3 Foshay Road, Dudley, MA 01571

Choosing or Specifying an RF Switch Matrix

Introduction

This application note describes the different types of switch matrix products available for IF, RF and microwave switching. It also discusses many of the important tradeoffs that should be considered when selecting a switch matrix to fit your application.

Many installations and facilities use switch matrix products throughout the entire signal chain. This paper is focused on those portions of the system where the signals being switched are (typically) modulated carriers. The type of modulation is not important to our discussion, since the switch matrix only has to pass the signal without degrading it. We are interested in the portions of systems passing RF signals that are already modulated (on the transmit side) or not yet demodulated on the receive side. Signals received from an antenna or down-converted by a LNB, for instance, are part of our discussion, as is multiple carrier multiplexes being prepared for transmission.

For these discussions we exclude switches designed for baseband or digital signals such as RS-232, ECL, RGB video, audio and similar signals. Although some of these signal types can cover many MHz of bandwidth, they require switch matrix designs that are substantially different than those of an RF, IF or microwave switch.

For the purposes of this paper, we will assume that the application is a satellite communication ground station. The process of selecting a switch can be applied equally well to other application areas, but it is convenient to use one of the most common applications, because it provides a variety of possible switching configurations for our discussion.

1.1 Key Definitions

A glossary of terms is included at the end of this document. In this document, we will need to distinguish the switch matrix from an individual switch element inside the matrix. We will refer to the switch matrix as a whole, by the term “matrix”. We will refer to an individual switch element inside a matrix, as a “crosspoint”. In this way, we will minimize the use of the term “switch” with its potentially ambiguous meaning.

Unless we specify otherwise, when the following discussions refer to “RF” signals, we mean any frequency, so that IF, RF and microwave signals are all being considered.

1.2 Organization

Section 1 of this paper will encourage you to read with your own application in mind. Some key questions are listed, which will allow you to make your own tradeoffs as you read.

Section 2 covers some basic switch matrix terminology and looks at block diagrams of contrasted architectures.

Section 3 describes the different types of physical switching elements (crosspoints) that are found in switch matrix products. Those technologies are compared in a number of categories and generalizations are made regarding the suitability of each for various application requirements.

Section 4 looks at the switch matrix as a complete system. The discussion focuses on tradeoffs in such areas as physical packaging, method of control, maintainability, and other system level issues.

1.3 Know Your Application

The first thing that will affect your selection is the matrix’s location in the signal chain sequence. Location implies things about signal characteristics and its importance to the overall system.

- Will it be located in the receive or the transmit side of the signal chain?
- Will the signals be very low levels or will they have been amplified to very high levels when they reach the matrix?
- Will the signals passing through the matrix be revenue generating (or otherwise critical to you or your customer) or will they be monitor and test signals used for maintenance?

The second thing that will influence your decision is the anticipated size of the switch matrix.

- How many inputs and outputs are required?
- How much physical space is available for the matrix?
- Is significant expandability required? Can a small percentage of spare inputs and outputs be adequate to cover future needs for the life of the matrix?

2 Switch Configuration Tradeoffs

Like many engineering problems, there are multiple solutions to the problem of switching signals. Several typical configurations have been used in the industry, each with its own merits. Vendors will optimize their own designs, of course, but in general, the respective product offerings can be grouped by a few characteristics of their internal design.

For our discussions, we will take an example switch matrix that has 4 inputs and 4 outputs. This matrix would be designated as a 4X4 configuration. A switch matrix can have any number of inputs and outputs, designated as NxM. A 4X4 configuration is unusually small, but it reduces the complexity of the drawings.

2.1 Directionality

Switches can be unidirectional or bi-directional. Few RF applications require true bi-directionality in the matrix, because the upstream and downstream equipment typically only moves the signal in one direction. If the switch has gain elements (amplifiers) it typically cannot be bi-directional.

Almost all bi-directional switches have a single input or single output. So they are 1xN or Nx1 configurations. For our purpose here, we will not consider bi-directional switches since they are rare and typically very small.

2.2 Blocking vs. Non-Blocking

A primary differentiator for switch matrices is whether they have a “blocking” or a “non-blocking” architecture. Blocking is the situation where the user wants to connect a particular input to a particular output, but the switch matrix is not able to do so, because that input is already committed to another output. So the request cannot be satisfied *without changing other connections or outputs*.

Some systems tolerate rearrangement of existing channels, so a blocking characteristic is acceptable. The public switched telephone system is one example. In commercial or military satellite environments, the momentary interruptions associated with any path re-arrangement would be unacceptable.

The figure below shows a simple 4x4 blocking switch matrix. The heavy line shows that input 1 is connected to output1. The inputs all pass through a single path in the diagram.

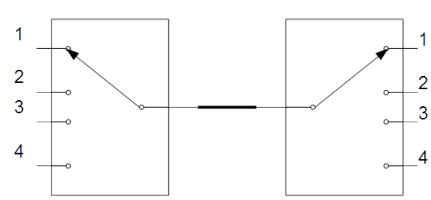

Figure 2-1 : 4x4 Blocking Switch Matrix

So far, this switch has accomplished 1 connection. But if the user of Output 2 needed to look at input 1, the switch could not service the request.

Output 2 cannot be connected to Input 1 unless Output 1 relinquishes its own connection to Input 1. This configuration can only connect an input to (at most) one output at a time. We say this configuration has a fan-out of 1, and this matrix design is not “full-fanout”

A blocking switch matrix can often be completely passive, since the only losses are two switch contacts and some interconnecting cable. It is possible to have total losses under 2 dB for a 16x16 blocking matrix. Lower loss also yields better bandwidth, better flatness, lower noise, lower distortion and higher power handling.

A non-blocking switch is one where no request for a connection is blocked. This requires that each input signal be handled in such a way that it can be sent to many destinations simultaneously. If we replace the input 1-to-4 switches with 1-to-4 power dividers, we can overcome the blocking situation. The figure below shows a full fanout, non blocking 4x4 switch matrix. The power divider is shown as a transformer with multiple secondaries. Each output of the power divider is an independent copy of the input, though it is at a reduced power level (the incoming RF power was divided equally among the N outputs)

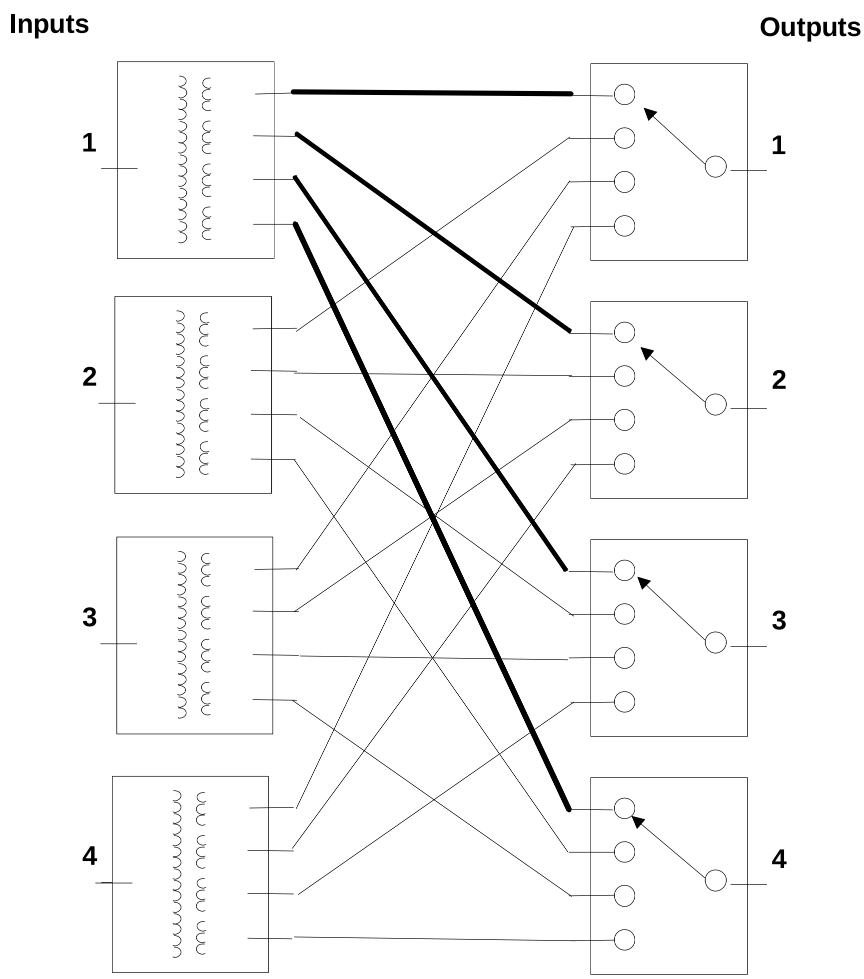

Figure 2-2 : 4x4 Full Fan-Out Non-Blocking Switch Matrix

The total number of interconnects between the power dividers and the switches is the same as in the blocking switch. For an NxM switch matrix, each input power divider will need M outputs, and each output selection switch will need N inputs.

