

NATURAL ENERGY AND STRESS FOCUSING PHENOMENA

Seth Putterman and Raffi Budakian

Physics Department, University of California, Los Angeles, CA 90095

ABSTRACT

Friction, triboelectrification and sonoluminescence are phenomena which indicate nature's propensity to focus energy and stress in continuous media driven into off-equilibrium motion. In cavitation luminescence sound energy can focus by 12 orders of magnitude to make light. In tribological processes quantum mechanics focuses applied stress by over a factor of 1 million to account for stick-slip friction. Observation of these effects makes use of new techniques that have been developed to measure charge transfer and bonding between sliding surfaces.

INTRODUCTION

The tendency of nature to spontaneously focus energy in the off-equilibrium motion of continuous media is spectacular. Figure 1 shows the spectrum of light that is emitted from a fluid as it is forced to flow through a Venturi tube at speeds which lead to a pressure drop of about an atmosphere[1]. Although the velocity of the room temperature water is measured in meters/second the emitted light is broad band out to energies of 6eV which means that the light is emitted from hot spots with temperatures of about 50,000K or higher. This light is due to the formation of bubbles which focus energy as they collapse supersonically[2]. A single such bubble can be trapped in a resonant sound field so as to emit one flash of light with each cycle of sound and as shown in Figure 1 its spectrum is similar to that of the Venturi tube flow. The pulsating motion of the bubble focuses the diffuse sound energy density by 12 orders of magnitude to make flashes of light whose width is measured in picoseconds. Cavitating bubbles have actually found use in surgical procedures where they form on the end of a vibrating tip. The energy focused by this cloud of bubbles can effectively cut through tissue[3]. The spectrum of a cloud of transient bubbles is also displayed in Figure 1. At high frequencies [over 1MHz] the bubbles formed in a high amplitude sound field collapse down to radii on the order of tens of nanometers[4] at which point light is emitted. This can be called nanoluminescence.

