ZLAN-111

Applications of the ZL50073 Designing Large Switching Matrices

Application Note

Contents

July 2004

1.0 Summary

This Application Note describes how to create large TDM switching matrices using Zarlink's 32 K TDM switches, one of the highest-density TDM switches in the market.

Zarlink has a few different flavors of 32 K TDM switches to meet customers' various requirements. In this application note, the most feature-rich device, ZL50073, will be used as an example. The other 32 K switches can be used similarly.

Two methods of building large switching matrices are covered in this application note - the square matrices and the three-stage matrices. They each have their own applications and design considerations, which will be discussed in details in this document. In general, the square configuration finds more applications in lower-capacity matrices, in which the matrices' capacities are no more than 4 times the capacity of the element switching device. In higher-capacity matrices, the three-stage configuration is used more often.

2.0 Introduction

There are various switching system architectures. For TDM systems, distributed and centralized architectures are most commonly used.

In a distributed switching system, the switching functionality is distributed to multiple logical points in the system. In most cases, these logical points are circuit boards with different functionalities. They are connected to each other, either directly or indirectly, within the system. They each have their own switching device(s) to control how they send and receive the traffic. This architecture is more flexible, scalable, and less vulnerable to single point of failure, but it requires more complex control software and has overall inferior switching performance (more delay, higher blocking probability, etc.).

In a centralized switching system, the switching functionality is centralized onto one logical point, which is usually a circuit board with a single switching device or a matrix of multiple switching devices. This architecture is very easy to control and have in general better switching performance. It is, however, less flexible, less scalable and more vulnerable to single point of failure.

In centralized switching systems, the best practice is to choose a single switching device that can fulfill the capacity requirement of the system. This will keep the cost and control complexity to a minimum. In cases where very large switching capacity is required, multiple switching devices may need to be used. These devices must be put in proper "matrix" configurations to minimize cost and maximize performance. Two of the most efficient configurations are the square matrices and the three-stage matrices

The ZL50073 is a non-blocking large TDM digital switch with 32,768 x 32,768 channel switching capacity operating at 16.384 Mbps, 32.768 Mbps or 65.536 Mbps. Its capacity reduces to 16,384 x 16,384 channel when operating at 8.192 Mbps. It has up to 128 input and 128 output streams. Data throughput delay is only 1 frame in variable delay mode and 2 frames in constant delay mode. With the high switching capacity, large number of input/output streams and low throughput delay, it is an ideal building block for the construction of cost effective and high-density switching matrices.

When designing a switching matrix, the followings are the most important considerations:

1. Size and cost of the matrix
2. Blocking probability
3. Control software complexity
4. Data throughput delay

In the following sections, the two methods of designing large switching matrices, the square matrices and the three-stage matrices, will be outlined, with examples shown using the ZL50073 as the building block. Their respective benefits and disadvantages will be analyzed along the above mentioned factors.

3.0 Definitions

Blocking and non-blocking are terms used to describe the ability of networks to establish new connections. The definitions of blocking, non-blocking and some other related terms are given in this section.

3.1 Blocking Configurations

A network is considered blocking if it supports a connection set that will prevent some additional connections from being established between idle ports, even with the rearrangement of the existing connections. Figure 1 illustrates a blocking network with 4 input streams running at 8 Mbps (512 channels) and only 2 output streams running at 8 Mbps (256 channels). Only 256 channels from the input streams can be switched to the output.

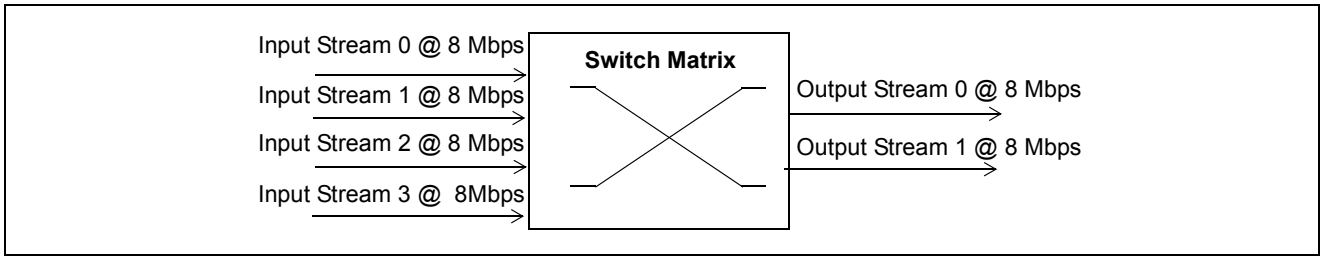

Figure 1 - Blocking Switch Matrix

3.2 Non-Blocking Configuration

There are a few different types of non-blocking network. Two of them will be discussed here:

