21st International Communications Satellite Systems Conference and Exhibit
AIAA 2003-2422

SLOT ALLOCATION IN A TDMA SATELLITE SYSTEM: SIMULATED ANNEALING APPROACH

T. Toker*, E. Altman*, J. Galtier†,*, C. Touati*, I. Buret‡, B. Fabre‡, C. Guiraud‡

We consider in this paper the uplink slot assignment problem in a multi-spot geostationary satellite. Radio interference impose constraints on the slots that can simultaneously be assigned in different cells that have the same frequency. The problem is shown to be an NP-complete one, which motivates us to search for a heuristic solution approach. We describe here a heuristic solution based on simulated annealing. We further investigate how to improve the performance of the simulated annealing and the rate of convergence of the annealing. Numerical experimentations are provided to test our proposed improvements.

1 Introduction

We consider a multi-spot geostationary satellite system for which a manager assigns satellite TDMA slots to service providers (operators) located at various zones in various cells. We consider uplink transmission. A slot cannot be assigned simultaneously to more than one zone in a cell. Cells of different colors (frequencies) do not interfere, but cells of the same color do, and a slot can be assigned to an operator in a given zone only if the interference it experiences is below a given threshold. Slot assignment is static but can be changed once per hour (due to changes in demands, on the one hand, and to changes in atmospheric conditions, on the other hand). Our goal is to maximize the goodput of the system (the throughput of useful information), taking into account the coding ratio of a slot that may vary from a zone to another. Casting the problem into coloring of graphs shows that it is NP-complete to maximize the goodput, and we therefore propose a heuristic approach based on simulated annealing[4,5] and analyze its performance.

The structure of the paper is as follows. In the next section we describe the slot allocation problem. In Section 3 we describe the simulated annealing algorithm and point out on some elements used to improve its performance. In Section 4 we describe other improvements in the system performance related to carrier allocation in the frequency domain and to a two-stage implementation of the simulated algorithm. We then present in Section 5 numerical results; we compare various implementations of the simulated annealing and compare our optimized version to other approaches. The paper ends with a conclusion section.

2 Problem Description

Our design of a multi-spot TDMA system consists of two phases. In the first radio planning phase, carriers are assigned to operators that are located in various zones of various spots. The carrier allocation takes into account the minimum throughput requirements of operators and the overhead due to coding rate which vary from one zone to another due to various atmospheric conditions. The allocation process aims at finding a fair solution with optimality properties. This phase was described in.[1] The solution is the input to the second phase of slot allocation which we present in this paper.

Since different carrier types have different throughputs, the duration of a slot varies with the c arrier type, which renders the slot allocation an even more complex problem.

A simplified version of our problem can be modeled as a so called "k-colorable induced subgraph" problem where one considers a graph G=(V,E) consisting of finitely many nodes and directional links. A valid coloring of the graph consists of coloring nodes such that no nodes with a common link have the same color. We look for a subset of nodes $V' \subset V$ and edges $E' \subset E$ such that the induced subgraph is k-colorable, i.e., there is a coloring for the subgraph (V', E') of cardinality at most k. The problem consists in finding such a graph with the maximal number of nodes. This problem turns to be NP-complete*. Our problem is in fact even more complex since both arcs and nodes have weights (nodes weights are related to the coding ratio and arc weights to the amount of interference), and exclusion constraints are more complex. This is in contrast to slot allocation in

*INRIA BP93, 06902 Sophia Antipolis Cedex, France
†France Telecom R&D, 06921 Sophia-Antipolis Cedex, France
‡ Alcatel Space Industries, Toulouse, France

*see http://www.nada.kth.se/~viggo/wwwcompendium/-node34.html and[6]

Satellite Switched TDMA systems that have polynomial solutions since they correspond to coloring problems with a simple bi-partite graph topology.[2] We therefore used a heuristic simulated annealing optimization approach.

Since a new allocation should be generated every hour, solutions are needed within a few minutes. We worked with eight spots, three zones per spot, four colors and six types of carriers. The total number of carriers of different types per spot was generated at random, and the average sum of number of carriers was around 100 per spot.