Full fanout, non-blocking matrices are very common. The price paid for the non-blocking feature is the signal path losses in the power dividers.

But most applications require a non-blocking architecture. The non-blocking matrix usually requires amplifiers to recover the signal attenuation of the power dividers. This means higher power consumption, some loss in signal fidelity, and increased complexity. But with proper design, the performance can remain excellent over a wide dynamic range.

2.3 Fan-In vs. Fan-Out

The non-blocking switch matrix block diagram in Figure 2-3 shows that it is possible for an output to connect to only one input at a time. However, any input can be connected simultaneously to multiple outputs, because of the power dividers. This type of matrix is called a fan-out matrix. It is the most common configuration. It is used mostly in receive-side applications, where a signal may need to be routed to multiple destinations simultaneously.

A fan-in configuration is a mirror image of the fan-out. The figure below shows a 4x4 fan-in non-blocking switch matrix. Output 1 is connected so that it provides the summation of all 4 input signals. The Fan-in configuration finds its usefulness in uplink paths, where multiple carriers (at non-overlapping frequencies) may be multiplexed together by the power combining property of the power divider.

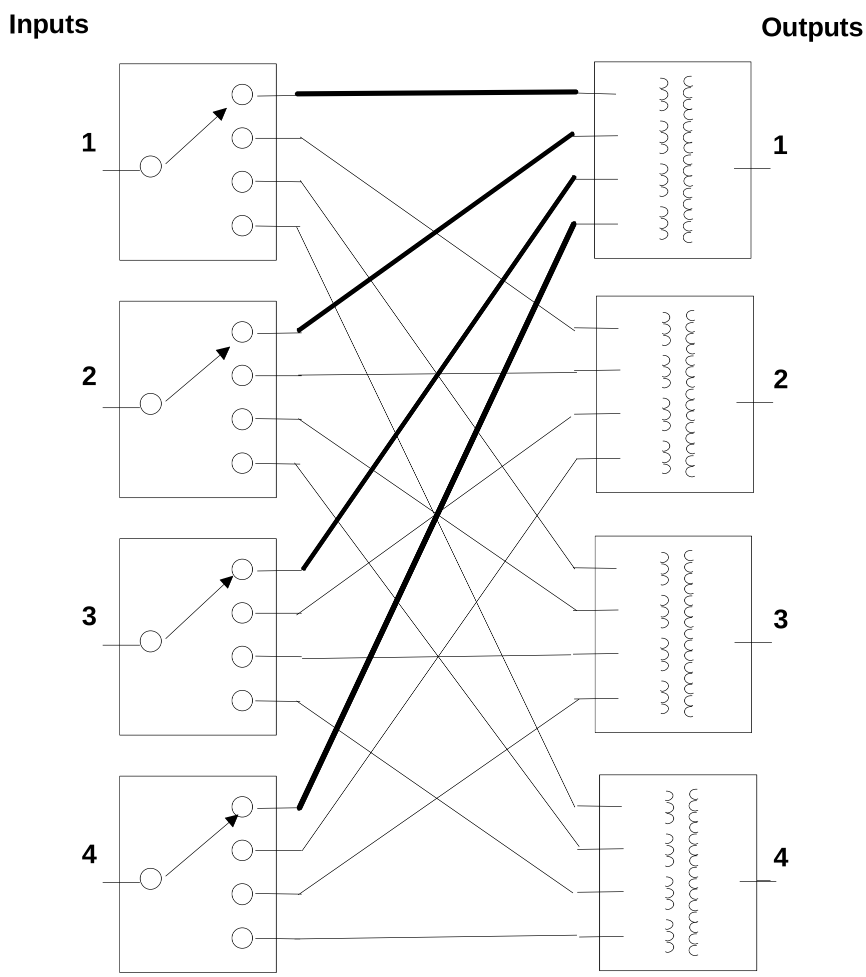

Figure 2-3 : 4x4 Full Fan-In Non-Blocking Switch Matrix

2.4 Single vs. Multi-level

Switching occurs in multiple stages within a large matrix. It is usually not possible to build a single multi-pole switch that has enough capacity for all the signals. Instead, switching components are cascaded until the full capability is realized. For instance, semiconductor crosspoints are limited by the number of pins that can be mounted on a single integrated circuit. The practical limit at high frequencies is about 8 inputs. So a 32 way switch will require multiple 8 pole switches to achieve its required capability.

A large switch matrix can be built up from smaller complete matrices. This is typically how vendors provide their largest switch offerings. So a matrix switch typically has a core switch size and a core power divider size that become the building blocks for all matrices in a product line.

The 4x4 matrix in our examples above is a single level switch matrix. All the power dividing occurs before any switching occurs. If you were to trace any connection path, you would encounter a total power division ratio of M before encountering any switches. Then you would encounter N selection switches without encountering any more power division.

A multi-level switch breaks the matrix into smaller independent non-blocking full fanout matrixes. When tracing a path from an input to an output, the signal would be divided by less than M, then selected by a switch that was less than N wide, then power divided some more, then down-selected some more, until the total power division was at least M and the total selection was at least N. A diagram will make this clearer.