1. A network is considered *strictly non-blocking* if any connection between two idle ports can be established immediately without interference to any arbitrary existing connections. In other words, the new connection can be established even if rearrangement of the existing connections is not allowed.
2. A network is considered *rearrangeably non-blocking* if any connection between two idle ports may be temporarily blocked, but can be established if one or more existing connections are rerouted or rearranged. In other words, it is always possible to add a new connection linking two idle ports by rearranging the existing connections.

In the rest of this application note, the term non-blocking will refer to strictly non-blocking unless otherwise stated.

3.3 Blocking Probability

When a connection needs to be established between an idle input channel and an idle output channel, the blocking probability is the probability that there is no path available in the switching network to route this connection.

3.4 Traffic Density

Traffic density is the average proportion of channels in use out of all the available channels at any given time during normal operation of a switching network.

4.0 Square Switching Matrix

Square switching matrices are mostly used if the size of the matrix is no more than 4 times the size of the fundamental switching element. They are strictly non-blocking, provided that the building blocks are also strictly non-blocking, which is the case for the ZL50073. Figure 2 illustrates the construction of a 64 K x 64 K channel switch by putting four ZL50073s in a square matrix configuration. For this configuration to work, the switching elements must have the ability to tri-state the output on a per-channel basis. This is because all the output streams are connected to more than one switch. For all the switches connected to the same stream, there should only be one switch driving a particular channel timeslot, while the other switches tristate their outputs during that timeslot. Otherwise, bus contention will happen, which may cause data corruption and even damages to the devices.

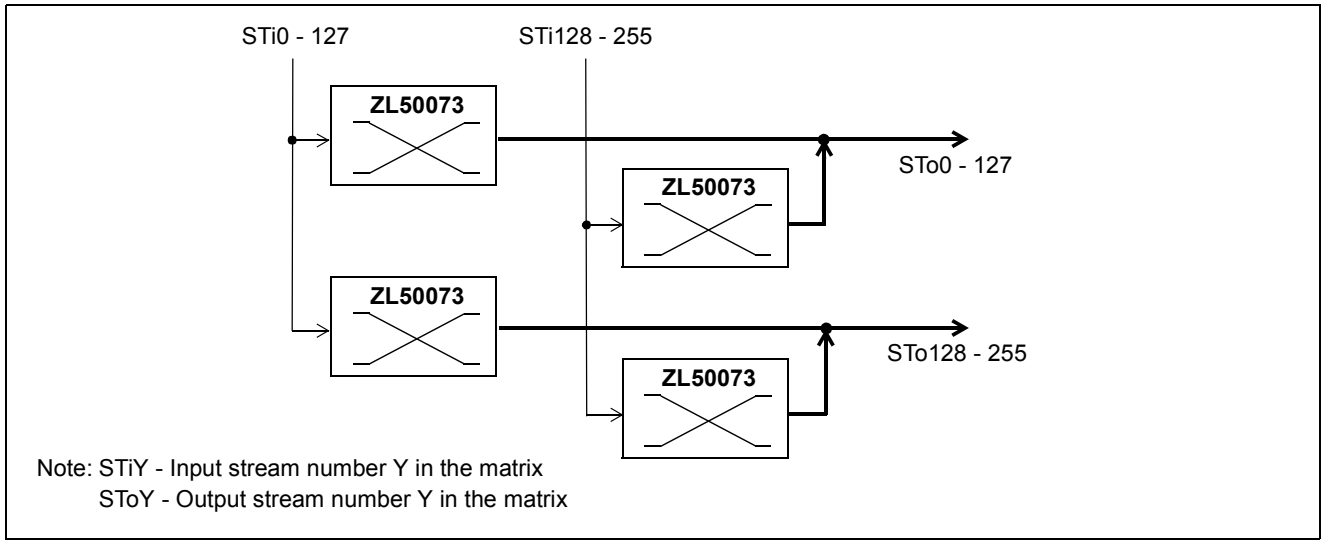

Figure 2 - 64 K x 64 K Non-Blocking Square Switching Matrix

The number of switching elements required can be calculated as follows:

$$\text{Total Number of Switching Elements} = \left(\frac{\text{Total channel capacity of desired switch}}{\text{Switch element capacity}} \right)^2$$

