 By David Marsh, Contributing Editor

Off-the-shelf components

make light work of machine vision

MAKING SYSTEMS RECOGNISE OBJECTS AND THEIR PROPERTIES IS CENTRAL TO MANY OF TODAY'S AUTOMATED PROCESSES. ONCE THE DOMAIN OF ARTIFICIAL-INTELLIGENCE EXPERTS, VISION SYSTEMS AND SOFTWARE PACKAGES ENABLE INDUSTRIAL-STRENGTH APPLICATIONS IN MINIMAL TIME AND WITHOUT SPECIALIST KNOWLEDGE.

With its origins in defence applications, the automotive industry adopted machine vision for parts inspection during the 1980s. Since that time, the technique has become a cornerstone of modern manufacturing processes. Videocameras and software are far better than human inspectors at extracting accurate information from repetitive image sequences—and far cheaper in the long term. But although the production cost and efficiency benefits are beyond dispute, many system engineers are reluctant to approach machine-vision projects. It's easy to recognise that successful vision systems require technologies including cameras, optics, and lighting control to augment familiar electronics and control-system knowledge. And 20 years of intensive development brings machine-vision applications within the reach of nonspecialists.

Even so, you need some knowledge before you embark on your first project. Fortunately, several short cuts can help get you started. The vibrant Internet-expert community is only too keen to share its knowledge: A simple "machine-vision" search at www.google.com returns some 880,000 hits. You can find information ranging from everyday industrial control topics to exotica, such as inspecting the Chernobyl nuclear-disaster site or analysing the surface of Mars. A key link is to the Automated Imaging Association (www.machinevisiononline.org), a US organisation dedicated to the technology and its applications. Machine vision also receives a great deal of academic input. Institutions, such as the computer-science department at the UK's University of Wales at Cardiff (www.cs.cf.ac.uk/department/staff/bruce.batchelor.shtml), specialise in research for industrial applications and publish much useful information. And if software is your passion, you can learn and help develop a common object-oriented environment by downloading software for the Image Understanding Environment project, which is sponsored by DARPA (the US Defense Advanced Research Projects Agency), at www.aai.com:80/AAI/IUE/IUE.html.

It's important to note from the outset that machine vision is a huge topic, and this feature can only scratch at its surface. Worse, no two machine-vision applications are the same. But, in general, application profiles fall into one of four categories that industry insiders describe with the acronym GIGI (gauging, inspection, guidance, and identification). "Gauging" refers to measurement; inspection is a visual go/no-go test; guidance is feedback, such as for a machine handler; and identification involves symbolic recognition, such as reading a bar code. And, as with so many other areas of systems engineering, machine-vision practitioners universally recognise the benefits of open architectures and common components. Although you can specify turnkey systems based on bespoke hardware, PC-based hardware and software increasingly replaces vendor-specific products. Specialist vendors, such as Cognex, provide a range of system components that you can quickly configure to suit applications from stand-alone imaging stations to full-blown industrial networks. Many other vendors supply cameras, boards, and software that you can build into a system from the ground up (see **sidebar** "For more information").

Advanced machine-vision systems recognise and inspect pc-board assemblies (courtesy Teradyne Assembly Test Division).

UNDERSTAND THE PROBLEM

A machine-vision system essentially comprises a camera/digitiser combination that converts the target image into a pixel array for software processing (**Figure 1**). In today's architectures, a frame-grabber card provides the interface between the camera and a PC. Software can perform geometric and statistical analyses to identify an object's presence, orientation, and size. Software processing can reduce noise before enhancement to extract the image's differentiating features. (Image-feature extraction differentiates machine-vision from image processing, a discipline that seeks to modify images.) But notice that no substitute exists for a clear, high-contrast image that a properly set camera-and-lighting system resolves. When you have a clear image, intelligent data reduction can group features such as multiple edges to arrive at a representation for comparison with a reference-image template.

