WHITE PAPER

Fundamentals of Image Processing Systems

Example of an image processing system: Equipment for inspection of glass substrates for thin film solar industry

Image processing systems handle tasks like measuring and counting vehicles or products, calculating their weight or volume, sorting goods at high speeds based on pre-defined characteristics and automatically extracting small but decisive pieces of information from giant volumes of data. They help experts interpret images by filtering, optimizing, supplementing and making them available for quick retrieval. They are tireless and give consistent effort.

But how should this kind of system be designed? What steps are necessary, what factors must be taken into account and what options are available?

As a former manufacturer of image processing systems, we have extensive experience in this field that we hope to pass on to you through a series of white papers. The topics in this series are oriented toward the steps required for conceptualization: Selection of camera, lens and light source, evaluation of image quality and selection of PC hardware and software. Last but not least we discuss optimal usage of camera functions.

This first paper is intended to serve as a starting point and touches briefly on each of the aforementioned topics to provide initial orientation. More detailed information will come in the papers to follow.

Content

1. Image Processing: One Technical Area, Numerous Facets 1
2. Know What's Necessary: Requirements Definition 2
3. Resolution and Sensor 3
4. Color or Monochrome? 4
5. What Else Does the Camera Offer: Camera Function and Image Quality 4
6. The Camera's Eye: Benchmarks and Imaging Quality 4
7. Frequently Overlooked: The Lighting Setup 5
8. PC Hardware 6
9. Software 6
10. An Apple Grower Needs Our Help 7
11. Summary 7

Image Processing: One Technical Area, Numerous Facets

Image processing systems are used for a variety of applications, including manufacturing, medicine, traffic monitoring and surveillance.

After the lighting setup and lens, the digital camera itself is the next most important link in the image processing chain. It uses a specific protocol to communicate with a computer, which then processes the image data and which can be used to modify the settings on the camera.

The requirements for an industrial image processing camera vary from application to application, but are in all cases markedly different from those you might be familiar with from your own consumer-grade digital camera. The key requirements for example are related to sensitivity, dynamic range, noise level, frame rate and resolution.

Example of an image processing system: Basler L6 System for controlling of printed images for CD/DVD labels

To achieve their full potential, cameras must be properly supported by their peripherals as well. As such, proper optics and a lighting setup for inspection tasks are also of crucial importance. The interface used to transmit the camera data is also very important.

Yet what exactly is involved with these inspection tasks? Let‘s take a look at the requirements definition.

Camera bar on an image processing system (Basler system for inspection of glass substrate for TFT displays)

Know What‘s Necessary: Requirements Definition

If you want to design an image processing system, then you first need to pose the question as to what exactly the system needs to achieve under which conditions.

This first step may seem so obvious as to merit no mention, and yet it is performed nowhere near as often as it should be. Take the time at the start to consider precisely the system you desire. This simple step saves time and money later.

Is the system

- only intended to display images of the object being examined and, for example, use magnification or specific lighting to make visible product characteristics not visible to the human eye?
- calculating objective product characteristics such as dimensions and volumes?
- intended to review correct positioning, such as for automated pick-and-place systems?
- determining characteristics based upon which the product will be classified into specific product classes?

In this context, another important question involves the step after the image processing. Should a person, assisted by the system, examine the workpiece? In this case, speed is unlikely to play a major role and decisions will ultimately be made by the human.

It‘s a very different situation if automated sorting into classes or rejection of defective parts is to be conducted. As the degree of automation for the final process rises, it becomes more crucial that the image processing is based on objectively verifiable specifications that can be defined in the software. Automated inspection and ongoing processing also necessitate a stronger degree of real-time compatibility by the system than situations where images are to be post-inspected by a human.

Second, it is important to know precisely which objects are being examined and which level of detail is required to handle the image processing task properly. If for example I want to measure an object, I must ensure that the object can be detected completely within the image and that the resolution of the image is better than the error threshold I‘m trying to detect. If I want to check whether an apple deviates by more than 1 cm in diameter from an ideal size, then my resolution within the image has to be significantly less than 1 cm.

