

Building a High Performance Camera for Wood Inspection Using VITA1300 Image Sensor and the XEM5010 FPGA Integration Module

Automatic visual inspection has gradually found its way into the forest industry. There are plenty of inspection tasks in sawmills as well as in the secondary wood industry, e.g. production of furniture and construction components, such as windows, beams, etc. Regardless of where in the production chain the systems are used, the primary task is to find various defects on the lumber surface, such as knots, cracks, stain, resin pockets, and dimension faults, etc.

Using the “Tracheid Effect” for Wood Inspection

A quantum leap in wood inspection technology was made through the discovery of the so called “tracheid effect,” illustrated in Figure 1. Today, this method is used by most suppliers of wood inspection equipment.

The tracheid effect works like this: When a wood surface is illuminated by a concentrated light source, e.g. a laser, the wood is somewhat translucent and some light will thereby spread onto the surface layer, creating a larger illuminated spot than the extension of the influence light itself. The light will mainly spread along the wood fibers, called tracheids in softwood, giving an elliptical spot as shown in Figure 1.


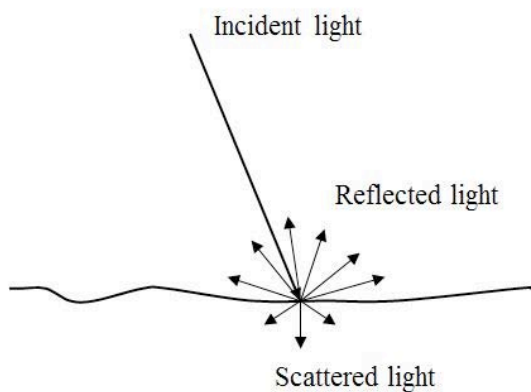


Figure 1. The Tracheid effect

The key to the method is the fact that the light scattering will be very different, both the amount and direction, in defective wood. For instance, a knot will scatter much less light than the surrounding clear wood. If we now illuminate the lumber surface with a line laser projector, the result will be as shown in Figure 2, where it can be clearly seen that the line looks narrower within the knot.




Figure 2. Light scattering in a knot

Assume now that we have a setup as shown in Figure 3, which basically is the traditional setup for 3D measurements by triangulation. This is in itself an important function in most wood inspection systems since, in addition to the detection of defects like knots, we need to measure lumber dimension and also detect so called wane edges and other geometrical faults. In order to make a compact system, these functions can be combined in the same setup with the measurement of the Tracheid effect.




Figure 3. Camera and laser setup

Basically, as the lumber passes through the inspection system, we will obtain a sequence of images from the camera showing the laser line projected onto the surface. From each such image we will extract the position of the line, column by column, giving the 3D profile of the wood surface. The result, when the entire board has passed, will be a “profile image” of the surface. This is standard 3D imaging and there are several known ways to calculate the line position, such as the center of gravity of the acquired light within the column.

If we now, in addition to the position of the laser line, also measure the amount of light scattering, e.g. by measuring the width of the projected line at a certain threshold level, the result, when the board has passed, will be a “tracheid image” or “scatter image.” We can also measure the reflectance of the light in the center of the laser line, giving, when the board has passed, a traditional

grayscale image of the surface where the laser merely is used as illumination. Figure 4 shows the grayscale image and the scatter image respectively of a board. Note the much higher contrast in the scatter image. One doesn't need to be an expert in computer image processing to understand that it is easier to detect the knots in this image.




Figure 4. Grayscale (top) and Tracheid or Scatter image (bottom) of a piece of wood

We start with a sequence of 2D images. From each such image we extract three different 1D lines of data, profile, scatter (tracheid) and grayscale. These lines, sequentially acquired as the boards pass through the system, are used to build three new different 2D images to be used in subsequent processing.

Increasing Camera Speed, to Meet Industrial Requirements, Results in Too Much Data

The challenge is to do this imaging at sufficient speed. If an ordinary video camera is used, we will get 50 frames/s resulting in 50 lines/s in the resulting images. If the lumber is moving at 2 m/s this will correspond to a lengthwise resolution of 4 cm. This is far from meeting industrial requirements, which are typically around 1 mm corresponding to 2000 frames/s. There are certainly high-speed cameras on the market that can do this but, besides being expensive, they would generate huge amounts of data to be transferred to the host and to be processed. There are also numerous “smart cameras” with embedded processing capabilities, but few of them can handle the high frame rates needed in this application.

The Solution—a High-Speed Sensor into FPGA Processing Engine

The solution is to use a high-speed sensor, with high bandwidth, into an FPGA processing engine located within the camera itself, to do the first level of processing. The output from the processing is the three result images—profile, scatter and grayscale, which require much less bandwidth.

The image sensor chosen for this project was the VITA 1300 from Cypress Semiconductors, now marketed by On Semiconductors. This CMOS sensor has 1280x1024 pixel resolution and is capable of

acquiring 150 frames/s at full resolution. Proportionally higher frame rates can be achieved using smaller regions of interest (ROI). In the wood inspection application, we only need high resolution crosswise. Typical ROI size can be 1280x64 pixels giving 2400 frames/s. The ROI height is mainly a compromise between acquisition speed and height resolution in the 3D profiling.

To avoid developing the complete camera, a readymade FPGA integration module, the XEM5010 from [Opal Kelly](#), shown in Figure 5, was used. The [XEM5010](#) is an integration module based on a XC5VLX50 Xilinx Virtex-5 FPGA. In addition to a high gate-count, high-performance FPGA, the XEM5010 includes the [FrontPanel](#) API that utilizes the high transfer rate of USB 2.0 for easy communication between the sensor and the PC.