But rather than jumping straight into choosing hardware and software, experience from system integrators suggests that your first task is to fully understand the process that you're trying to monitor. Start out by considering what the vision system must accomplish. Are you trying to measure a single parameter or more fully inspect an object? Do these tasks interrelate? For example, does one station read a bar code before issuing an instruction to another station? Does the task happen in a real-time process or off-line, as in taking a sample for SPC (statistical-process-control) purposes? Do you need networking connectivity, remote control, or both? And what object throughput do you need? Additional information, such as the number of quality-control steps, measurement tolerances, the typical pass/fail ratio, and the nature of common defects, complement basic details, such as an object's shape and size. Engineering drawings of the object with a full analysis of what you're looking for, together with good and bad samples of that target, help to clarify the system's task. Also, be prepared to acknowledge that the project's scope may extend beyond vision-system selection. For example, you may need a handling system to accurately position objects on a production line.

For any project of real consequence, you should make process data available to management or for use in other process functions, such as further assembly or test stages. Doing so requires network connectivity and, possibly, an independent data-management function, such as an SPC software package. SPC tools help optimise process efficiency by fine-tuning measurement parameters. This step is typically continuous and involves evaluating a balance between rejecting too many good samples and not detecting marginally bad ones. Importantly, the system should be comprehensible to the end users, who must trust its results and know when the system is performing inadequately. Ideally, end users should also be able to adjust or modify the system to improve accuracy and make setup changes to suit different production runs. These factors stress the man-machine interface, an area in which GUIs provide a familiar environment for task control.

Making suitable camera, frame-grabber, and lighting-system choices involves steps that deeply interrelate and must couple with a software selection that provides the facilities your application demands. First, determining how to accurately resolve target images involves key considerations, such as image size and resolution, the camera's field of view, the number of images required for a representative picture, and the image-acquisition speed. Lighting and optics are new territory for most system engineers, but vision-system integrators advise that knowledge of as few as three lighting techniques can solve most problems. These major lighting scenarios are front lighting, backlighting, and structured lighting (see **sidebar** "Illuminate for robustness").

Lenses generally divide into conventional constructions and the telecentric designs that machine-vision applications often favour (**Table 1**). The **table**'s author, Light Works' principal Spencer Luster, stresses that his analysis is general.

As with most machine-vision topics, he notes, "There are exceptions across the board." However, a telecentric lens collimates the major light rays from an object

AT A GLANCE

▷ No two machine-vision applications are the same.

▷ Options divide into off-the-shelf systems and do-it-yourself configurations.

▷ Prepare to understand cameras, lighting, optics, and software functions.

▷ Acknowledge that all component choices deeply interrelate.

▷ Get your lighting right first.

to provide parallel photon paths at the expense of a very narrow field of view. Thus, telecentric designs minimise perspective errors to provide near-constant magnification across a range of working distances, also minimising magnification errors with moving objects. The industry-standard lens mount is a C type that uses a 1-in.-diameter, 32-thread-per-in. screw mount with a fixed back focus at 0.690 in. from the back-mounting flange. Knowing the focal length of a lens, you can calculate the range and depth of focus for a given magnification using Newton's lens equations. See the discussion at the Light Works resource pages (www.lw4u.com/LightId001.html) for a representative example of this procedure.

CAMERAS PRESENT BEWILDERING CHOICE

Choosing a camera is a seemingly bewildering task with selections that range from conventional area-scan videocameras to purpose-built digital devices. Examples of special-purpose devices include line-scan cameras that scan a single video line to rapidly detect an object's shape; infrared cameras that can view objects in poor lighting conditions; and X-ray cameras that can check solder integrity beneath a BGA package. Camera vendors include such familiar names as Hitachi, Panasonic, and Sony, as well as specialists including Adimec, Basler, Dalsa, and Pulnix. Consequently, interface board and software vendors, such as Data Translation and National Instruments, include selection tools within their Web pages that ensure compatibility between the camera and the frame-grabber card. Similarly, camera vendors can specify frame-grabber hardware that they've tested for compatibility with their systems. Also, most vendors will evaluate as-yet-untested hardware combinations for you.

