

Bandwidth Calculation
for Pico Tetra and Pico Tymo
Documentation Addendum

January 5, 2006

WARNING

EURESYS S.A. shall retain all rights, title and interest in the hardware or the software, documentation and trademarks of EURESYS S.A.

All the names of companies and products mentioned in the documentation may be the trademarks of their respective owners.

The licensing, use, leasing, loaning, translation, reproduction, copying or modification of the hardware or the software, marks or documentation of EURESYS S.A. contained in this documentation, is not allowed without prior notice.

EURESYS S.A. may modify the product specifications or change the information given in this documentation at any time, in its discretion, and without prior notice.

EURESYS S.A. shall not be liable for any loss of or damage to revenues, profits, goodwill, data, information systems or other special, incidental, indirect, consequential or punitive damages of any kind arising in connection with the use of the hardware or the software of EURESYS S.A. or resulting of omissions or errors in this documentation

Contents

1. Introduction	3
2. Acquisition modes	4
2.1. Streaming	4
2.2. Switching	4
<i>General</i>	4
<i>Definitions</i>	5
2.3. Acquisition performance	6
<i>Operation conditions</i>	6
<i>Streaming mode</i>	7
<i>Switching mode</i>	7
3. On-board data transfer validation	9
3.1. Throughput evaluation	9
3.2. Bandwidth evaluation	10
3.3. Bandwidth constraint	10
3.4. Working around the bandwidth constraint	11
3.5. Examples	11
4. Out-board data transfer evaluation	11
4.1. Combined throughput	11
4.2. Practical PCI bandwidth	12
4.3. Available PCI bandwidth	12
4.4. PCI constraint	12
4.5. Chip-set architecture	12
4.6. PCI latency effects	13
5. Assessing spoiled images	14
5.1. Spoiled images	14
5.2. Performance curves	14
<i>Piccolo Tetra</i>	14
<i>Piccolo Tetra – Synchronized video</i>	15
5.3. Conclusion	15

1. Introduction

Video surveillance systems generate large amount of data. These data, mainly images coming from surveillance cameras, are gathered in the frame grabber and sent for processing inside a PC through a PCI bus.

In systems with multiple cameras, the PCI bus can quickly become a transfer bottleneck. It is very important to correctly configure a video surveillance system in order to guarantee a smooth transfer of data throughout the PC.

This document exposes the different elements necessary either to size your system according to your requirement or to adapt your application to your system.

The following calculation techniques are applicable to any system but the figures given in this document are specific to Pico Tetra and Pico Tymo boards.

2. Acquisition modes

2.1. Streaming

A streaming application aims at recording a video sequence during a substantial amount of time.

The video flow representing this sequence is processed by one of the four on-board digitizers, which cannot serve another video signal in the same time.

For this kind of application, one video source is selected for the duration of the sequence recording. The selector in front of the digitizer can be seen as a facility to choose the video cable through which the video sequence is issued.

On Picolo Tetra, if only one video source is intended for streaming, no additional bracket is needed. When several video sources are successively chosen for streaming, additional brackets are needed.

Picolo Tymo is equipped with a high-density connector on its bracket and does not need additional brackets.

The time required for the digitizer to settle from one source to the next one is usually not an issue, as it is short compared to the sequence recording duration.

As long as the data transfer capability of the host PC is not exceeded, the streaming application will deliver to the system memory a regular video flow that can reach real-time speed and/or full resolution.

2.2. Switching

General

A switching application means that a single digitizer is due to capture images from more than one video source in a cyclic way. One image is captured from a source, then the digitizer quickly settles to another source, and another image is captured. The process is repeated until all active multiplexed sources are scanned.

For Picolo Tetra and Picolo Tymo, the switching process should be considered for one digitizer fed by 2, 3 or 4 inputs. The other digitizers can be used independently, either in a streaming or switching mode.