The placement of the object — its situation — is also of interest.

- Are the objects in the image clearly separated from one another?
- Is this a sheet product, like paper rolls?
- Is this a bulk good that must first be separated so that the individual components can be identified in the image?
- Is the object moving while the picture is being taken, and if yes, how quickly (such as on a conveyor belt in a sorting system or during the actual production process)?
- Is the object stationary at the moment of recording (for example, has it been stored someplace by an automated filling machine)?

Basler system for inspection of plastic sheets

Another set of framework conditions are the total field of view, i.e. the width and length of the area to be covered. This is generally defined as the width of a conveyor belt or product itself, or by the volume of the objects being inspected. In addition, consideration must be given to the question of whether the available physical space represent restrictions, such as the minimum and maximum distance between camera and object, whether other concerns arise through moisture, increased or decreased ambient temperatures or maximum energy consumption.

Finally, the overall system must be taken in account. The major concern here is whether the image processing system is part of a larger system, such as a production line or sorting system, and how the interfaces to this are constituted or whether this is an offline device. This information is important for achieving an image process system optimally designed for its respective purpose.

Resolution and Sensor

What camera is required for an application? The answers to the questions above can be used to derive the necessary resolution and sensor size for the camera.

But what exactly does resolution mean? In classic photography, resolution describes the minimal distance in real life between two points or lines such that they can be identified as distinct within the image.

When discussing digital cameras, it's often said that they have a resolution of 2 megapixels, for example. This, however, refers solely to the number of pixels on the sensor — not the resolution. The actual resolution is determined by the total system comprised of camera, lens and geometry, i.e. distance between camera and object. Higher resolutions do in fact require a greater number of pixels, so the pixel count can be quite relevant. It indicates the maximum possible resolution under optimal conditions.

In industrial image processing, the benchmark for resolution of the total system is frequently stated in terms of how much of an area or which dimensions of the object being inspected are depicted on a single pixel.

This size is frequently describes as the ‚pixel resolution'. If, for example, a square measuring 0.5 mm per side is depicted on a camera pixel, then this is described as a pixel resolution of 0.5 mm.

For a specific detail to be available for inspection, the resolution must be better than the size of the details. If the resolution is equal to the size of the detail, then the detail is reproduced as a whole pixel and no structure from the detail will be identifiable. The pixel resolution is initially a calculated size. What must also be taken into consideration is the maximum resolution that the optics in front of the camera can deliver. In other words: a poor lens will make the actual resolution for the camera worse than the theoretical pixel resolution.

The field of view, meaning the entire area later available for review in the image captured by the camera, should cover at least the size of the object. If, for example, the application involves measurement with dimensional tolerances in the range of 10mm, then the resolution must be significantly better than the dimensional tolerance, so 0.5mm per pixel. If an object is some 200 mm large, then the sensor must possess at least 200 mm/(0.5 mm/pixel) = 400 pixels in one direction, i.e. the range of VGA resolution.

The requirements can further rise if, instead of deviations from dimensional tolerances, deviations in shape are specified, or if multiple objects are to be inspected simultaneously at the same resolution. In that case the requirements for pixel counts on the camera sensor can easily increase many times over.

