Redundancy
Pays for Itself

By Oliver Herzberger, Manager, Systems Engineering at DEV Systemtechnik

Investment in redundant system design saves long term operational costs.

Oliver Herzberger
Manager Systems Engineering
DEV Systemtechnik GmbH & Co. KG

Oliver Herzberger graduated in 2001 at the University of Applied Science in Communications Engineering. He started in the R&D Department of DEV Systemtechnik during his Diploma period.

First working as an RF (radio frequency) developer responsible for the development and innovation of RF technology within DEV devices, he changed in 2007 to Systems Engineering and is now the Manager of the Systems Engineering Department. Due to his solid expertise in solutions design of complete RF systems from satellite to receiver, he is now responsible for the technical project support of the DEV sales team and internal and external technical training. He is also responsible for the project management of international broadcasting projects such as solution design for the RF part of the main head-end in Mumbai (India), RF integration in a large teleport project in Turin/Italy and for various challenging projects, including military applications.

DEV Systemtechnik GmbH & Co. KG is a key integrator for the transmission and distribution of radio frequency signals such as those required in head-end stations. The company develops and produces products and systems for the entire signal transmission path. The portfolio from DEV comprises fully coordinated product solutions from antenna to receiver. All products are "Made in Germany" and meet demanding requirements on system availability, reliability and controllability.

Satellite ground stations, CATV head-ends and playout centres have to cope with demanding availability requirements. Unplanned shutdowns may cause enormous follow-up costs and contract penalties. Redundancy for important system functions and minimisation of the number of system components are proven measures to achieve and sustain availability rates of more than 99 per cent.

Designing system redundancy simply entails the provision of more components or subsystems than is actually necessary. Redundancy is a proven method to enhance system reliability and minimise the risk of failure. It causes additional complexity depending upon system availability requirements, in turn, one can expect lower operational costs and follow-up penalties in the case of failure.

Redundancy is not everybody's darling. The main arguments against it are that modern components and subsystems are reliable enough to do the job without redundant design. Furthermore, it is said to be costly in terms of money and rack space, wasted power and increased programming and configuration efforts for monitoring and control systems.

Its advocates mostly point to the the possible interruption-free operation of systems as the main asset. Higher initial investments are offset by lower operational costs. Redundant systems are closed systems which work autonomously. Only alarm-handling is necessary to incorporate them into a Monitoring and Control (M&C) system. There is no extra intelligence needed within the M&C system. Redundancy enables an undisturbed operation with scheduled maintenance intervals. No instantaneous response actions are necessary, not even in the case of failure.

Time vs. money

How redundancy influences system availability can be shown from a very simple example, a modulator subsystem in a head-end running 365 days per year around the clock. Its statistical mean time between failures is said to be 18 years. Without redundancy, failure of the modulator means at least a one-day operational disruption for repair or exchange of the defect parts. Correlating these figures leads to a statistical steady-state unavailability of 3.3 hours downtime per year.

Investing in a second modulator and redundancy switch brings the same configuration to a statistical downtime of 51 seconds per year, an improvement in system availability of a factor of 234.

What this means in terms of money depends upon the total system environment. With appropriate service contracts, exchange or repair time can be reduced to hours instead of days, but such contracts tend to be costly. Furthermore, the modulator in our example is only one link in a chain of elements with multiplying failure probability. But, as a rule of thumb, it can be stated that initial investment in redundancy pays for itself in the long term.

Antenna redundancy

The visually dominant system component of ground stations and head-ends is the parabolic antenna. It receives the weak satellite signals in the 10 GHz band, bundles them and feeds the LNB in the antenna focus. The LNB amplifies the signal and converts it to the L-band around 1.5 GHz, allowing for easier handling and transport via copper coax or fibre glass cables to the receivers, distribution and switch matrices.

Antenna redundancy usually is of the n + 1 type, providing one redundancy antenna for all antennas in operation. If one of the operational antennas fails, the redundancy antenna takes over, positioned manually by the system administrator or automatically. Fig. 1 shows the principle of a 4 + 1 antenna redundancy.