To characterize the switching performance of one Picolo Tetra or Picolo Tymo digitizer, the following reasonable assumptions are made:

- The inputs sourcing the digitizer assume the same video standard (PAL or NTSC).
- They all operate in field or frame mode.
- The parity is globally "locked" or "free" for all sourcing inputs.
- All sourcing inputs are intended for the same capture rate.

It may be useful to consider the switching mode as grabbing snapshots, or photos, from several cameras scanned in sequence as quickly as possible. By opposition, the streaming mode consists in shooting video from a single camera.

Definitions

Field or frame

The switching application should choose to operate in the field or frame mode.

The field mode means that the captured photos come from a single video field, 16.7 or 20 ms in duration according to the standard. The field photo can exhibit the full horizontal resolution, or be downscaled to a smaller width.

The following formats belong to the "field" category:

▼ FIELD IMAGES

	NTSC	PAL	Comment
Full resolution	640H x 240V	768H x 288V	
Half resolution	320H x 240V	384H x 288V	Square pixels
Broadcast resolution	720H x 244V	720H x 288V	
SIF	360H x 242V	360H x 288V	Square pixels

The frame mode means that two consecutive fields are processed to build the photo. A set of two fields is called a frame, and lasts 33.3 or 40 ms according to the standard. Several horizontal resolutions can be selected.

The following formats belong to the "frame" category:

▼ FRAME IMAGES

	NTSC	PAL	Comment
Full resolution	640H x 480V	768H x 576V	Square pixels
Broadcast resolution	720H x 488V	720H x 576V	

Parity locked or free

The switching application should choose to operate in the parity locked or free mode.

The parity lock mode means that all captured photos from a same camera will have a determined parity. In the field mode case, all photos will come from the odd field, or from the even field, according to a programmable selection. In the frame mode case, the application can select the odd-even or even-odd structure for all the captured photos from a specific camera.

The parity free mode means that the field or frame photos are acquired without synchronizing to the odd or even parity. This modes enables a slightly faster switching rate because there is no delay for parity synchronization.

However, because of the one-line gap between the odd and even fields, the parity free mode may result in a visual tiny up-and-down image jump. The parity lock feature removes this artifact.

Capture rate

When considering an individual camera, one photo from the video flow will be captured periodically and delivered to the PC memory.

The periodicity of this image capture process is the capture rate.

The higher the number of cameras involved in the scanning process sourcing the digitizer, the lower the capture rate.

In this paper, we consider the situation where the application does not set a special priority to any of the cameras feeding the digitizer. This means that the capture rate is the same for all involved cameras. In this case, the capture rate is also the period of the whole cycling process scanning the sourcing cameras.

Load factor

It is customary to quote the performance of a capture board in term of fps, or field per seconds. The field lasts 16.7 ms in NTSC and 20 ms in PAL.

For a given camera, the capture rate in fps is twice the above-defined capture rate when frames (pairs of fields) are acquired.

The load factor for the digitizer is the sum of the capture rates in fps of all switched cameras.

The load factor of the whole capture board is the sum of the load factors of all on-board digitizers.

The maximum load factor for a digitizer is achieved when it is operated in real time from a single source. It amounts to 60 fps in NTSC or 50 fps in PAL.

A capture board is sometimes characterized by the maximum load factor. For Picolo Tetra and Picolo Tymo, this figure amounts to 240 fps in NTSC or 200 fps in PAL. This performance is only achieved when all four digitizers are operated in the streaming mode. As soon as switching is involved, the performance is below this maximum value.

2.3. Acquisition performance

Operation conditions

The following tables state the expected performance for an ideal PC able to unconditionally sustain the data rate generated by the video capturing activity.

Actual PCs may exhibit uncontrolled internal bus traffic sporadically reducing the performance to a smaller figure.

The load factor for the whole Picolo Tetra or Picolo Tymo board is computed with the assumption that all four digitizers operate under identical conditions.

