Active and Passive Range Sensing for Robotics

Martial Hebert
hebert@ri.cmu.edu

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Abstract

In this paper, we present a brief survey of the technologies currently available for range sensing, of their use in robotics applications, and of emerging technologies for future systems. The paper is organized by type of sensing: laser range finders, triangulation range finders, and passive stereo. A separate section focuses on current development in the area of non-scanning sensors, a critical area to achieve range sensing performance comparable to that of conventional cameras. The presentation of the different technologies is based on many recent examples from robotics research.

1.0 Introduction

The use of three-dimensional information is paramount in robotics, from detecting and avoiding obstacles in a three-dimensional workspace, to recognizing objects, to mapping environments. Although simple solutions such as probes and light beams are sufficient in many industrial applications, robotics problems typically require a much higher sensing rate, longer range, and high spatial resolution that cannot be delivered through those conventional means.

Development of range sensing in robotics has used virtually all the possible modalities. Laser range finders were considered as far back as [24], while passive techniques based on stereo vision have been under development for decades. One of the most widely used types of sensing in robotics was ultrasonic sensing owing to its applicability to a number of problems in mobile robotics.

More advanced range sensors, however, are not as widespread as one would hope. The main reason is that the requirements of workspace, speed, and accuracy often led to systems that were too costly and bulky to be of practical use beyond research experiments. Several advances have changed that picture. First of all, in the area of active sensing, advances in solid-state technology and electronics integration have rendered possible the design of compact and relatively affordable active range sensors whose performance is compatible with the requirements of robotics. Furthermore, current development on non-scanning range

sensors as well as of very high speed sensors have opened up new exciting possibilities. Second, in the area of passive range sensing, the tremendous increase in computing power and integration that occurred in the past few years has made passive stereo ranging a viable option for real-time robotic systems.

This paper is a brief survey of the technologies currently available for range sensing, of their use in robotics applications, and of emerging technologies for future systems. The paper is organized by type of sensing: laser range finders, triangulation range finders, and passive stereo. A separate section focuses on current development in the area of non-scanning sensors, a critical area in achieving range sensing performance comparable to that of conventional cameras¹.

Rather than listing sensors and applications with their specs, the paper explores the range sensing technologies through many examples, both from the applications standpoint and from the sensor technology standpoint. Those examples are not intended to constitute an exhaustive list but to provide snapshots of typical activities in the area of range sensing. In particular, many examples are from commercial prototypes, or even products in some cases. Those examples were chosen as typical examples of high-end systems for the robotics applications of the future. That selection is not the result of a systematic quantitative evaluation or of an exhaustive list of available systems. Whenever possible, the source of the information is listed for all those examples.

Finally, the emphasis here is on technologies with demonstrated potential in robotics applications. Except for a brief mention of radar, other ranging techniques are not considered. The presentation is further restricted by emphasizing imaging, or scanning systems that are most relevant to robotics applications.

1. The paper aims at presenting a broad survey of the technologies available and under development. As a result, it is not possible to include detailed performance tables for each device. Instead, a comprehensive bibliography is included at the end of the paper.

2.0 Range Sensing and Robotics

Because much of the development effort in range sensing are concerned with industrial applications such as metrology, surveying, or inspection, it is important to identify the requirements that are specific to robotics applications. Broadly speaking, the first distinguishing feature of robotics applications is that they are dynamic, that is, either the robot is in motion or the robot observes a dynamically changing environment. The impact of this requirement on sensing technology is that data acquisition speed is critical and that many of the advanced concepts developed for the metrology and surveying areas cannot be transferred to robotics.

The second consideration in robotics is the physical packaging of the sensor system. Sensors are typically integrated in robot platforms. Current trends in robotics, e.g., personal robotics, are pushing for smaller robots with more limited power storage. As a result, many of the most advanced ranging systems cannot be transferred to robotics unless significant advances are made in the area of sensor integration.

Finally, operation in difficult environmental conditions, such as low-light, dust, fog, etc. is critical in many field applications; thus, new developments in robustness are needed before many of the proposed technologies can be used in robotics settings.

In the remainder of the paper, we emphasize those three axis of speed, integration, and robustness through examples of typical existing ranging systems in robotics and of on-going research in addressing their limitations.

