# Routing in the SITA Network

## Part 1: do you remember the best routes tables?

I joined SITA in June 1978 as an engineer to work in the Technical Planning Team then managed by Alfred Weyers, and part of the Technical Studies department with was headed by Geoges Giraudbit.

I had experience and skills in the planning of large public telecom network having worked in France Telecom's reseach arm CNET, but the SITA network was a completely new world for me, from a technical and human point of view. As I was eager to learn, this was a terrific experience for me. The working relationship with Alfred was quite pleasant as we had a lot of freedom and autonomy to do our job in our own way. I was lucky to have the opportunity to quickly have my hands on actual issues and problems faced by the network. The feeling of having a direct bearing on how the network was engineered, planned and operated by inventing and designing solutions to enable a smooth development of the network has brought me a lot of satisfaction.

On a distributed network like the HLN, the data packets have to follow a path from the origin HLS to the destination HLS and this path must adapt to events on the network such as a circuit or a node (HLS) failure and it is the role of the routing function to select the most appropriate path for each traffic flow from a given origin to a given destination. These routes were stored in tables in the HLSs, known as best routes tables and each HLS would select in its best routes tables the most apropriate route for each destination depending on the status of the network (i.e. on which circuits or other nodes were down). Then this HLS would switch the traffic, i.e. put the packets on the outgoing queue of the circuit corresponding to that route.

At the time I had joined SITA, coping with traffic growth was already an issue. Circuits (and modem technology) offered a limited capacity, usually 2400 bits per second (bps) or 4800 bps and at best 9600 bps. When a traffic flow for a given destination was approachig or even exceeding the circuit capacity, long transmission delays and dropping of traffic occured, thus resulting in poor data quality and complaints from the airlines. Annie Javerzac who was in charge of planning the topology (i.e, the use of circuits) of the HLN and ofproducing the best routes tables, and Alfred Wyers were worried. Hence had come the idea of split-routing : dividing the traffic flow from an origin node to a destination node in two distinct traffic flow, so that each of them would be smaller that the circuit capacity, and each would use a different circuit. There were discussions, but nobody was in charge of specifying this split-routing and nobody, in the Technical studies department, seemed to volunteer, a fact which surprised me. Alfred proposed me to be in charge of that project. I knew nothing about the routing in the HLN and about the best routes tables, but as I liked challenges and was unconscious of the complexity of the rules governing the best routes tables, I gladly accepted.

The switching function in the HLN was simple, as it was based on the destination HLS only, but this mechanism resulted in complex functional requirements for the best routes tables. Actually once I understood these logical constraints on the routing tables, designing the split-routing was not that difficult and within a month I issued the specification of the split-routing function to be implemented in the HLN.

The specification itself was rather short. Besides understanding the routing in the HLN, this initial work had another important benefit for me: in order to fully understand the problem, I did not only discuss with my colleagues from the planning team, but with the software development team, and in particular with Jean-Pierre Ladam, and the HLN operations team and ist manager, namely Gérard Gallot. This relationship with Jean-Pierre and with Gérard has always been excellent over the many years I worked on routing issues and very precious for my work, because I could make sure that the solutions would „fly“ from an HLN software point of view (of which I was not a specialist) and were meeting the operational requirements (we will see an example of such issues later in this article).

Although Jean-Pierre and Gérard were from a very different background from mine, the mutual respect and openness that we shared, was the foundation for this excellent relationship, which helped me a lot in winning the support of their departments when I was coming with new solutions or projects. Of course this not to say that my relations wiith other colleagues from these departments were not good, in many cases they were good indeed. However with my background and personality of a telecom engineer from a renowned engineering school, I was perceived by many from both the software development department and the operations department as an „egghead“, and I runned the risk that the solutions proposed would not be trusted and be dismissed on the ground that they were „missing the point“. And this has been a very enriching experience for me, and which I had not found in the previous companies I hadbeen with, to work in such a diverse environment.

Once the specifications for split-routing were accepted and then quickly implemented in the HLSs, the best routes tables were designed and loaded in the HLN, and it worked! I was really exhilirated by this success.

What it actually meant is that if the routing from e.g. LON to HKG was split, then the LON HLS would have two different best route tables to route the traffic to HKG (but in the opposite directin it could well be that HKG only had one best route table to reach LON). Hence there were now more routing tables to be prepared on a regular basis. On one side this was solving the issue, but as a drawback, Annie had now more work to do (i.e. more tables to prepare).