Streaming mode

▼ FIELD MODE

	Number of Sources per Digitizer	Capture Rate		Rate per Camera (fps)		Load Factor (fps)			
		PAL	NTSC	PAL	NTSC	per Digitizer		per Board	
						PAL	NTSC	PAL	NTSC
Parity Free	1	50	60	50	60	50	60	200	240
Parity Lock	1	25	30	25	30	25	30	100	120

▼ FRAME MODE

	Number of Sources per Digitizer	Capture Rate		Rate per Camera (fps)		Load Factor (fps)			
		PAL	NTSC	PAL	NTSC	per Digitizer		per Board	
						PAL	NTSC	PAL	NTSC
Parity Free	1	25	30	50	60	50	60	200	240
Parity Lock	1	25	30	50	60	50	60	200	240

Switching mode

▼ FIELD MODE

	Number of Sources per Digitizer	Capture Rate		Rate per Camera (fps)		Load Factor (fps)			
		PAL	NTSC	PAL	NTSC	per Digitizer		per Board	
						PAL	NTSC	PAL	NTSC
Parity Free	2	7.1	8.6	7.1	8.6	14.3	17.1	57	69
	3	5	6	5	6	15	18	60	72
	4	3.8	4.6	3.8	4.6	15.4	18.5	62	74
Parity Lock	2	6.3	7.5	6.3	7.5	12.5	15	50	60
	3	3.6	4.3	3.6	4.3	10.7	12.9	43	51
	4	2.8	3.3	2.8	3.3	11.1	13.3	44	53

▼ FRAME MODE

	Number of Sources per Digitizer	Capture Rate		Rate per Camera (fps)		Load Factor (fps)			
		PAL	NTSC	PAL	NTSC	per Digitizer		per Board	
						PAL	NTSC	PAL	NTSC
Parity Free	2	5.6	6.7	11.1	13.3	22.2	26.7	89	107
	3	3.8	4.6	7.7	9.2	23.1	27.7	92	111
	4	2.9	3.5	5.9	7.1	23.5	28.2	94	113
Parity Lock	2	5	6	10	12	20	24	80	96
	3	3.6	4.3	7.1	8.6	21.4	25.7	86	103
	4	2.8	3.3	5.6	6.7	22.2	26.7	89	107

3. On-board data transfer validation

3.1. Throughput evaluation

Each digitizer delivers a digital data flow that has to propagate through the host interface (PCI bus) towards the PC memory.

In order to evaluate the system performance in a particular application, the system integrator should be able to qualify the peak data throughput requirements.

The first step consists in determining the data throughput generated by one particular digitizer. This depends on:

- the image format being digitized,
- the pixel format.

The following table provides a tabular way for the user to quantify the digitizer Peak Throughput, measured in Megabytes per second (MB/s).

▼ PEAK THROUGHPUT (MB/s)

		Video Standard Image Format (H x V)					
		NTSC 320 x 240	NTSC/PAL 360 x 244 360 x 288	PAL 384 x 288	NTSC 640 x 240 640 x 480	NTSC/PAL 720 x 244 720 x 288 720 x 488 720 x 576	PAL 768 x 288 768 x 576
		Pixel Format	Y8 Monochrome	RGB16, YUV 4:2:2	RGB24	RGB32	
		6.14	6.75	7.38	12.27	13.5	14.75
		12.27	13.5	14.75	24.55	27	29.5
		18.4	20.25	22.13	36.82	40.5	44.25
		24.55	27	29.5	49.09	54	59

As a general rule, it is highly recommended to choose a combination of pixel and image format compatible with the application requirement and leading to the minimum throughput.

In particular, it is advisable to transfer color images with the YUV 4:2:2 or RGB16 pixel format.

It is worth to mention that the YUV 4:2:2 pixel format does not degrade the image quality, because the composite video signal delivered by the camera is originally in the YUV format.

Additionally, many software compression algorithms (JPEG, MPEG) need the video information in the YUV format. Using the YUV 4:2:2 format simultaneously lowers the bandwidth requirement, and decreases the CPU workload by suppressing the need for transforming RGB to YUV. Accelerated display like DirectX works also with YUV.