3.0 Scanning Laser Range Finders

Essentially, three basic technologies are used in active laser ranging. AM-CW lasers use the difference of phase $\Delta\Phi$ between emitted and received beam; time of flight lasers measure the travel time ΔT of a pulse; and FM-CW use the frequency shift of a frequency modulated laser for measuring range. Although those are somewhat generalizations, a rough comparison of technologies can be done as follows: the AM-CW lasers are faster and perform best at close to medium range (e.g., 50m range). They are typically more sensitive to ambient natural light, and therefore more suitable for indoor use. TOF scanners can perform at long range and are best suited for mobile robot application in outdoor settings; however, their acquisition rate is typically lower. FM-CW sensors can be considerably more accurate, but at a cost of more complex design and more brittle packaging. Except for a few experimental exceptions, the successes in robotic applications to date have been obtained with AM-CW or TOF sensors; therefore, FM sensors will not be discussed further.

Although the existing systems in this class of laser range finders tend to be costly, many turnkey systems are now available, a number of which have been demonstrated in robotics research systems. We consider below some examples of systems relevant to robotics applications. We limit ourselves to those sensors that use a mechanical scanning system. (Non-mechanical scanning will be briefly discussed in a separate section.)

A major step toward the widespread use of laser range sensors in robotics has been the introduction of affordable and relatively compact range scanners. Those sensors permit the acquisition of considerably more accurate and denser data than was possible with sonars. In the area of indoor mobile robots in particular, this has been coupled with the recent breakthroughs in concurrent mapping and robot localization [36,8,14] which opened new areas of research and application. A popular illustrative example is the SICK laser scanner [<http://www.sickoptics.com/>] which uses a TOF ranging laser with a 50m range with a resolution of ± 50 mm. The scanner provides a single line scan with 180° coverage. Figure 1 shows the scanner mounted on a RWI robot and a typical result of accurate mapping of a building. This example is taken from the research initiated in [36,14].

Figure 1 Typical example of global mapping using a single-scan laser range finder; (a) SICK laser scanner mounted on a mobile platform; (b) final map of building; (c) initial map from odometry only (from Thrun.)

Because they are based on the detection of a pulse rather than on the integration of a continuous signal which rapidly decays with range, TOF systems can be built to achieve much longer ranges. In the context of robotics, this means that we can now consider new applications in which 3-D data previously could not be directly used, such as AUVs. The Riegl system is a typical example of a TOF system [www.riegl.co.at] operating in the near-infrared -- similar systems exist, for example [32]. The laser can use either single or multiple pulse detection. One-axis or two-axis scanning can be achieved with the basic system. Such a system can easily achieve maximum ranges of one to several hundred meters. The advertised accuracy is ± 5 cm. It should be noted, for such a TOF system, the accuracy does not decrease with range as rapidly as with the AM-CW systems.

Such approaches to laser ranging systems are used in robotic applications involving long-range mapping and navigation. One example is the use of an AUV, an autonomous helicopter, for autonomous mapping of large areas [21]. Figure 2 shows the map obtained by accumulating a large number of scans acquired while flying over an area.

Figure 2 Application of laser ranging to an AUV; (left) autonomous helicopter with scanner; (right) map constructed from repeated scanning of an area.

A major objection to the use of laser sensors is their sensitivity to environmental conditions, e.g., scattering due to fog or dust. A major development in TOF technology over the past few years is the introduction of “last pulse measuring” techniques. As shown in Figure 3, the basic concept is to detect the last echo received from a pulse, or a train of pulses, so that the echo from the surface that is the further away will be used for range computation. When operating under hostile environmental conditions, such as fog, dust, or smoke, this technique guarantees that the range to the target is returned instead of the range produced by the scattering from the medium.

Figure 3 Last pulse measuring technique.

Figure 4 3-D scan of a truck for a robotic excavation application (from [34])

Although this idea has been around for a long time, operational systems are now available. The implication for robotics is that field applications that were not previously possible can now be developed, thus extending the reach of robotics applications. For example, Figure 4 shows the set of points acquired from such a laser scanner. The scene contains a truck and a dirt pile and is typical of the environment used for robotic field applications [34]. Although the environment could generate a great deal of spurious measurement from scattering from the medium, the data is correct and the shape of the object can be easily recovered.