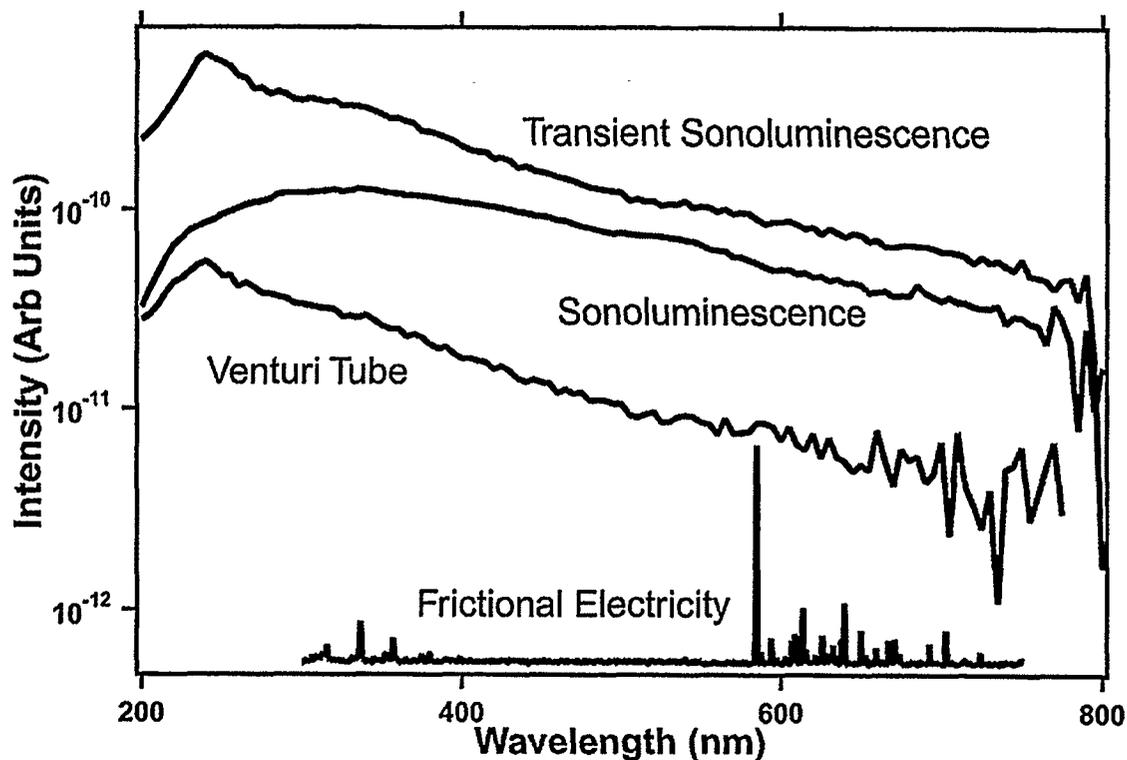


FIGURE 1: SPECTRA OF VARIOUS OFF-EQUILIBRIUM MOTIONS

Energy focusing is not limited to cavitation. One of the earliest [Picard 1676] investigations of luminescence dealt with the light emitted by mercury sloshing around inside of a barometer. In a controlled laboratory experiment[5] one can rotate a container of mercury around a horizontal axis and observe light emission from the region of the meniscus with the unaided eye. For the purpose of taking the spectrum shown in Figure 1 neon gas was added to the system. In the case of "rubbing electricity" the spectrum appears to come out as lines. But excitation of the neon line requires electrons with energies above 20ev. So the relative motion of mercury against glass at about 1mm/second leads to the repetitive acceleration of electrons to about 1% the speed of light. These fast electrons collide with the neon atoms to make the light. The mechanism of this effect is still unclear[6].

EXPERIMENTS AND DATA

In the experiments where mercury rubs against glass one observes that the mercury sticks to the glass between discharges. It thus undergoes a kind of stick-slip motion. In order to determine whether these observations can be generalized so as to gain insight into the long standing problem of friction[7], the charge transfer between rubbing surfaces has been measured[8]. Figures 2A,B show the apparatus used to measure the force of friction between surfaces in relative motion as well as the charge transfer. In this case the tip is gold and the

substrate is a dielectric [PMMA]. Charge deposited onto the substrate as a result of a stick slip friction event will lead to a time varying voltage when chopped with a conductor as shown the figure.

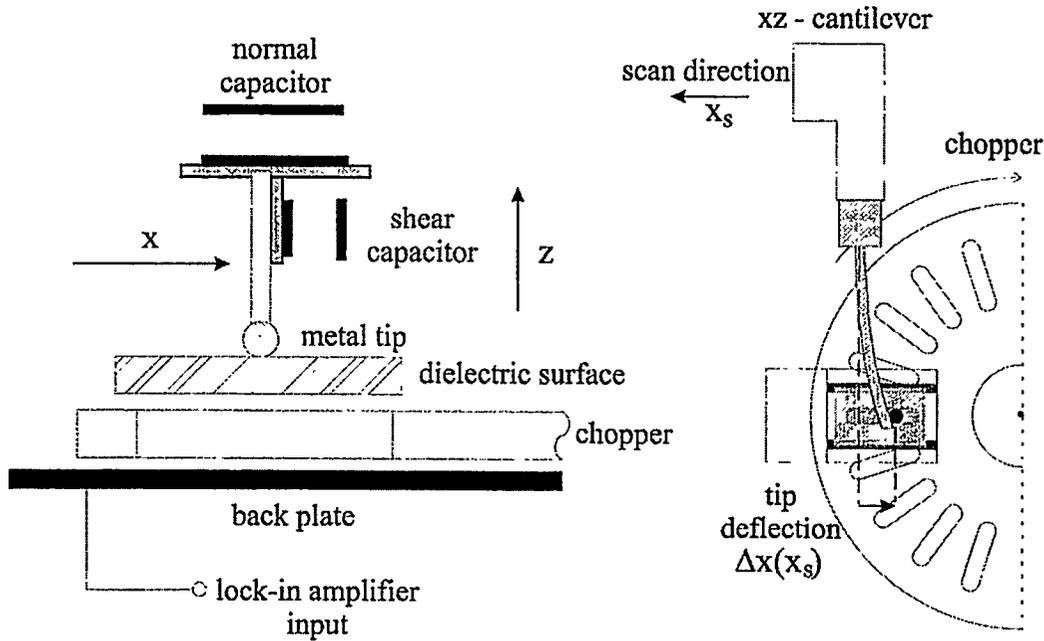


FIGURE 2A - APPARATUS TO MEASURE STICK-SLIP FRICTION AND CHARGE TRANSFER

Figure 3 shows the force of friction as a function of distance scanned as well as the total charge laid down by the gold tip. Note that each slip event in Figure 3 reveals a jump in the accumulated charge and furthermore there exists a scale factor $.48eV/\text{\AA}$ which can be used to collapse these two measurements onto one curve as shown in Figure 3B.