3.2. Bandwidth evaluation

All digitizers deliver their data through a common embedded bus. This bus is shared between the four digitizers on the board. When digitizers compete for the bus, the digitizers suffer from delays and reduced bus bandwidth available per digitizer.

The tables below give the bandwidth per digitizer on Picolo Tetra and Picolo Tymo in practical conditions.

▼ BANDWIDTH PER DIGITIZER FOR BEST FLUIDITY (MB/s)

Board	Format	Operation with ...		
		2 digitizers	3 digitizers	4 digitizers
Tetra	Packed	39	36	21
	Planar	28	22	18

▼ USABLE BANDWIDTH PER DIGITIZER (MB/s)

Board	Format	Operation with ...		
		2 digitizers	3 digitizers	4 digitizers
Tetra	Packed	43	40	25
	Planar	37	29	27

3.3. Bandwidth constraint

When only one digitizer is in use, the on-board bandwidth is not a constraint.

Assuming that the digitizers in use on Picolo Tetra or Picolo Tymo share the same image format and pixel format, the following condition must be verified:

$$\text{Throughput} < \text{Bandwidth}$$

This rule is applied per digitizer.

Satisfying the constraint in the “usable” case means that images are occasionally spoiled because of bandwidth limitations.

According to the frequency of the phenomenon, it may be acceptable for the application or not. The driver conveniently reports the spoiled images and allows the application to take appropriate measures. This is described later in this document.

The statistical nature of the problem prevents to fully guarantee a zero-defect operation. Satisfying the constraint in the “best fluidity” case means that less than 1% of images are lost due to bandwidth limitation.

3.4. Working around the bandwidth constraint

When the constraint is not satisfied, the system designer will relieve it in one of the following ways:

- Reducing the amount of digitizer simultaneously in use.
- Choosing an image format yielding fewer pixels per line.
- Choosing a less demanding pixel format.

Additionally, slowing down the acquisition rate reduces the failure probability. You can achieve this by artificially lowering the acquisition rate through MultiCam control. Note that the switching acquisition mode inherently reduces this rate.

In any case, you should qualify the system performance in representative conditions. A set of typical performance curves is provided later in this document.

3.5. Examples

Standard & Format	Throughput (MB/s)	Digitizers in use	Piccolo Tetra
NTSC 320H RGB16	12.27	4	< 21 MB/s Best fluidity
PAL 768H YUV 4:2:2	29.5	3	< 36 MB/s Best fluidity
NTSC 720H YUV 4:2:2	27	4	> 21 MB/s Not usable
NTSC 640H RGB24	36.82	4	> 21 MB/s Not usable
PAL 384H RGB24	22.13	4	> 21 MB/s Usable
PAL 768H RGB24	44.25	4	> 21 MB/s Not usable

4. Out-board data transfer evaluation

The next assessment step consists in figuring out whether the data generated within the board will safely cross the Piccolo Tetra or Piccolo Tymo board boundary through the PCI bus to the PC memory.

4.1. Combined throughput

The combined throughput is the summed throughput from all used digitizers sharing the same PCI bus in the system.

4.2. Practical PCI bandwidth

The PCI performance is often characterized by its practical bandwidth.

Bus variant	Bus width (bits)	Bus speed (MHz)	Practical bandwidth (MB/s)
Conventional PCI	32	33	90
Wide PCI	32	66	180
	64	33	
	64	66	360
PCI-X	64	100	540
	64	133	720

Refer to the application note “PCI bus Variation” included in the MultiCam documentation.

4.3. Available PCI bandwidth

You should however keep in mind that your acquisition application will share this bandwidth with other agents in the system.

In other words, the available PCI bandwidth is always lower than the practical bandwidth.

4.4. PCI constraint

The following condition must be verified:

$$\text{Combined throughput} < \text{Available bandwidth}$$

The available bandwidth is the practical bandwidth reduced by the amount of non-acquisition-related traffic.

4.5. Chip-set architecture

As a general rule, there is a requirement for the host PCI bus to be as available as possible for peripheral data transfers generated by add-on boards.

