Distributed sensing and probing with multiple search agents: toward system-level landmine detection solutions*

Donald E. Franklin†, Andrew B. Kahng and M. Anthony Lewis

Commotion Laboratory, UCLA CS Department
Los Angeles, CA 90024-1596

†U. S. Army Night Vision and Electronic Sensors Directorate
Ft. Belvoir, VA 22060

ABSTRACT

The problem of landmine detection has been studied for decades. Mine detection systems have typically been developed by first identifying a sensor technology, then testing on particular manmade testbeds, then deploying the sensor on a vehicle or man-portable device. Despite much effort, current systems still exhibit gaps between existing and desired capability, e.g., in terms of rate of advance, detection rate, and false alarm rate within demonstration testbeds. In this paper, we propose a new system-level approach to landmine detection. We argue that "the landmine detection problem" cannot be attacked in a piecewise fashion: *system-level solutions* must *simultaneously* consider functional requirements, sensor technologies, models of sensors, the method of sensor application, and the platforms from which sensors are applied. This perspective allows us to shift our focus from the previous emphasis on novel sensor technology, and to go somewhat beyond traditional doctrines governing standoff or man-portable detection. We first propose a new theory of geometric sensing and probing in the mine detection context. Specifically, we propose new formulations of "object identification by probing" which correspond to various sensing modalities. We demonstrate that multiple agents can achieve probe classes that are not serializable for emulation by a single probe agent. With this in mind, our second main contribution lies in proposing a new paradigm for landmine detection, based on (i) close-in observation with simple spectra, and (ii) small, inexpensive, networkable robotic sensing platforms which can act in a cooperative fashion to implement powerful multi-agent probing strategies.

Keywords: landmine detection, geometric probing, autonomous robots, multi-agent search, sensor platforms, multi-spectral sensing

1 FOUNDATIONS

The problem of landmine detection demands multi-spectral sensing as well as sophisticated sensor fusion, signal processing, imaging, and image analysis. Landmine detection has been studied for many decades; reviews such as [28] cite attempted technologies ranging from biological and chemical detectors (including nitrogen-philic bacteria, trained rats and dogs) to various types of physics-based imaging (including radar, thermal IR, X-ray photon backscattering, and acoustic).[1] Mine detection systems have often been developed by (i) identifying a suite

*This work was supported in part by NSF Young Investigator award MIP-9257982 and grant MIP-9223740. ABK and MAL co-direct the UCLA Commotion (Cooperative Motion) Laboratory, which is supported by NSF CISE CDA-9303148. Address correspondence to: Andrew B. Kahng, UCLA Computer Science Department, 3731 Boelter Hall, Los Angeles, CA 90024-1596. E-mail: *abk@cs.ucla.edu*.

[1] We treat "landmine detection" as fundamentally a problem in object (or anomaly) detection. This is for two reasons. First, landmine *discrimination* from arbitrary forms of clutter is in some sense intractable (in the extreme, one might ask whether the system can discriminate live landmines from defused landmines), and many would characterize an acceptable countermine system as one which reliably detects buried anomalies. Second, higher-level countermine tasks such as statistical pattern recognition of scatter

of promising sensor technologies, (ii) building a sensor which integrates these technologies, (iii) demonstrating the sensor on particular manmade testbeds, and (iv) deploying the sensor on a vehicle or a man-portable device. Despite decades of effort, there remains a gap between existing and desired detection capabilities, particularly for such performance parameters as "rate of advance", "detection rate", "false alarm (false-positive) rate", etc.[2]

In this paper, we ask whether the present gap in detection capability stems from a missing synergy between (i) identification of the detection system requirements and (ii) design of the system itself. We argue that future work in landmine detection should make a strong shift in perspective toward *system-level solutions* that simultaneously consider such ingredients as: system functional requirements (in particular, those relating to geometric aspects of the detection task), sensor technologies, models of sensor designs, methodology for sensor application, and attributes of sensor platforms. The system-level perspective allows a shift in focus away from novel sensor technology as the source of a "magic bullet". Also, by formulating detection requirements from basic principles, we can identify new formalisms to describe the detection process, as well as new approaches to detection that lie outside current standard doctrines. Our discussion begins with two motivating observations.

1.1 Motivating observation 1: Geometric contraints on sensor application

Our first observation is that the actual functional requirements for a fielded detection system can vary greatly, especially in light of *geometric constraints* on the sensor application.

- There is no "canonical" landmine detection scenario. "Landmine detection" has often been equated with "detection of subsoil targets such that both false-positive returns (false alarms) and false-negative returns (missed detections) are minimized", but this equation can be misleading when transformed into "performance criteria". For example, in demining activities around the world (i.e., clearing former battlegrounds, former areas of low-intensity conflict, or former ordnance test ranges), users are typically not very sensitive to false-positive rate or rate of advance. Yet, landmine (or unexploded ordnance) detection systems must often meet performance criteria for such parameters.[3]

- Even assuming a particular landmine detection scenario, the functional requirements for a detector can still vary, according to constraints on the detector's application. We believe that the most important of these constraints are essentially *geometric* in nature.