3 Efficient Simulated Annealing

Slots, chosen at random, are assigned sequentially. A slot can be assigned if the sum of the interference it receives is below a given threshold, and if with the extra interference that its assignment generates, the total interference experienced by any other already assigned transmission at the same slot remains below a given threshold.

The interference created by a slot assigned to zone i of a given spot l over a zone j of a spot k depends only on i, l and k. In other words, all zones of spot k experience similar interference from a slot assigned to some zone i of a spot l (note that the interference occurs at the receiving antennas in the satellite). We denote this interference $I(l, i, k)$. These interferences as well as the antenna gains $G(i, l)$, corresponding to zone i in a spot l, are given in a table (whose entries can be obtained by computing the link budgets) that is an entry to the simulated annealing program.

At each step of the simulated annealing,a random spot and zone, a random carrier, and a random slot are chosen. We call this choice a candidate assignment. If the slot is not yet assigned then we perform the candidate assignment if and only if the above interference conditions are satisfied.

Next, we consider the case that the chosen slot is already assigned to another zone of the spot. Then there are three possibilities:
P1. maintain the existing assignment,
P2. use our candidate assignment instead of the existing one,
P3. cancel the existing assignment without performing an alternative assignment of that slot.

To decide which of the three to use at each step, we introduce a performance measure $\mathcal{G}$, which is a weighted sum of the number of allocated slots, where the weight is related to the coding rate. The weighted sum corresponds to the goodput of the system. The decision is then done as follows.

- If with the candidate allocation, the value of $\mathcal{G}$ is strictly greater than under the existing allocation, then we use the candidate allocation (option P2), provided that the interference conditions are satisfied.
- If the value of $\mathcal{G}$ is the same under the existing and the candidate allocations, then we switch to the candidate allocation (option P2) if the interference conditions are satisfied and moreover, if it creates less total amount of interference than the existing assignment. Otherwise we maintain the existing assignment (option P1). In this step we thus follow a hierarchical optimization approach by introducing a second (lexicographic) criterion, of minimizing the created interference.
- If the $\mathcal{G}$ corresponding to the candidate allocation is strictly smaller than the existing one (but it is still feasible in terms of interference) then we maintain the current allocation (option P1) with high probability, denoted by P, which is a function of the difference $\Delta\mathcal{G}$ and of a so called temperature T parameter. P decreases to zero as $T \to 0$, and is given by

  $$P(T) = \exp\left(-\frac{|\Delta\mathcal{G}|}{T}\right). \tag{1}$$

  With probability $P/2$, however, the existing allocation is canceled (option P3); with option P1 is also chosen with probability $P/2$. These choices (that seem nonoptimal) allow to prevent the convergence to a local minimum). As the simulated annealing progresses, T is gradually decreased, so that the probability of choosing "nonoptimal actions" vanish.

The algorithm ends if either one of the following conditions hold:

- A given time limit is exceeded,
- Some bound on the number of iterations is exceeded,
- During a given period of n consecutive iterations, the gain $\mathcal{G}$ is unchanged, where n is some fixed threshold.

Next we describe the way to update T. We tested both linear and logarithmic rates for cooling T. The linear cooling has the form

$$T_{n+1} = ReductionFactor \times T_n \tag{2}$$

where $ReductionFactor$ is some constant, whereas the logarithmic cooling has the form

$$T_n = \frac{C}{\ln(n)}. \tag{3}$$

The linear rate of cooling is often used as it is much faster than the logarithmic. But asymptotic convergence to a global optimum is not guaranteed, where-as it is for logarithmic rate of cooling.[3]

We tried to improve the performance of the simulated annealing in various ways. The first improvement was by using the hierarchical optimization mentioned above. A second direction of improvement was in optimizing the choice of the convergence rate.

4 Carrier Allocation In The Frequency Plan

We further improve the system performance by some actions not directly related to the simulated annealing described in the previous section.