To calculate the scaled curves we introduce the displacement Δx_i of the tip on the i^{th} slip event which is resisted by a force, $\Delta F_i = k\Delta x_i$ where k is the lateral spring constant of the cantilever [47N/m]. The total distance x_s scanned by the stage obeys:

$$x_s = \sum_{i < s} \Delta x_i + \Delta x(x_s) \quad (1)$$

where $\Delta x(x_s)$ is the deflection of the tip from its position of mechanical equilibrium after being scanned a distance x_s . Setting $\Delta F_i = \alpha \Delta N_i$ and summing over all events yields

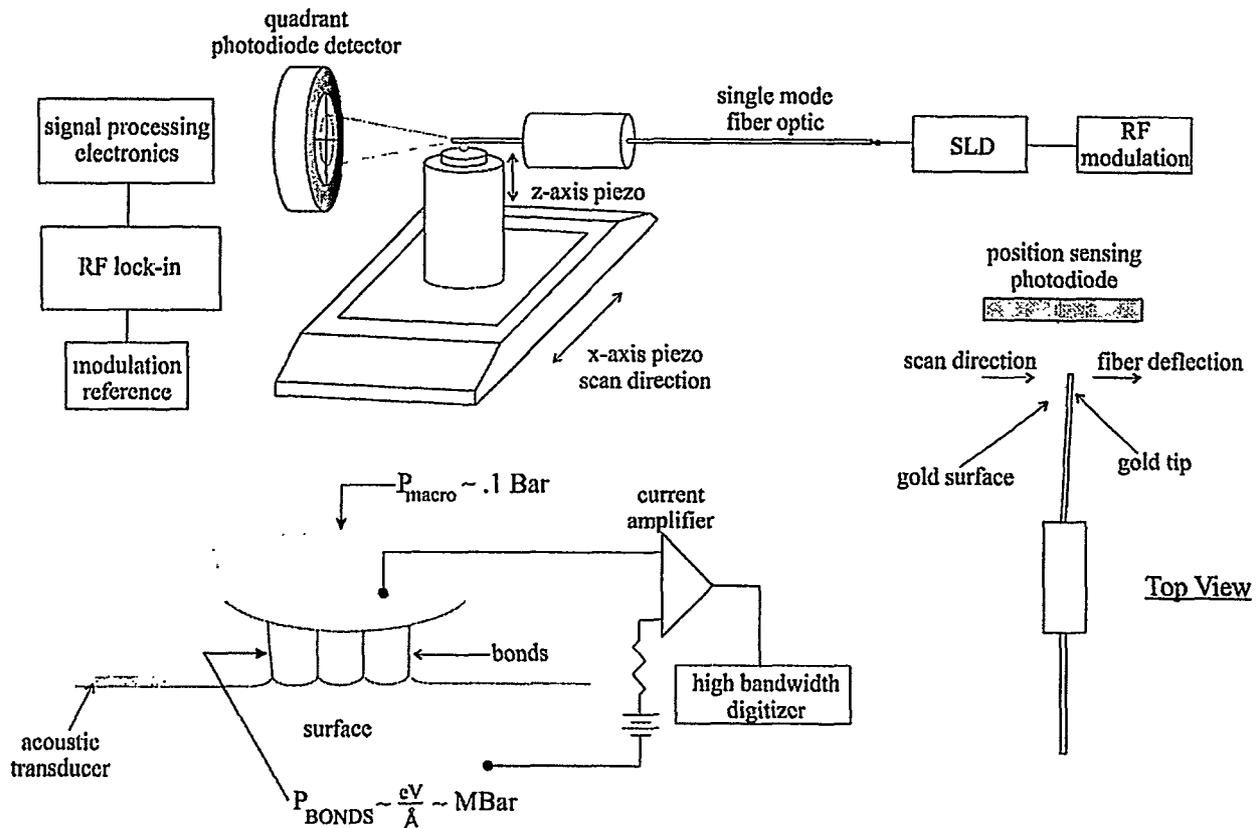


FIGURE 2B - APPARATUS TO MEASURE BOND STRENGTH

$$k(x_s - \Delta x(x_s)) = \alpha N(x_s) \quad (2)$$

where ΔN_i is the charge transferred on the i^{th} event, $N(x_s)$ is the total charge deposited in scanning x_s and α is the average lateral force per charge. The force curve in Fig. 3A, which is equal to $k\Delta x(x_s)$, has been combined with the imposed x_s to obtain a scaled force curve equal to the left hand side of Equation (2). This is plotted in Figure 3B along with $\alpha N(x_s)$, using the value $0.48 \text{ eV}/\text{\AA}$ for α . The energy and length scale that combine to determine α are characteristic of a single molecular bond, which suggests that in these experiments friction arises from bunches of bonds that form between and grab to rest two surfaces in relative motion. The charge transfer is a marker for the number of bonds ruptured at a particular slip event. The force of a single bond is remarkably large as an $\text{eV}/\text{\AA}$ (the natural scale for α) is 1 nN so that 10^8 bonds/ mm^2 , which is typical of our findings, corresponds to a force of 1 mN. This corresponds to an average macroscopic stress of about an atmosphere, and a focused stress at each bond of a Mbar ($\sim 1 \text{ eV}/\text{\AA}^3$).

By use of a liquid crystal film under the substrate the location of the charges can be instantaneously imaged with a resolution of $10^{-12} \text{ Coulomb}/\text{mm}^2$. Future generations of this charge microscope should reach resolutions of $10^{-14} \text{ Coulomb}/\text{mm}^2$. This device shows that the charge transfer is indeed concentrated at the slip locations [8].