  Example 1: If the landmines are surface-laid (rather than buried) then a 2-dimensional (rather than 2 1/2-dimensional) sensing task applies. Example 2: Depending on the type of trigger mechanism (e.g., based on pressure or inductance), it may or may not be possible to place a sensor directly above a potential landmine target (e.g., on a rotorcraft or a hand-held detector). Figures 1(a) and 1(b) contrast these two possibilities for (say) a radar: (a) when the transmitter and receiver cannot be above the target, the signal return is reduced, particularly with certain target shapes; on the other hand, (b) when the transmitter and received can be directly above the target, a stronger return is likely. Example 3: In current doctrine, the detector ideally will not be placed above a potential landmine target – or worse yet, above uncleared ground. Clearly, such constraints on the *accessibility* of the region under interrogation will affect the utility of a given detector. Returning to the same radar example, if the receiver is located symmetrically to the transmitter with respect to the target (Figure 1(c)), then a strong return becomes possible. Figure 1(c) also represents an instance of *spatial distribution* of the sensor (see below).

Thus, the geometric constraints tell us the ways in which one can "poke at" the region under interrogation.[4] A main precept in our work is that these constraints dramatically affect the system-level use of sensor technology.

minefields must still address the underlying object (individual landmine) detection problem.

[2] These are among the traditional exit criteria for DOD technology demonstrations involving landmine detection.

[3] Even whether landmine detection should be equated with subsoil target detection is very much an open issue. For example, if "subsoil" means "*emplaced* within the ground" then developers and contracting agencies should not casually separate the two aspects of mine identification (material identification) and emplacement detection (detecting the soil anomaly caused by the emplacement process). The most successful approaches will clearly address both these aspects – but present technologies either center on only one aspect or do not identify which aspect is more important given the underlying sensors.

[4] Note that we are not concerned with environmental factors (moisture, soil granularity, incidence and type of clutter, etc.), which can certainly affect the utility of sensors. Our observations are independent of the environment, and deal only with the geometry of the detection task.

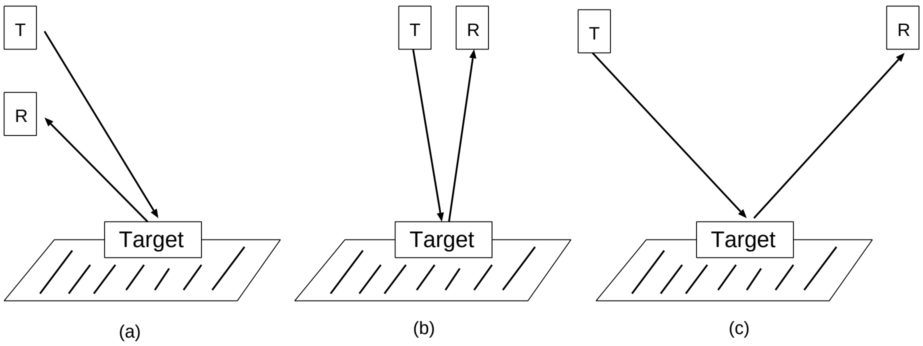

Figure 1: Sensor utility varies with constraints on sensor location. (a) When the sensor (transmitter and receiver) is constrained to be "on the uncleared side" of the target, the return signal may be weak. (b) When the sensor can be placed above the target, a strong return is possible. (c) When the sensor cannot be placed above the target, but can be spatially distributed about the target location, again a strong return is possible.

1.2 Motivating observation 2: It's not the sensor

A second motivating observation is, "It's not the sensor (it's the system)". We believe that *current sensor suites can yield enough information to acceptably solve the detection task* (certainly, the negation of this belief would be a distressing state of affairs!). Of course, there is a "variable selection" problem to be solved: systems designers must determine the best combination of many available spectra and signal processing techniques. However, fundamental obstacles at present seem to be a lack of coupling between detailed physical models of sensing and image classification approaches; a concomitant lack of validated synthetic sensor images; etc. – *not* any lack of a "perfect sensor".

From this observation follows a simple corollary: "Noise is not the problem (speed is the problem)". In other words, because existing sensors are sufficient, simply "poking at" a region long enough will reliably detect landmine targets; the challenge is to reduce the time spent in applying these existing sensors. It is instructive to note the existence of demonstrated technologies for landmine detection where the image resolution gives only a few pixels across the target (e.g., airborne infrared detection of surface-laid mines [18,20]). This suggests the possibility of a *system-level* approach to increasing the speed of detection, by quickly identifying appropriate "pixels" to query within the region under interrogation. We believe that important ingredients of such a system-level solution include: (i) adaptive application of sensors (e.g., changing the sensing strategy based on recent sensor returns [36]), (ii) intelligence in the sensor (e.g., image compression before downlink to an off-board image classifier in a distributed, networked dismounted battle scenario), and (iii) sequencing and hierarchy in the application of sensors (e.g., optimizing the paths of the platforms on which sensors are deployed, or using a preliminary coarse scan to identify locations that merit close examination).