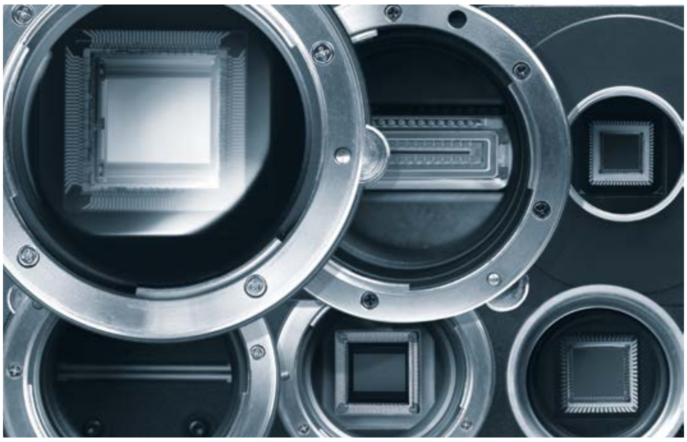

View of the sensor of various cameras (top: area scan, line scan, area scan sensor; bottom: line scan, area scan, area scan sensor).

The finer the resolution or larger the inspected range, the greater the number of pixels required. Multiple cameras may be needed to inspect a large area at high resolution. In many cases, it's more cost effective to deploy multiple cameras with standard lenses than one camera with a special lens capable of handling the difficult imaging conditions involved.

The sensor size and field of view can be calculated together into a reproduction scale that later plays an important role in selecting the lens.

A line scan camera is often the best choice for moving objects or sheet/rolled objects. Line scan cameras only record the image for one line, meaning one row of pixels. The object under the camera continues to move past, so the recordings of the lines are then patched together into a total image by the computer hardware backend for the system. For objects moving at high speeds, the camera interface must also be able to keep up with the speeds involved. Line scan camera are often also used when special handling requirements or trigger conditions for the camera are present, or when a special framegrabber card is required.

In most cases, however, an area scan camera with a cost-effective standard interface, such as GigE or USB 3.0, is sufficient. Here too fast cameras are available in a broad range of resolutions and sensitivities. Many tasks, including highly challenging ones, can be successfully handled using these cameras.

Color or Monochrome?

The next question is whether you'll need a monochrome or color camera. In general, it can be reasonably stated that color cameras are not a necessity, but that viewers do often find it easier to look at images in color. If realistic color reproduction is required, then white illumination must be used in combination with the color camera. If characteristics regarding the color can be detected (such as red flecks on the object), then in many but not all cases color reproduction is required. In many cases these characteristics also appear on the b&w images of a monochrome camera through the use of colored illumination. Experiments on good samples are necessary to determine whether this is possible. If color is not relevant, then a monochrome camera should be selected, since the color filter on color cameras inherently makes them less sensitive than b&w cameras.

For more complex inspection tasks, it may also be necessary to deploy multiple cameras, such as to record very different characteristics requiring different lighting or optics.

What Else Does the Camera Offer: Camera Function and Image Quality

When selecting the camera, pixel counts should not be the only selection factor. Image quality and camera functions matter too.

Beyond the resolution, three factors are essential for describing the image quality of a digital camera

- Light sensitivity: The greater the light sensitivity, the less additional light that is needed. Because lighting is frequently a low-cost component, high-sensitivity cameras are typically deployed only when supplemental light sources cannot be deployed due to external factors.
  Strong light sensitivity can also be required if the camera has to work with short exposure times, such as with highframe rates or to ‚freeze' movements.
- Dynamic range: A large dynamic range allows for good images in variable lighting conditions and shows details for scenes in which very dark and light areas occur. Cameras with this feature are frequently advertised as HDR or WDR (high or wide dynamic range).
- Signal-to-noise ratio: A camera's noise level should be low in relation to the discernible signal. This property is particularly important in low light situations.

Among the camera functions, speed is among the most important. This is typically stated as fps or frames per second and indicates the maximum number of images that can be recorded per second.

Other functions include automatic exposure times and gain adjustments and automatic white balance. For image processing systems it's frequently necessary for the camera to possess an external, real-time compatible trigger input to allow for synchronization of images or handling at high speeds with other parts of the system.

Later white papers will explore other image quality characteristics, including color fidelity and functions.