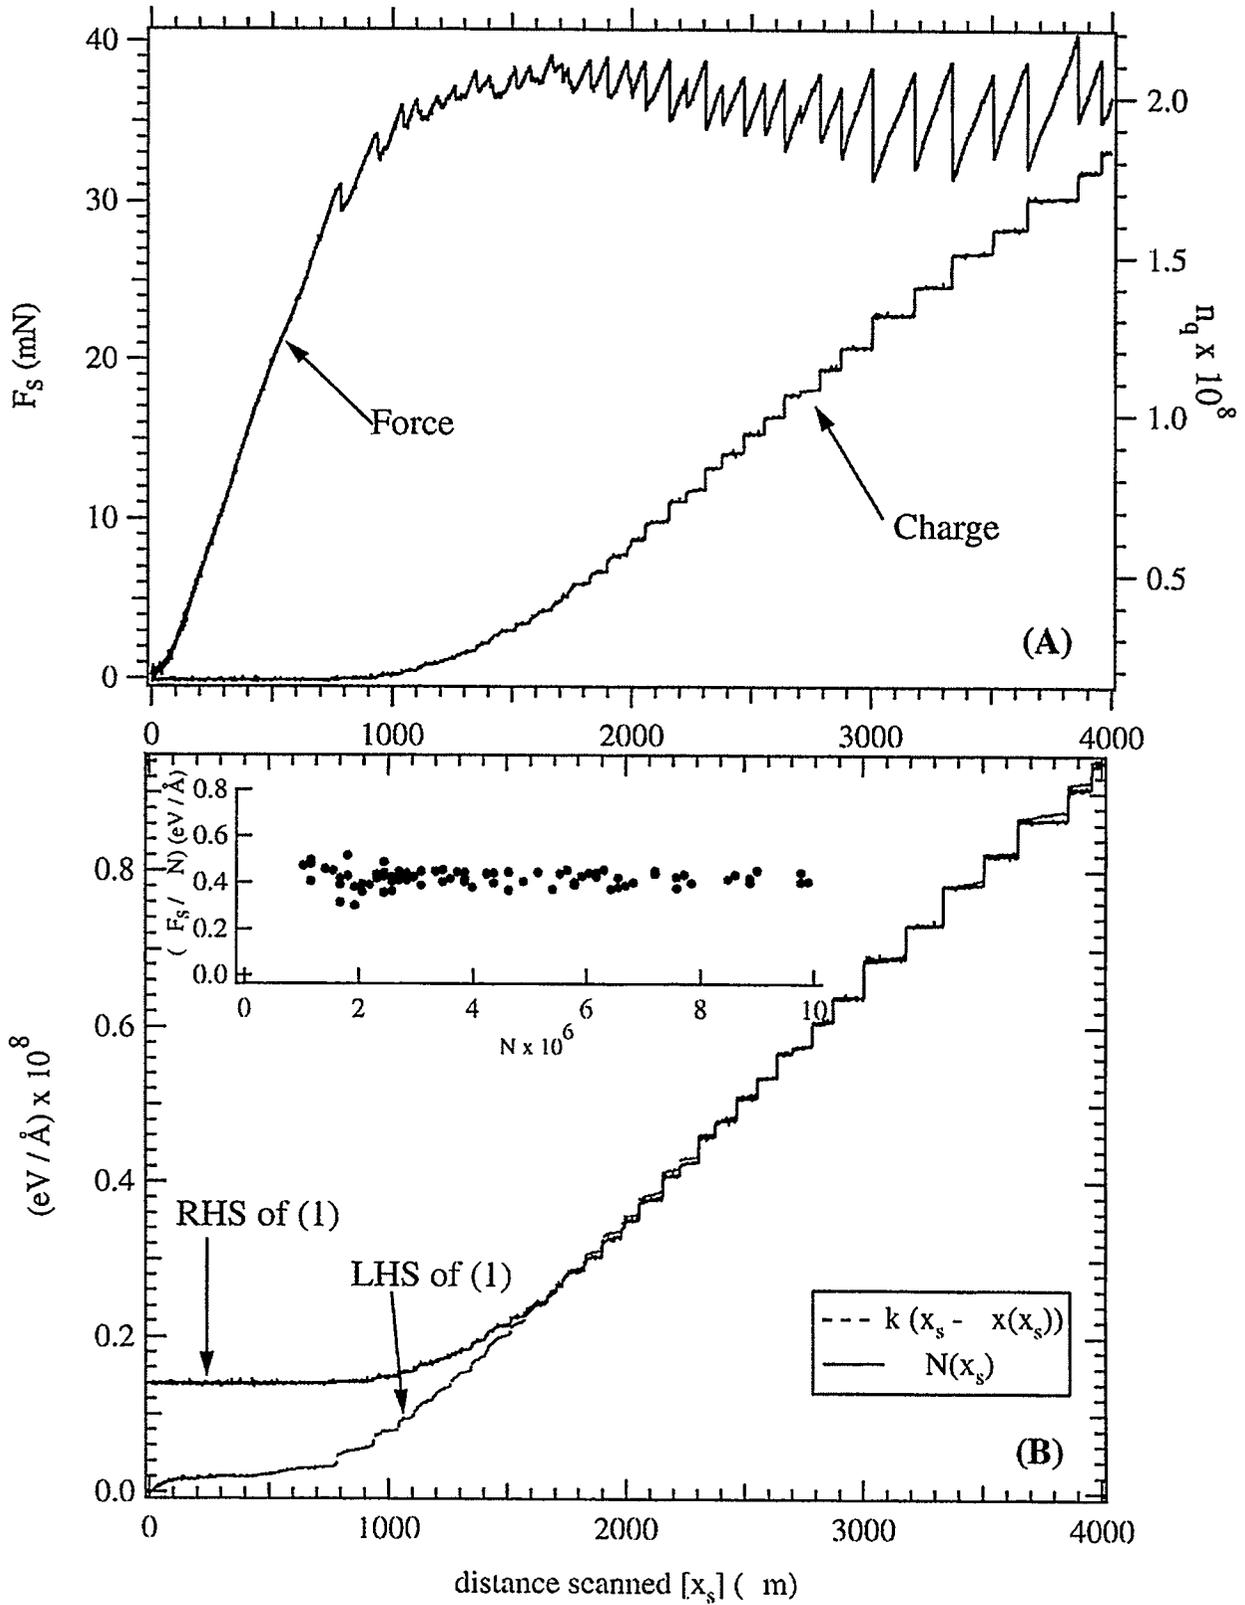


FIGURE 3 - STICK-SLIP FORCES AND CHARGE TRANSFER

For a metal surface sliding over another metal surface it is not obvious how to measure the charge transferred as a result of stick-slip events. In this case one can measure the elasticity of the bonds which form when the surfaces come into contact. The apparatus is shown in Figure 2. A 100micron gold ball hangs from the end of a fiber optic which also serves as cantilever. The natural frequency of oscillation of the ball on the fiber is 1,600Hz. As a 1mm gold ball is brought into contact with the suspended 100micron ball one observes that the resonant frequency for motion parallel to the interface stiffens and that the balls attract each other. This behavior is shown in Figure 4. Attraction between the surfaces is a process that can take off on its own[9]. For the parameters presented here the centers of the spheres can approach by an additional 40nm.

This behavior can also be interpreted in terms of bonds that form between surfaces in contact. As some bonds form the local attractive stresses increase so as to favor the formation of additional bonds until the elastic forces set up by the distortion of the surface balance the process. The observed forces of stick-slip friction for gold on gold agree with the bonding forces documented in Figure 4.