1.3 Toward a theory of distributed sensing and probing

In the remainder of this paper, we develop an approach to landmine detection that has as its basis a geometric, system-level perspective. While no detection system has yet been built following our precepts, our work does provide a framework within which we may address the landmine detection problem from first principles. First, we begin by formulating landmine detection as a type of geometric search in 2 dimensions or 2 1/2 dimensions (i.e., a region with bounded depth). Second, we propose a theory of sensing and probing for detection. The key idea is that any available sensor, together with its application to interrogate the ground, can be viewed as a type of *probe* operation. Such probes can be probabilistic, spatially limited, etc. as necessary to model the underlying physical phenomena of sensing. For a given object detection task, we may establish lower and upper bounds on the number of (given types of) probe operations (or other related information requirements) that are necessary to accomplish the task. Similarly, we may analyze the effort required to apply these probe operations if we are given a model of the sensor platform. (Typically, we consider optimal sequencing of required probe operations.) After these elements have been established, we may pose the question of achieving a prescribed certainty of target detection using minimum resources. We may assume that detection quality increases monotonically with resources expended; thus, we are typically interested in finding cost-performance tradeoff points as illustrated in Figure 2.[5] Third, we give a brief taxonomy of probe operations and their correspondence with example sensing technologies or modalities. Here, we make yet another enabling observation, namely, that certain probe types which assume multiple, spatially distributed sensing platforms can be fundamentally more powerful than any probe achievable using a single sensor platform. In other words, the spatial presence afforded by multiple, distributed sensors cannot be emulated by any single sensor via "serialized parallelism". The problem of efficient sensor application generalizes to, e.g., flavors of the multiple traveling salesman problem. Finally, from our studies within the probing framework and observation of recent trends in the field, we argue that future research in landmine detection should embrace the concept of autonomous, distributed, cooperating, robotic agents that perform close-in detection using simple spectra. This vision allows us to go beyond traditional doctrines of standoff or man-portable landmine detection, and to explore fresh new solution architectures. We conclude the paper with a short review of the UCLA Commotion (Cooperative Motion) Laboratory's efforts toward (i) development of capable, low-cost autonomous robotic platforms which can be used in countermine applications, and (ii) a paradigm for user control of networked, remote taskable robotic hardware.

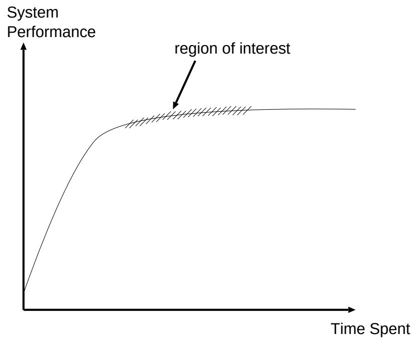

Figure 2: With any system, we may assume that detection quality increases monotonically with time spent in applying sensors, processing sensor returns, etc. Thus, we are typically interested in achieving system performance in the region of interest shown.

2 A GEOMETRIC PROBING FRAMEWORK FOR LANDMINE DETECTION

We view landmine detection as a form of *geometric search*, which may be posed as "Detect all objects of interest within a given 2-dimensional or 2 1/2-dimensional region". Often, the detection must be made with error rate bounded by some prescribed constant. The low-dimensional problem embedding reflects the fact that landmines are flat, convex objects which are emplaced within a bounded distance of the surface, which suggests the application of techniques from computational geometry and motion planning. In practice, the search problem may be simplified by known upper and lower bounds on the object size (e.g., objects with diameter greater than two feet, or less than 2 inches, are probably not landmines), by the fact that there will usually be only a few possible mine types in a given region, and by symmetries and convexity of common mine types.

[5] We believe that a benefit of our approach is that it highlights *algorithmic* speedups for detection, rather than *hardware* speedups; in some sense, the latter are the most obvious and most expensive means of enhancing performance, and thus should be the methods of last resort.

2.1 Sensors as probes

We propose to model sensors as *probe types*, each of which has a given constant cost per application. Beyond the following examples, the reader is also referred to the body of computational geometric work by Skiena [33–36], which established many key results for basic probe classes including several of those mentioned below.[6] Note that by "sensor", we include any sensing operation about which we can reason, e.g., a prone soldier systematically poking around a flag with a wooden stick is a "sensor". Classic probe types include:

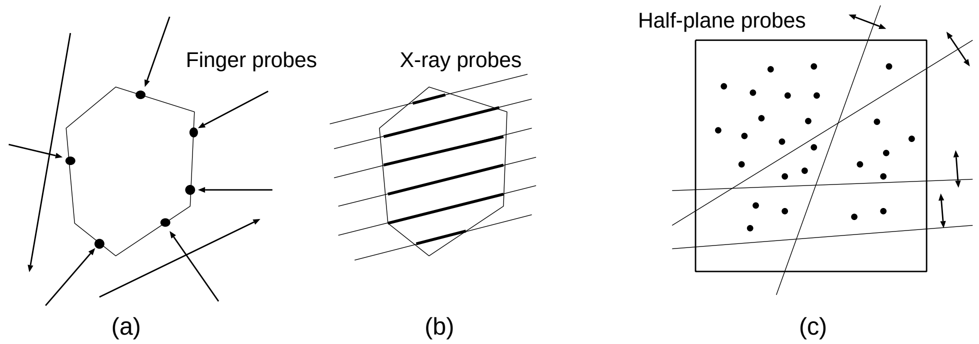

Figure 3: Illustration of three basic probe types formulated by Skiena.