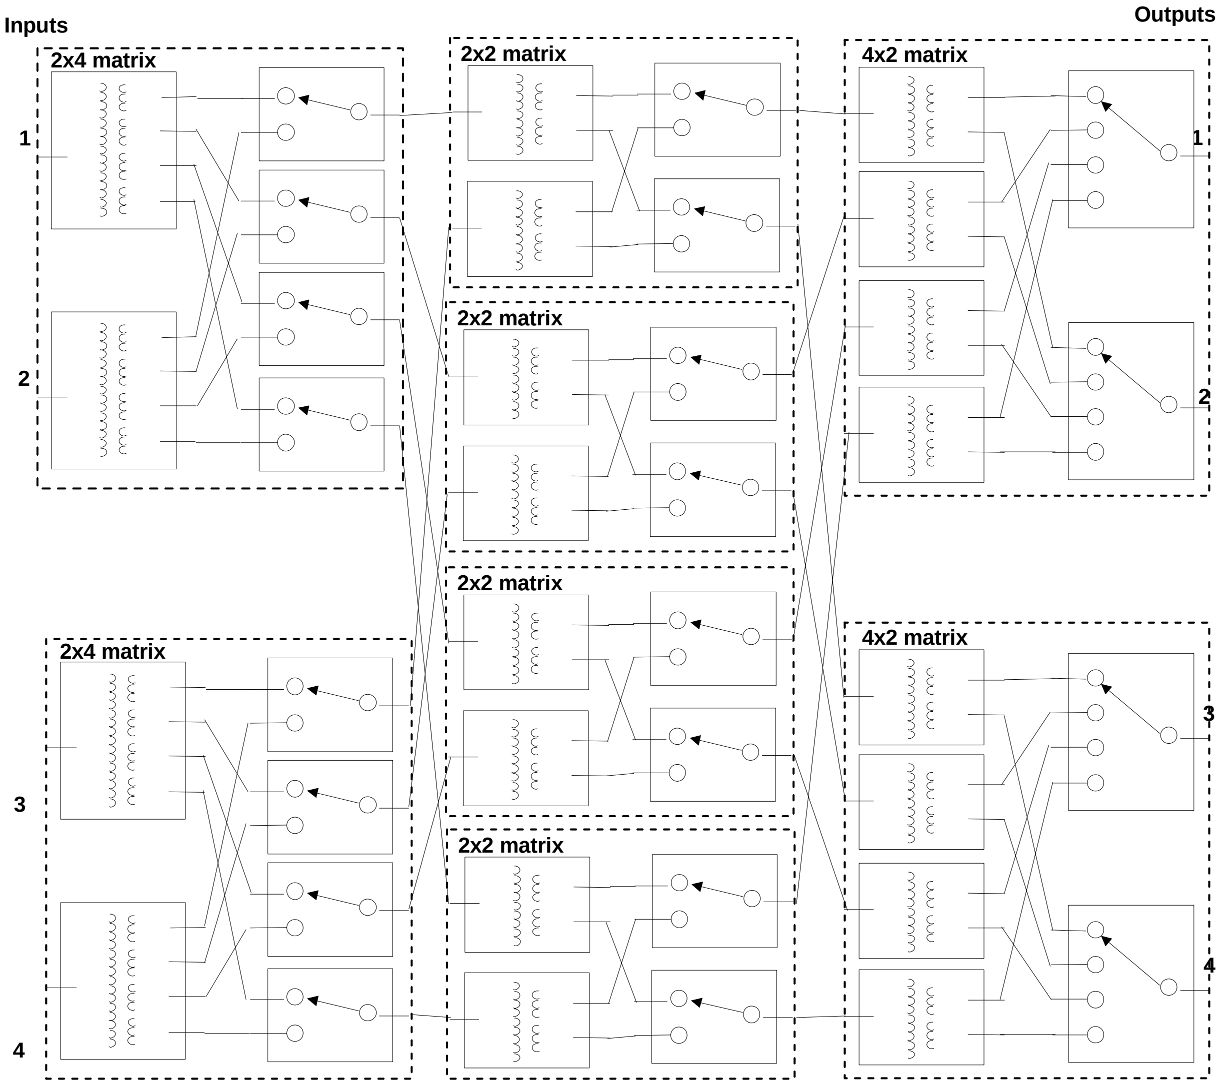

Figure 2-4 : 4x4 3-Level Non Blocking Matrix

A 4x4 matrix is too small to achieve any benefit from a 3-level architecture, but for simplicity we use it anyway. In fact, there is a significant cost penalty associated with 3 level switches, until the total matrix size reaches about 32x32.

The 4x4 is shown as interconnections of smaller NxM matrices. A 2X3 matrix is used at the input, a 2X2 matrix is used in the center column, and a 3X2 matrix is used at the output. This configuration meets the mathematical criteria for a minimal size 3-level 4x4 non blocking switch. All output requests can always be satisfied, without rearranging any other paths. [for a discussion of the mathematical criteria for 3-level non-blocking switches, see Benes, reference 1).

Compare the 3 level switch to the single level switch matrix. The single matrix required 16 interconnecting paths between power dividers and switches. The 3-level version requires 48. The total power division ratio (power loss) for the single level switch was 4 (6 dB). For the 3-level switch, that loss ratio is 12 (11 dB).

The benefit to the 3 level switch are two-fold. In very large switches, the total switch crosspoints can be reduced to less than NxM. This results in substantial cost savings in a large switch. The second benefit is that there are redundant paths between any input and output pair. In critical applications, this redundancy may be worth the extra complexity.

3 Switching Technology Tradeoffs

3.1 Crosspoint Types

There are 5 major types of electronic switches that are applicable to the RF/IF environment. Two are solid state devices, and three are electromechanical devices. The electromechanical devices consist of coaxial relays, reed relays, and microelectromechanical systems (or MEMS). As of 2016, some MEMS are available, but yet constantly changing. As a result, no data will be presented. The solids state devices are GaAs MESFETs and PIN diodes. CrossPoint Technologies can provide any of these switching elements in a matrix. The most common is the GaAs FET, followed by the microwave coaxial relay. The reed relay is popular at lower frequencies, and the PIN diode is requested in only a few applications.

The sections that follow will give a comparative overview of these switching technologies. Most of the comparisons are shown graphically, but the numeric scales should be treated as approximations only. Manufacturers are constantly extending the performance of their components. The limits are intended to show relative strengths, rather than specific performance limits.

In the graphs and discussions that follow, only the crosspoint components themselves are under consideration. The impact of amplifiers, power dividers, cables, etc are not considered unless explicitly stated. (Section 4 takes up the system-wide issues)

3.2 Frequency

The electromechanical devices have the potential to pass DC and very low frequencies. However, if amplifiers are used in the matrix, the low frequency capability is usually lost. The dotted line in the reed relay shows that some relay vendors claim usable performance beyond 1 GHz.

The GaAs FET and PIN diode both operate above 10 GHz. GaAs FET distortion tends to increase dramatically in the lower MHz region, though GaAs is still usable in some low frequency applications. The ease of biasing and low power (see below) are often more important than low distortion, especially when signal levels are < -10 dBm.

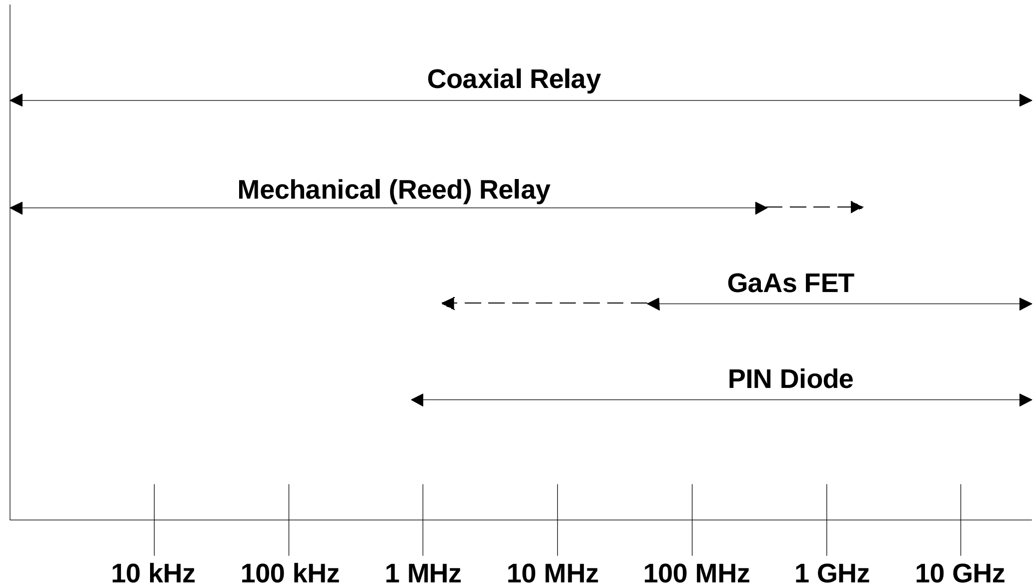

Figure 3-1 : Frequency Ranges

3.3 Flatness and Bandwidth

All the devices are very broadband. Reed relays exhaust usable range near 1 GHz, the others can operate much higher.

The bigger contributors to loss of flatness tend to be system issues, such as cable slope (especially in large coaxial matrices where cable lengths are very long) and amplifiers. Typical COTS matrix products are very broadband.

Because the matrix does not have filtering, phase accuracy and group delay tend to be excellent. Adding a switch matrix to a satellite signal chain does not typically influence the overall phase linearity.

3.4 Isolation

The chart below shows the range of isolations of a single switch element, against the frequency range. Isolation of all devices degrades with increasing frequency. Again, MEMS are excluded due to the constantly improving products. It should be noted they have isolation comparable to electromechanical solutions. The coaxial switch is substantially better than all the other solutions. This means that a coaxial relay matrix only requires a single switch to achieve the total isolation. The other technologies usually have to put several devices in series to build up enough isolation for the application.