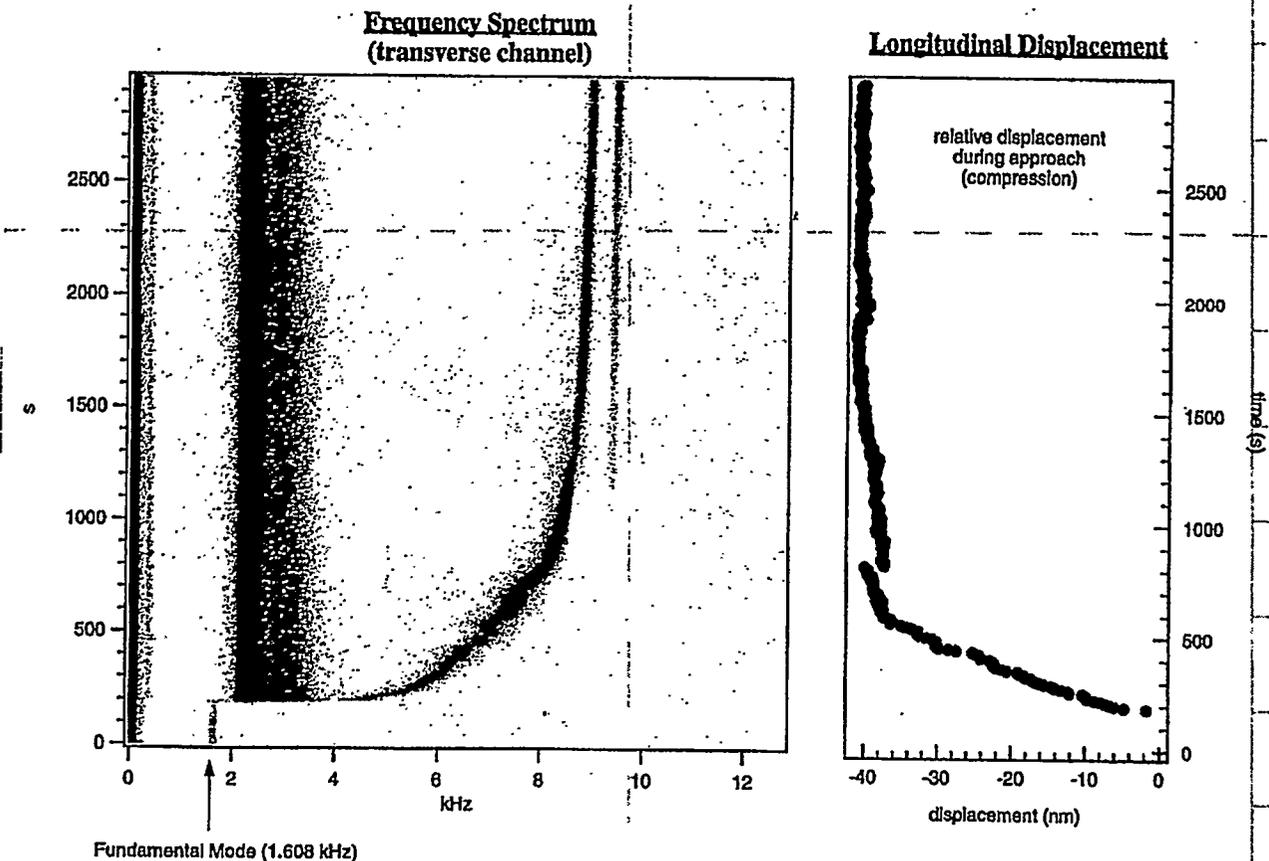


FIGURE 4 - STRENGTH OF SPONTANEOUS BONDING

CONCLUSION

In conclusion macroscopic measurements of surfaces in relative motion suggest that triboelectrification and friction have a common origin: namely bond formation and rupture. The quantum processes that lead to bonding also lead to highly focused local stresses that constitute an essential consideration for engineering at the nanoscale.

ACKNOWLEDGMENT

Research supported by the U.S. Department of Energy.

REFERENCES

1. K.R. WENINGER, C.G. CAMARA and S.J.PUTTERMAN 'Energy Focusing in a Converging Flow' *Phys. Rev. Lett.* 83, 2081 (1999).
2. S.J. PUTTERMAN and K.R. WENINGER 'Sonoluminescence' *Ann. Rev. Fluid Mech.* 32, 445 (2000).
3. K.R. WENINGER, C.G. CAMARA and S.J.PUTTERMAN 'Physical Acoustics of Ultrasound Assisted Lipoplasty' *Clinics in Plastic Surgery* 26, 463 (1999).
4. K.R. WENINGER, C.G. CAMARA and S.J.PUTTERMAN, Unpublished.
5. R. BUDAKIAN, K. WENINGER, R.A. HILLER, S.J. PUTTERMAN 'Picosecond Discharges and Stick-Slip Friction at a Moving Meniscus of Mercury on Glass' 391, 266 (1998).
6. B.D. TERRIS, et al 'Contact Electrification Using Force Microscopy' *Phys. Rev. Lett.* 63, 2669 (1989).
7. B.N.J. PERSSON, 'Sliding Friction' (Springer, Berlin 1998).
8. R.BUDAKIAN and S.J. PUTTERMAN 'Correlation between Charge Transfer and Stick-Slip Friction at a Metal-Insulator Interface' to appear *Phys. Rev. Lett.*; 'Real Time Imaging of Two-Dimensional Charge on Dielectric Surfaces' *Rev. Sci. Inst.* 71, 444 (2000).
9. K.L. JOHNSON 'Contact Mechanics' (Cambridge University press, Cambridge, 1977).

Instabilities and Defect Chaos in Models for Rotating Non-Boussinesq Convection

Hermann Riecke, Blas Echebarria, Vadim Moroz, and Filip Sain*

Department of Engineering Sciences and Applied Mathematics
Northwestern University, Evanston, IL 60208, USA

Abstract