A monochrome output suits most vision-system applications, but if you need colour-recognition capability, analogue videocameras still offer the widest selection. Their composite-video output signal format requires an ADC on the frame grabber to process 768×576-pixel native-format frames at the maximum 25-Hz field rate. Such cameras typically use CCD (charge-coupled-device) sensors as their imaging elements, but CMOS sensors are now beginning to make inroads into the CCD's market. CCDs are available in various sizes and resolutions including popular 1/3- and 1/2-in. types that provide cameras, such as Sony's XC-E series, with 752×582-pixel spatial resolution. In television-system terms, this resolution equates to 560 horizontal lines. Notice that overall size does not equate to the size of the CCD's active area, and most examples use a 4-to-3 aspect ratio. A traditional frame-transfer CCD comprises a sensor array of light-sensitive photo sites and an adjacent field-storage array that opaque material shields. The sensor array integrates incoming light for the exposure time, after which a clocking scheme transfers the accumulated charge to the storage array. The storage array connects to an output multiplexer that clocks out the image. In the meantime, the sensor array integrates the next image. But rather than employing two separate arrays, an interline-transfer CCD arranges its storage elements adjacent to each active photo site to allow a simultaneous transfer of all the pixels in the frame. This strategy speeds image transfer, but the greater light leakage into the shielded storage elements potentially provides slightly lower spatial resolution than a frame-transfer CCD of equivalent size.

A machine-vision system rarely needs television-system compatibility, because it can use dedicated PC monitors, further freeing camera choices. One common variation is the progressive-scan camera, which transfers its entire image at once and without any interlacing. With features such as direct digital outputs that dispense with external ADCs, this design is rapidly gaining popularity. A serial-digital-stream output also exhibits far greater noise immunity than its analogue-signal counterpart, an important consideration in many industrial environments. Adimec's 1000m camera uses a 2/3-in. interline-transfer

CCD to provide a 1004×1004-sq-pixel image. An RS-232 interface provides full image-acquisition control, including the ability to scan a number of image lines that can start at an arbitrary line number. You can capture data asynchronously or under trigger control and transfer the data immediately or after a programmable delay time. The delay-time feature is useful for frame-grabber boards with multiple input channel, or "taps." The camera's onboard ADC processes as many as 50 images/sec and outputs 8- or 10-bit data via the newly standardised Camera Link interface.

The Camera Link standard builds on 1996's EIA-644 specification for LVDS (low-voltage differential signaling) and provides specifications for a standard 26-pin MDR (miniature-delta-ribbon) connector from 3M; National Semiconductor supplies the interface-chip set. The camera-to-frame-grabber data-transfer format comprises 24 data bits that divide into three 8-bit serial ports to suit monochrome and colour cameras of various resolutions (**Figure 2**). Another port transmits four status bits: data valid, frame valid, and line valid, plus one bit reserved for future use. There's also a dedicated deserialiser clock channel for a total of 10 transmission lines. The frame-grabber-to-camera link provides four

ILLUMINATE FOR ROBUSTNESS

Lighting is so critical to machine vision that it's the number-one contributor to process robustness. Applications such as cellphone inspection stress the lighting system with the need to resolve symbols on miniature screens that lie underneath reflective plastic surfaces. The techniques at your disposal comprise front lighting, backlighting, and structured lighting. You may deploy these techniques alone or in combination. You may also mix and match lighting sources to illuminate different parts of an object. Lighting sources include halogen-filament lamps, for bright directional light; fluorescent lamps for diffuse, nondirectional light; lasers for structured lighting; strobe lamps for stopping motion; and high-brightness white and coloured LEDs for multiple duties. Other tools include shaped light sources, such as fluorescent-ring lamps, and point-delivery systems, such as fibre-optic light guides that illuminate selected areas. Ring lamps are also available that fit around the camera's C-mount lens to provide another source of local illumination.