The Camera's Eye: Benchmarks and Imaging Quality

Good optics are expensive. In many cases, standard lenses are more than enough for the application at hand. To determine what's really needed, several parameters must be defined such as

- lens mount
- pixel size
- sensor size (with the measurement in inches referring to the historical diameter of TV pickup tube. If the data sheets from the camera manufacturer refer to a one-inch sensor, this corresponds to a sensor measuring 16 mm across the diagonal)
- magnification, with the ratio of the diagonal measurement of the sensor (image size) here corresponding to the diagonal measurement of the object field to be examined (item size). This also corresponds to the ratio of the size of the individual pixels divided by the pixel resolution (the pixel resolution is the edge length of a square on the object being examined such that it fills precisely one pixel on the camera sensor.

- focal width of the lens, from which the magnification and the distance between lens and camera are determined
- maximum aperture or "speed" of the lens designation its ability to gather light.

This information is helpful for examining the data sheets from lens makers for an initial assessment of whether a cost-effective standard lens will suffice or whether a lens in the higher price category is needed.

It is important to take into account the lens characteristics, including distortion, resolution (described by the MTF curve), chromatic aberration and the lens's optimal spectral range. There are for example special objects for near infrared, extreme wide angle lenses (fish-eye) or telecentric lenses designed especially for length measurements. These special lens characteristics often come at a high price, however.

Tests and critical examinations of sample shots are essential for clearing up open questions related to tough applications or specifications omitted from the data sheets.

Frequently Overlooked: The Lighting Setup

Without good lighting you'll never see anything; this

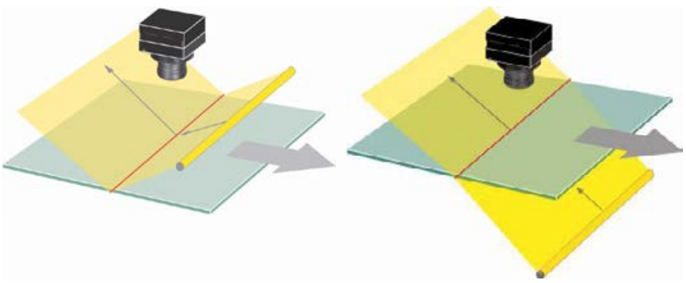

Two sample illumination setups: Dark field illumination in reflection (left) and in transmission (right)

truism applies especially for image processing systems.

High inspection speeds often require sensitive cameras and fast lenses. In many cases it's easier to modify or optimize the illumination to produce the same boost in image brightness. The following options for achieving higher image brightness are always available:

- In some cases it is enough to raise the ambient light using standard commercial light sources. In simple terms: make sure it's bright enough in there. But you should also check that the illumination is receiving the right energy supply. Under the wrong circumstances, power drawing on alternating current from the standard power mains can mesh poorly with the image recording frequency, creating beats and pulsating image brightness. Illumination for image processing is therefore frequently conducted using direct-current or high-frequency AC.
- Using light shaping, such as through lenses or mirrors, the existing light can be bundled and cast on the area of interest at a high level of light intensity.
- Flash photography using a suitable light source allows for very bright illumination of a temporary, not permanent, nature, to be synchronized with the capture of the image. That's because many light sources can be operated briefly at a much higher power intake level (and thus higher brightness) that would destroy the light source under steady-state operation.

But illumination brightness isn't the only important factor. The path that the light takes, striking the object and then into the camera, also matters. Special types of lighting are often used to make specific characteristics or properties visible. Do the desired details or characteristics appear better

- on a light field or dark field,
- with targeted or diffuse illumination,
- with reflected or backlighting?

One example from photography that will be familiar to most people: ambient light is typically diffuse, but often isn't bright enough. Using a directional flash brings plenty of light — but also undesirable reflections on smooth surfaces in the image that can overwhelm the details being sought. These effects can be desirable in image processing, where low-reflecting, straight surfaces are exposed to strong light intensity. Diffuse light is better for objects with many surfaces reflecting in numerous directions.

An image is viewed in reflected light, while a stained glass window only reveals its splendor through backlighting.