TOF scanners are inherently limited in speed (typically to a few thousand samples per second) and therefore are limited in the spatial density of the data that can be acquired in a scan. AM-CW technology can provide a much higher sampling rate at the cost of the maximum range, which is typically limited to tens of meters, assuming eye safety limitations. Such scanners are mature and have been used in manufacturing applications. For example the Perceptron scanner can measure 256x256 images at 2Hz in workspaces of 10m [7]. Further technology development in this area concentrates on applications that require very high resolution and fast scan rate, for example, autonomous high-resolution mapping. One example of such a development effort is the scanner described in [10] which can acquire 150Ksamples per second at a maximum range of 57m. The range accuracy, i.e., 3σ , is 34mm on reflective targets. An example data set is shown in Figure 5. The high accuracy is obtained by using two lasers offset in modulation frequency. Such a high-end ranging technology does combine high resolution and acquisition speed, which makes it potentially suitable for a variety of robotics applications. A major obstacle, however, is that the rate at which a scene can be scanned is considerably lower than the acquisition rate because of the limitations of the scanning technology. (This issue will be addressed in a separate section.)

Figure 5 Typical example of data from a high-end AM-CW scanner: range and intensity (left) and 3-D view (right)

4.0 Active Triangulation

Historically, active triangulation is one of the first range imaging approaches used in robotics. Recall the basic principle: a stripe (or a single spot) of light is projected

onto the scene. A sensor, usually a CCD camera, views the scene. The depth to the point on the object is found as $z_o = \frac{B}{x_o/f + \tan \alpha}$, where B is the baseline separation of the laser and the camera optical centers, f is the focal length of the camera lens, x_o is the image location within a row where the laser stripe is detected, and α is the projection angle of the laser in respect to the z axis. In order to ease the detection of the laser in the image, the illumination conditions are usually adjusted so that the projected laser generates the brightest features in the scene.

A fairly wide selection of commercial scanners based on this principle is available. Figure 6 shows one example, the Minolta VIVID scanner which acquires a complete scan in 0.6 second. A complete scan produces a 200x200 range image and a 400x400 color image in an operating range of 0.6 m to 2.5 m. The various incarnations of the basic triangulation technology differ in the beam forming technique, the focusing technique used to ensure sharpness of the stripe in the image, e.g., the Minolta uses an auto-focus camera to get maximum sharpness on the target, and in the implementation of the scanning mechanism.

Figure 6 Typical data from a triangulation light-stripe sensor -- the Minolta VIVID.

The attractive aspect of the active triangulation approach is the simplicity of its implementation. However, although the basic triangulation technology as described above has been successful in object modeling and recognition applications, it has several major drawbacks that need to be addressed in the context of robotics applications. We now turn to those robotics-specific issues and solutions currently being investigated.

The first problem is that the resolution of a triangulation system is directly related to the baseline of the system. The resolution relative to depth of field typically needed to robotics applications can be achieved only by increasing the baseline, thus increasing the occurrence of occlusions and missing data, and making the size of the system impractical for on-robot embarked systems, except for very short range applications. The situation is complicated by the fact that the projected beam must stay in focus over a wide depth of field, consistent with the workspace typically required in robotics systems.

There has been considerable effort in addressing this problem. One of the most successful attempts is the synchronized scanning approach, originally introduced in [25]. In this approach, both the emitted beam and the receiver optics are scanned so that the projected beam remains in focus over a large depth of field, and high resolution is maintained despite a short physical baseline. The original idea has been extended to include color by using an RGB laser. The technique can be modified to allow for random access scanning, thus enabling active 3-D target tracking [2]. This example shows the applicability of high-resolution color/range scanners to robotics [3].

A second problem is the frame rate of the scanners. With the basic triangulation approach as described so far, images must be acquired and processed for each position of the beam. Although this can be done efficiently through proper design of the electronics, the basic principle does not allow frame-rate data acquisition. An obvious solution to this problem, proposed many years ago, is to project a structured pattern over the entire scene and recover depth through triangulation by processing only a few -- or even one -- image. Early examples of this technique include the OGIS camera [28]. Although this is an obvious and simple approach, it is dramatically limited in practice for robotics because of its sensitivity to ambient light and its limited range, which stems from the fact that it is difficult to project enough power in the scene at long range. For short range applications in controlled lighting, the projected pattern approach continues to receive some attention.

A more promising approach is that of time-to-range: Suppose that the laser is continuously swept across the scene, say from right to left. Each pixel in the sensor has its own line of sight and "sees" the laser stripe only once as it sweeps by. Based on the observation that, if we can record the time t_o at which a particular pixel at location x_o sees the laser, the range is calculated as $z_o = \frac{B}{x_o/f + \tan \omega t_o}$, where ω is the angular velocity of the mirror.