- *Single-point* probes model beam-forming and sounding sensors (e.g., radar) which return a vector of values for each point in the region. The size of the vector can be large (e.g., amplitude and phase returns as the radar is scanned over multiple frequencies), and values can be binary (e.g., when thresholding or bandpass filtering is applied).
- When scanned across a region with maximum resolution, a single-point probe will yield an *image* of the region (i.e., sensor returns at the integer lattice points).[7] If the single-point probe is scanned along a line, the effect is that of a swept vehicle-mounted sensor or a hand-held detector that is swung in an arc. If we additionally threshold the sensor return, we obtain a so-called *X-ray* probe: such a probe returns the intersections of the line and all objects of interest (Figure 3(b)).
- If a single-point probe is scanned along a line and stopped when it first contacts an object, we obtain a so-called *finger probe* (Figure 3(a)). Often, it is assumed that the finger probe can return both the point of contact and the normal to the object boundary at the point of contact; this is a reasonable model for human tactile sensing (e.g., the soldier probing with a stick).
- Finally, when the scan line of a probe is itself moving forward along the normal to the scan line, the finger and X-ray probes respectively generalize to the *hyperplane* and half-plane probes. The hyperplane probe returns the first point of contact between a line moving parallel to itself and an object; the half-plane probe returns the sum of all sensor readings over the area traversed by the moving scan line (Figure 3(c)).

These and other probe types can be extended in various ways to achieve greater fidelity to actual sensing modalities. For example, [21] has proposed *finite* and *dissipative* variants which capture a sensor's range limitation and decreasing accuracy over long distances, as well as *probabilistic* variants which return the correct answer with some probability $1 - \epsilon$.

[6] Skiena treats object types including convex polygons, collections of convex polygons, star-shaped and simple polygons, or simple point sets. Probing objectives in his work include *Determination* (give upper and lower bounds on the number of probes needed to completely determine a particular object), *Verification* (given an object description, bound the number of probes needed to verify the description), and *Feature Determination* (given some object feature such as orientation or convexity, bound the number of probes needed to determine this feature).

[7] Notice that if the scan is coarsened by a factor of k in each axis (basically, looking at every k^{th} coordinate value), we will require a factor of k^2 fewer probe operations to image the region. Associated image processing operations (e.g., computing matrices of autocorrelations) will typically enjoy similar speedups with the coarse scan. Compared with custom hardware, this is a trivial and often more enabling type of system speedup.

2.2 On spatially distributed probe types

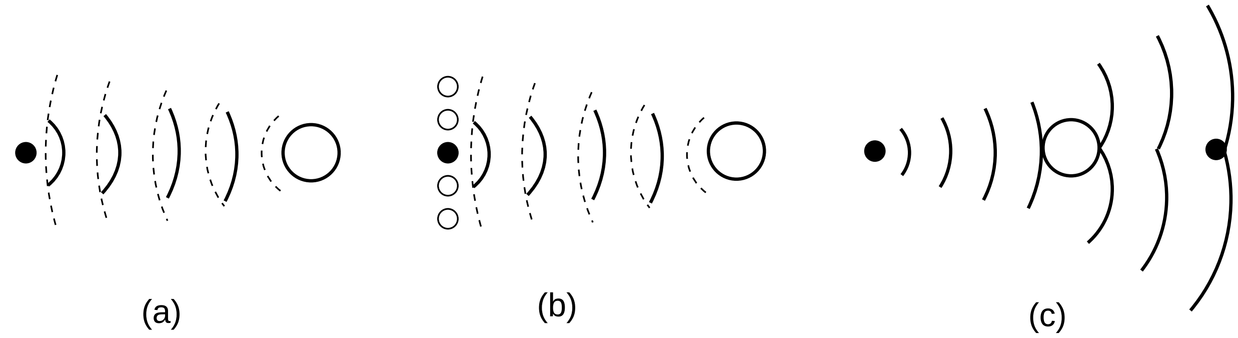

Figure 4: Three configurations of acoustic-based sensors, showing the power of distributed locations, i.e., *spatial presence*, in the sensor. (The drawings are from the perspective of an observer located above the ground.)

Figure 1(c) hinted that *spatial distribution* can increase the power of a sensor. Here, we give a second example which more crisply illustrates the fundamental advantage of spatial distribution. Consider an acoustic subsoil detection approach, where a spark tube containing a column of air or water can be used to "thump" the surface of the earth. As shown in Figure 4(a), surface waves will encounter a discontinuity at the target (hollow circle) and bounce back to the sensor (solid dot). The time of flight yields the distance to the target. Figure 4(b) illustrates that *location* of the target is possible only when the sensor moves (so that triangulation is possible), or when multiple receivers form a receiving array. This location strategy can be complicated when there are multiple targets in the region. Figure 4(c) shows that when we allow sensors to be spatially distributed, a fundamentally different *occlusive* probe type arises. Specifically, when the time of flight is greater than expected between the source and receiver, we may infer the presence of an obstacle between the two locations. The key observation is that it is difficult (or impossible, depending on assumptions) to "serialize" the occlusive probe, i.e., emulate it via a sequence of probe operations using a single sensor.[8] That spatial presence yields strictly more powerful probe types motivates our concept of distributed multi-agent mine detection in Section 3 below.