We are interested in the types of spatially and temporally complex dynamics that may arise in systems in which patterns with *hexagonal* planform become unstable. Motivated by the Küppers-Lortz instability of rotating Rayleigh-Bénard convection, which leads to persistent domain chaos, we study the effect of rotation on convection with a hexagonal planform as it arises for sufficiently strongly temperature-dependent fluid parameters. Using coupled Ginzburg-Landau equations and order-parameter models of the Swift-Hohenberg type, we find that the steady hexagons that arise near onset can become unstable via steady and oscillatory, short- and long-wave instabilities. There are regimes in which the instability gives rise to persistent dynamics in which the Fourier spectrum of the disordered pattern effectively rotates with time. Within weakly nonlinear theory, the oscillating hexagons that arise farther away from threshold are described by the single complex Ginzburg-Landau equation (CGLE) and *generically* exhibit bistability between periodic oscillations and *defect chaos*. They appear to be the first system that should allow detailed quantitative comparison between experiments and the CGLE in one of its regimes of complex dynamics. As a first step towards a quantitative comparison we derive the relevant coupled Ginzburg-Landau equations directly from the Navier-Stokes equations.

1 Introduction

Among the current challenges in the understanding of chaotic systems are systems that behave chaotically in time but are also disordered in space. Such spatio-temporal chaos can arise in systems ranging from waves on fluid surfaces to electrical excitation waves in hearts undergoing fibrillation, a serious heart condition. The understanding of low-dimensional chaos of systems with few active degrees of freedom has reached a high level and has allowed substantial progress in its control as well as in harnessing its features for applications, e.g. in communications or in mixing.