In front lighting, an overhead source provides light that the object reflects back to the camera. Varying the angle between the light source and the object changes the nature of the illumination—from a bright field with the source directly overhead to a dark field at low angles. A directional bright field highlights edges by creating shadows that maximise the contrast ratio to improve object recognition but may also create glare that locally "blinds" the camera; a dark field emphasises the background detail. Adding a diffuser between the light source and the object smoothes illumination and minimises highlights or shadows, reducing glare from smooth surfaces but lowering the contrast ratio. Backlighting is familiar from display technologies in which a light source illuminates an object's edges from behind. In machine-vision applications, the backlight shines at the camera, and the object partially blocks transmission. This arrangement especially suits applications that measure object parameters. Structured lighting refers to modifying the light beam to form a shape, such as a line, that you can also use to locate edges and measure distances.

Also consider the effect of ambient light on the process. Some applications may require an enclosure to prevent ambient light from creating unwanted shadows or glare. Other applications may require multiple light sources and even multiple cameras mounted at different angles to provide sufficient resolution. One example is Teradyne's AOI (automated-optical-inspection) system for pc-board assemblies that must unambiguously resolve details such as tombstoned components and solder bridges in a fast-moving production environment (**Figure A**). Other techniques include adjusting the wavelength of the light source to resolve details such as the characters on a coloured label. For example, red light sharpens contrast by making any red components blend into the background. A backlight may further enhance the image to ease gauging or optical-character-recognition applications. Of course, you also need to match the light source to the camera's characteristics and specify a device that offers high sensitivity to the wavelength that you plan to use. For this reason, many camera vendors offer variations on a base product to accommodate custom lighting applications.

Figure A

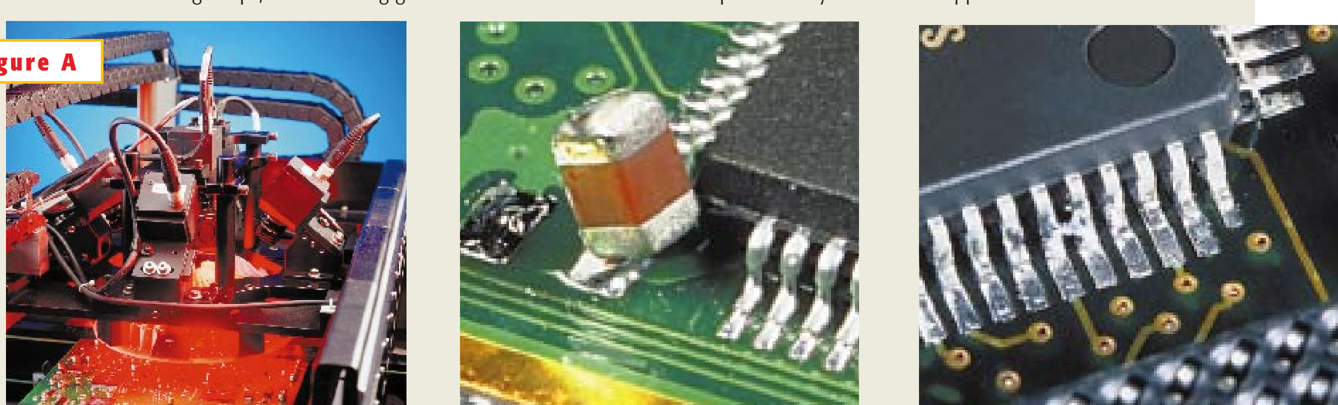

Teradyne's automated-optical-inspection system (left) employs multiple cameras and light sources to unambiguously resolve tombstoned components (center) and solder bridges (right) in a high-speed production environment.

general-purpose command-and-control channels that use the same protocol as a PC's serial port. Camera vendors are free to define their own use for these lines. The standard lets vendors build on this basic configuration by adding another 26-pin connector that can carry as many as five data ports, eight control signals, and two clock signals. This configuration provides a degree of future-proofing for coping with high-speed, high-resolution digital cameras. Vendors can also include RS-232 links and other digital-I/O lines for specific control functions, including backward compatibility. You can download the Camera Link standard at www.machinevisiononline.org.

