

Application Note

Number 12/2000

Updated: 2000-05-05

CELLULAR SYNCHRONIZATION NETWORKS FOR TELECOM APPLICATIONS BASED ON THE GPS AND ON SDH/SONET NETWORKS



1. INTRODUCTION

The most widely used method for distributing synchronisation in telecommunication networks uses SDH (Synchronous Digital Hierarchy) or SONET (Synchronous Optical Network) transport networks as the synchronisation carrier. Traditionally a central Primary Reference Clock (PRC) generates a synchronisation reference signal with high frequency accuracy. The synchronisation signal is then distributed to all nodes of the network via a tree of slave clocks connected by SDH/SONET traffic links. This paper presents an alternative method based on the Global Positioning System (GPS) for coarse distribution to a number of sub-networks, and on the SDH/SONET network for fine distribution to all nodes within the sub-networks. The method consists in partitioning the network at hand into sub-networks of predefined internal structure, and to interconnect the sub-networks according to well-defined rules. These rules determine the inter-sub-network or 'external' structure. The obtained sub-networks are called "cells", the resulting synchronisation network is called a 'cellular synchronisation network'. The chosen internal and external structures result in some interesting properties, which will be discussed in detail.

2. CLASSICAL SYNCHRONIZATION NETWORKS

As stated above, synchronisation is usually distributed throughout a telecommunications network using the traffic links of an SDH or a SONET transport network. Often a Primary Reference Clock installed in a centrally located node generates the frequency reference. The PRC usually consists of a redundant set of atomic Cesium clocks and/or GPS-receivers. The SDH/SONET network provides functions for the distribution of synchronisation from the PRC to all nodes of the network via a tree of slave clocks interconnected by SDH/SONET traffic links. There is also a simple messaging protocol for communicating a reference source quality information from clock to clock throughout the synchronisation distribution tree. This protocol is called Synchronisation Status Message (SSM). SSM is also useful for automatic protection switching in case of clock or synchronisation link failures.

Let us consider a linear chain of SDH/SONET nodes with their slave clocks. Synchronisation reference signals are fed to the chain in at least two nodes, typically at the two ends of the chain. The direction in which synchronisation is distributed inside the chain depends on priority table settings in the SDH/SONET equipment. Let us assume that under normal conditions, synchronisation is distributed from the west end node, which is also a synchronisation feeding point, to the east. In case of a link failure, all nodes east of the failed link are cut off in terms of synchronisation. Within the cut off half of the chain, the synchronisation flow will automatically change direction. After the automatic reconfiguration, the synchronisation will flow from the east end node, which is the second synchronisation feeding point, to all nodes east of the failed link. This automatic protection switching makes use of the SSM and of the priority-based reference selection function.

These classical synchronisation distribution techniques are illustrated schematically in figure 1. The upper part of the figure shows a portion of the network that is synchronised by a tree of master-slave clocks. The lower part of the figure shows one instance of a linear chain of clocks, which uses the SSM-controlled automatic reconfiguration mechanism. The numbers next to the arrows (= synchronisation links) indicate the priorities assigned to the synchronisation reference inputs. The dashed arrows inside the circles (= clocks) shows the synchronisation flow under normal operating conditions: in absence of failures each clock selects the priority 1 reference input.

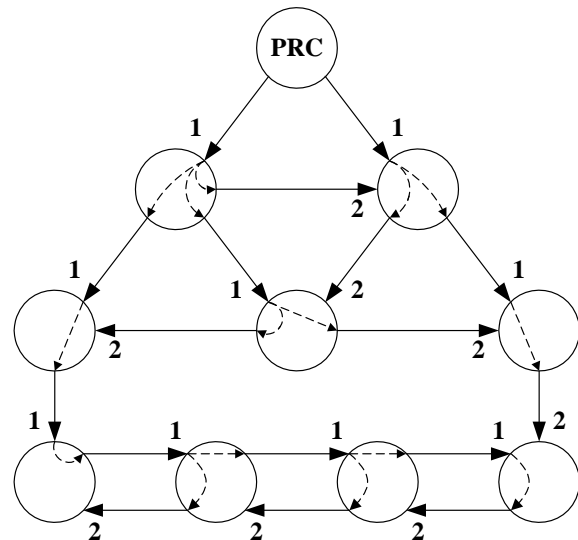


Figure 1: Classical synchronization network.