The disadvantage of this configuration is the large number of switches required as the switching capacity grows, as shown in Table 1. When the capacity doubles (e.g., from 64 K x 64 K to 128 K x 128 K), the number of switching element quadruples. A 128 K x 128 K channel switch constructed using square matrix configuration is illustrated in Figure 3. This matrix requires 16 devices instead of 4 in the 64 K x 64 K matrix. Similarly, a 256 K x 256 K square matrix will require 64 devices. The matrix then becomes very expensive and complicated to implement.

Matrix Capacity	Number of 32 K Devices in a Square Matrix Configuration
64 K x 64 K	4
128 K x 128 K	16
256 K x 256 K	64
512 K x 512 K	256

Table 1 - Size of Square Switching Matrix

The main advantage of the square matrix configuration is the simplicity of the control software. In this configuration, there is only one path between an input stream and an output stream. In other words, only one switch element can make a connection between a particular input stream and a particular output stream. The control software can therefore determine the connections very easily. It only needs to make sure no more than one switch element drives the same channel timeslot.

Another advantage is the strictly non-blocking nature. In a square matrix configuration, connections can always be made from any idle input channel to any idle output channel without interference to any existing connection.

A square matrix does not introduce extra data throughput delay. No matter how big the matrix is, data always experience the delay of only one switching element.

