77] Fiorini, S. and Wilson, R.J.
Edge-Colourings of Graphs.
Pitman, London, 1977.

[FSCS 89] *UHF Demand Assignment Multiple Access (DAMA) System TD-1271B/U Multiplexer Program Performance Specification*
Department of Defense Index of Specifications and Standards FSCS-212-16, Revision D, 1989.

[Garey 79] Garey, M. R. and Johnson, D. S.
Computers and Intractability.
Freeman, San Francisco, 1979.

[Goldberg 84] Goldberg, M.K.
Edge-Coloring of Multigraphs: Recoloring Technique.
J. of Graph Theory 8:123-137, 1984.

[Inukai 79] Inukai, T.
An Efficient SS/TDMA Time Slot Assignment Algorithm.
IEEE Transactions on Communications 27(10):1449-1455, October, 1979.

[McDiarmid 72] McDiarmid, C.J.H.
The Solution of a Time-tabling Problem.
J. Inst. Maths. 9:23-34, 1972.

[Mobile 89] *The Communications Ground Segment to Provide Mobile Satellite Services in North America, Hawaii and Puerto Rico/Virgin Islands: Functional Requirements*
Draft edition, Telesat Mobile Inc. and American Mobile Satellite Corporation, 1989.

[Papadimitriou 82]
Papadimitriou, C. H. and Steiglitz, K.
Combinatorial Optimization.
Prentice-Hall, New Jersey, 1982.

[Schneider ar] Schneider, Hans and Li, Chi-Kwong.
A Note on the Time Slot Assignment Problem for TDMA Systems.
IEEE Transactions on Communications , To Appear.

[Shannon 49] Shannon, C.E.
A Theorem on Coloring the Lines of a Network.
J. Math. Phys. 28:148-151, 1949.

[Vizing 65] Vizing, V.G.
The Chromatic Class of a Multigraph (in Russian).
Cybernetics 3:32-41, 1965.