The user will advantageously get a good knowledge of the topology of the busses running inside his PC.

For instance, some chip-sets share the traffic running over an internal 64-bit PCI bus with the traffic consumed by the AGP graphic board. In fact, in this particular architecture, the AGP bus is issued from a bridge hooked to the 64-bit PCI bus. It has been shown in this case that any significant graphic activity reduces the bandwidth available to the add-on devices.

4.6. PCI latency effects

The PCI bus architecture allows for several agents to exchange large amount of data through a set of copper tracks running on the motherboards. However, the PCI protocol does not offer a mechanism to *guarantee* the bandwidth delivered from an agent to another. On the contrary, a priority mechanism attempts to fairly satisfy all transactions. Practically speaking, Picolo Tetra or Picolo Tymo issuing video related data may not be allowed to deliver it properly if another intervening traffic is running.

This traffic saturation results in so-called PCI latencies. Each digitizer located on-board issues a request for bus ownership when a few data have been decoded and are ready for transfer. Before the bus arbiter allows this digitizer to effectively deliver the data, a latency time is spent. This time is used by other agents to perform their own transactions. During that time, the digitizer accumulates incoming video data into an internal buffer, called a FIFO (First-In-First-Out memory).

Unfortunately, this FIFO buffer is rather small, meaning that too long a latency will cause the buffer to overflow.

Globally, this does not stop the transfer process. Instead, a number of data are simply discarded and not transferred towards the PC memory, causing an anomaly in the transferred image. The missing data affects a section of the image ranging from a part of a line to several lines.

In summary, high PCI traffic sporadically induces a PCI latency that causes a FIFO overflow inside the digitizer, resulting into a defective image.

The FIFO overflows statistically appear at a rate strongly dependent on the required throughput and on the internal PC activity.

5. Assessing spoiled images

5.1. Spoiled images

An image is considered as “spoiled” as soon as a FIFO overflow phenomenon occurs during its transfer.

You will benefit from the fact that the MultiCam driver marks the spoiled images. The application can ignore those images. Or the application can use them because the FIFO overflows effect is acceptable.

5.2. Performance curves

The following curves show the rate associated to the occurrence of missed images per Pico Tetra and Pico Tymo digitizer.

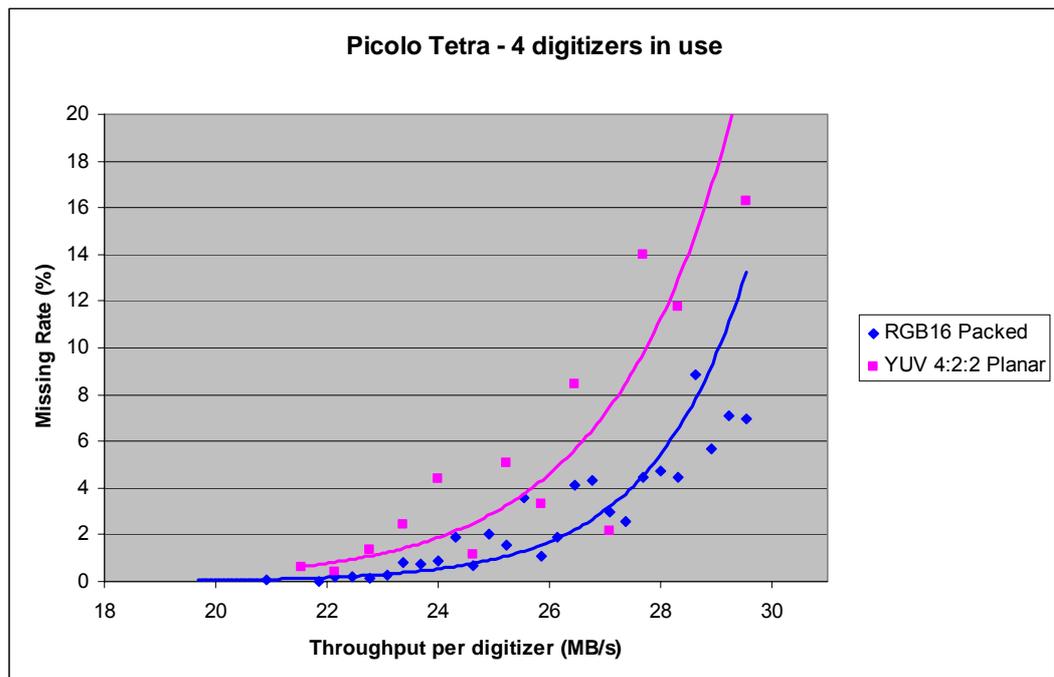
This rate is quoted in percent. A 1% rate means that on average one image every one-hundredth is missing in the sequence. This represents one image every 3 to 4 seconds in the video sequence.