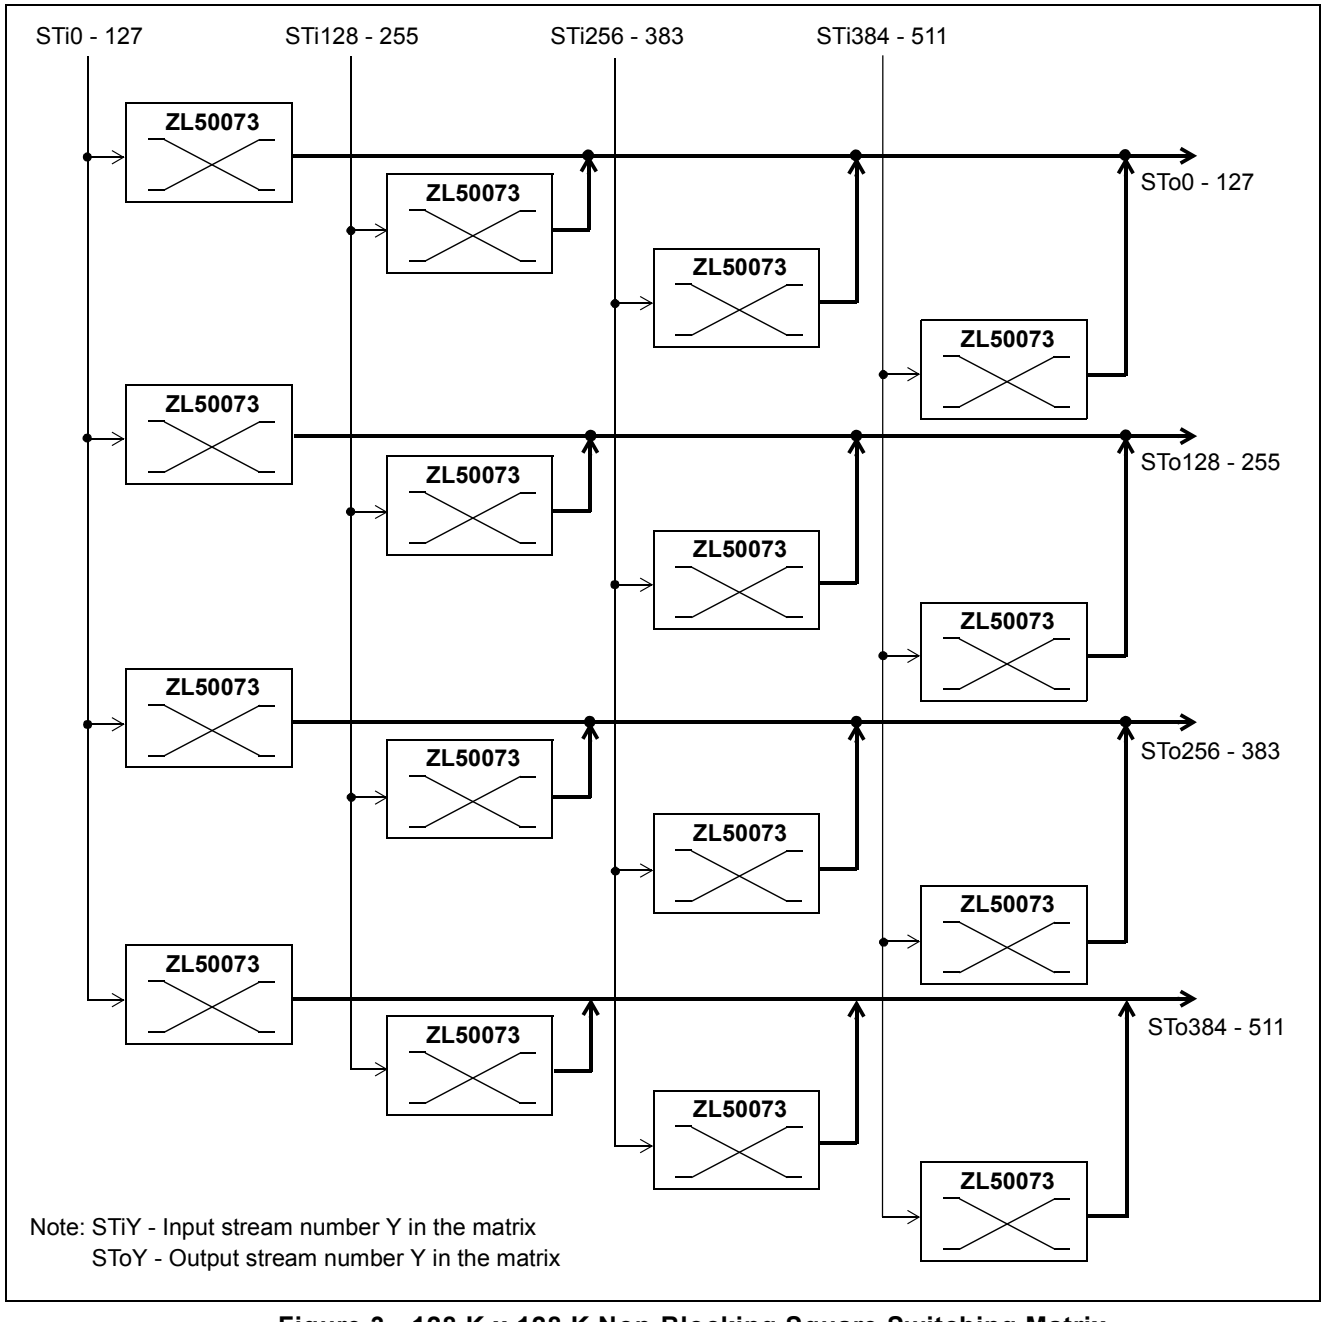

Figure 3 - 128 K x 128 K Non-Blocking Square Switching Matrix

5.0 Three-Stage Switching Matrix

Any square matrix that is more than four times the size of the switch element becomes too large to be practical (e.g., 512 K x 512 K requires 256 ZL50073s). To reduce the number of switching elements, a three-stage switching matrix can be implemented. Figure 4 shows a typical three-stage switching structure.