2.3 Issues in geometric probing for detection

Having modeled sensors as probe types, it is natural to seek useful complexity bounds and algorithmic solutions for the problems of object detection and discrimination. Key questions are:

- How many probes of a given type are needed to recognize or discriminate given classes of objects?
- Which probes should be applied, and in what sequence, for maximum possible efficiency?
- How should the sensor platform move through the region to maximize efficiency, safety, and other objectives?

While it is beyond our present scope to discuss these issues in detail, we now provide a brief list of pointers.

First, "static analysis" has been a theme of works by Skiena, Edelsbrunner, and others in the computational geometry literature [1,15,32–36]: the typical result is that a $O(n)$ probes are necessary and sufficient to identify targets that are convex n-gons. Similar results hold for finite, dissipating, and probabilistic probes types [21], with finite range or the need to probe repeatedly (to achieve a given certainty) yielding straightforward increases

[8] The occlusive probe in some sense has a "stronger need" for sensor locations to be known (e.g., by GPS), or for positions to be dead-reckoned so that sudden changes in time of flight can be observed as the source and receiver traverse the region. Observe that the occlusive probe can be coerced into the framework of the X-ray probe (i.e., it is the "or" of the return from the X-ray probe when there is exactly one object). However, the X-ray probe is fundamentally a model of near-field sensing (via a swept sensor), while the occlusive probe is fundamentally a model of far-field sensing.

in the number of probe operations required.[9] For occlusive probes, the complexity of detection will also depend on the minimum size of any target and the minimum allowed inter-target spacing.

Second, given that a brute-force probing of a target has complexity proportional to the size of the target (rather than to the size of the target's description, e.g., as a polygon with n sides), intelligent application of probes offers significant potential speedup. Typical results are heuristics for sequencing of probe operations, often cast as a traveling salesman optimization [22]. More interesting formulations are on-line, or "dynamic", i.e., the probe operations are adaptively chosen and sequenced based on the results of preceding operations.

Finally, little is known about the application of multiple sensors in optimally probing a given region. It is conceivable that "cooperative speedups" [13] are possible, i.e., that k sensors operating simultaneously can be expected to clear a region in less than $1/k$-th the time required by a single sensor. Our current research efforts center on provably good strategies for this sort of cooperative probing.[10] It should be noted that the subject of "cooperative mobile robotics" has been actively studied in recent years; see [9] for a comprehensive review. Of some interest are works on multi-agent "foraging" [37,2,3,17,25,14,6]; Figure 5 shows a cartoon of a potential-field based strategy where robots are repelled by each other (so as to disperse over the region), and are attracted to unexplored candidate landmine targets. Other related works in the robotics literature deal with distributed movement into "formation" (e.g., if the sensors are to evenly divide up the coverage of a region, or if the platforms must implement a dragnet that can be moved over the region) [29,38,11,12,39].

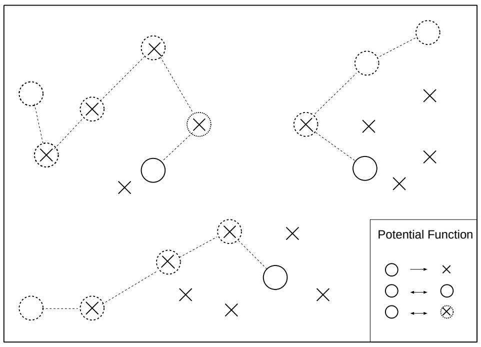

Figure 5: Cartoon of multi-agent "foraging" to examine candidate landmine targets, via a potential-field approach. The three agents are represented by circles; target locations are X's, and locations that have already been explored are circled X's. The potential function heuristically repels the agents from each other and from explored locations, while attracting agents to unexplored locations.

3 TOWARD AUTONOMOUS ROBOTIC SENSING PLATFORMS FOR COUNTERMINE

We have tried to argue that the science of landmine detection can benefit from (i) escaping the preoccupation with novel sensor technology, and adopting a system-level perspective, (ii) considering fundamental requirements for detection, e.g., by modeling sensing modalities as probes, and (iii) considering the concept of spatially distributed sensors, which are strictly more powerful than traditional monolithic sensor implementations. The system-level perspective ("speed is the issue") suggests hierarchical solutions, e.g., using a rotorcraft-mounted

[9] If the range of a finite or dissipating probe is d, then the number of probe operations typically blows up by a factor of L/d or L^2/d^2, where L is the longest dimension of the region under interrogation. With probabilistic probes, there is typically a constant-factor increase in the number of probe operations for any required level of certainty.

[10] Problem-specific assumptions will affect the optimal probing strategy, chiefly because information requirements for object detection will change. Example issues: multiple targets can cause interference with certain types of probes; the architecture of information transfer between sensors (e.g., whether location information is distributed, bandwidth of the communication topology, etc.) affects the type of cooperation that is possible; and testbed parameters (multiplicity of landmine types, shape of region under interrogation, statistical knowledge of minefield pattern, etc.) also affect the search strategy.

wide-area sensor to identify soil anomalies, then using close-in detectors only to resolve these anomalies. The system-level perspective also allows us to obtain interesting inferences from recent trends in the detection literature. The notion of spatial distribution suggests more powerful distributed sensors which can also yield parallel speedups of the detection process. Finally, probing models allow us to establish tight bounds on the complexity of detection for any given sensing modality.