4.1 The Frequency Allocation

As input, we had been given the number of different types of carriers assigned to various zones, as well as their location in the frequency plan. This location was chosen in our tests randomly. We then tried to check whether the frequency location could be chosen in a way that would give a better performance. We did not try to include the carrier allocation in the frequency plan as part of our combinatorial optimization, since already the optimization of the slot allocation turned to be quite complex with a huge number of possible configurations. Nevertheless, we proposed an alternative way of allocating the carrier according to their bandwidth: in half of spots of a given color, we placed the carriers in an increasing order of bandwidth, and in the other half, in a decreasing order. This allows us to have less interference between carriers of large bandwidth.

4.2 A Two-Step Simulated Annealing

A second approach for improving the performance of the allocation is to perform a two-phase simulated annealing instead of one; each phase is allocated half of the time limit we initially had for the whole simulation. At the first step, the simulated annealing is restricted to assign slots to the carriers having the largest bandwidth only. Then the second simulated annealing step is performed for assigning the rest of the carriers.

In each phase we restart cooling the T parameter from scratch. For comparison purposes, in our experiments below we tested this two-step annealing also in a context in which it is not related to the size of the carriers. In other words, we tested also the "classical" simulated annealing in which the approach of Section 4.1 is not applied. In that case, the only difference between the classical simulated annealing is that after half of the dedicated time, we do not continue to cool T, but restart from a higher temperature using as an initial configuration the output of the first phase.

5 Results

Since a new allocation should be computed once per hour, this poses a constraint on the duration of the simulated annealing: we were limited by a total duration of 1min to perform this task. We therefore performed the annealing experimentations during 1min, 4min and 10min. All experiments were performed on a DELL INSPIRON 2500 computer with a Pentium III 1GHz processor and a 256 MB SDRAM. (We thus beleive that on a faster computer, the performance we obtained for 4min or 10min could be achieved). The annealing is programmed in a 1500 lines Java program.

We first ran a much simpler simulated annealing with no constraints on interference so as to obtain a reference upper bound on the performance that can be achieved. We then made many experimentations with simulated annealing, each time with another limit of allowed interference, and plot the obtained goodput normalized by the one obtained by the upper bound. The results are presented in Fig. 1, 2 and 3 which consider simulated annealing run during 1min, 4 min and 10min, respectively. The best performance in all three tests is obtained with our optimized simulated annealing. We compare it to four other schemes three of which have only one single simulated annealing part (consisting of optimizing all carriers type at once) some schemes use logarithmic and some use linear cooling etc.

We present various simulation results with different approaches and compare them to an "optimized approach": it uses the approaches of Sections 4.1 and 4.2, it uses the hierarchical optimization approach, and has an optimized logarithmic cooling rate.

Fig. 1-3 show that our optimized simulated annealing improves the solutions by around 50 % for small interference thresholds, e.g. -30 dB and 600 % for large interference thresholds e.g. $+50$ dB.

As for the duration of the annealing, we see that our optimized algorithm obtains an improvement of around 60% (for all the tested range of interference thresholds) when increasing the duration from 1min to 10min. At least for low thresholds, it is clear from the figures that we cannot hope to gain much more with even longer durations than 10min since we obtain with our algorithm around 96% of the reference upper bound for 10min. We note that for other implementations that we tested (not the optimized one), the improvement in performance when increasing the duration to 10min is even larger.