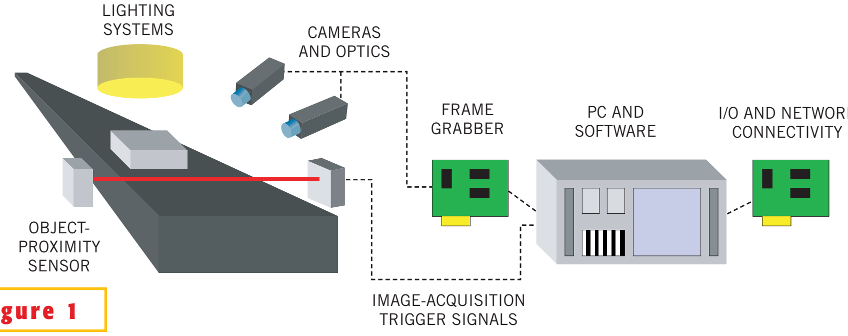

Figure 1

A basic machine-vision application comprises a camera and its lighting system together with communication links to a PC that runs the image-analysis software.

FRAME GRABBERS GO 100% DIGITAL

There's a wealth of frame-grabber cards to choose from vendors such as Coreco Imaging, Data Translation, and National Instruments, reflecting the variety of cameras in use. Analogue-input frame grabbers vary in channel count to accommodate formats including single-channel CCIR monochrome and PAL composite signals; dual-channel S-video, which separates the luminance (Y) and chrominance (C) components for improved colour rendition; and three-channel inputs for the RGB video streams that PC monitors use. A recent example is Coreco's Bandit-II CV, a general-purpose analogue-input frame-grabber and video-accelerator card that handles maximum display resolutions of 1600×1200 pixels at maximum refresh rates of 100 Hz. Available to fit a PC's AGP/PCI buses, the card captures video from as many as six colour or monochrome composite-videocameras or two S-video cameras.

Notice that even a low-end monochrome camera that outputs 8-bit-resolution VGA-format data (640×480 pixels) requires some 307 kbytes of storage for each frame. Thus, data flow can become a critical consideration. Accordingly, the Bandit-II CV card carries as much as 16 Mbytes of dual-port video

TABLE 1—ANALYSIS OF LENSES*

Lens type	CCTV (most C/CS mount)	35-mm photographic camera	Photographic enlarger	Telecentric	"Zoom" microscope
Magnification range	~1/∞ to 0.10X	~1/∞ to 2×	~0.025 to 2× (reverse these lenses for ~2 to 10X)	~0.05 to 5×	~1/∞ to 20×
Angular field of view	To ~100°	To ~200°		Typically less than 1°	Less than 1° (true telecentric) to ~30°
Optical performance	Poor to good	Good to excellent	Good to excellent	Fair to excellent; some low-cost lenses have high distortion	Good to excellent
Cost	$50 to $250	$150 to $5000	$150 to $1000	$250 to $4000	$500 to $2000
Typical applications	Defect detection, part recognition, pattern matching, OCR and OCV, surveillance	Defect detection, part recognition, pattern matching, OCR and OCV, flat-field gauging, high-resolution and critical applications	Flat-field gauging, defect detection, part recognition, pattern matching, OCR and OCV, high-resolution and critical applications	Gauging of thick objects, gauging of objects with variable depth, refractive-defect detection	High-resolution gauging and defect detection; alignment; long-standoff, high-magnification inspection
Other comments	Usually poor for gauging due to distortion; usually poor for critical or high-resolution defect detection	Designed to image well onto a flat surface, usually good optical abberation correction; wide-angle lenses (~28 mm fl and less) are poor for gauging; many designs available, including "micro" lenses as well as very fast designs (with small f/number for good light collection)	Designed for excellent flat-field performance with low distortion and good to excellent optical-abberation correction; being used more in machine vision as a high quality "workhorse"; usually not designed for high-light-collection –f/4 or larger	Designed for constant-view angle over their field and large depth of field; usually excellent for gauging; true telecentric lenses cannot view fields larger than their diameters; telecentric does not necessarily mean low distortion	Becoming more popular for machine vision; several modular systems allow quick changeover for different magnifications or standoff distances; you can build in light sources along with motorized focus and aperture control