The measurements have been conducted in the streaming condition, which is the most severe. All four on-board digitizers are operating with the same data format.

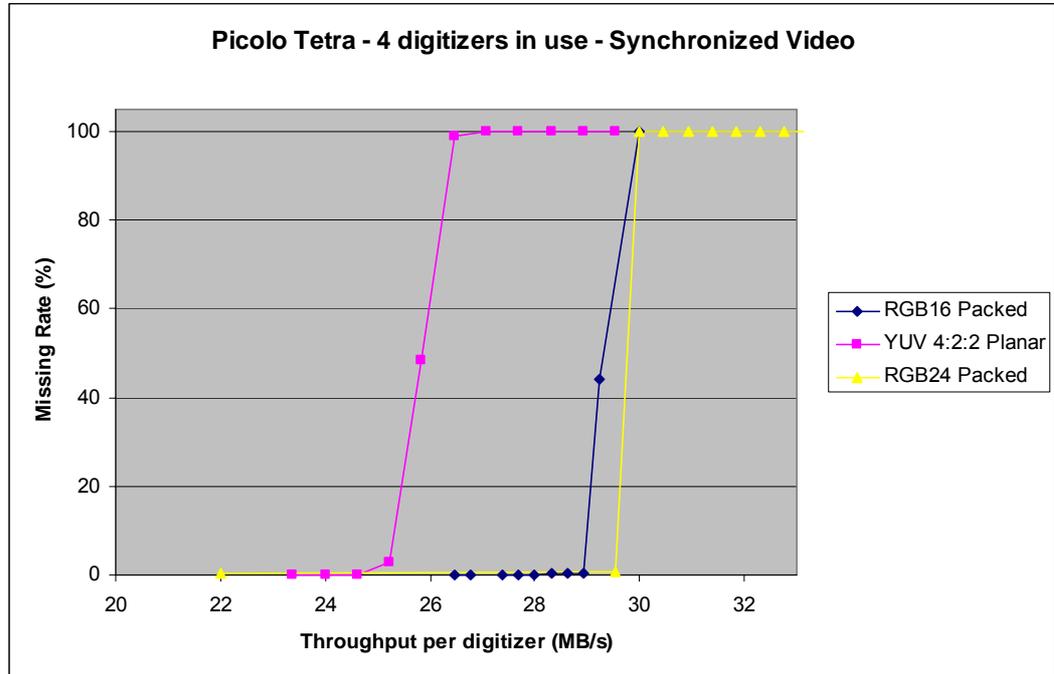
The curve for each pixel format is obtained in testing the missing rate for a range of line width from 320 to 768.

The curves are representative for off-the-shelf server class PC systems with wide PCI busses (64 bits, 66 MHz).

Picolo Tetra



Picolo Tetra – Synchronized video



5.3. Conclusion

The above curves provide an objective way to appreciate the level of performance of a video surveillance system using multiple video channels in real time.

They draw the attention of the system integrator to the factors to consider for reaching the top-level performance for the combination of PC and capture board.

Real-time full-size streaming of color video images can be achieved with a low rate of missed captures.

It is expected that the observed missing rates will decrease substantially in the future as more and more wide-PCI and PCI-X motherboards become available, with better performance.

Copyright Euresys s.a.
2006