We conclude this paper with a somewhat speculative proposal of "distributed, autonomous, robotic sensing platforms" for landmine detection. Such a proposition is timely for the following reasons. (1) Recent work allows such robotic platforms to have high capability with respect to sensing/vision, on-board signal processing, and networking; the "foraging for landmines" task is comparatively simple, and feasible, compared to other autonomous applications. The wireless networking allows operator control and is compatible with future "networked battlefield" concepts. Note that robotic platforms can be made compact and easily deployable, and that a multiple-robot system is inherently less vulnerable than a monolithic system to single points of failure. (2) Recent work in landmine detection has reverted to simpler spectra (e.g., visual, IR (thermal contrast), and moisture probing (determining soil anomalies from disrupted capillary action)), with greater resolution. Such sensing is easily mounted on present robotic platforms. (3) The availability of inexpensive robotic platforms erases many of the cost-benefit tradeoffs which have led to current doctrine constraining close-in and standoff mine detection; clearly, this allows entirely new detection strategies. As examples: having a sensor accidentally trigger a landmine becomes more acceptable; remoteness of human operators allows close proximity to targets (and higher imaging resolution), landmine neutralization can be incorporated into a hierarchical system organization ("detecting" sensors can guide a cheap "walking bomblet" robot to neutralize a landmine); and "rate of advance" requirements can be satisfied simply by applying parallelism. (4) Arguably, distributed robotic solutions are the only means of addressing the huge scope of unexploded ordnance and demining tasks that must be performed worldwide. (5) Perhaps most important, the availability of robotic sensing platforms will be enabling to the modeling and simulation of fielded detection systems. Because the dynamics of a slow-moving robot can be well-characterized[11], accurate synthetic image databases can be obtained for sensor development and evaluation. For purposes of, e.g., modeling the performance of detection systems in dismounted battle (say, for the DIS thrust), such a capability will be invaluable.

At the UCLA Commotion (Cooperative Motion) Laboratory, several projects are under way that study issues pertaining to future robotic platforms for distributed remote sensing and actuation. One theme of our work is aimed at determining appropriate paradigms for remote operation of semi-autonomous or autonomous taskable hardware. To this end, we are currently implementing a robotics laboratory testbed that is accessible and operable over the Internet [10]. The system provides remote users with robust taskability of multiple mobile robots in an unmapped environment. It is currently implemented on ten wheeled mobile robot bases hosting Unix workstations, and our near-term plans include extension to aerial and legged robots. Our work has revealed important issues in interface design for remote tasking on mobile robots: we have found that standard network access must be seamless, that the environment must be robust (i.e. pre-emptively multi-tasked with memory protection), and that the interface must be familiar to the user. To accomplish this functionality, we have implemented the system with industry standard hardware (which provides access to mass-market pricing, standard software environments and tools, device drivers, etc.) and accessible front-end software (i.e. a World Wide Web interface).

More specifically, our project consists of ten R-3 robots from IS Robotics, Inc., each equipped with infrared proximity sensors, bump detectors, force-sensing grippers, differential drive, and radio modems. There is also an off-board vision system that localizes the robots to within a three centimeter tolerance; in a fielded scenario this can be replaced with highly integrated GPS. To study remote taskability, each of the robots has been outfitted with Linux running on 486/DX2 processors running at 66 MHz (see Figure 6). A programmer's interface has been developed using a shared-memory paradigm, allowing users to control the robots using any language that allows file access. Thus, tasks can be automated using a graphical user interface, conventional programming techniques, or both. Using PPP (Point to Point Protocol), the robots are available over the Internet. Note that this funcionality is available in a package that is less than fifteen inches tall including unix workstation, sensors, communications, and embedded robot controller. Total development cost (at retail, off-the-shelf pricing) for each platform is under fifteen thousand dollars.

In addition, our laboratory is exploring two aerial platforms for sensors that must be deployed from above. Our current aerial robot (Figure 7) is based on a X-Cell 60 model helicopter that has been modified to accept a high-

[11] As opposed to, say, the dynamics of the detector at the end of the pole that is held by the hand on the swinging arm of the walking soldier (who has indeterminate physique, mental state, fatigue, etc.

Figure 6: Photograph of an IS Robotics R-3 robot, outfitted with on-board Linux workstation and wireless networking in a small (approximately 5 inches on a side) module. A 10-robot system is used to study paradigms for user interaction with remote, taskable sensing/actuation platforms.

power electric motor. The sensors consist of a vertical gyro, three rate gyros, an ultrasonic altimeter, a flux-gate compass with level sensor, and three cameras (two for x-y location determination and one for target acquisition). Computatational needs are satified by an on-board 486/DX4 processor running at 100 MHz. Framegrabbers, data aquisition modules, and PWM synthesis modules are connected as standard PC/104 bus devices. The hierarchical controller has a top-level sequencer that controls the vehicle by activating behaviors in order to achieve a sequence of user-specified subgoals [26]. (These subgoals could, for instance, take the form of a script.) Again, by exploiting standard peripherals and operating environments, we have limited cost for the aerial robot to under eight thousand dollars while minimizing development time.