And the creation of the best route table was indeed becoming a burden, due to the constant increase in the traffic over the network. A project had been launched to automate the production of these best routes tables. A colleague from the technical planning team, Jacques van der Steen had already been working on this project. Jacques had developed a computer program to automatically generate the best route tables. Actually he had adapted an algorithm which had been developed by a well known guru in the area of modelling and designing packet-switched networks, Leonard Keinrock. Kleinrock had published a book in two volumes on the modelling of compter networks, the title of which was „Queing Systems“ ; this book was the bible for network modelling and computing network delays (later L Kleinrock also promoted himself as one of the father of packet-switching but this fact has been contested, in particular by the two recognised inventors of this technology, the Polish born American Paul Baran and the British Donald Davies, who published papers and lead experiments in 1965). Kleinrock had also proposed an algorithm for optimising the routing in distributed packet-switched networks, the „Extremal Flow (EF) algorithm which was well-known among specialists of network optimization. Jacques had a mathematical background and was also very clever in computer programming (whereas I was not at ease at writing computer programs). He had modified the EF

algorithm (to get rid of the so-called multiple-path routing that this algoprithm was generating but which was not possible on real networks like the SITA network). The program was working well and was generating best routes tables. Adapting this program to the new split-routing functionality was not an issue. Before this program could be used however , the results had also to be validated by the Operations department. And here came the bad suprise: the Operations team had found that the best routes tables generated by this program did not comply with the functional requirement that imposed that if all routes from HLS A to HLS B were down, then necessarily all routes from HLS B to HLS A had to be down also. In other words HLS B could not see HLS A as still reachable if HLS A saw HLS B as isolated. This requirement was necessary for the HLN protocol then in use, namely the P1000 protocol, to work properly in all situations. Actually this condition implied that the best routes table from A to B and the best routs table from B to A needed to include the same set of routes, although possibly in a different order. This seems to be a rather trivial condition to meet, but because of other functional requirements on these best routes tables, it was actually a difficult condition to meet. And the routing tables generated by the program based on the EF algorithm, were not meeting this condition. As a result, the use of this program to automatically generate the best routes tables had to be dropped, crashing the then strong hope of the technical team of an immediate relief by having the production of these damned routes tables automated. The team was shocked but we had to accept this outcome.

My initial work on split routing had helped me understand the issue at stake. Alfred asked me to work on it. I could understand the problem but finding a solution was another matter! Besides, the work relationship with Jacques was far from perfect. We had different personnalities. He was an extravert , talking easily with whoever, whereas I was rather an introvert, and in addition there was some competition between us as weh ad similar responsibilities and I had joined the team latter. Alfred was embarassed by the tension created by this competition and was trying to soothe the relationship, but on the other hand Georges thought that this was just sound competition and should stay like that. On my side, starting to work on the automated production of the routing tables, I was not feeling at ease.

The technical problem at stake was difficult. No known method seemed to meet the conditions that were imposed on the routing tables (this was indeed a SITA specific problem). I was investigating and discussing with the people who had worked on the routing tables (I remember a discussion with Chedly Redjeb who explained to me how he was proceeeding to create these tables when the responsibility to produce them was still with the operations department), but this was not providing any clue for an automated solution. Actually producing these tables by hand was a heavy and difficult task which involved many attempts before being able to complete the work. But still at the end of the day we could do it. However I was not finding a way to replicate that work with an automated program, and I was not aware of any published algorithm which could do the trick. I was puzzled and feeling disheartened. I remember one day I went to share my worries with Giovanni Strigari who was then deputy Director of the Technical Studies department, and he straightforwardly answered that “if we could produce these tables by hand, then there should be no difficulty in writing a program which would just be replicating what we were doing by ourselves”. I did not expect such an answer from Giovanni who was (and is still, I believe, although I have not met him since long!) a very nice person and I vaguely felt that there was a trick in this answer but I could not clearly see what the trick was. Now, with the hindsight of nearly forty years of experience the use of computer programs, I see clearly where the trick was.

Actually, automating the production of the best route tables was an Articial Intelligence (AI) problem. And we all know now that automating some tasks which seem to be governed by relatively simple rules may be dauntingly difficult. For instance, it took more than fifty years and a huge increase in computer power and immense progress in computer software, as well as a lot of work, to design a program which could beat human champion at the Go game, although the rules of this game are relatively simple. And the production of best routes tables was definitely an artificial intelligence problem because no straightforward algorithm could deliver the desired result. One of the difficulty of AI is that there are a lot of things that we do, but actually we do not know how we are doing them. For instance, taking a trivial example, when we wake-up in the morning, how do we decide which leg reaches out to the floor first?Usually we are not aware of it although it is always the same leg. If a computer program has to take this decision, the rule governing this decision has to be explicitly specified. Coming back to the answer from Giovanni about computer-generated routing tables it frustrated that this answer did not acknowledge the difficulty of the task, but on the other hand, Giovanni's reaction triggered a new lash of energy and determination in my mind to find a solution!