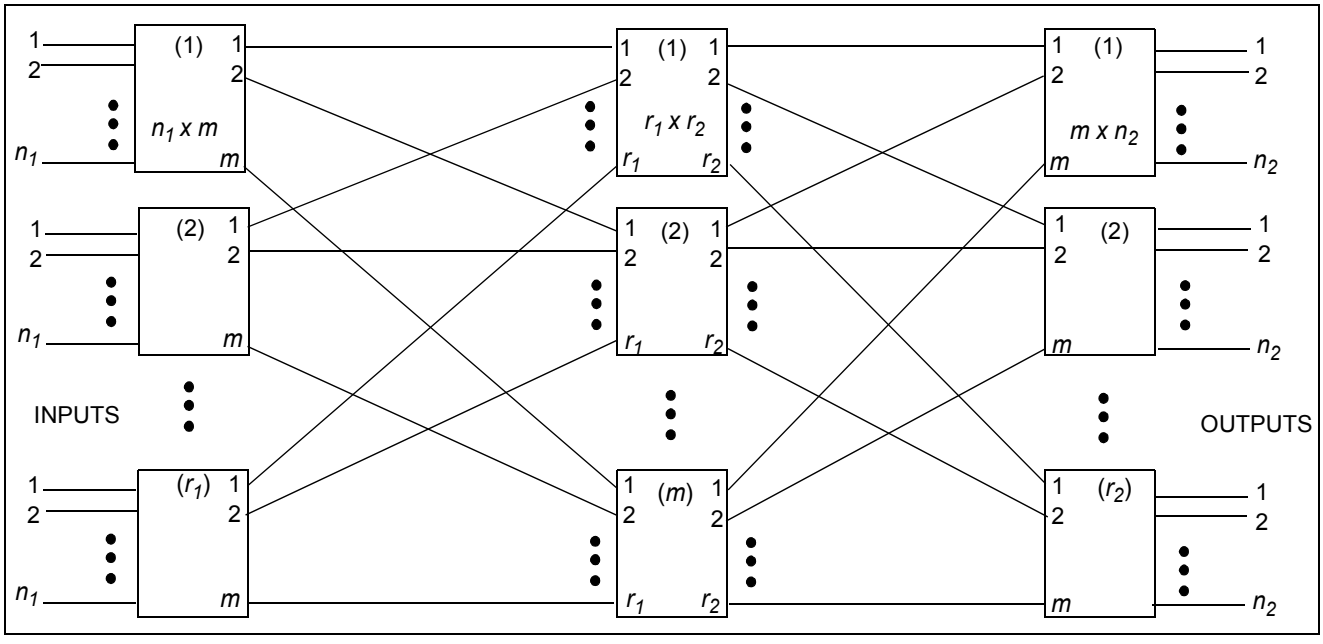

Figure 4 - General Structure of a Three-Stage Switching Matrix

In a symmetrical matrix where $r_1 = m = r_2$, the number of switching elements required can be calculated as follows:

$$\text{Number of Switching Elements per Stage} = \frac{\text{Total channel capacity of desired switch}}{\text{Switch element capacity}}$$

$$\text{Total Number of Switching Elements} = \frac{\text{Total channel capacity of desired switch}}{\text{Switch element capacity}} \times 3$$

The number of elements per stage is limited to the number of streams. The maximum number of elements per stage using the ZL50073 is 128 and a total switch count of 384. This would result in a 4 M x 4 M channel switch. Figure 5 shows the implementation of a 256 K x 256 K three-stage matrix using the ZL50073s. The benefit of this type of switching matrix is the small number of switches required as compared to the square matrix. A few matrix sizes are shown in Table 2.

Matrix Capacity	Number of 32 K Devices in a Three-Stage Matrix Configuration
64 K x 64 K	N/A
128 K x 128 K	12
256 K x 256 K	24
512 K x 512 K	48