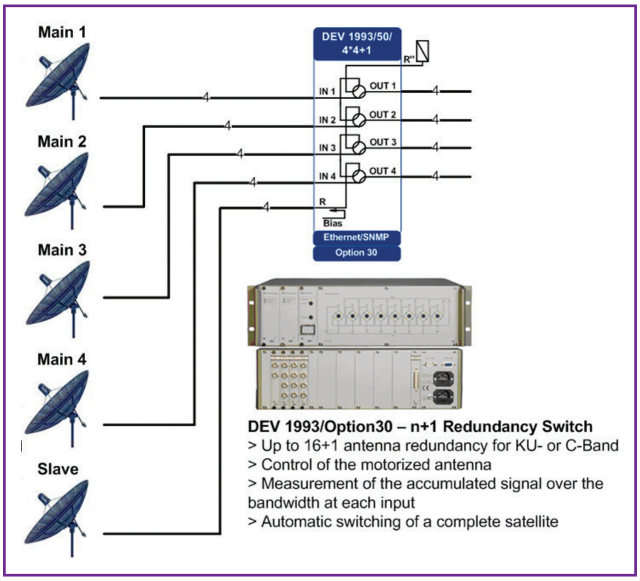

Fig. 1: Principle of a 4 + 1 antenna redundancy

Fig. 2: Motorised redundancy antenna at QTel in Doha, Qatar

A practical example is the installation of a motorised redundancy antenna at QTel in Doha, Qatar (Fig. 2). It is moved by an AKS 200 Antenna Control Unit (ACU) and controlled by a DEV 1993 Option 30 antenna monitoring and control system, memorising the position data of up to 99 satellites. The DEV 1993 redundancy switch enables the automatic positioning and switching of a motorised redundancy antenna. Antenna functions are monitored remotely via a web interface and communication with the network management system is based on the SNMP protocol.

Critical applications may justify a 1+1 antenna redundancy, providing one redundancy antenna for each operational antenna. The principle of this configuration is shown in Fig. 3.

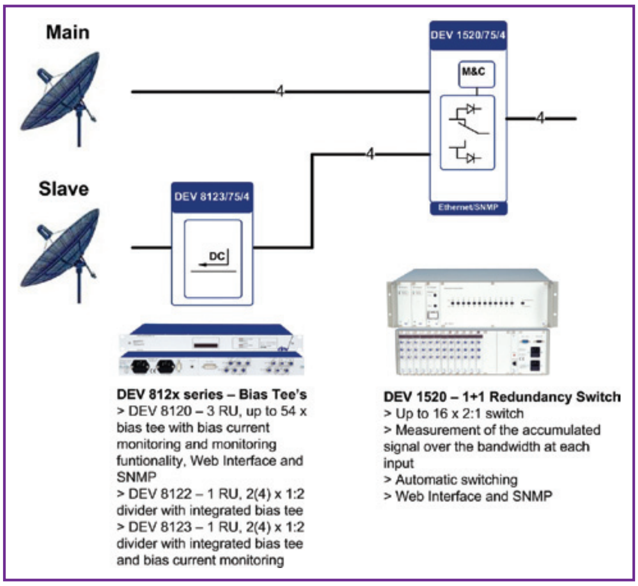

Fig. 3: 1 + 1 antenna redundancy, achieved with a bias tee of the DEV 812x product family and a DEV 1520 redundancy switch. Total signal intensity over bandwidth is monitored. It switches automatically in the case of failure.

Path redundancy

Low transmission losses over long distances, negligible frequency-dependent losses in the transmission band, galvanic isolation between transmitter and receiver, insensitivity to electro-magnetic interference and high interception security are the advantages of opto-electronic transmission compared to classical coax cabling. Glass fibre cabling turns out to be an economically attractive solution, particularly for new installations.

Of course signal paths can be designed redundantly, too. Dependent upon availability requirements, a 1+1 redundancy, offering one redundant signal path per base signal path is possible, or n +1 redundancy with one redundant signal path per n base signal paths. As each antenna normally delivers four signals (Upper band, Lower band, H & V polarity), n most often is 4 or a multiple of it. The principle of optical path redundancy is shown in Fig. 4.