For detailed information on these synchronisation techniques refers to [1] to [5].

3. CELLULAR SYNCHRONIZATION NETWORKS

3.1. General approach

The cellular synchronisation network concept was developed with the intention of synchronising telecommunications network to the GPS. Its architecture was designed with the following criteria in mind:

- The GPS is the sole primary reference source (no atomic Cesium clocks)
- The synchronisation distribution to each node is protected (redundancy) against link failures, clock failures, GPS-receiver failures, and against local corruption of the GPS radio signal
- The protection mechanism is automatic (no human intervention required)

These criteria reflect a concern about the possibility of failing GPS reception due to radio interference or intentional jamming. Because of the stringent availability requirements in telecommunications networks, a synchronisation distribution network must provide protection against failure modes with non-negligible probability of occurrence.

The chosen architecture consists of an aggregation of synchronisation sub-networks called 'cells'. There exists

a limited number of standard cell structures called cell patterns. The cells are interconnected to form a synchronisation distribution network that covers the entire SDH/SONET transport network. The way cells are interconnected follows a set of simple rules.

3.2. The cell's internal structure

There are two standard cell structures called pattern A and pattern B. Figure 2 shows the typical internal structure of a pattern A cell.




Figure 2: Cell pattern A.

It consists of a linear chain of clocks using the SSM protocol for automatic protection switching. At both ends of the chain there is a GPS-receiver which delivers a synchronisation signal to the node's co-located clock. This synchronisation signal is derived from the GPS system time. The GPS-receiver's output complies with the specifications of ITU-T Recommendation G.811 (see [6]). The clocks at the two ends of the chain provide outgoing synchronisation links for interconnection with other cells. Figure 2 also shows the direction of synchronisation flow under normal operating conditions (i.e. no failures). Figure 3 shows what happens in case of failure of the GPS-receiver in the west end node: the automatic reconfiguration mechanism causes the GPS-receiver in the east end node to become the active primary reference source.




Figure 3: Failure of a GPS-receiver in an A-type cell.

Figure 4 shows the automatic reconfiguration caused by a link failure. After the automatic reconfiguration, all the clocks east of the failed link will get synchronisation from the GPS-receiver located in the east end node. The failure of a clock inside the chain has similar consequences.




Figure 4: Link failure in an A-type cell.

Cell pattern B resembles pattern A in the sense that it also consists of a linear chain of clocks using SSM-based automatic reconfiguration. Unlike pattern A, pattern B does not contain any GPS-receiver. Instead, there are incoming synchronisation links connected to the chain's western and eastern end clocks. This is shown in figure 5. A B-type cell receives synchronisation from other cells to which it is connected via the just mentioned incoming synchronisation links. This will be shown in section 3.3. In case of failures, a pattern B cell essentially behaves the same way as a pattern A cell.




Figure 5: Cell pattern B.

The maximum length of a chain, i.e. the maximum number of clocks in a chain, is limited by considerations regarding the maximum amount of jitter and wander that is allowed at any point in the network. The jitter and wander network limits for synchronisation interfaces are specified in ITU-T Recommendation G.823, chapter 6 (see [7]). The maximum chain length for the two standard cell patterns will become clear once the interconnection rules are defined. Therefore we will resume the discussion about the maximum chain length at the end of section 3.3.

3.3. The interconnection rules

A cellular synchronisation distribution network consists of a set of A- and B-type cells interconnected according to the following rules:

- 1) For all B-type cells of the network: each of the two synchronisation inputs must be connected to a synchronisation output of an A-type cell; moreover these inputs must be connected to different synchronisation outputs (however, the two synchronisation inputs of a B-type cell may be connected to the two synchronisation inputs of a single A-type cell).
- 2) An A-type cell's synchronisation output may be connected to one or many synchronisation inputs of B-type cells.

Rule 1) prohibits connecting the two synchronisation inputs of a B-type cell to one and the same synchronisation output of an A-type cell. However, the two synchronisation inputs of a B-type cell may be connected to the two distinct synchronisation outputs of one and the same A-type cell.




Figure 6: Cellular synchronization network.