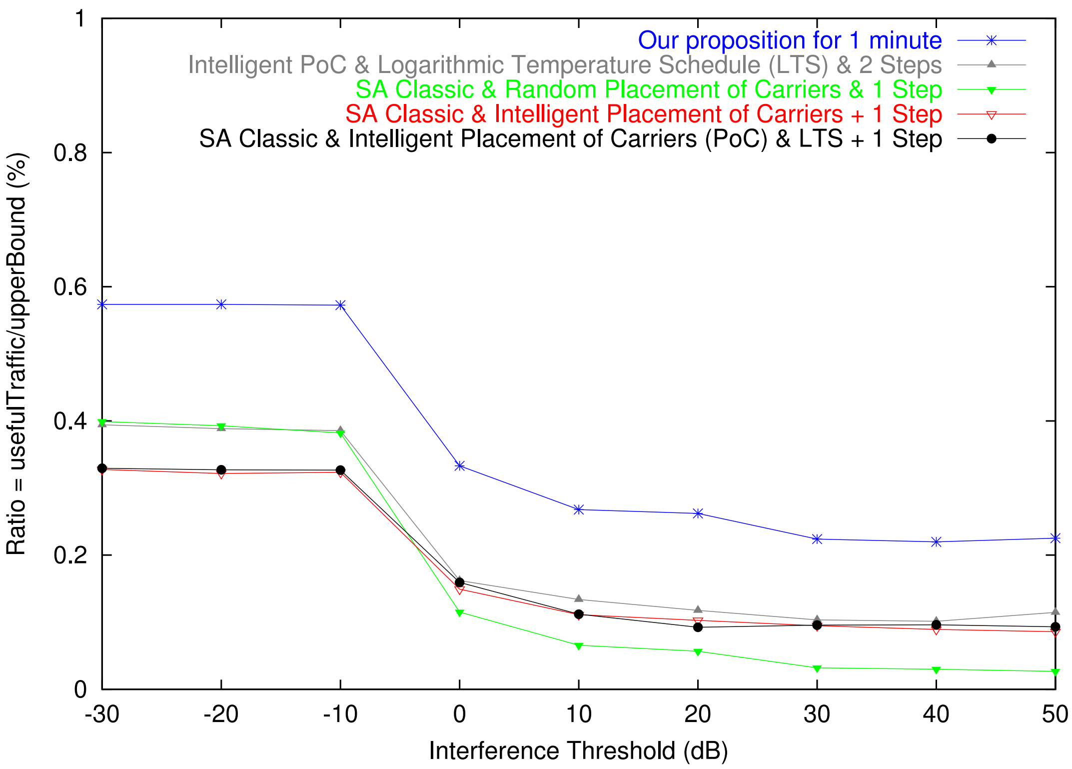

Fig. 1 Comparison for one minute

6 Conclusion

We have proposed in this paper a simulated annealing approach for the multi-beam geostationary uplink slot allocation problem. We proposed a hierarchical optimization approach that tries not only to maximize the throughput, but also prefers allocations that generate less interference. We optimized the simulated annealing with respect to the rate of convergence, and integrated some other system improvements concerning carrier allocations. These improvements were further implemented in a two-step simulated annealing approach which we compared numerically to other approaches. We showed that a substantial gain is obtained.

References

[1]E. Altman, J. Galtier, C. Touati, "Fair Bandwidth Allocation between Providers in a Geostationary Satellite Network", In the *Proceedings of the 21st International Communications Satellite Systems Conference* (AIAA-ICSSC), Yokohama, Japan, April 2003. Also INRIA Research Report RR. 4421, march 2002.

[2]S. Chalasani and A. Varma, "Efficient time-slot assignment algorithms for SS/TDMA systems with variable-bandwidth beams", *IEEE Transactions on Communications*, 42(2/3/4), pp. 1359-1370, Feb. 1994.

[3]B. Hajek. "Cooling schedules for optimal annealing", *Mathematics of Operations Research* Vol. 13, No. 2, May 1988.

[4]S. Kirkpatrick, C.D. Gelatt Jr., M.P. Vecchi, "Optimization by Simulated Annealing", *Science*, 220(4598), pp. 671-679, 1983.

[5]N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller. "Equation of State Calculations by Fast Computing Machines", *The Journal of Chemical Physics*, Vol. 21, No. 6, June 1953.

[6]B. Toft, "Coloring, stable sets and perfect graphs", Chapter 4 in of Graham, Grötschel, Lovász (eds.), Handbook of combinatorics, North Holland 1995 vol1., 233-288.