I was spending many hours, whether at work, travelling in the tube in Paris, or at home thinking about how to tackle this problem. Actuallyafter some time of thinking, I had realised that the difficulties I was facing summed-up to the two following problems:

- Translate the constraints on the best routes tables into rules that could be easy to understand
- Find a way to easily implement an artificial intelligence approach, i.e. programming an heuristic resulting in most cases in an acceptable solution, rather than programming an algorithm that would provide the « optimal solution », as this latter approach seemed too ambitious and out of reach.

Two circumstances helped me address these issues. First of all, prior to joining SITA, in my work at CENT, I have had the opportunity to study and apply the the mathematical discipline of graph theory (actually I had worked with a guru of this theory, Michel Minoux, with whom I had publishe a short paper on the financial modelling of network investment planning). Applying graph theory to my problem , I was able for the first time to translate the constraints applying to routesinto mathematical rules : basically the result was that I could not consider these routes one by one, but had to include in the tables whole blocks of routes sharing constraints, in order to meet the functional requirements. Once I could establish thisrule, it was indeed a very significant progress in my quest for a solution.

The second positive circumstance came from the programming language that we were using at SITA on our IBM computer, namely PL/1. Before joining SITA, I had been programming in Fortran. But for my problem, PL/1 had, by chance, a very significant advantage over Fortran, as it allowed recursion. Recursion is something that we as humans have difficulty to use for dealing with a problem, but it is something very easy for computers (this fact was noted in a recent article in the French edition of the scientific magazine « Science »). I realized at some point that use of recursion would allow me to easily program the « trial and error » approach I had to use, rather that an approach based on theoretical model, to construct the best routes tables. I had noted myself that when constructing these tables manually, we would often encouter deadlocks in constructing these tables and in order to get rid of these deadlocks, we had to change an earlier choice which had seemed good but which

was actually creating this deadlock. This was, in a nutshell, the difficult part to program, and this is where the use of the recursion (or program re-entry feature as it is sometime called) would help tremendously in programming the heuristic I had now in mind.

I was ready to start programming the heuristic I had now designed, but as stated earlier, I did not like much programming. Fortunately I have been helped by a an intern, a young lady who was very good at programming. So the programming was a teamwork which worked very well and we did not have too many errors to correct (and this was fortunate because recursive routines are difficult to analyse and debug). I launched the development of three computer programs which were given the names of famous French clothing brands : Dior, Cardin and Patou. Dior was the one which generated the best route tables and it worked fine (the other programs were adressing the issues of finding the best topology, i.e. which circuits to use to route the traffic, and taking into account the fact that the HLN being a worldwide network, whe had to take into account three daily peaks of traffic corresponding to Europe, Asia and te Americas). The best routes tables generated by Dior were validated by the Operations team and the program was put at use by Annie, thus relieving the technical planning team from manually handling a task was was getting heavier and more difficult as the network was rapidly growing.

I have to confess that I was very proud of the results having solved what certainly has been the most technically difficult problem I have had to deal with in my career. This lead me to present the results of this work (which also included a graph theory result) at the IEEE Global Telecommunications Conference in San Diego, CA, in Nov-Dec 1983 (remember the work had started in 1978!). The IEEE conferences and publications were the most prestigious and advanced forum at that for telecommunication matters,and I was very happy to have the opportunity to present my work at this conference :

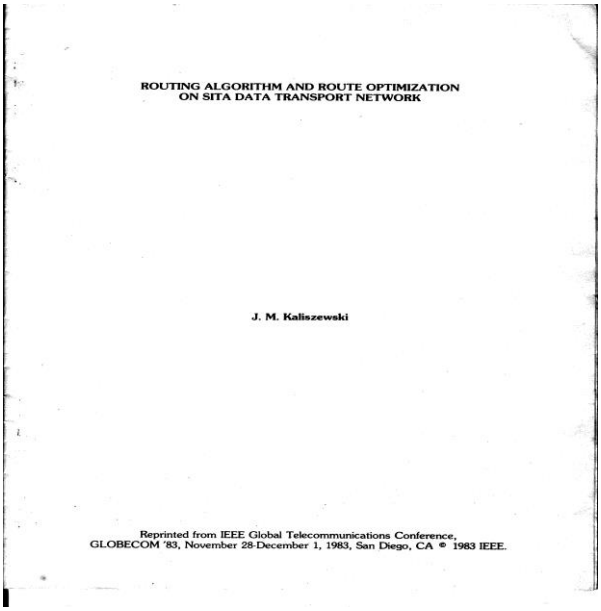

ROUTING ALGORITHM AND ROUTE OPTIMIZATION
ON SITA DATA TRANSPORT NETWORK

J. M. Kaliszewski

Reprinted from IEEE Global Telecommunications Conference,
GLOBECOM '83, November 28-December 1, 1983, San Diego, CA © 1983 IEEE.