Light and dark field illumination is typically less of an issue in daily life. For light field photography we look directly into the light reflected back from an object, corresponding to our normal vision. For dark field illumination, targeted light is cast at an angle so that it is not directly reflected into the camera. Only light that meets specific parameters — typically flaws in the object that refract the light — are then visible to the camera.

As previously mentioned, consideration should also be given to the wavelength of the lighting being using. Monochrome light sources frequently produce sharper reproductions since there are no chromatic effects on the lens. That said, it can be difficult to achieve the proper intensity of monochrome lighting. Working with an extremely small-band light (laser) can also lead to interference phenomena. Care must also be taken that the camera being used is sufficiently sensitive in the selected wavelength range.

Another benefit is that for light at a specific wavelength is that it can be used to bring properties to the forefront that are otherwise invisible, and thus can be weighed against other disadvantages.

In designing the lighting, a large number of lighting geometries and lamps with highly varied characteristics and properties can be selected. Halogen, halide and gas discharge lamps, LEDs, fluorescent tubes and many other lamps are available for the construction of a lighting system.

The system's optics responsible for shaping the illumination beam path can also take advantage of a variety of tools to help identify special defects or characteristics in the objects being inspected.

The framework conditions must not be forgotten here, especially in terms of energy consumption, life span, triggering options and cooling, which can on a case-by-case basis make the difference between a good solution and insufficient one.

PC Hardware

The required hardware is determined based on the task and the required processing speed.

While simple tasks can be handled using PC hardware and standard image processing packages, complex and fast image processing tasks can in some cases require specialized hardware.

Special framegrabber cards are required for specific fast camera interfaces such as Camera Link. These frequently handle the pre-assessment, extending the system's computational power. This provides an indication of just how complex a job really is. A fully automated system that extracts and evaluates characteristics and then passes on instructions to a sorting machine must, for example, be real-time compatible and requires significantly more computational power than a system that detects characteristics for post-processing by a human in offline mode.

For smaller devices intended for resale to customers, an embedded system is appropriate. Special hardware is sometimes needed, but always has the disadvantage of being very dependent and is frequently tied to high costs and lower user satisfaction. Affordable consumer devices by contrast frequently have the disadvantage that models are quickly discontinued and spare parts become difficult to find, meaning the successor has to be designed to spec to fit the system. In terms of operating system, the eventual user of the system is important for determining how powerful it needs to be.

Software

Software is typically required to evaluated the images. Most cameras come with applications used to display images and configure the camera. That's sufficient for the first-time setup. Special applications and image processing tasks require special software, either bought in or developed.

Many cameras work with standards (such as GenICam, and USB3 Vision) and can be operated using generic drivers and ports. Others have manufacturer-specific drivers integrated into a separate piece of software, often with sample software code. These are offered as Software Development Kits (SDK), including for example Basler pylon, a camera software suite that the manufacturer (Basler) provides as a free download from its website.

For the actual image processing a separate, specialized program can be written, including development and implementation of an image process algorithm. This is generally only economically feasible for commercial providers of systems or software who can then obtain patent protection for their developments.

Solutions are already available for many standard image processing tasks solutions, including free or commercially available software libraries (such as Halcon) or programs like LabView as components.

An Apple Grower Needs Our Help

Let's now take a look at a more practical case study.

Specifically: inspecting apples. A 300 mm wide conveyor belt transports apples at a speed of 200 mm/s. 3 apples are typically transported next to one another.

The apple grower wishes to use an image processing system to measure the size of the apples down to a half-centimeter.

Resolution
To ensure that 5 mm can still be comfortably resolved, this length should not be depicted on less than 20 pixels. Each pixel is thus responsible for a distance of 5 mm/20 = 0.25 mm.

For the entire resolution of the conveyor best to be captured at this so-called pixel resolution, we need 300 mm/(0.25 mm/pixel) = 1200 pixels.