OCR=optical character recognition.
OCV=optical character verification.
Courtesy Spencer Luster of Light Works

RAM and transfers data via onboard DMA (direct-memory access) to minimise host-processor loading. You can select image storage of 8, 16, or 32 bits per pixel in 4:2:2, 4:2:0, or RGB formats. Functions include scaling and zooming by arbitrary factors and video overlays on live video. The card's display controller has dual-head outputs that can simultaneously drive a high-resolution PC monitor and a TV monitor. This feature allows you to record video on a standard recorder while viewing full-detail images. Software support is via a range of Coreco's software-development tools for Windows NT/2000.

If you need digital capture together with Camera Link compatibility, consider National Instruments' NI 1428 frame grabber (**Figure 3**). This PCI card has bus-master capability that can sustain 100-Mbyte/sec transfers across the PC's bus and comes with 16 Mbytes of SDRAM for use as an image buffer. A multitap data formatter accepts the multiple channels that many digital cameras produce to increase the camera's frame rate. The formatter can reorder the data from as many as four taps to place the pixels in the correct on-screen position. The card includes two 64k×16-bit look-up-table memories that you can also configure as four 256×8-bit look-up tables to perform gamma correction or—more usefully in a machine-vision context—contrast enhancement or another nonlinear function. A 68-pin SCSI connector complements the Camera Link connector, adding data and other lines for use as fast triggers or general-purpose digital I/O. The card also supports the RTSI bus, which transfers real-time synchronisation information between compatible cards, such as National Instruments' image- and data-acquisition hardware. An acquisition-, scaling-, and ROI (region-of-interest)- control block monitors incoming video signals and routes active pixels to the data formatter or into SDRAM. The card can perform ROI data reduction on any video frame and line number for output to the system bus; pixel- and line-scaling operations transfer multiples of pixels or lines into onboard SDRAM. The NI 1428 costs $1495 and is available now.