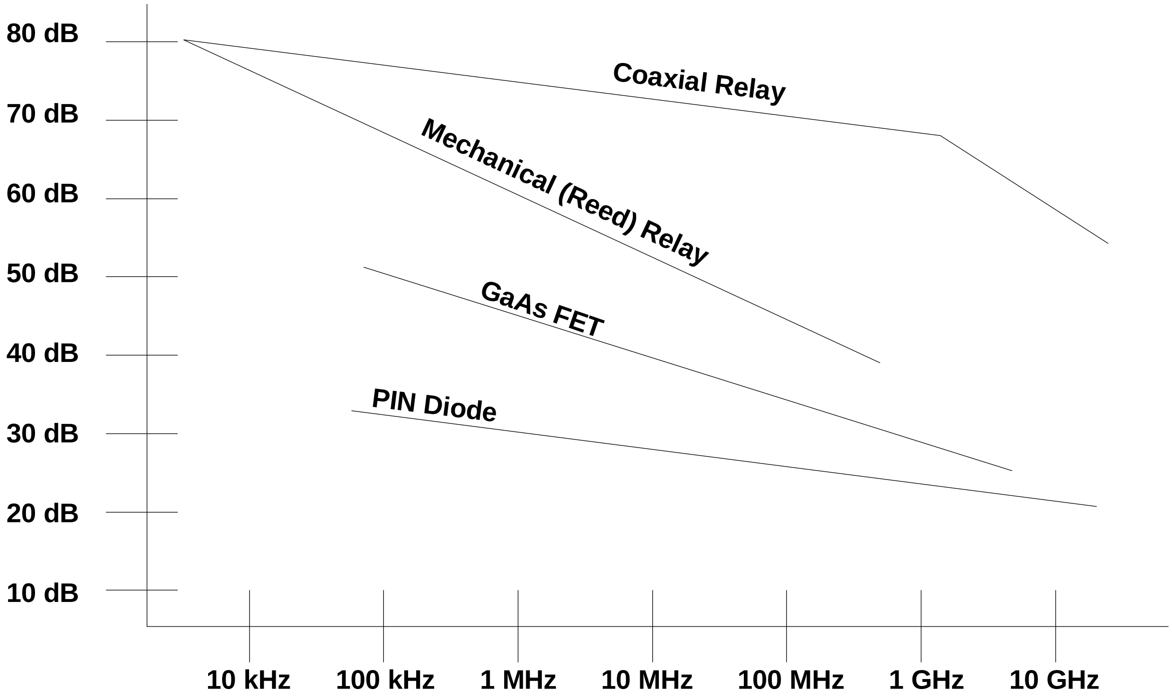

Figure 3-2 : Single device Off-Isolation

3.5 Signal Power Handling

Signal power level is important to the solid state devices. If the signal levels are too high, distortion will result. The electromechanical devices can handle much higher CW power. However MEMS are a unique case. Because of the majority of customers' needs to switch live signals, MEMS can not always be considered as alternative crosspoint elements like other electromechanical solutions. Most ground station applications with signal levels below +10 dBm range can be served by any of the technologies.

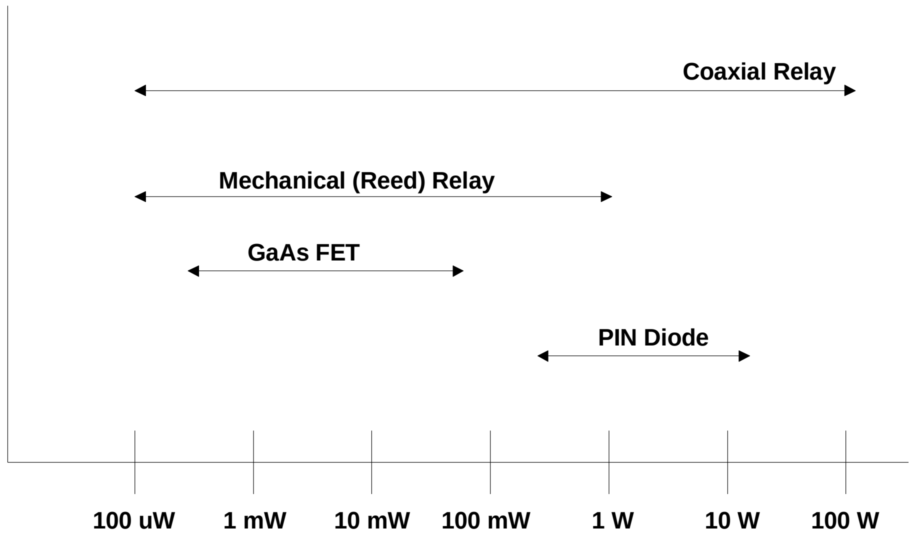

Figure 3-3 : Signal Power Handling

3.6 DC Power

The amount of DC power required to operate a single stage device is compared below. The coaxial relay is shown in two configurations. The standard coil arrangement requires power continuously, while the magnetic latching configuration can change state with a pulse. Power is not running continuously no matter what state the switch is commanded to. A large coaxial relay matrix will benefit from using latching relays, though there is increased complexity in the drivers and auxiliary contacts are required to make sure the latch pulses achieved the desired connections.

PIN diodes require significant bias currents – tens to hundreds of milliamps per diode. A PIN switch with 60 dB isolation will require 3 or more diodes. PINs are not cost effective in general purpose matrices because of the power dissipation issues.

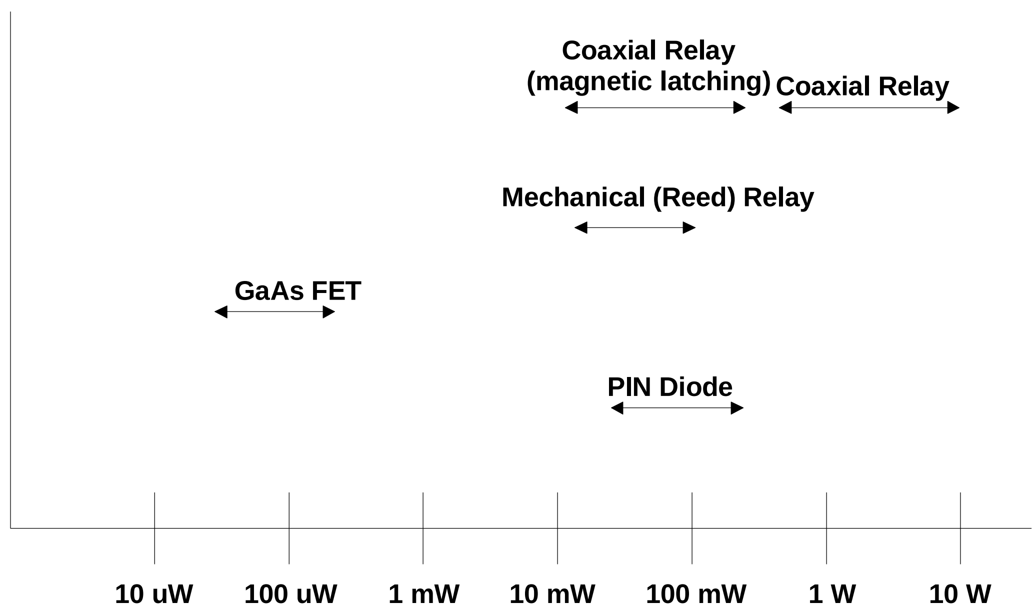

Figure 3-4 : DC Power Consumption

3.7 Switching Speed

The solid state devices are faster, as one would expect. This comparison is for the device itself, not the overall matrix switching speed. This graph does not account for driver rise times, control delays, etc. PIN diodes are suited to high speed, high RF power applications. For lower power applications, GaAsFETs are a better choice. They retain the speed without the power consumption.