Let's presume we opt for a Basler ace camera with 1200 x 1000 pixels. From the data sheet we can determine that each pixel of this camera has an edge length of 3.75 µm.

For the reproduction scale, this produces 0.00375 mm/0.25 mm = 0.015.

If we use a standard lens focal length of between 18 mm and 35 mm, a distance of 2.5 m between camera and conveyor belt is sufficient.

Maximum Exposure Time

There are multiple limitations for the maximum exposure time. In this case, the limitation results from the movement-related blurriness of the object.

The blurriness from motion should not be greater than 1–2 pixels on the sensor. To move by 2 pixels in the direction of transport, 2*0.25 mm = 0.5 m, the band needs only 2.5 ms (namely 0.5/200 m/s).

This travel time produces the maximum exposure time, i.e., during an exposure of 2.5 ms, the object has moved 2 pixel.

The exposure time is thus very short, meaning an external light source must be used.

If no other details about the apple are of interest, backlighting could be used to put the edges of the apple in sharp contrast, although the apple itself would be dark.

Is the Camera's Frame Rate Sufficient?

An apple is approx. 10 cm large. How often do I have to capture the image to ensure that the apple is completely in the frame?

The camera possesses over 1000 pixel in the belt transport direction. That means at a pixel resolution of 0.25 mm a field of view of 25 cm. To ensure that the apple appears at least once completely in the image, it's necessary to take a photo for each 15 cm the belt moves forward ((Field of View) - (Size of Object)). To account for the size differences among apples, the actual figure is closer to every 10 cm.

For 10 cm the belt requires 0.5 s. The resulting refresh rate is thus 1/0.5 s = 2 fps. In other words, the image needs to be captured only twice per second. In terms of refresh rates, this is a very low bar to clear and any standard commercial camera can handle it.

Summary

The conceptual design of an image processing system involves numerous considerations related to all components within that system, from the camera and its optics to the illumination and downstream PC hardware and software.

Yet this task can be managed effectively step by step by clarifying in advance how the system will be used and which factors are most important.

In the white papers on image processing still to follow, we'll be going into detail on aspects that are only touched upon in passing above. The next white paper in the series will explore the question of selecting the right camera.

Author

Dr. Britta Niederjohann
Product Manager

Dr. Britta Niederjohann has worked at Basler AG since 2005. Her broad knowledge of the field of image processing systems was built up over the course of her years with the company and her broad spectrum of tasks within various product areas.

She initially encountered many of the different application ranges in her role as a developer: Optical Media Inspection (inspection of CDs and DVDs), MicroCrack Inspection (inspection of wafers for the solar industry), and inspection of foodstuffs for sorting systems

She later worked as a project engineer, gathering practical experience on-site by providing support during the installation of various SENSIC glass inspection systems. This included systems to inspect glass for thin-film photovoltaic applications as well as TFT displays and color filters. Years of providing close support for customers provided Britta Niederjohann with extensive know-how in various applications.

Ms. Niederjohann is now product manager for key accounts in the camera area. She is able to apply her knowledge of image processes to benefit her customers. It‘s important for her to understand the customer‘s systems and application cases to determine the best specific camera solution for each customer and assist in the selection of further components.

Contact

Dr. Britta Niederjohann – Product Manager

Tel. +49 4102 463 476
Fax +49 4102 463 46476
Email: britta.niederjohann@baslerweb.com

Basler AG
An der Strusbek 60-62
22926 Ahrensburg

Basler AG

Basler is a leading developer and manufacturer of high-quality digital cameras for applications in manufacturing, video surveillance, medicine and traffic monitoring. Product development is led by the demands of industry. The cameras offer simple integration, compact sizes and an outstanding price/performance ratio. Basler offers over 25 years of experience in the image processing industry. The company is home to more than 400 employees at its corporate headquarters in Ahrensburg, Germany and its subsidiaries in the USA, Singapore, Taiwan, China and Korea.