In other words, we can continuously scan the scene and the range at each pixel will be reported as soon as the pixel is triggered by the laser stripe, that is, as soon as the pixel sees a maximum of intensity. The acquisition rate in this case is limited only by the maximum speed of the mechanical scanning apparatus and by the minimum integration time at each pixel. In both cases, the theoretical speed is far greater than the speed achievable through conventional means. Most importantly, the speed does not depend on the image size since the "computation" is performed simultaneously at each pixel.

The main challenge in this approach is that each sensing element of the receiver must be equipped with additional triggering electronics, thus causing a serious integration problem. Current VLSI integration technology does make

such an approach viable.

Figure 8 Principle of the time-to-range approach.

Theoretical speeds of 1000 frames/sec. are ultimately targeted for such time-to-range triangulation systems. While early attempts [41,29,9] did not match those goals, they demonstrated the feasibility of such an approach and 10 frames/sec acquisition rate independently of image size. Application to 3-D object tracking for robotics were reported in [33]. More recently, taking advantage of advances in VLSI integration of computational sensors, 100 frames/sec sensors were demonstrated in [4] -- several other efforts are described in the same proceedings. Similar approaches are being investigated for specific robotic applications, e.g., fast range sensing for obstacle detection for mobile robots [13]. Although the obstacles, in particular the integration required for full-frame range sensing, remain formidable, this technology is a promising and attractive path toward fast range sensing for dynamic robotic applications.

5.0 "Scannerless Scanning"

One of the key limitations of today's range sensors is the inability to acquire dense data at a sufficiently high rate. The limitation comes from the mechanical constraints on the scanning mechanism; these prevent the motion of the deflection system from keeping up with the actual range measurement rate. Ideally, one would like a sensor in which the scanning is performed electronically or without any scanning at all. This would also address the issue of sensor size which is still a major obstacle to their practical use.

Several concepts have been proposed for non-mechanical scanning. One example is the use of acoustico-optical devices AO for beam steering. AOs are crystals placed in front of the optical path of the beam with the property that their optical characteristics change when an acoustical signal is applied to it. By properly modulating the signal, the

crystal deflects the beam according to a pre-programmed pattern, thus effectively scanning the scene without moving parts. An extensive literature exists on AOs which have been used for a variety of applications [37]. Research on beam deflection using AOs was originally motivated by applications to inspection, image formation, and 3-D displays [30,23]. Recent efforts have concentrated on extending both the range and angular deflection of those devices. A comparison of current techniques can be found in [35].

A key issue in this approach is the loss of power incurred as the beam is deflected by the AO device. Additional complex optical effects affecting focusing and diffraction of the beam are additional challenges under investigation. Other approaches based on solid-state techniques borrowed from the data storage industry do not suffer from those obstacles. However, they still do have limitations in the maximum angle that can be scanned as well as other issues with the allowable transmitted power.

Figure 9 Acoustico-optical scanning.

Another approach is eliminate scanning altogether. The basic idea is to send a laser pulse that covers the entire field of view in the scene. The returned pulse is then observed by an array of receivers. The key is that the returned pulse will be received at different times for receivers corresponding to points at different ranges. The problem is to then efficiently detect the time difference at each receiver simultaneously. Although the principle is straightforward, the technical challenge is to be able to detect such small time intervals over a large array efficiently. One early example of such a device was implemented by Sandia Natl. Lab [1] - incarnations of this technology are sometimes referred to as "flash lasers."

Practical implementations include range gating of the received pulse through fast electronic shuttering, and off-line processing of the time data. It is also possible to compute the time of flight directly at the receiver by handling the charge created during sensing at each pixel of a CCD array [22]. Another benefit of this approach is the simultaneous acquisition of range and color data since the receivers are co-located. Examples of recent implementation of this technology can be found in [26,31,11]. Quantitative evaluation techniques for those devices are proposed in [27].

Progress reported in the literature on such devices is promising. Several difficult technical challenges needed to be

overcome, however. High resolution relative to depth of field is difficult to achieve. In fact, some of the earlier examples of this technology were for measuring ranges much longer than those normally used in robotics applications; short range applications are more challenging. The electronics required for the range measurements is also an issue and will require more integration before being cost-effective and practical.

Although the deployment of those technologies in truly turnkey systems for robotics is still in the future, those lines of development open up the exciting possibility of having true range cameras. For example, the Zcam device from 3DVSystems [38] provides full-frame range/color images at 30Hz with a resolution of up to 10 bits in depth.

Figure 10 Range channel (left) and color channel (right) from a sequence acquired using the zcam range camera (the images are slightly out of phase because of the capture program used for generating the illustration.)