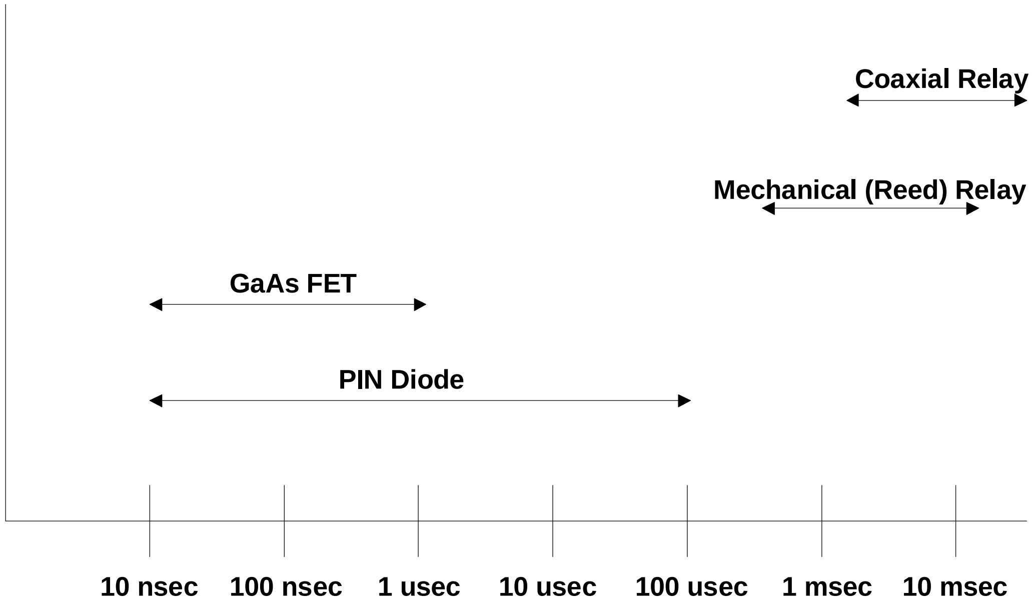

Figure 3-5 : Switching Speed

3.8 Size

This chart shows the range of physical volume of the devices themselves, scaled to a single crosspoint. Driver circuitry is not considered. PIN diode driver circuitry is much larger than GaAsFET driver circuitry.

For a large matrix, the coaxial relay consumes more volume for the switch elements, but it also will have enormous cabling volumes, because coaxial relays are typically built in a round configuration. This form factor is incompatible with a plug-in row/column format available with PCB mountable switch elements.

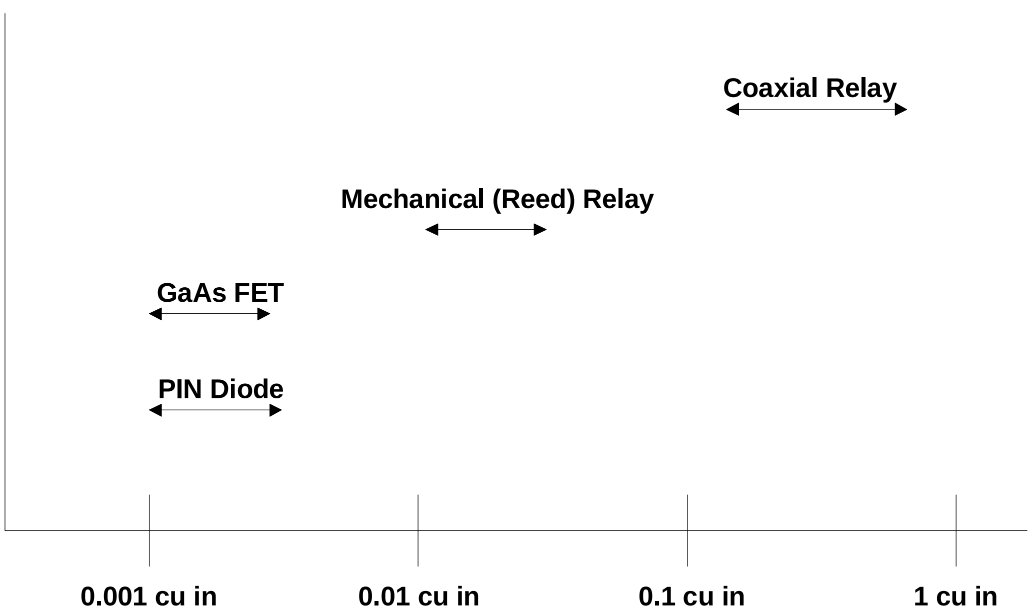

Figure 3-6 : Switch Physical Volume per crosspoint

3.9 Lifespan

Lifespan of the solid state devices is very large, the scale numbers are arbitrary. On the other hand, electromechanical devices have specified life spans (operating cycles). Coaxial relays are typically specified at 1 million operations. The usual end of life condition is electrical performance degradation (resistance rise), not mechanical failure. For matrices with only occasional path changes, electromechanical devices can provide many years of trouble-free service.

When considering MTBF, remember that a solid state switch may use a number of devices in series to achieve isolation as well as to make a multi-way switch. But the coaxial switch is a single multi-pole device. So the difference between a single device is less significant when considering the entire system. But solid state switches will still give much longer service life than electromechanical devices.

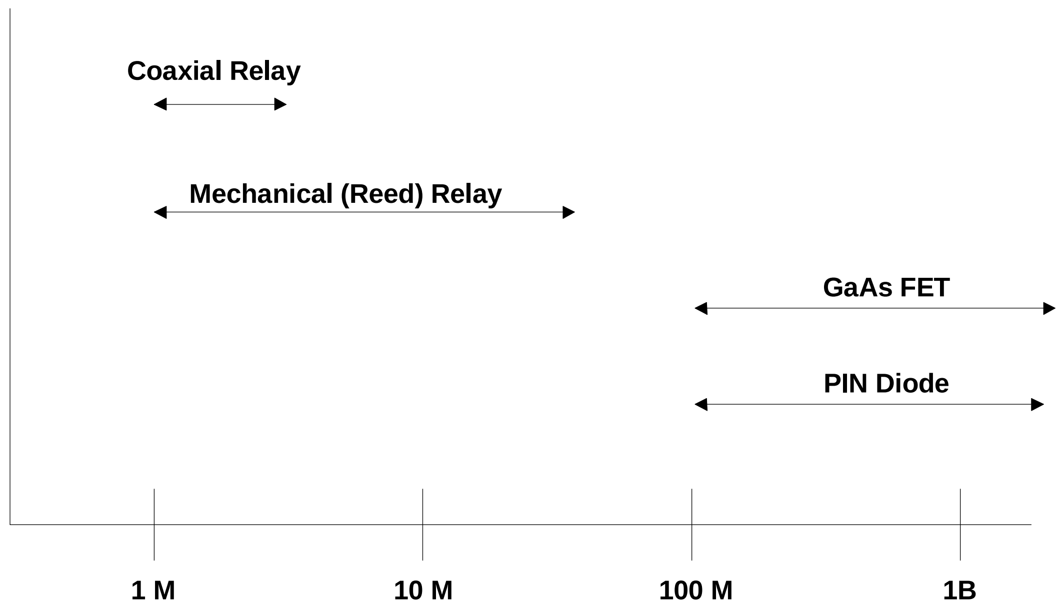

Figure 3-7 : Device Operations

4 System-Level Tradeoffs

In this section, we discuss system-wide factors that influence design and specification decisions. The 4X4 non-blocking configuration is embellished in the drawing below, to show one way in which amplifiers are incorporated into a matrix. In this case, a single amplifier provides all the make-up gain to overcome all the losses from power divider, interconnects and switch insertion loss.

This configuration is shown because amplifiers tend to be the limiting factor in system performance. They also can be a significant contributor to MTBF. Distributing the gain throughout the signal path is another alternative. But that approach makes it harder to match channel to channel (differential) gains. CrossPoint has matrix products with amplifiers that are individually hot swappable from the rear panel. The amplifier modules can also be customized for noise or distortion performance, and can include adjustable attenuators or filters if required.