Table 2 - Size of Three-Stage Switching Matrix

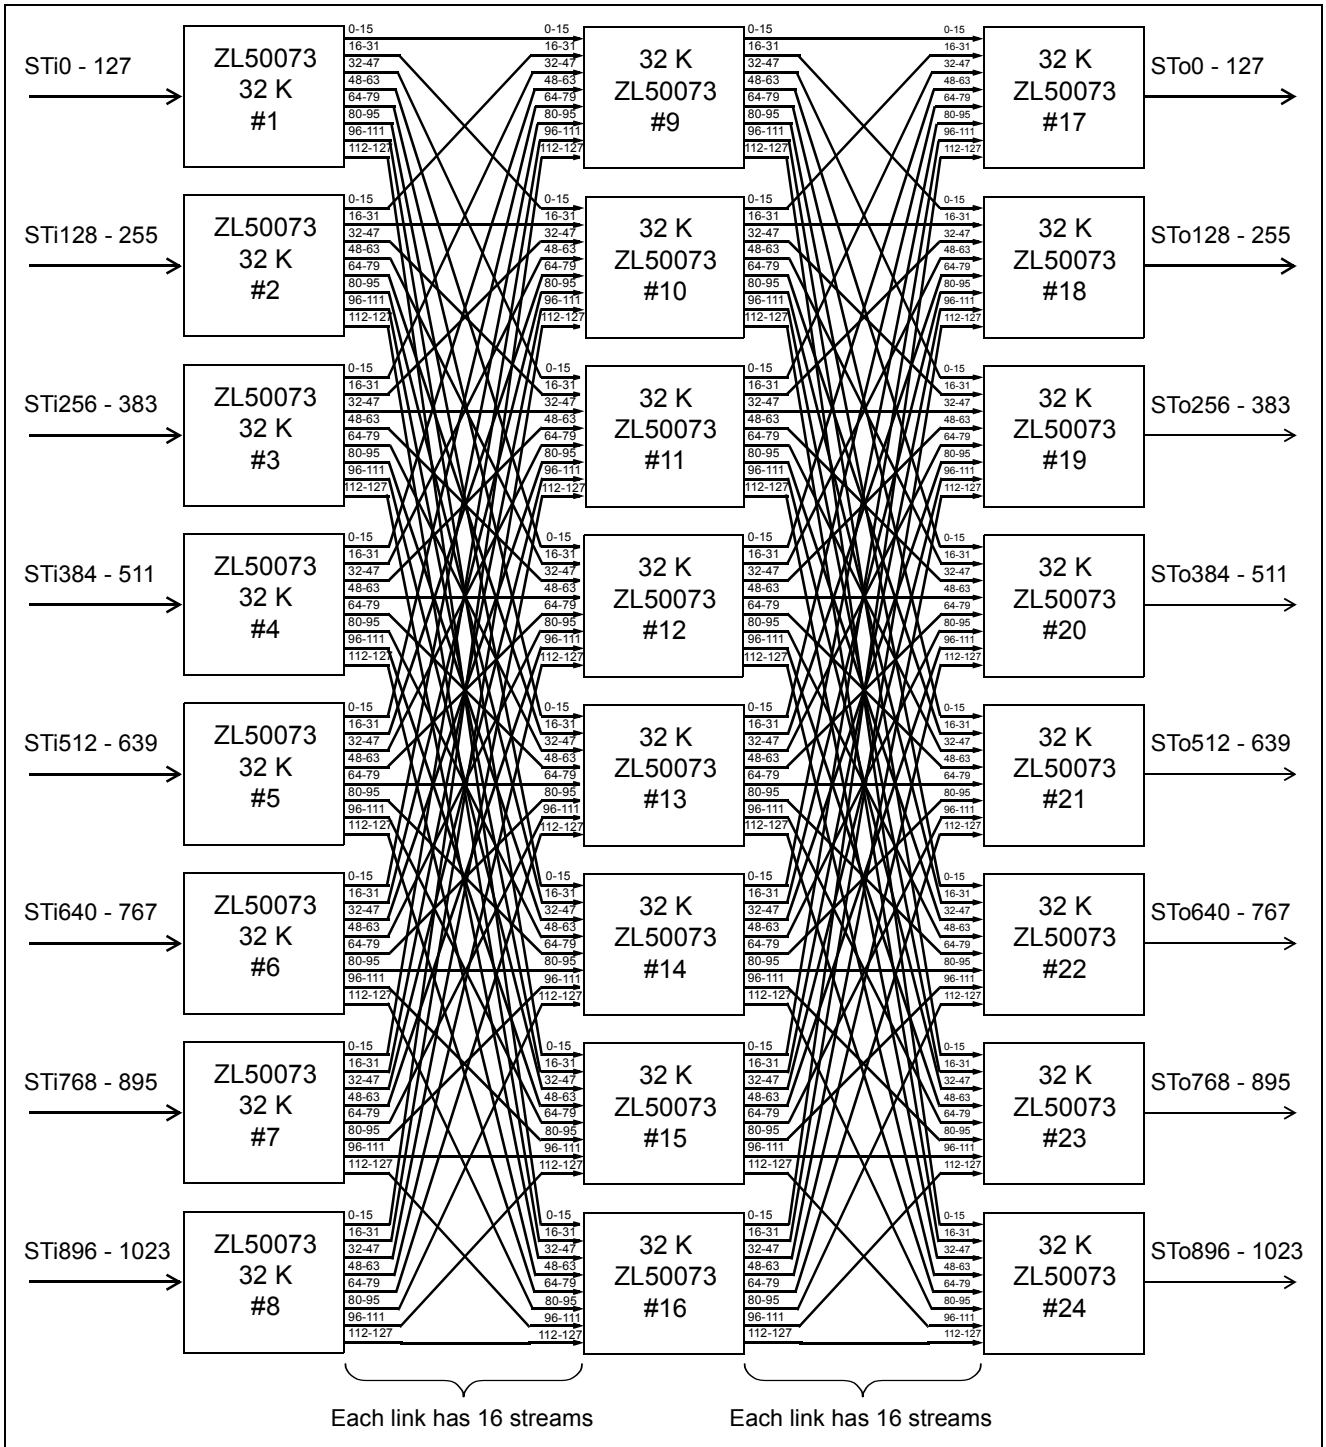

Note: All streams are running at 16 Mbps with 256 channels per stream
STiY - Input stream number Y in the matrix
SToY - Output stream number Y in the matrix

Figure 5 - 256 K x 256 K Three-Stage Matrix

A three-stage matrix is rearrangeably non-blocking, but not strictly non-blocking. However, even if rearrangement of the existing connections is not allowed, the blocking probability is very low. Blocking probability is a function of the traffic density in a three-stage matrix. It can be approximated using a statistical model with the assumption that traffic is evenly distributed and no multicasting is used. Table 3 shows the blocking probability of the 256 K x 256 K matrix shown in Figure 5 for a few different traffic density values. For traffic density below 0.9995, the blocking probability is well below 1.0×10^{-6}.

Number of Idle Channels	262	157	131	26	0
Traffic Density	0.999	0.9994	0.9995	0.9999	1
Blocking Probability	5.19×10^{-14}	5.11×10^{-8}	1.47×10^{-6}	0.187	0.882

Table 3 - Blocking Probability of a 256 K x 256 K Three-Stage Matrix Constructed from ZL50073s

Other than the finite amount of blocking probability, the three-stage matrix also has the disadvantage of requiring more complex control software. Unlike the square matrix, a three-stage matrix has many possible paths to route an input channel to an output channel. When the traffic density is high, the software may have to go through many iterations to find an available path. It may also need to rearrange the network to allow all traffic to go through.