Although the cost and size of this type of device still place it out of reach of robotics work -- it was designed for the film industry-- it is obvious that the development of such sensors will have a substantial impact on robotics applications.

6.0 Passive Stereo

Passive stereo is also a triangulation technique with the same geometric characteristics as active triangulation, except that the range is computed by triangulation between the locations of matching pixels in images rather than between a known source and an observed pixel. Correspondences between pixels are established by searching through the image and using correlation or sum of square differences measures to compare local neighborhood. The raw output of a stereo system is an image of the disparity -- or, equivalently, inverse range -- between images at each pixel. The literature on stereo vision is too massive to be discussed here; we limit ourselves to those considerations that bear directly on the robotics applications of stereo.

Although passive stereo vision is one of the oldest research topics in the Computer Vision community, its use in robotics was limited by the large amount of computation required, basically the equivalent of dozens of correla-

tion operations at each pixel. The rapid advent of computing power has generated a revolution in the way stereo systems are implemented and used. This trend has accelerated owing to the introduction of parallel image computing hardware, such as MMX, in standard hardware. Only a few years ago, stereo systems could barely generate a few frames per second at reduced resolution, even when using custom hardware or special-purpose image processors such as Datacube boards. Today, a large number of operational systems operate at near-frame rate on conventional PC hardware.

Beyond the computation issues, stereo suffers from limitations due to the triangulation geometry. Specifically, the trade-off between high resolution (large baseline) and reduced ambiguity in matching (small baseline), which also existed for active triangulating systems, is exacerbated in the case of passive systems. Stereo systems with three or more cameras do address this problem. The effect of using shorter baselines is compensated for by integrating matching results from multiple cameras. Practical trinocular stereo systems were made possible not only because of the increased computational power but also because of a better understanding of the geometry of the stereo problem, including development in calibration, a result of fundamental research in the Computer Vision community.

Examples of real-time stereo systems now abound in robotics. For example, the Point Grey system runs on a conventional PC and is able to compute one million pixels per second by taking advantage of the MMX architecture [www.ptgrey.com]. If a reduced frame size is used, 30 frames/sec can be achieved. Figure 11 shows an example of data from that system in an outdoor scene. Range resolution and matching accuracy are achieved by using a trinocular system. Another example is the SVM system from SRI [12], which also uses computation on a conventional PC to deliver 320x240 range images at 12Hz with 16 disparity levels. This system has been used in a number of applications for navigation, localization, and mapping. Interestingly, the acquisition rate is fast enough to do motion detection and people tracking [6], which are greatly simplified by using range information.

As those examples show, recent developments in stereo systems place the possibility of real-time range acquisition on the horizon. Perhaps more important than sensing speed is the fact that, because they rely only on conventional cameras, the physical packaging of stereo systems can be made fairly small. For example, Figure 12 shows the layout of the Point Grey trinocular system which is a little over 100mm on the side. Rapid decrease in packaging size is amplified by the advent of progressive scan board cameras as standard commercial equipment.

Figure 11 Example data from the Point Grey trinocular stereo system.

Figure 12 Compact layout of stereo system

The impact of size reduction on robotics applications is considerable. As we have seen, laser and structured light systems are still too large for many robotic systems and require further integration. In contrast, the small stereo systems can now be easily integrated in the small robots that are becoming increasingly important. For example, [18] describes a man-packable robot from IS Robotics, equipped with a stereo system from JPL (Figure 13.) This stereo system is able to generate 120x128 range images at 5Hz. The range images are used for obstacle avoidance in outdoor, unstructured terrain at travel speeds close to 1m/s.

This example is typical of the impact that the recent progress in stereo has had on robotics systems: The stereo package, including computing on a standard PC104 CPU, is small enough and of sufficiently low power to provide ranging capabilities where laser systems would still be too large for rugged terrain driving. As a result, capabilities for rugged terrain driving on a small robot can be demonstrated; those would not have been possible a few years ago, and will lead to new developments in future small-scale robotic systems [39].

Figure 13 The JPL stereo system on the Urbie robot

The use of stereo for navigation is becoming widespread, both for short range application, as shown above, and also

for longer range, higher speed driving. For example, the work on unmanned ground vehicles reported in [19,20] involves the demonstration of mobile robots with target speeds of 20mph operating in unstructured environments. In that case, ranging and obstacle detection must take place at ranges of tens of meters.