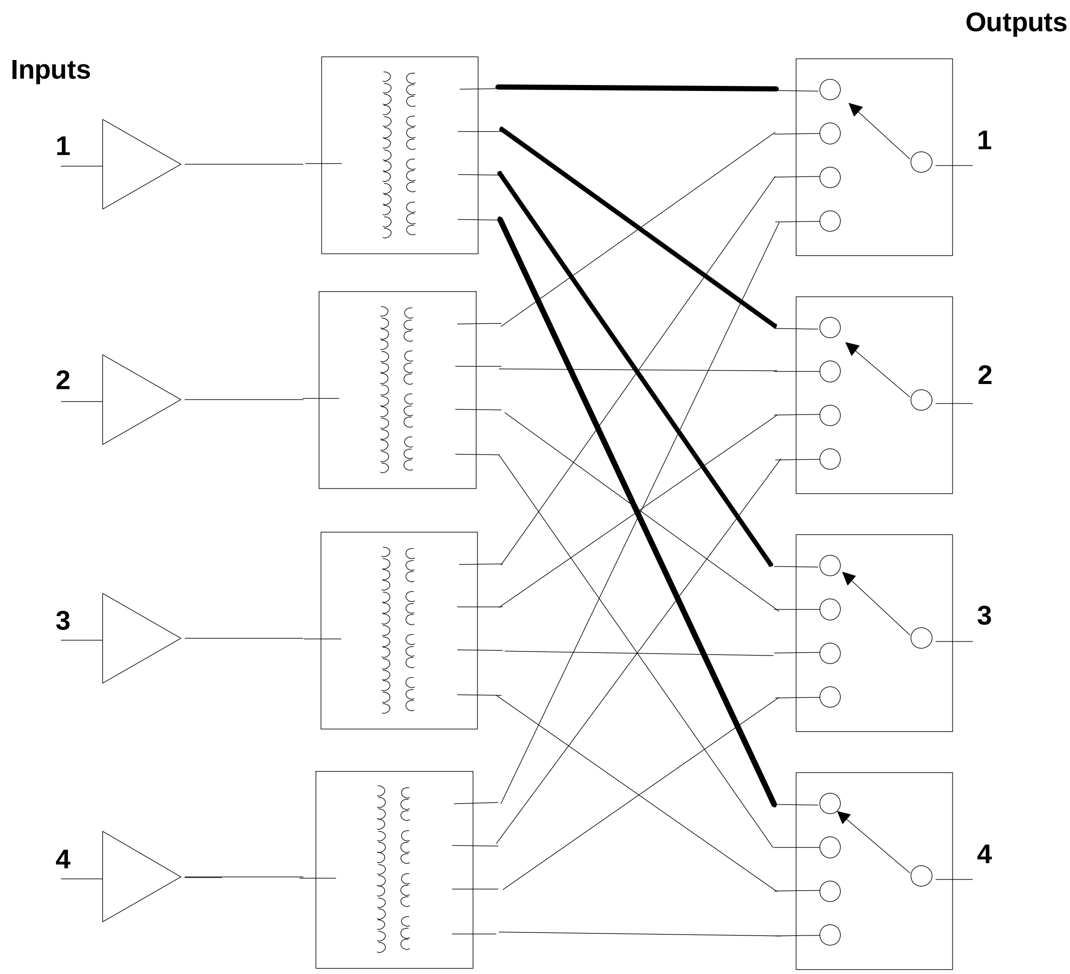

Figure 4-1 : Amplified 4X4 Non Blocking Matrix

4.1 Signal Level, Gain and Dynamic Range

We have already mentioned that the power dividers introduce loss that must be counteracted for most applications. Large matrix installations will also have significant loss in the switches and the interconnecting cables. The typical switch matrix application requires unity gain (0 dB). Adjustable gain or attenuators may also be provided to allow flexibility. When the matrix is used in the receive side, the system noise figure is typically fixed before the signal arrives at the switch matrix. So the matrix does not usually need an extremely low NF. Values in the range 6-12 dB are often adequate.

As with any RF system, the gain and noise performance are tradeoffs against power handling and distortion. Increasing gain improves noise but degrades distortion and compression performance. Specifying intermodulation distortion and output compression points that are adequate for the task, but not over-specified will keep costs down.

Compression point and intermod distortion points are very dependent on amplifier transistor bias points. Increasing the compression point by 10 dB will raise the DC power of an amplifier by about a factor of 10. In a GaAs FET matrix, the amplifier power tends to be a large part of the total power consumption, since the GaAsFET's themselves operate with virtually no power consumption.

When possible, allow the signal to be attenuated through the matrix. This increases noise figure, but generally pays dividends in relaxing the IP3 requirements. And many systems can tolerate the noise figure increase without serious detriment to the signal to noise ratio presented to the demodulator.

4.2 Reliability

In military and commercial satellite applications, high-value signals are continuously passing through the switch matrix. Continuous operation is the norm. Therefore, reliability is a very important criterion. When the matrix is used for maintenance and test, some reduction in reliability may be tolerable to reduce cost.

4.2.1 Power Supply Reliability

The most common feature to improve system reliability is redundant DC power supplies. The failure rate of the AC/DC power supply tends to be one of the largest contributors to final MTBF. A typical AC/DC supply can have a specified MTBF of 100,000 hours. By using redundancy, that MTBF value is effectively squared, which tends to remove the effect of a power supply failure from the MTBF computation.

One important aspect of redundancy is to ensure that the circuit elements that sum the power supply outputs are either replaceable or testable. In a typical low cost scheme, series diodes are used in the outputs of each supply. If the series diode fails open, that supply cannot provide any current when the other supply dies. But monitoring that condition in a live circuit is difficult. In cases where MOSFET’s provide this function, the situation is similar. When swapping out a failed power supply, the best systems will also swap the summing elements, to ensure that all the potential failures are removed. Parallel diodes are another technique to address some of the risk of diode failure.

The ease of swapping the supply is discussed below under Maintainability

4.2.2 Switching Path Reliability

A single level non blocking switch has no redundant switching paths. A three level switch will have some amount of redundancy, depending on the actual design. Having alternate paths during failures can protect critical outputs. But the firmware complexity grows to accommodate this fact. It also means that either the three level switch must determine the path failure on its own (Built In Test Equipment - BITE) or there must be a manual method of selecting alternate paths when a failure is spotted by external equipment.

The additional BITE features increase the cost of the three level switch. BITE to test paths typically involves signal generation, which is a potential interferer for adjacent channels. Power sensors are required for all inputs and outputs. Additional series switches are included to keep the internally generated signals from appearing on inputs or outputs during testing.

4.3 Environmental

Most applications are benign rack mounting situations, with good forced air in the racks. However, sometimes the switch matrix must be transportable or even mobile. In those cases, the environment needs to be considered.

4.3.1 Cooling

Cooling requirements increase with the density of the packaging. Most matrices have forced air internally. As more crosspoints are packaged per PCB, heat density rises. With GaAs FETs, this is usually not significant. The amplifiers are the primary heat source in a GaAs FET switch.

As third order intercept and 1 dB compression requirements rise, so does power consumption. In transmit applications, where the signal levels may already be fairly high when they arrive at the matrix, consider external amplifiers after the matrix. This will move heat out of the matrix and probably ease the maintenance situation. Reduce the input level at the matrix to achieve real cost savings and size reduction.

4.3.2 Shock and Vibration

A mobile matrix will usually require solid state devices. Mechanical relays are susceptible to vibration and shock. Magnetic coaxial relays are especially vulnerable, as there is no holding current to keep the contacts closed. The contacts are bi-stable, with magnets used to hold the contacts in their various states. If the contact bounces free, it could become latched to an alternate position. When magnetic relays are used, auxiliary contacts are always required. Sensing circuitry or software must monitor the auxiliaries to make sure they remain in the required position.

4.4 Control

There are several popular control methods in the industry. The most common physical interface is asynchronous serial communication, such as RS-232, RS-422 or RS-485. Ethernet control is also popular. For systems with extreme speed or synchronization requirements, a customized parallel interface may be required.