Another disadvantage is the larger delay through the matrix. All the traffic has to go through 3 switches in a three-stage matrix instead of 1 switch in a square matrix.

6.0 Conclusions

Two methods of constructing large TDM switching matrices from smaller switching elements were discussed. Examples were shown using the ZL50073 32 K switches. Both methods have their pros and cons. Table 4 is a comparison between the two methods.

		Square Matrix	**Three-Stage Matrix**
Matrix Size (Number of 32K Switches)	64 K x 64 K	4	N/A
	128 K x 128 K	16	12
	256 K x 256 K	64	24
	512 K x 512 K	256	48
Blocking Probability		zero	very small but finite
Software Complexity		low	high
Data Delay		1 switch	3 switches

Table 4 - Comparison between the Square Matrix and the Three-Stage Matrix

The square matrix has all the advantages except for its high switch element count when the matrix capacity becomes large. In general, when the matrix capacity is no more than 4 times the switching element's capacity, the square matrix is recommended. If larger capacity is required, the three-stage matrix becomes more practical.

In designing a three-stage matrix, considerations must be taken to reduce the blocking probability, for example by limiting the traffic density below a certain level. It is also important to allocate enough software resources to control the matrix. The extra data delay introduced should also be taken into account.

For more information about all Zarlink products
visit our Web Site at

www.zarlink.com

TECHNICAL DOCUMENTATION - NOT FOR RESALE