Our next generation aerial robot, which is a ducted fan prototype, fulfills several needs that a helicopter cannot. First, since a ducted fan shrouds its propellers with a nacelle, risk to operators is minimized and collisions with, e.g., outdoor structures can be tolerated. A prototype four-blade design with a Kevlar nacelle is under development. By differentially controlling the thrust of the four blades, the machine can be positioned in three-dimensional space without the need for complicated articulated control surfaces. In simulation studies, we have begun to explore cooperative pushing and manipulation of objects with groups of ducted fan vehicles. The projected cost of the ducted fan prototype is on the order of ten thousand dollars; current prototypes are approximately two feet in diameter.

Finally, for sensors that do not require aerial deployment, legged robotic vehicles provide platforms that are

Figure 7: Photograph of the aerial robot with all on-board sensing and vision processing.

capable of navigating very rough terrain. The UCLA/USC walking machine, "Geo" (Figure 8), is a reconfigurable walking machine designed to evaluate the walking control mechanism in animals. The machine features a flexible spine and a reconfigurable number of legs (i.e. the robot is capable of being configured as a quadruped or hexapod). The flexible spine can emulate the side-to-side flexing during walking of amphibians as well as the dorsal-ventral flexing of mammals. Each leg has four degrees of freedom and can be outstretched for amphibian walking or tucked under the body for mammalian walking. Each joint is instrumented with position feedback and strain gauges, and each foot has miniature load cells to measure forces between the robot and the ground. In quadruped mode the robot has 19 degrees of freedom, while in hexapod mode the robot has 30 degrees of freedom. The current-generation of robot (shown in the photograph) can be anywhere from six to eight inches long, and is four inches tall; it is capable of supporting on-board computation. Fabrication of a single walking machine, including machining, sensing and actuation, costs approximately six thousand dollars. Future generations will incorporate the strengths of the current controller, while increasing the complexity and payload capacity of the vehicle to accomodate fielded operation.

4 SUMMARY

We have proposed a new theory of geometric sensing and probing within the mine detection context. Examples show that multiple agents can achieve probe types that are not serializable for emulation by any single probe agent; this is fundamentally due to the *spatial presence* of multiple agents. Given these preliminaries, we also propose a new paradigm for close-in landmine detection using simple spectra and inexpensive, networkable robotic sensing platforms which can implement cooperative landmine detection strategies. Our current work creates autonomous robotic sensing platforms which can form the basis of fielded countermine systems. Closely related work is aimed at developing new paradigms for user interaction with, and control of, remote taskable hardware.

Figure 8: Photograph of the UCLA/USC walking machine, "Geo". The robot is shown in its quadrupedal configuration.

5 REFERENCES

[1] E. M. Arkin, M. T. Goodrich, J. S. B. Mitchell, D. Mount, C. D. Piatko, and S. S. Skiena, "Point Probe Decision Trees for Geometric Concept Classes", *Proc. Algorithms and Data Structures Workshop*, 1993, pp. 95-106.

[2] R. C. Arkin, "Cooperation Without Communication: Multiagent Schema-based Robot Navigation," *Journal of Robotic Systems*, vol. 9, no. 3, 1992, pp. 351–364.

[3] R. C. Arkin and J. Hobbs, "Dimensions of Communication and Social Organization in Multi-agent Robotic Systems," *Proc. Simulation of Adaptive Behavior*, 1993.

[4] M. R. Azimi-Sadjadi, D. E. Poole, S. Sheedvash, K. D. Sherbondy, et. al., "Detection and Classification of Buried Dielectric Anomalies Using a Separated Aperture Sensor and a Neural Network Discriminator", *IEEE Transactions on Instrumentation and Measurement*, Feb. 1992, pp. 137-143.

[5] M. R. Azimi-Sadjadi and S. A. Stricker, "Detection and Classification of Buried Dielectric Anomalies Using Neural Networks-Further Results", *IEEE Transactions on Instrumentation and Measurement*, Feb. 1994, pp. 34-39.

[6] R. Beckers, O. E. Holland, and J. L. Deneubourg, "From Local Actions to Global Tasks: Stigmergy and Collective Robotics," *Proc. A-Life IV*, MIT Press, 1994.

[7] K. D. Boese and A. B. Kahng, *manuscript in progress*, 1995.

[8] R. A. Brooks and A. M. Flynn, "Fast, Cheap and Out of Control: A Robot Invasion of the Solar System", *Journal of the British Interplanetary Society*, v. 42, no. 10, Oct. 1989, pp. 478-485.

[9] Y. U. Cao, A. S. Fukunaga, A. B. Kahng and F. Meng, "Cooperative Mobile Robotics: Antecedents and Directions", *Proc. IEEE/RSJ Symp. on Intelligent Robotics and Systems*, Pittsburgh, August 1995, to appear.

[10] Y. U. Cao, T. W. Chen, M. Harris, A. B. Kahng, M. A. Lewis and A. D. Stechert, "A Remote Robotics Laboratory on the Internet", *Proc. INET-95*, Honolulu, June 1995, to appear.

[11] Q. Chen and J. Y. S. Luh, "Coordination and Control of a Group of Small Mobile Robots," *Proc. IEEE Intl. Conf. on Robotics and Automation*, 1994 pp. 2315–2320.

[12] Q. Chen and J. Y. S. Luh, "Distributed Motion Coordination of Multiple Robots," *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 1994, pp. 1493–1500.