Asynchronous serial interfaces are ubiquitous. They are cheap and conversions readily exist between the various formats listed above. Virtually every microprocessor that a vendor would like to use will support this interface. The personal computer industry has abandoned asynchronous serial interfaces, in favor of USB. But USB has not appeared as an industry standard on rack mount instruments. The destinations are typically too far from the computer to make this interface feasible.

Ethernet is the interface that is taking over in rack installations. The costs of internal hardware and software have continued to drop so that most vendors will now provide Ethernet interfaces. Although the internal microprocessor must be more sophisticated, the ease of transferring data makes Ethernet a preferred interface.

There is wide variation in the control speed of a matrix. Ethernet has the potential to get the data to the matrix in much less time than asynchronous serial data, but that transfer time is typically not included in the control time specification. And Ethernet is non-deterministic - it is not possible to predict when a message will arrive at its destination on a busy network. Once the message arrives, there is a processing delay before that command gets executed and the RF switch path is changed. The total delay depends on internal software as well as the internal signaling methods used within the matrix.

CrossPoint matrices switch in under 250 msec from the time a full command is received (125 msec typ). The GaAsFETs switch in microseconds, but the data must arrive at the FETs before that can happen. There is a tear-down operation that removes any existing path to the specified output, then a connection operation that makes the new path available. Removing the old path is necessary to ensure maximum isolation (minimum crosstalk). The larger the matrix, the longer the entire process may take.

4.5 Maintainability

Electronic systems have become extremely reliable, but a failure will eventually occur. In an environment where the device is running 24/7 and processing valuable signals, the ability to diagnose and repair the matrix becomes very important.

4.5.1 Built In Test (BIT)

The first line of defense is the ability to determine that failure has occurred, or is immanent. In the discussion of the 3 level switch, we pointed out that redundant paths are only useful if there is some way to determine that you need to use one. Many single level matrices have built in test capability of some sort. The amount of BIT included must be weighed against the cost of not having it at all. In many cases, a combination of manual and automatic test is the most cost effective. If failure probability is low, and repair time is low, it might be simpler to use a non-automated method to find and fix signal path problems.

Testing the non-signal components such as power supplies, digital controllers, and the like, is commonly done. These kinds of BIT are not costly, and should be present as a matter of course. Power monitoring in a system without redundant supplies is vulnerable to losing the very power that is required to report the failure. Power supply hold-up time must be sufficient to get a message out before the processor stops operating. When redundant supplies are installed, replacing a failed supply can be scheduled based on the risks of a second power supply failure while the first supply remains down. Control system failures will require immediate attention, unless the switch configuration is completely static and remains that way during the failure period.

Switch matrices, particularly single level matrices, are very simple architecturally. No matter how many inputs it contains, all inputs circuitry is identical. The same is true for the outputs. In many installations, having a spare input that connects to a signal generator, and a spare output that connects to a spectrum analyzer or power meter, will be sufficient to allow fast troubleshooting. When a signal fails to appear at the downstream equipment, the test generator can be connected to that questionable output to verify the output circuitry. The questionable input can be connected to the spectrum analyzer to verify the upstream signal is received and the input circuitry is operating properly. One of those two connections will demonstrate the failure of the surrounding system.

To increase the automation of the troubleshooting, RF BIT systems can be included. Power detectors at inputs and outputs can validate incoming signals from upstream equipment, and can verify successful routing to the intended output. To localize the problem, however, there must be additional resources beside power detectors. The paths must be sub-divided in some manner to allow portions of the path to be verified while other portions are removed from consideration. This may mean spare channels dedicated to BIT, it could mean transfer switches at strategic locations, or other techniques. Once the software has localized the problem, it can notify the system operator and identify the parts required to repair the problem.

The simple approach is to have all this software as an offline diagnostic. The BIT software is given complete control of the switch matrix, and no user signals are being passed through. Diagnostics become a scheduled event, during matrix downtime. Investing in this kind of diagnostic that can only operate when the switch matrix is in a maintenance mode is often not cost effective. Once the matrix is offline, verification of all RF paths in a semi-automatic manner is straightforward and easily performed by a technician. Verifying a 32x32 matrix manually takes under an hour for all crosspoints, with a network analyzer or signal generator and spectrum analyzer. Once the switch matrix is offline, the major inconvenience has already occurred – the end user's signals have been disconnected.

To avoid the downtime, customers may specify real time signal path BIT. In this situation, any unused paths are being monitored on a periodic or continuous basis. Any new connection path is verified before confirmation is reported. If the system detects a failure when establishing a new path, the system operator is notified. Notice that this form of BIT is not monitoring established paths with its own test signal. Once the user's signal is flowing through the switch matrix, no test signals can be added to that path. Using power monitors at input and output allows verification that the channel gain is as expected. If the output power monitor shows loss of signal, while the input monitor shows power arriving, a fault can be declared. But to perform further fault localization means disconnecting the user signal at input and output. Here again, the BIT software can only notify of a fault after the loss of signal. So the inconvenience to the user has already occurred. Only insofar as the automated BIT speeds the troubleshooting cycle can it be considered valuable. It is not valuable in preventing faults, or routing around them (except in a 3 level switch matrix).

When specifying BIT, be sure that the entire maintenance philosophy and process is considered. Then determine if the expense of the additional BIT will materially improve performance.

4.5.2 Modularity

How a matrix is physically constructed plays a big part in its maintainability and overall cost. Matrix designers have long recognized that matrices represent a 3-dimensional interconnect problem. And every matrix vendor has built versions that are physical row and column architectures. An internal photo of the CrossPoint 16x16 L band matrix shows the 16 to 1 output switches running vertically. Underneath, running horizontally, are 1 to 16 power dividers. By this plug-in architecture, all the interconnects can be made without cables.

Figure 4-2 : 16x16 Switch Matrix Internal View

As parts continue to shrink, matrices are appearing with planar construction. This approach puts power dividers and switches on the same PCB. The overall size of the matrix shrinks, but the increment of modularity is larger. When a failure occurs, the repair does not affect just the failed channel. It affects a group of channels. And although the failed card may be hot swappable, its replacement will interrupt multiple channels.

In a 3 level matrix, a similar problem will exist. Replacing modules usually requires replacing a complete sub-matrix. Therefore, more than one input or output is affected. Replacing the mid matrix can theoretically interrupt every output, depending on the state of the switch.

For many operators, these modularity considerations will require scheduled downtime (or scheduled replacement) to minimize the impact on end users. If the application requires repairs to be isolated to only the failed channels, planar designs will not meet the objective.

4.5.3 Hot Swapping

Hot swappable (redundant) power supplies are available from many vendors. There is no system impact to a failed supply and no impact from a replacement. Hot swapping RF assemblies is often required for matrices that operate round the clock.

The system should identify where the failure has occurred (LRU), to a very high percentage of accuracy. No BIT system can be 100% accurate in its diagnosis. The suspect part should be removable without turning off power to any other elements, removing obstructions, etc. Front panel access is ideal, but not always practical. CrossPoint mounts matrix chasses on slides to allow access after removing covers.

Once replaced, the system may accept the new part automatically. Other times a command is required to initiate it.

4.6 Expansion

Expandability is one of the most requested and least used features in specifying a switch matrix. Often the anticipated expansion becomes a major cost driver. No one wants to buy a matrix and then throw large parts of it away when expansion is required. But it is a rare installation that actually uses the expandability that they envisioned. By the time expansion is required, the entire facility is often ready for major upgrades, and the existing matrix may not meet the new performance objectives. Before specifying expansion, determine if a small number of spare channels can meet the likely needs of the future, rather than provisioning to enlarge the entire configuration at a later date.

If significant expansion is required, consider that expansion of outputs requires wider power division and more gain in the front end. Adding more inputs requires wider selection switches at the outputs, but probably no more gain. Of the two, increasing outputs tends to be more costly, because of its impact on the input structure.