Figure 2

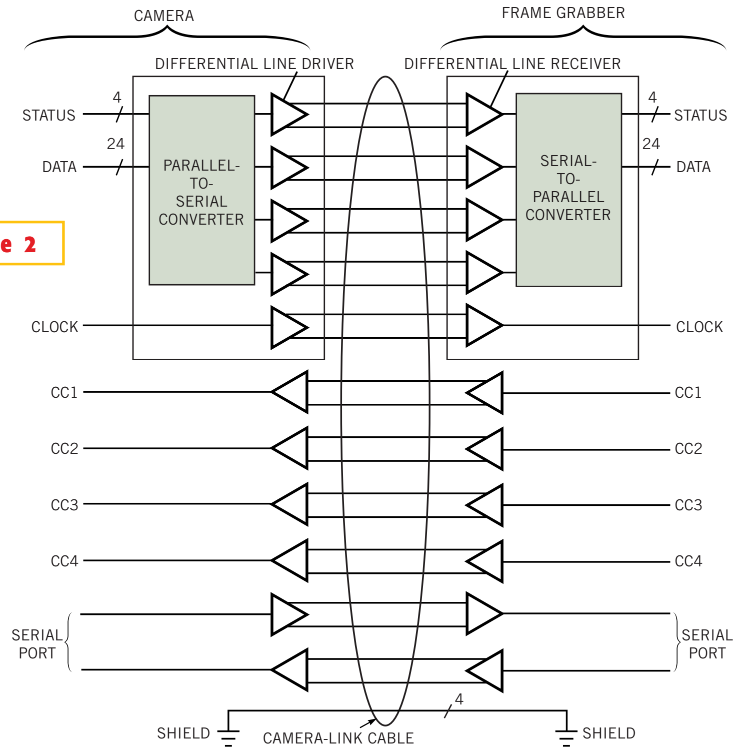

The Camera Link standard rationalises the camera-to-frame-grabber connection but still allows latitude for dedicated I/O signals and future growth.

SOFTWARE DRIVES DECISIONS

Software choices provide another baffling array of options to consider. But from the outset, your choice of frame-grabber hardware typically dictates the

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to www.edn-info.com. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN Europe*.

Adimec
www.adimec.com
Enter No. 412

Amerinex Applied Imaging
www.aai.com
Enter No. 413

Basler Vision Technologies
www.baslerweb.com
Enter No. 414

Cognex
www.cognex.com
Enter No. 415

Coreco Imaging
www.imaging.com
Enter No. 416

Dalsa
www.dalsa.com
Enter No. 417

Data Translation
www.datatranslation.co.uk
Enter No. 418

Hitachi Denshi America
www.hdal.com
Enter No. 419

Light Works
www.lw4u.com
Enter No. 420

Microsoft
www.microsoft.com
Enter No. 421

National Instruments
www.ni.com
Enter No. 422

National Semiconductor
www.national.com
Enter No. 423

Panasonic
www.panasonic.com
Enter No. 424

Pulnix
www.pulnix.eu.com
Enter No. 425

Sony Broadcast & Professional
http://bpgprod.sel.sony.com
Enter No. 426

Teradyne
www.teradyne.com
Enter No. 427

3M
www.3m.com
Enter No. 428

SUPER INFO NUMBER
For more information on the products available from all of the vendors listed in this box, enter no. 429 at www.edn-info.com.

Figure 3

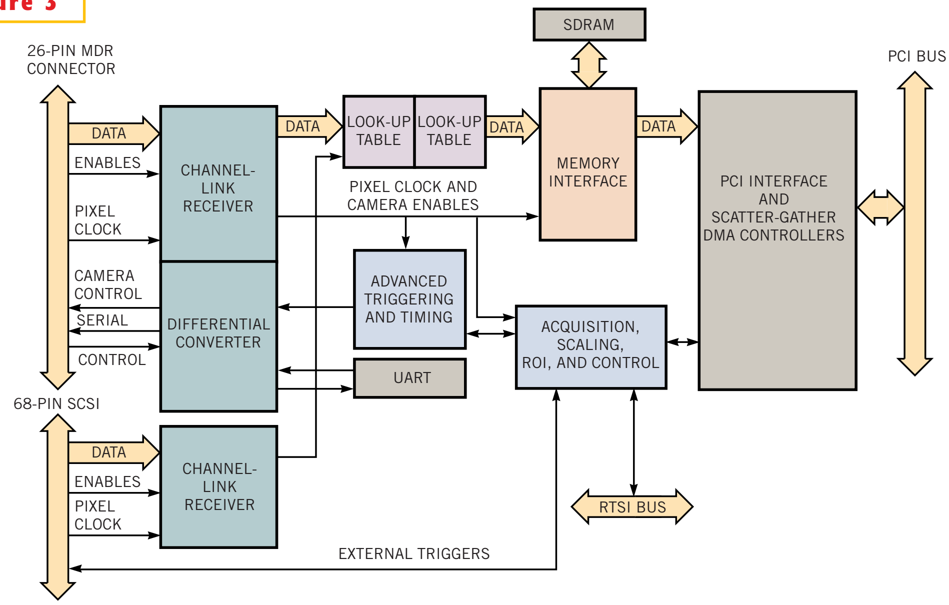

National Instruments' NI 1428 frame-grabber card is Camera Link-compatible and fits PCI-bus slots to sustain 100-Mbyte/sec data transfers.