Stereo has even been applied to applications in which ranging at long distances. e.g., 100m, is required. For example, the system described in [40] is able to reliably detect 14cm obstacles at over 100m for on-road obstacle detection. This system, illustrated in Figure 14, computes stereo disparity after rectification of the image with respect to the anticipated ground plane. A similar technique, known as focused stereo [5], had been used before but required specialized hardware for real-time performance [17]. At the current rate of progress in stereo, it is clear that the same real-time algorithm for vehicle guidance applications will soon be operational on conventional processors [15].

As mentioned before, another dimension of many robotics applications is robustness to environmental conditions. In that respect, additional developments [20] include the use of conventional infra red cameras or of light intensifiers for night range sensing, another example of how the use of stereo permits the conventional use of sensing hardware. Those developments open the door to continuous operation of stereo systems.

Figure 14 Typical example of real-time stereo processing from obstacle detection (from [40])

Owing to the use of multiple cameras to increase resolution, it is now possible to use stereo for mapping applications which require higher range resolution than that required by pure navigation applications. One example is reported in [16] in which a track robot is equipped with a trinocular stereo system which is used for building maps of an unstructured environment -- the application is the mapping of the inside of the Chernobyl reactor. The robot, stereo system, and an example map built from multiple range images are shown in Figure 15. This example is sig-

nificant because active ranging is typically not a viable option in applications such as robotic remediation of nuclear sites. Thus, the advances in stereo ranging may open up a new arena for more capable robotic systems.

Figure 15 Typical example of autonomous mapping from stereo; robot and trinocular head (left) and reconstructed map of the environment (right).

Finally, stereo has advanced from research work in Computer Vision to a practical solution to many robotic problems through algorithmic advances, increase in computing power, size reduction in computing and sensing. Fast range imaging with an accuracy compatible with navigation and mapping application is now possible, enabling the implementation of systems not possible with the larger, more costly laser systems.

Despite all those successes, stereo still has substantial drawbacks which are current topics of development. Since it relies on ambient illumination, stereo is much more sensitive to environmental conditions, in particular for outdoor applications. More limiting is the fact that the range resolution and accuracy of stereo are still far below what can be achieved with active techniques, such as TOF. As a result, despite cost and size, active scanners still remain preferred solution for applications in which accuracy matters most. Because the only way to alleviate this problem is to increase image resolution, another leap in computing power will probably be necessary before stereo closes that gap.

7.0 Conclusion

Advances in sensor design, electronics integration, and computing power have allowed the development of new solutions for range sensing systems. Performance is dramatically improving on all fronts: acquisition speed, accuracy, packaging size, and robustness. This progress is having an impact on robotics, not just in improved performance of existing robotic approaches, but also in new opportunities and lines of research. For example, affordable and accurate laser scanners now permit the development of practical approaches to 3-D mapping, an advanced which has spurred a flurry of activity in this area. Similar systems with long-range capabilities open up the possibil-

ity of using 3-D mapping and navigation techniques for AUVs. New techniques for robust operation of range sensors in hostile environmental conditions have substantially increased the possibilities in field robotics. The rapid advances in stereo systems over the past few years offer the opportunity for using range imaging on small-size robots.

Despite all those successes, the field is still quite a way from the "dream" range sensor for robotics: short to long range, better than frame rate acquisition speed, high resolution, color data, small size comparable to video cameras and low power draw. Some of the research directions indicated above attempt to bring such a dream closer to reality. For example, active research in scannerless approaches and to on-chip integration of ranging functions will be critical in moving closer to a true range imager. Research and development in stereo is progressing quickly toward general-purpose range imagers.

Acknowledgements

The author wishes to thank the following individuals and institutions for providing illustration and data: Sebastian Thrun (Fig. 1), Omead Amidi (Fig. 2), Tony Stentz (Fig. 3 & 4), Dirk Langer, Z+F Inc. (Fig. 5), Z-Cam Inc. (Fig. 10), Point Grey Research (Fig. 11 & 12), Larry Matthies (Fig. 13), Scott Thayer (Fig. 15).