If the matrix uses a plug-in architecture, leaving unpopulated slots for future expansion is the best choice. However, this means that you are often purchasing the larger configuration, with certain components deleted. The cost will probably be closer to the larger matrix than the smaller one you could have started with.

In the case of major expansion (doubling inputs or outputs, for instance), there is significant additional expense. An NxM matrix grows geometrically. Taking a 16x16 matrix to a 32x32 will require 4 times as many crosspoints.

One typical way of approaching this expansion is shown in the figure below. The original 16x16 matrix is shaded. All the rest of the equipment is new. In addition to the three new 16x16 switch matrices, there are multicouplers on the front. A multicoupler is an amplified power divider. Here the power division is only a factor of 2, but gain is required unless the 16x16 matrix has additional gain available. Gain is also required because of the new losses in all the cables and the output combiner switches. On the output side, there are two way selector switches to give each output access to all 32 inputs. If the 16x16 matrix is rack mounted, the total rack space grows more than 4x.

RF performance will likely be hurt by this expansion, unless changes are made to the amplifiers inside the 16x16 matrix. This is because the new multicoupler amplifier provides more gain ahead of the original amplifier, raising the equivalent input level at that point. (because some excess gain is required for the additional cable and output losses). The resultant third order intercept points degrade fairly rapidly unless amplifiers are carefully matched. So if the same RF performance is required for the 32x32, as for the 16x16, modifications to the original matrix will probably be required.

But if the 32x32 chassis were purchased, with only 16 channels installed, there would be no additional rack commitments, no external cables, and the RF performance would not be degraded by the addition of a second amplifier. The multicoupler and output combiners would have been accommodated in the chassis, without the expense of two (or more) additional chasses. In the case of CrossPoint products, sixteen of the 32 way power dividers would be installed, with 16 of their outputs terminated until those channels were populated. The other 16 power dividers could be left out until required. Although this approach is considerably more expensive than a simpler 16x16, if the probability of expansion is real, the eventual cost savings and size reduction would be enormous.

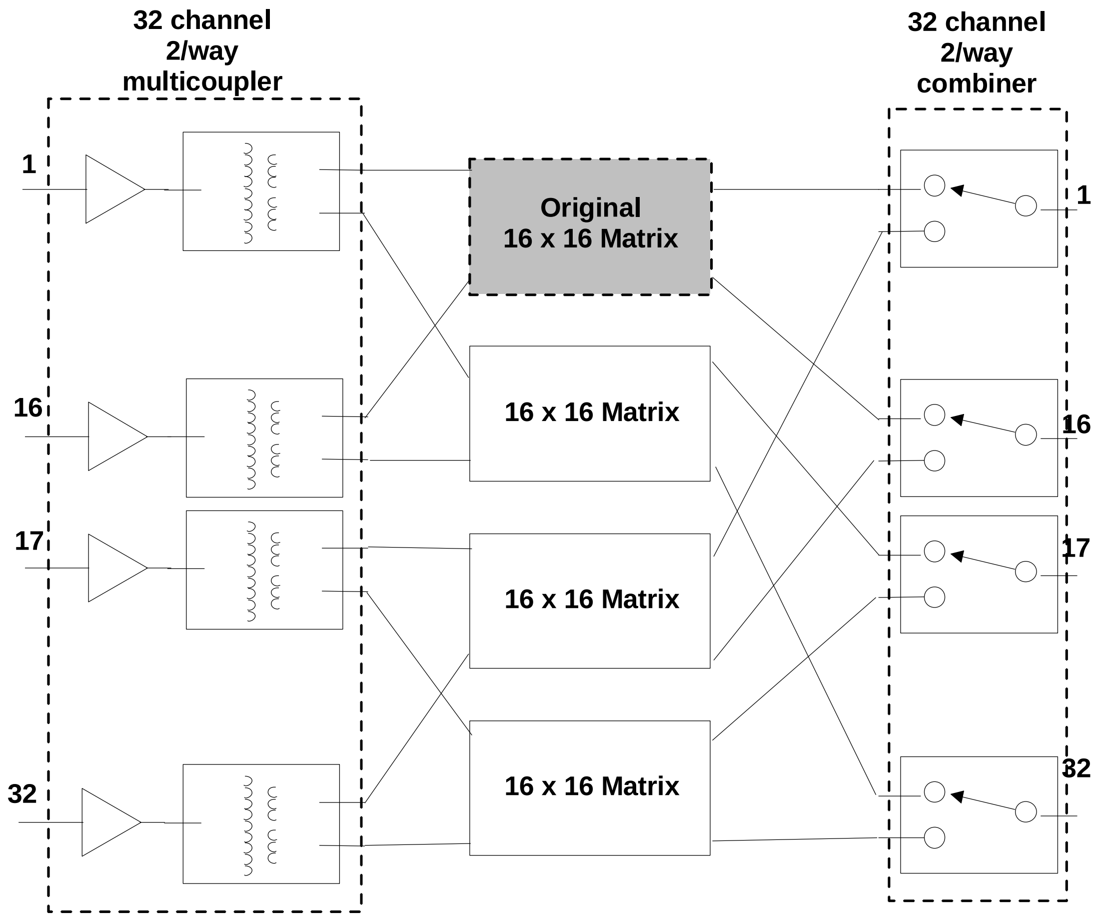

Figure 4-3 : Matrix Expansion from 16x16 to 32x32

5 Glossary

Blocking – Use the word to indicate there are ways the switch matrix is being used that does not require all the ports to be available all the time. An example of a blocking matrix is shown in 2.2. It is a switch matrix with 4 inputs and 4 outputs, but only enough hardware to support 1 path at any given time. Conversely a switch matrix with 4 inputs, 4 outputs, and hardware supporting 4 paths at any given time can be called non-blocking.

Conference Matrix – This is used to describe the capability in a unidirectional switch matrix to provide both fan-in and fan-out capability in one switch matrix. Note, while it is possible to create a passive switch matrix to use this feature bidirectionally, its usefulness is dependent on overcoming the signal loss.

Crossbar – This term is dated, however is still used in some circles. It is carried over from telephone company origins. Schematically this was a switch matrix that can connect any one path from any input port to any output port. It can be thought of as two perpendicular bars with a crosspoint at the intersection, however this literal topology is not adequate to handle many common RF frequencies of today because of the inherent stubs of the design. Nevertheless, the term is still used to describe a non-blocking style matrix because it is the schematic equivalent.

Fan-In – This is used to describe the capabilitiy in a unidirectional switch matrix to connect multiple input ports to a single output port.. This is useful in transmit (uplink) applications where multiple signals can be sent to one antenna. Note, while it is possible to create a passive switch matrix to use this feature bidirectionally (as fan-in, or fan-out), its usefulness is dependent on overcoming the signal loss.

Fan-Out – This is used to describe the capability in a unidirectional switch matrix to connect multiple output port to a single input port. This is useful in receive (downlink) applications where multiple receivers can be attached to one antenna feed. Note, while it is possible to create a passive switch matrix to use this feature bidirectionally (as fan-in, or fan-out), its usefulness is dependent on overcoming the signal loss.

Non-Blocking – Some switch matrixes are configured such that certain paths may be unusable given specific crosspoint configurations. A non-blocking switch matrix describes the ability to switch a path from a given input port to a given output port (for example input 1 to output 1), while not blocking another path whose ports are different (for example input 2 to output 2). In this example a non-blocking switch matrix can connect 1 to 1 and 2 to 2 simultaneously. In a blocking switch if 1 to 1 is connected, it is possible 2 to 2 may be “blocked” due to the switch topology.

Path – This describes the signal connection through a crosspoint from an input port to an output port.

6 References

V.E. Benes, *Mathematical Theory of Connecting Networks and Telephone Traffic*, 1965, Academic Press. A foundational work on N-level switches and non-blocking criteria

W.E. Doherty, Jr, R.D. Joos, *The PIN Diode Designer’s Handbook*, 1998, MicroSemi Engineering Department. Available from www.microsemi.com

CrossPoint Technologies
A Division of I.F. Engineering Corp