Figure 5. XEM5010 FPGA USB 2.0 module

The 10-bit pixel data output from the sensor comes in four parallel LVDS serial channels plus a sync channel running in DDR mode at a clock frequency of 310 MHz giving 620 Mbit/s peak transfer rate for each channel. In total this gives $4 \cdot 620 / 10 = 248$ Mpixel/s. The image sensor was mounted on a separate PCB, shown in Figure 6, of the same size as the XEM5010 (85mm x 61mm) with a connector on the bottom side mating directly to one of the I/O-connectors of the FPGA module.




Figure 6. The image sensor board

The Virtex 5-based module was selected because it gives the ability to directly interface to the DDR LVDS channels through the use of the IODELAY and ISERDES primitives. The interface needs to be trained using a training bit pattern output from the sensor. The sensor PCB is a four-layer board with the LVDS routes implemented as impedance controlled microstrip lines due to the high frequencies involved.

The processing pipeline in the FPGA inputs image data from the sensor and calculates the laser line position, and the amount of light scattering around the line, column by column. The image comes in line by line from the sensor, requiring a line buffer holding intermediate results implemented using block RAM in the FPGA. Since the four channels output data from different image columns, they can be processed in parallel by instantiating four identical processing pipelines in the FPGA. The pipelines run at the quite moderate 62 MHz pixel clock from the LVDS deserialization, which reduces FPGA heat dissipation.

The final image data from the processing engines are then stored in a large output FIFO. To save FPGA block RAM for line buffering, processing lookup tables, etc., the external 36 MByte synchronous SRAM on the XEM5010 was used for this purpose, running at 124 MHz to facilitate one read and one write operation for each pixel cycle. Data from the FIFO is then transferred to the host through the USB interface using the block-throttled pipe primitive in the Front Panel API. The three images are generated at a rate of 2400 lines/s each. The scatter and grayscale images use 8-bit pixels while the profile image uses 16-bits, i.e. we have in total $4 * 2400 * 1280 = 12$ Mbyte/s which is within the capability of USB 2.0. The image data is then saved, again in a large FIFO buffer in the host RAM, before it is processed.

There is also an image mode sending the raw image data from the sensor directly to the host. This is mainly for setup and calibration purposes and can't be done at full speed, since the sensor data rate is more than the USB channel can handle. There is an option though, for testing purposes, to dump raw pixel data from the sensor into the XEM5010 SDRAM, which can write the data at full sensor speed. Data from the SDRAM can then be run through the processing pipelines or downloaded to the host for offline testing and algorithm prototyping. The size of the SDRAM is 256 MB giving room for 3200 frames, or 3.2 m lumber, given the above mentioned frame size and feeding speed. Note though that enabling the SDRAM will significantly increase the heat dissipation from the FPGA making cooling an issue. Without the SDRAM a passive heatsink is sufficient in the closed camera box.

In addition to the processing pipelines, the XEM5010 FPGA contains an SPI interface for controlling the sensor through a set of internal registers that define exposure parameters, ROI size and position, etc. Even if a 64-line window is sufficient for the processing, the window must be repositioned on the sensor to cope for varying lumber dimensions. In order to also handle movement of the board within the field of view, there is a tracking logic in the FPGA that constantly updates ROI position based on the calculated line position. Since the tracking is entirely performed in the FPGA, this can be done in a very tight feedback loop with a latency of only one frame.

Finally, there are a number of digital I/O signals with interface logic available at a connector on the sensor board. Some of these signals are connected to two quadrature counters, implemented in the

FPGA, to keep track of the lumber movement using rotary encoders. Other signals can be connected via opto-couplers to photocells and other external signals. The idea is that encoder positions, photocells, etc. are tagged to the data stream directly in the camera. By doing so, the host will be revealed from all real-time event handling. It just needs to process image data in a FIFO manner while all events, such as detecting the beginning and end of a board, are encoded into the data stream.



Figure 7. The camera

The Camera and Industrial Enclosure

Physically, the sensor and XEM5010 FPGA boards are enclosed in a 10 x 7.5 x 5 cm aluminum box, see Figure 7. The C-mount lens is mounted on an adjustable adapter so that the lens can be angled relative to the sensor in order to satisfy the so called Scheimpflug condition. This will ensure that the laser line always stays in focus, regardless of the distance to the object. Finally, the camera and the laser are enclosed in a sealed box, due to the dusty environment in a sawmill, as shown in Figure 8. A complete system is then made up using four such boxes, one for each side of the lumber, shown in Figure 9. This system uses a separate white LED illumination in addition to the lasers, as shown in Figure 10. The reason is that there are some drawbacks using the laser as illumination for traditional grayscale imaging. The LED grayscale data can either be obtained from a separate (single-line) ROI, with appropriate filters placed directly on the sensor, or be time-multiplexed by turning the lasers and LEDs on and off for alternating images controlled by FPGA I/O signals. In the latter case, the resulting output lines are only generated one for each second sensor frame, so this procedure puts no additional burden on the host.




Figure 8. Camera and laser enclosed in a box




Figure 9. A complete system consisting of four camera/laser boxes and white LED illumination for grayscale data acquisition.



Figure 10. A piece of wood illuminated by the laser line and the white LEDs

Advanced, Low-Cost, Smart Camera for Wood Inspection using FPGA

To conclude, an advanced special purpose smart camera can be built at a relatively low cost, with a relatively small development effort, by using the XEM5010 FPGA integration module. To design the same camera from scratch would be a much bigger project. The camera today utilizes around 40% of the Virtex 5 FPGA resources, so there is room for expansion and both the external SRAM and SDRAM of the XEM5010 come in handy for this design. The Front Panel framework, with its API primitives to communicate with the FPGA, also cuts development time significantly. Finally, despite the high sensor output data rate, the USB 2.0 interface in the XEM5010 is fast enough for this application, eliminating the need for expensive frame grabbers.