References

- [1] J.P. Anthes et al. Non-scanned RADAR imaging and applications. *Proc. SPIE. Vol. 1936*, pp. 11-21, 1993.
- [2] J.A. Berardin, F. Blais, M. Rioux, M. Cournooyer, D. Laurin, and S. MacLean. Short and medium range 3D sensing for space applications. *Proc. SPIE Aerosense 97, Visual Information Processing VI*. Vol. 3074, pp. 29-46, Orlando, FL, 1997.
- [3] F. Blais, M. Rouix, J. Domey, and J.A. Berardin. On the implementation of the BIRIS range sensor for applications in robotic control and autonomous systems. *Proc. of the Canadian Conference on Electrical and Computer Engineering*, Vol. 1, pp. 37.1.1, Ottawa, Ont. 1990.
- [4] V. Brajovic. A Triangulation-Based VLSI Range Mapper. *In Proc. ISA/Expo'98*, 1998.
- [5] P. Burt, L. Wixson, G. Salgian. Electronically directed "focal stereo". *In Proc. of the Fifth International Conference on Computer Vision*. pp. 94-101, Cambridge, 1995.
- [6] C. Eveland, K. Konolige, R.C. Bolles. Background modeling for segmentation of video-rate stereo sequences. *In Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1998.
- [7] P. Forno. 3-D robot guidance with the PERCEPTOR LASAR System. *In Proc. SPIE vol. 2271. Industrial Applications of Laser Radar*, 1994.
- [8] D. Fox, W. Burgard, and S. Thrun. Markov localization and applications. *In Modeling and Planning for Sensor-Based Intelligent Robot Systems*. Springer Verlag, Berlin, 1999.