software-package permutations. Thus, you need to consider both elements to build a suitable system. Machine-vision software includes several fundamental operations that combine geometrical and statistical techniques that include blob (binary-large-object) analysis, edge detection, histogram analysis, morphology, pattern matching, and threshold comparisons. Most applications require these operations on only the ROI, reducing the amount of data to process. Simple quantitative techniques calculate the average pixel intensity or standard deviation of pixel intensities within the ROI to report, for example, the presence of a component on a pc board. Alternatively, a histogram might bin the frame's entire pixel count into each of the 256 values in an 8-bit monochrome image. This step filters an image's characteristics before further processing, such as blob analysis. Blob analysis finds an object's shape by reducing each pixel to a zero or a one value according to a predetermined threshold. By scaling the resulting shape by a representative factor, you can then establish more than 50 parameters, including an object's area, centre, and perimeter.

Morphology changes shapes to improve software-assessment reliability. For example, an image may contain reflections that overlap object perimeters. A morphology function can reduce the overlap area to make it possible to count the number of objects present. But both blob analysis and morphology rely upon high contrast levels to locate an object's edges, meaning that the relevant pixels in the image are at or near a threshold. To locate fuzzier edges within lighting gradients, edge-detection algorithms typically examine the pixels in a line across the object to establish areas of rapidly changing values. Calculating the derivative of a continuous algebraic function locates inflexion points that represent the positive-to-negative or negative-to-positive slope changes that occur at the object's edges. Finding the background, edge information, and shape of a known-good object provides a template for a pattern-matching algorithm, which typically correlates the grey-scale information within the ROI with a template to establish a match.

Now at Version 3 and suiting all modern Windows variants, Data Translation's Vision Foundry package supports monochrome and colour images to 32-bit format (**Figure 4**). The software supports the company's range of frame-grabber boards and provides a point-and-click scripting tool that allows you to build an application by connecting library functions. The scripting tool supports formulas and variables together with nested conditional statements, such as "do-while" and compound "ifs." Library tools include bar-code reading, blob analysis, contour classification, gauging, histograms, image classification, line profile, and optical character recognition. Image-preprocessing routines include arithmetic, filter, morphology, and threshold functions. There's also a set of display tools and system utilities. ROI functions include ellipse, freehand, freehand line, point, polyline, polyfreehand line, rectangle, and straight line. You can declare an unlimited number of ROIs per image and attach them to an image for faster processing. The software dynamically allocates ROIs, which coordinate data from other tools can reposition. Alternatively, you can call Vision Foundry APIs from Microsoft's Visual C++ to create a custom application. And if the package's library functions and APIs can't fulfil your application's requirements, you can create custom algorithms in Visual C++ and GUIs by using the MFC (Microsoft Foundation Classes) interface. Prices for the runtime package start at $995, and it is available now.□

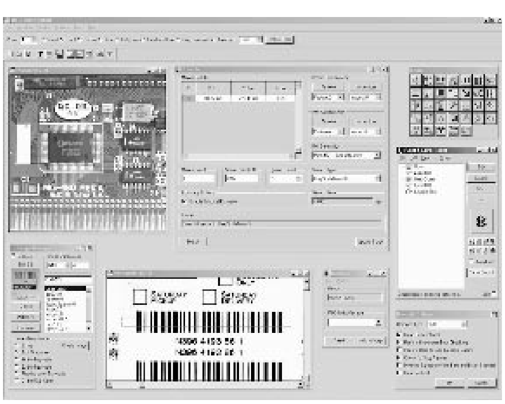

Figure 4 **Data Translation's Vision Foundry software provides an extensible set of image-analysis tools that you can configure using a point-and-click scripting tool.**

You can reach Contributing Editor David Marsh at forncett@compuserve.com.