**Figure 1** : first page of my article for the conference

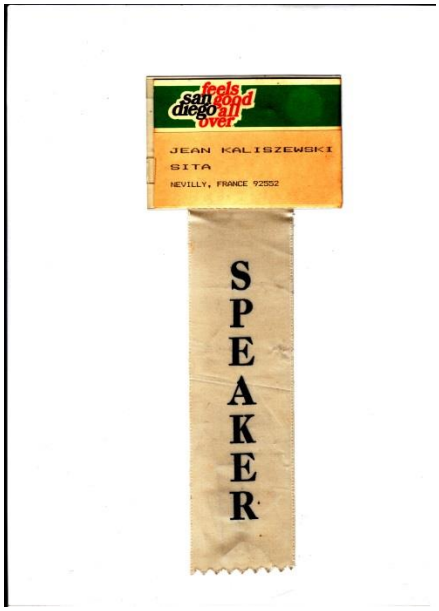


**Figure 2** : my speaker badge

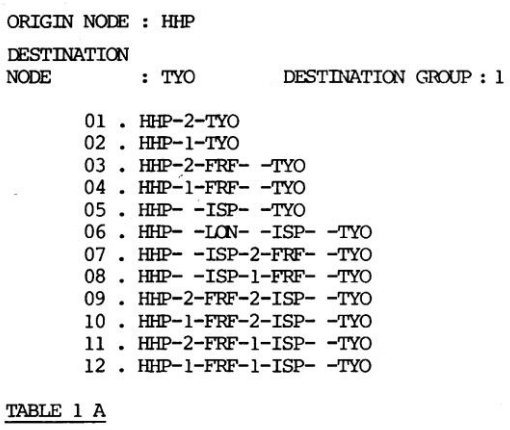

```
ORIGIN NODE : HHP
DESTINATION
NODE        : TYO        DESTINATION GROUP : 1

      01 . HHP-2-TYO
      02 . HHP-1-TYO
      03 . HHP-2-FRF- -TYO
      04 . HHP-1-FRF- -TYO
      05 . HHP- -ISP- -TYO
      06 . HHP- -LON- -ISP- -TYO
      07 . HHP- -ISP-2-FRF- -TYO
      08 . HHP- -ISP-1-FRF- -TYO
      09 . HHP-2-FRF-2-ISP- -TYO
      10 . HHP-1-FRF-2-ISP- -TYO
      11 . HHP-2-FRF-1-ISP- -TYO
      12 . HHP-1-FRF-1-ISP- -TYO

TABLE 1 A
```

**Figure 3** : A best routes table generated by DIOR

I had some difficulties when travelling to San Diego (I will tell about this in another article).

I remember proudly wearing my speaker badge at the conference which attracted a huge audience. However there were many sessions running in parallel, and I was a little bit disappointed that a relatively small audience (less than one hundred) attended my presentation. With hindsight, I must say that my expectations were not realistic at that time.

This conference was also the opportunity to learn about advances and meet some celebrities in the field of computer networking and I remember in particular Professor Mischa Schwarz from Columbia University. A very nice person who became a friend. He had published a boook on the design on computer communication networks and who was interested in knowing more about the SITA network.

Annie and her team used the DIOR program satisfactorily for many years and I had to work on this program again to improve its friendliness and meet new requirements (such as the implementation of the DTN phase 1). However the logic of the program was quite complex, including the recursion part, and at some time I was happy not to have to work on it anymore (see part 2 of this article)!

The figure 3 above shows a best routes tables produced by DIOR. There are some interesting features in this table. First of all we see nodes with unusual three letter codes: HHP, FRF, ISP: these were the n front-end nodes (probably DISs) with HLSs in back-end. Also there are two links between HHP and TYO: HHP-1-TYO and HHP-2-TYO. The reason was that the circuit speed was insufficient to accomodate the traffic between these nodes and to interconnect them with twol circuits to be able to route the traffic. Trying to cope with the growth of the network traffic has been a difficult fight for many years on the SITA network (but a challenge that we were lucky to have!). And the mention „Destination group: 1“ is a reflection of the fact that we had split-routing, in which case there were several destination groups for the same destination.

When writing this story I went to read again my original article written for the conference (I had kept some reprimpts), and it took me some time to understand the methodology and findings I was then describing. But on the other side, I must say that a deep emotion is still there when I am thinking

about this time and accomplishment, as this has been a unique moment in my professional life, linking my personal history to SITA's history,and a lasting feeling of a very positive outcome.