- [9] A. Gruss, L.R. Carley and T. Kanade. *Integrated Sensor and Range-Finding Analog Signal Processor*. IEEE JSSC, vol. 26, No. 3, March 1991.
- [10] J. Hancock, D. Langer, M. Hebert, R. Sullivan, D. Ingimarson, E. Hoffman, M. Mettenleitner, C. Froehlich. Active Laser Radar for High Performance Measurements. In *Proc. IEEE International Conference on Robotics and Automation*. May 1998.
- [11] Y. Huimin, L. Zhukang. Scannerless laser three dimensional imaging method. In *Proc. of the SPIE*. vol.3558, p. 49-52. 1999.
- [12] K. Konolige. Small Vision Systems: Hardware and Implementation. *Eighth International Symposium on Robotics Research*. 1997.
- [13] H. Lamela, E. Garcia, A. De La Escalera, M. Salichs. A new laser triangulation processor for mobile robot applications: preliminary results. In *Proc. Intelligent Components for Vehicles*. 1998.
- [14] F. Lu and E. Millios. *Globally consistent range scan alignment for environment mapping*. Autonomous Robots, Vol. 4. 1997.
- [15] Q.T. Luong, J. Weber, D. Koller, J. Malik. An integrated stereo-based approach to automatic vehicle guidance. In *Proc. of the Fifth International Conference on Computer Vision*. pp. 94-101. Cambridge. 1995.
- [16] M. Maimone, L. Matthies, J. Osborn, E. Rollins, J. Teza, S. Thayer. A photo-realistic 3-D mapping system for extreme nuclear environments. In *Proc. IEEE/RSJ International Conference on Intelligent Robotic Systems*. Victoria, B.C. 1998.
- [17] R. Mandelbaum, M. Hansen, P. Burt, S. Baten. Vision for autonomous mobility: image processing on the VFE-200. *Proc. of the 1998 IEEE International Symposium on Intelligent Control*. 1998.
- [18] L. H. Matthies, C. Won, M. Hebert, L. E. Parker, G. Sukhatme. System design of a portable mobile robot for urban reconnaissance. *Proc. SPIE AeroSense Conference, Unmanned Ground Vehicle Technology workshop*. Orlando. 1999.
- [19] L. Matthies, A. Kelly, T. Litwin, G. Tharp. *Obstacle detection for Unmanned ground vehicles: A progress report*. Springer-Verlag. 1998.
- [20] L.H. Matthies, T. Litwin, K. Owens, A. Rankin, K. Murphy, D. Coombs, J. Gilsinn, H. Tsai, S. Legowik, M. Nashman, B. Yoshimi. Performance evaluation of UGV obstacle detection with CCD/FLIR stereo vision and LADAR. *Proceedings of the 1998 IEEE International Symposium on Intelligent Control*. 1998.
- [21] J.R. Miller, O. Amidi. 3-D Site Mapping with the CMU Autonomous Helicopter. In *Proc. of the 5th International Conference on Intelligent Autonomous Systems (IAS-5)*. June, 1998.
- [22] R. Miyagawa, T. Kanade. *CCD-Based Range-Finding Sensor*. IEEE Transactions on Electron Devices. Vol. 44, No. 10. Oct. 1997.
- [23] Y. Nishiyama, H. Tsukahara, Y. Oshima, F. Takahashi, T. Fuse, T. Nishino. Development of high-speed 3D inspection system for solder bumps. In *Proc. SPIE Vol. 3306. Machine Vision Applications in Industrial Inspection VI*. 1998.
- [24] D. Nitzan, A.E. Brain, R.O. Duda. *The measurement and use of registered reflectance and range data in scene analysis*. Proceedings of the IEEE, vol.65, no.2, p. 206-20. 1976.
- [25] M. Rioux, F. Blais, J.A. Beraldin and P. Boulanger. Range imaging sensor development at NRC laboratories. *Proceedings of the Workshop on Interpretation of 3-D Scenes*, Austin, TX. November 27-29. 1989.
- [26] J.T. Sackos, R.O. Nellums, S.M. Lebien, C.F. Diegert, C.F., J.W. Grantham, T. Monson. Low-cost, high-resolution, video-rate imaging optical radar. In *Proc. of the SPIE*. vol.3380. *Laser Radar Technology and Applications III*. 1998.
- [27] J.T. Sackos. Building accurate geometric models from abundant range imaging information. In *Proc. of the SPIE*. vol.3065. *Laser Radar Technology and Applications II*. 1997.
- [28] K. Sato, S. Inokuchi. *Range picture input system based on space-encoding*. Transactions of the Institute of Electronics and Communication Engineers of Japan, Part D, vol.J68D, no.3. 1985.
- [29] Y. Sato, K. Araki, and S. Parthasarathy. High speed rangefinder. *Proc. Optics, Illumination, and Image Sensing for Machine Vision II*, SPIE, vol. 850, pp. 184-188, 1987.
- [30] V.V. Savelyev, P.E. Tverdokhlebo, A.V. Trubetskoi, Y.A. Shchepetkin. *Three-dimensional image generation by means of a cascade high-speed acousto-optical deflector*. In *Optoelectronics, Instrumentation and Data Processing*, no.2. 1997.
- [31] W. Schroeder, E. Forger, S. Estable. *Scannerless laser range camera*. Sensor Review, vol.19, no.4, p. 285-91. 1999.
- [32] Schwartz ElectroOptics, Inc., Autosense Scanner. www.seord.com/comsen/autosen.htm
- [33] D. Simon, M. Hebert, T. Kanade. Real-time 3-D pose estimation using a high-speed range sensor. *Proc. International Conference on Robotics and Automation*. May 1994.
- [34] A. Stentz, J. Bares, S. Singh. A robotic excavator for autonomous truck loading. In *Proc. IEEE/RSJ International Conference on Intelligent Robotic Systems*. Victoria, B.C. 1998.
- [35] D.J. Svetkoff, D.B.T. Kilgus. Comparison of beam deflectors for ultra-high speed 3D imaging. *Proc. SPIE, Vol. 2909. Three-Dimensional Imaging and Laser-Based Systems for Metrology and Inspection II*. 1997.
- [36] S. Thrun, D. Fox, and W. Burgard. Probabilistic mapping of an environment by a mobile robot. In *Proc. of the IEEE International Conference on Robotics and Automation*. 1998.
- [37] C.S. Tsai. Recent progress on guided-wave acoustooptic devices and applications. In *Proc. First European Conference on Integrated Optics*. 1981.
- [38] www.3dvsystems.com/zcam_fs.html
- [39] C.R. Weisbin, J. Blitch, D. Lavery, E. Krotkov, C. Shoemaker, L.H. Matthies, G. Rodriguez. *Miniature robots for space and military missions*. IEEE Robotics & Automation Magazine, vol.6, no.3. 1999.
- [40] T. Williamson, C. Thorpe. A trinocular system for highway obstacle detection. *Proc. IEEE International Conference on Computer Vision and Pattern Recognition*. 1998.
- [41] A. Yokoyama, K. Sato, T. Yoshigahara, S. Inokuchi. Realtime Range Imaging using Adjustment-Free Photo-VLSI - Silicon Range Finder. *Proceedings IROS*. pp. 1751-1758, 1994.
- [42] L. Zhao, C. Thorpe. Qualitative and quantitative car tracking from a range image sequence. In *Proc. CVPR'98*. pp. 496-501. Santa Barbara. 1998.
- [43] L. Zhao, C. Thorpe. Qualitative and quantitative tracking from a range image sequence. *Proc. Computer Vision and Pattern Recognition*. 1998.