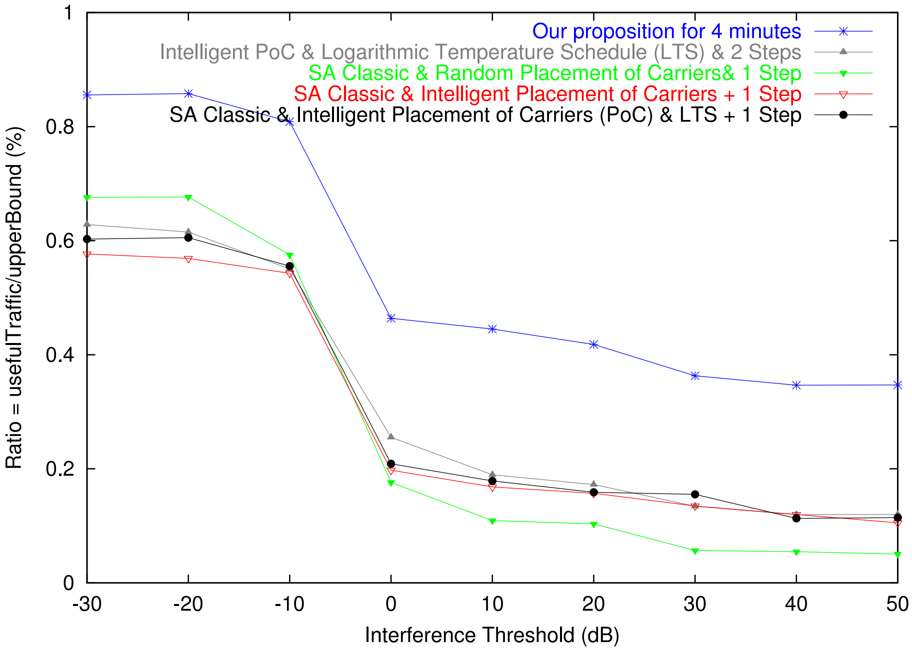

Fig. 2 Comparison for four minutes

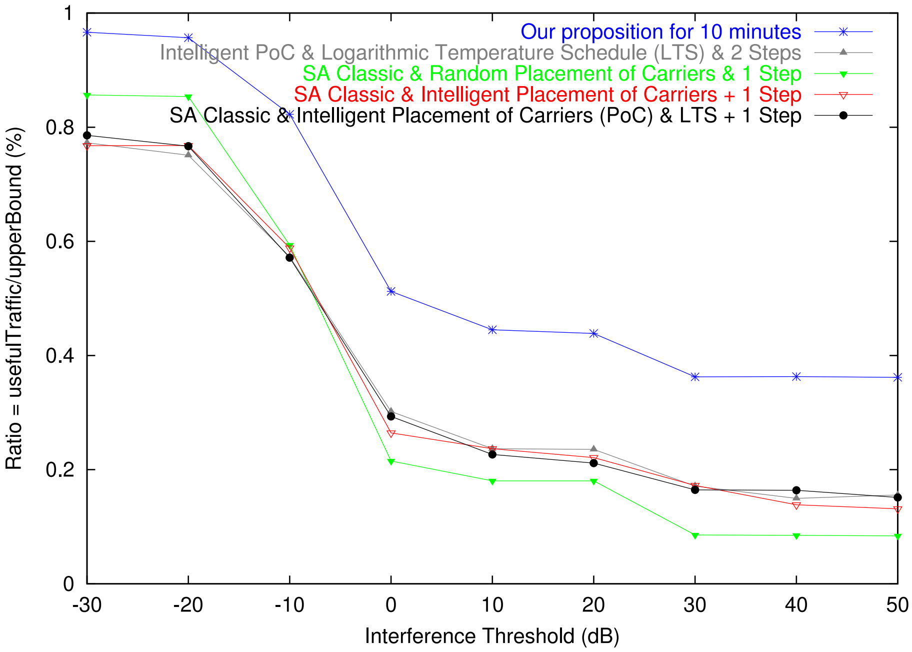

Fig. 3 Comparison for ten minutes