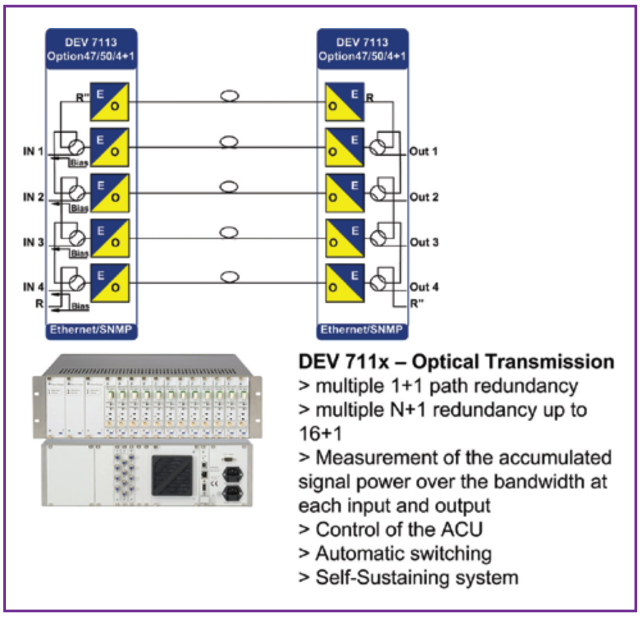

Fig. 4: 4 + 1 path redundancy with the DEV 7113 Optribution chassis. 20 slots with two transmitter/receivers each would allow for up to 8 x 4+1

With 1+1 redundancy options, the RF input signal is split and distributed to the main optical transmitter module and the redundant optical transmitter module ("hot stand-by"), thus the optical transmission is performed via two optical fibres. The two optical receiver modules convert the signal back to an RF signal which feeds the switch module in the receiver chassis.

There are four RF sensing units per signal channel which can be used for the surveillance of the transmission path. But, relevant for the autonomous switching between the main channel and the redundancy channel, are the two RF sensing units at the optical receiver modules.

If the RF sensing unit of the optical receiver module for the main channel detects that the signal is below the defined threshold, the RF sensing unit of the optical receiver module for the redundancy channel is checked to see whether this RF signal is above its defined threshold and the switch module is triggered to switch. Therefore, the RF output is fed by the redundant optical receiver module.

In n+1 redundancy configurations, the RF input signals pass through the transmitter redundancy switching section and the optical transmission will be via the corresponding optical transmitter receiver module pairs. After conversion back to the RF, the signals will feed the receiver redundancy switching section in the receiver chassis. In normal operation, an RF output is fed by the corresponding optical receiver module.

There are 2 * (n+1) RF sensing units for the complete configuration which can all be used for the surveillance of the n+1 transmission path. But relevant for autonomous switching are the RF sensing units of the optical receiver modules. If one of the RF sensing units of an optical receiver module detects that the signal is below the defined threshold for a channel but the signal level is above the threshold on the transmitter side, the switching of that channel to the redundancy channel in the receiver redundancy switching section and in the transmitter redundancy switching section is initiated. This is done autonomously. Detected are:-

- defective transmitters
- defective optical signal paths (glass fibre cables)
- defective optical receivers.

The RF signal (RF sensing) as well as the transmitted and received light is measured.

Minimising components

Finally the signal arrives in the building hosting the receivers. There may be nearly 100 of them, as in the QTel example, or many more. The task then is to connect the receivers correctly with the antennas. Until recently, the signal distribution system had to be implemented at RF after the optical receivers.

There is a now a device available which is not only able to transmit, receive and distribute (1:8, 16, 32, 64, 128) but also switch (4x8, 16, 32 and 4x64) optical signals, the DEV 7114 (Fig. 5). At the ANGA Cable 2011 trade fair in Cologne, Germany, internal redundancy options for the DEV 7114 were introduced giving a maximum of 6 x 1+1, 1 x 4+1 or 1 x 8+1.

Previously, this functionality could only be achieved with at least two system components, a DEV 7113 3RU chassis with 20 optical slots and a DEV 2190 4RU signal distribution system. Thus ”more functions in fewer devices“ is not only a marketing slogan but also a valuable contribution to the minimisation of device count, enhancing the availability of the total system.

Fig 5: Front view of the Intelligent Optribution Chassis DEV 7114, now available with internal redundancy options.