[13] S. H. Clearwater, B. A. Huberman and T. Hogg, "Cooperative Solution of Constraint Satisfaction Problems", *Science* 254(5035), 22 Nov. 1991, pp. 1181-1183.

[14] A. Drgoul and J. Ferber, "From Tom Thumb to the Dockers: Some Experiments with Foraging Robots," *Proc. Simulation of Adaptive Behavior*, 1993.

[15] H. Edelsbrunner and S. S. Skiena, "Probing Convex Polygons with X-Rays", *Society for Industrial and Applied Mathematics*, Vol. 17, No.5, 1988, pp. 870-882.

[16] W. L. Gans, R. G. Geyer and W. K. Klemperer, "Quantifying Standard Performance of Electromagnetic-Based Mine Detectors", *technical report* NISTIR-3982, National Institute of Standards and Technology, Oct. 1991.

[17] S. Goss and J. Deneubourg, "Harvesting by a Group of Robots," *Proc. European Conference on Artificial Life*, 1992.

[18] N. D. Grande, "Temperature Evaluated Mine Position Survey (Temps) Application of Dual-Band Infrared Methodology", *Proc. IRIC Passive Sensors*, No. 1, 1990, pp.365-384.

[19] G. G. Grinstein, M. T. Maybury, R. B. Mitchell, "Intelligent Virtual Interfaces for Telerobotics", *Cooperative Intelligent Robotics in Space III*, 1992, pp. 401-408.

[20] J. K. Hackett, E. Gold, D. T. Long, E. Cloud, H. Duvoisin, " Model-Based Approach to Real-Time Target Detection", *Automatic Object Recognition II*, 1992, pp. 269-276.

[21] A. B. Kahng, "Architectural and Algorithmic Results Pertaining to Object Recognition", *final report*, ARO Contract DAAL03-86-D-0001/DO2664, November 1991.

[22] E. L. Lawler, J. K. Lenstra, A. Rinnooy-Kan and D. Shmoys, *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*, Wiley, 1985.

[23] S. Lee and H. S. Hahn, "An Optimal Sensing Strategy of a Proximity Sensor System for Recognition and Localization for Polyhedral Objects", *Proc. IEEE Intl. Conf. on Robotics and Automation*, 1990, pp. 1666-1671.

[24] S. Lee and H. Hahn, "Recognition and Localization of 3-D Natural Quadric Objects based on Active Sensing", *Proc. IEEE Intl. Conf. on Robotics and Automation*, Apr. 1991, pp. 156-161.

[25] M. Lewis and G. Bekey, "The Behavioral Self-organization of Nanorobots Using Local Rules," *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 1992, pp. 1333–1338.

[26] M. A. Lewis, A.H. Fagg and G.A. Bekey,"The USC Autonomous Flying Vehicle: An Experiment in Real-Time Behavior-Based Control", *Proc. IEEE Intl. Conf. on Robotics and Automation*, vol. 2, May 1993, pp. 422-429.

[27] M. Mataric, *Interaction and Intelligent Behavior*. PhD thesis, MIT EECS Department, May 1994.

[28] R. V. Nolan, H. C. Egghart, L. Mittleman, R. L. Brooke, F. L. Foder and D. L. Gravitte, "MERADCOM Mine Detection Program: 1960 - 1980", *Report* 2294, MERADCOM Belvoir RDE Center, Mar. 1980.

[29] L. E. Parker, "Designing Control Laws for Cooperative Agent Teams," in *Proc. IEEE Intl. Conf. on Robotics and Automation*, vol. 3, May 1993, pp. 582–587.

[30] N. S. V. Rao, W. Wu, and C. W. Glover, "Matching Algorithms for Finite Points Sets in Plane", *Proc. IEEE SOUTHEASTCON*, 1991, pp. 1229-1232.

[31] L. S. Riggs and C. A. Amazeen, "Research Efforts with the Waveguide Beyond Cutoff or Separated Aperture Dielectric Anomaly Detection Scheme", *report*, Belvoir RD&E Center, Dec. 1989.

[32] K. S. Roberts, "Robot Active Touch Exploration: Constraints and Strategies", *Proc. IEEE Intl. Conf. on Robotics and Automation*, 1990, pp. 980-985.

[33] S. Skiena, "Geometric Probing", *Ph.D. Thesis* , CS Dept., University of Illinois, Apr. 1988.

[34] S. S. Skiena, "Problems in Geometric Probing", *Algorithmica*, No. 4, 1989, pp. 599-605.

[35] S. S. Skiena, "Probing Convex Polygons with Half-Planes", *Journal of Algorithms*, No. 12, 1991, pp. 359-374.

[36] S. S. Skiena, "Interactive Reconstruction via Geometric Probing", *Proc. of the IEEE*, Sep. 1992, pp. 1364-1383.

[37] L. Steels, "Cooperation Between Distributed Agents Through Self-organization," *European Workshop on Modeling Autonomous Agents in a Multi-Agent World*, 1990, pp. 175–195.

[38] K. Sugihara and I. Suzuki, "Distributed Motion Coordination of Multiple Mobile Robots," *Proc. IEEE International Symposium on Intelligent Control*, 1990.

[39] H. Yamaguchi and T. Arai, "Distributed and Autonomous Control Method for Generating Shape of Multiple Mobile Robot Group," *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 1994, pp. 800–807.