Figure 6 shows a simple example of a cellular synchronisation distribution network (or simply 'cellular synchronisation network'). Each ellipse represents a cell. The cells are numbered: C1, C2, etc. The letter within the parentheses indicates whether the cell is an A-type or a B-type cell.

Here we need to resume the discussion about the maximum length of the clock chains in the cells. ITU-T Recommendation G.803, chapter 8 (see [4]) describes the architecture of a classical SDH-based master-slave synchronisation distribution network. Section 8.2.4 states that, in a chain of clocks, there should be no more than 20 SDH equipment clocks between the PRC and the first so-called 'node clock' (a node clock is high performance slave-clock featuring a bandwidth that is three orders of magnitudes narrower than that of an ordinary SDH equipment clock; node clocks are used to attenuate accumulated jitter and wander). Since there are no node clocks in a cellular synchronisation network, the maximum of 20 SDH equipment clocks applies to the maximum number of clocks that may be found in the longest possible chain following a GPS-receiver. As a consequence of the interconnection rules described earlier, the longest possible chain of clocks following a GPS-receiver is given by the chaining of one A-type cell and one B-type cell. It follows that the rule from ITU-T Rec. G.803 is fulfilled, if the chain length in both A- and B-type cells is limited to 10 SDH equipment clocks. Of course, this limit can be extended by the deployment of narrow bandwidth node clocks, as described in ITU-T Recommendation G.803; but this approach is outside the scope and the motivation of this paper.

3.4. Extended architecture

It is possible to define a somewhat richer syntax for the cellular synchronisation network architecture.




Figure 7: Cell pattern A*.

Figures 7 to 9 give some indications on how this can be done. Figure 7 shows a cell pattern A* which is derived from pattern A. Compared to the latter, pattern A* features additional synchronisation outputs indicated by the boldface arrows. Similarly, figure 8 shows a cell pattern B* which is derived from pattern B. Compared to the latter, pattern B* has additional synchronisation inputs. It can be shown that the SSM-based automatic reconfiguration works also with pattern B*. These new cell patterns allow for more complex topologies to be covered by the cellular architecture. This is illustrated only schematically in figure 9. Compared to the example of figure 6, there is an additional synchronisation link (indicated by the boldface arrows) connecting cells C1 (A) and C6 (B). In this case the additional link provides additional protection against failures affecting the synchronisation distribution between C1 and C6, and within C6.




Figure 8: Cell pattern B*.




Figure 9: Extended cellular synchronisation architecture.

4. PROPERTIES

The chosen network structure results in some interesting network properties.

4.1. Protection against failures

An interesting feature is the fact that all nodes of a cellular synchronisation network get two reference signals coming from two GPS-receivers located in different sites via geographically separate routes. Thus a cellular synchronisation network is at least protected against any single failure of a link, a clock, or a GPS-receiver. Of course, there is also protection against many multiple failure cases. The protection against GPS-receiver failures includes the case of local corruption of the GPS radio signal caused by radio interference or intentional jamming. Simply installing two GPS-receivers and two antennas on that site would not cover this failure mode, since both receivers would capture the same corrupted radio signal.

4.2. Timing loop prevention

The topologies resulting from the described cellular architecture have the property of necessarily excluding timing loop formation. Timing loops occur, when a chain of slave-clocks forms a loop, so that the active reference input of any of the clocks is actually locked indirectly to the output of that same clock. The occurrence of a timing loop can have dramatic consequences. The clocks on the timing loop are not locked to a PRC anymore. Because of the dynamic characteristics of the loop, the loop's clock frequency diverges, thus causing catastrophic degradation of traffic performance (e.g. bit error rate) in the telecommunications network. Timing loops occur because of human error during the network design process or during equipment commissioning. Making sure that there is no possibility of timing loop formation in a classical redundant master-slave tree network is a complex task. One must not only check the primary synchronisation distribution paths, but also all possible path combinations resulting from protection switching events. Failures in the network will cause some of the clocks to switch to another reference input (automatic reference selection function). These switchover events modify the topology of the synchronisation distribution paths. The network designer must make sure that these topology changes never lead to the creation of a timing loop. Verifying that a network is exempt of timing loops is much easier with cellular synchronisation networks. The internal structure of the cells and the cell interconnection rules necessarily exclude timing loops. A cell either delivers synchronisation to other cells (cell patterns A and A*), or receives synchronisation from other cells (cell patterns B and B*). There are no cells that would simultaneously deliver and receive synchronisation to or from other cells. It is thus impossible to interconnect cells in such a way, that a timing loop is created. Inside a cell the creation of timing loops is prevented by the SSM protocol. The way SDH equipment clocks process the SSM is standardised in ITU-T Recommendation G.781 (see [5]). The standard clock behaviour prevents creation of timing loops in chains of SDH equipment clocks using SSM. Finally, the simplicity of the design rules minimises the risk of human error during the network design process.