Power supply redundancy

Depending upon the specified availability requirements, system critical components within devices are also designed with redundancy. For good reason, much attention is paid to the power supply. As an example Fig. 6 shows the rear view of the DEV 7114. There are two different mains plug receptacles, offering the possibility to supply the device from two different mains circuits.

Fig 6: Rear view of the DEV 7114, featuring two mains plug receptacles for redundant power supply.

By providing three separate internal power supply units, (Fig. 5 shows three horizontally-positioned power supply modules above the smaller vertically-arranged signal modules) power supply redundancy can achieved. If one of the power supplies fails, the others will continue in 1 + 1 redundancy mode. The defect unit can be changed in “hot-swap“ mode without interrupting operation. Power supply module redundancy and mains supply redundancy are independent.

Redundancy of amplifiers

Fig. 7 shows a common amplifier redundancy circuit. The incoming signal is split and fed to the inputs of two parallel amplifiers with a 90° phase shift for one of the signals. Afterwards, the two signals are united again. If one of the

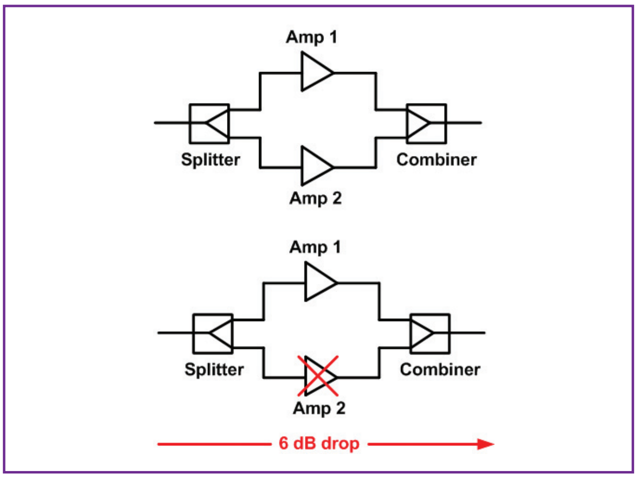

Fig. 7: Redundancy by simply paralleling two amplifiers.

amplifiers fails, there is still a correct signal from the combiner's output but with a signal level of 6dB less than when both amplifiers are working. This solution generates higher cost and system power. Furthermore, a spare part has to be kept in stock to regain the original signal level as quickly as possible.

An alternative solution is shown in Fig. 8. Here, amplifier redundancy is achieved at module level. A splitter parts the signal and puts two identical signals to two amplifier modules. The outputs of these modules are monitored by a redundancy switch. In the case of amplifier failure, the signal is switched from one module to the other within 3 ms. Rebuilding redundancy can be done without time pressure and operational interruption by changing modules.

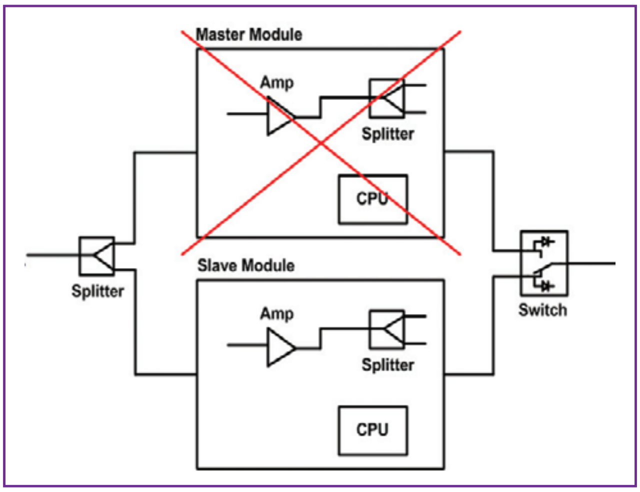

Fig. 8: Amplifier redundancy, the DEV way

Aims, costs and initiatives

The first question when dealing with redundancy is: What should actually happen if the system fails? Should the entire system be shut down completely or brought to a stable state, delivering error messages and waiting for remote diagnosis and maintenance? Should it try to become operational again, following pre-defined rules and schedules, or should it continue to work without noticeable interruption as in the case of mission-critical systems?

All of these scenarios raise technical and economical questions. More safety means more effort. Redundancy has many facets but the truth, at least for the long haul, is that redundancy pays for itself!