4.3. Simple and scaleable network design

One of the most beneficial advantages of cellular synchronisation networks over classical master-slave tree networks lies in the reduced complexity of the network planning process. Network planning includes the maintenance of the synchronisation network plan when the traffic network grows or when it undergoes topology, equipment, or configuration changes. Experience shows that adapting the synchronisation network to these modifications is a time-consuming task with classical master-slave tree designs. Sometimes it is necessary to redesign the entire tree topology. Also, these changes entail a lot of equipment reconfiguring. Adapting a cellular synchronisation network is less complex. All it takes is appending new cells to the existing design, or to change the internal design of the cells. Adding new cells to a cellular network does not affect the structure of the already existing network part. Changing the internal design of a cell is easy because of the cell's simple chain structure. Moreover, a change of the internal design of a cell does not affect the inter-cell topology. The only point that requires some attention is the length of the cell's clock chain, since there exists an upper limit (see section 3.3). But again, if the number of equipment clocks in a chain grows beyond the upper limit, all that is required is splitting the cell into several smaller cells. In short, one can say that cellular synchronisation networks are easily scaleable.

5. CONCLUSIONS

This application note presents a novel method for synchronising telecommunications networks. It is based on the Global Positioning System (GPS) for coarse distribution to a number of sub-networks, and on the SDH/SONET transport network for fine distribution to all nodes within the sub-networks. Two types of sub-networks with different internal structures are defined. A set of rules describes how these sub-networks, also called cells, can be combined in order to build a network that distributes synchronisation to the entire telecommunications network under consideration. This set of rules determines a class of inter-cell or external structures. The chosen internal and external structures result in some interesting network properties. Synchronisation networks of this type are scalable: adding new cells can expand them very easily, hence the name "cellular synchronisation networks". This kind of scalability greatly facilitates the maintenance of a synchronisation network plan during times when the telecommunications network is growing. Furthermore the chosen architecture has the property of excluding the formation of timing loops. Another interesting feature is the fact that all nodes of a cellular synchronisation network get two reference signals coming from two GPS receivers located in different sites via geographically separate routes. Thus the network provides protection against link failures, clock failures and corrupted GPS radio signals. These properties show that the cellular synchronisation network concept is an interesting alternative to the more traditional designs based on a tree of master-slave clocks locked to a central Primary Reference Clock. Cellular synchronisation networks are suitable for a wide range of telecommunications networks, both large and small.

6. REFERENCES

- [1] M. Sexton and A. Reid, Broadband Networking: ATM, SDH and SONET, Boston: Artech House, 1997, ch. 11, pp. 415-424.
- [2] St. Bregni, "A Historical Perspective on Telecommunications Network Synchronisation", IEEE Communications Magazine, vol. 36, pp. 158-166, June 1998.
- [3] S. Wong, "How to Synchronize Telecommunications Networks", presented at the 11th European Frequency and Time Forum, Neuchâtel, Switzerland, March 4-6, 1997.
- [4] ITU International Telecommunications Union, ITU-T Recommendation G.803: Architecture of transport networks based on the synchronous digital hierarchy (SDH), Geneva, 1997, ch. 8, pp. 14-25.
- [5] ITU International Telecommunications Union, ITU-T Recommendation G.781: Synchronisation layer functions, Geneva, 2000.
- [6] ITU International Telecommunications Union, ITU-T Recommendation G.811: Timing characteristics of primary reference clocks, Geneva, 2000.
- [7] ITU International Telecommunications Union, ITU-T Recommendation G.823: The control of jitter and wander within digital networks based on the 2048 kbit/s hierarchy, Geneva, 2000.