High-dimensional chaos is far less well understood. Spatio-temporal chaos in spatially extended dynamical systems falls into this category. One of its characteristic features is the extensivity of the chaotic attractor, i.e. the number of active modes grows linearly with the size of the system. This suggests that these systems may be thought of as interacting dynamical, possibly chaotic units. The identification of the relevant units is, however, far from obvious. In many pattern-forming systems, e.g. in Rayleigh-Bénard convection, natural candidates for building blocks of the chaotic dynamics are defects in the pattern. These can be point defects like holes or dips in the wave amplitude in one dimension [1, 2], or disclinations and dislocations in two dimensions [3, 4, 5, 6], or they can be line defects like domain walls separating, for instance, stripe patterns of different orientation [7, 8]. Whether the defects play in fact an active role in the overall dynamics has not been clarified in most systems. In simulations of the single complex Ginzburg-Landau equation (CGLE) it has been established that a certain fraction of the total fractal dimension of the attractor can be attributed

*Current address: The CNA Corporation, 4401 Ford. Ave, Alexandria, VA 22302

to the dislocations in the wave pattern with the remainder being due to the wave field between the defects [4]. In simulations of coupled Ginzburg-Landau equations the dynamics of the dislocations have been found to provide an intuitive understanding of the transition between a spatially ordered and a spatially disordered state, both of which are chaotic in time [6]. In experiments on waves in binary-mixture convection progress has been made in reconstructing the complete wave field from the field generated by the dislocations alone [5].

Experimentally, spatio-temporal chaos has been investigated in a number of systems, in particular in various fluid flows in layers with large-aspect ratio. In Rayleigh-Bénard convection at low Prandtl numbers quite complex textures have been found that are characterized by stripe patches of various orientations, spirals, disclinations, and dislocations [3]. This state does not arise very close to threshold and is therefore not a good candidate for theoretical approaches using weakly nonlinear reductions of the full Navier-Stokes equations. In the presence of rotation, roll convection can become unstable and chaotic immediately above threshold due to the Küppers-Lortz instability [9]. In this system the dynamics are characterized by domain walls separating stripes (rolls) or different orientation. The success of a description using Ginzburg-Landau equations is, however, limited since these equations break the isotropy of the system. In electroconvection of nematic liquid crystals a regime is found in which spatio-temporal chaos arises immediately above onset [10]. Since the system is anisotropic, coupled Ginzburg-Landau equations should be able to describe the dynamics quantitatively. Unfortunately, while most of the coefficients in these equations have been derived from the underlying fluid equations [11], one coupling coefficient is still missing.

On the theoretical side, spiral-defect chaos in convection has been treated by full numerical simulations of the Navier-Stokes equations [12] as well as model equations of the Swift-Hohenberg type [13]. The latter approach has also been used for the study of the domain chaos arising from the Küppers-Lortz instability [7]. Most theoretical investigations of spatio-temporal chaos have been concerned, however, with the complex Ginzburg-Landau equation (CGLE) which applies quite generally to systems undergoing an instability to homogeneous oscillations. Despite its simplicity the CGLE exhibits a number of different chaotic and ordered regimes [14, 15, 16]. Surprisingly, however, up to recently no system appears to have been identified that exhibits one of the complex regimes found theoretically in the two-dimensional CGLE and that is suitable for detailed quantitative comparison with experiments.

In this communication we enlarge on recent results on the dynamics of hexagonal patterns in the presence of rotation [17, 18, 19, 20, 21]. The motivation for this work is two-fold. Most of the previous studies of spatio-temporal chaos have been addressing regimes that arise from stripe-like states and therefore inherit some of their characteristics. It is to be expected that chaotic states that arise from instabilities of a hexagonal planform will differ in a qualitative way. Motivated by the results on the Küppers-Lortz instability, we conjectured that rotation may also trigger persistent dynamics in the hexagonal planform. This is indeed the case, albeit in a quite different way. In addition to inducing instabilities of the steady hexagons, the rotation introduces a completely new state, oscillating hexagons, that arises in a secondary Hopf bifurcation off the steady hexagons. Strikingly, these oscillating hexagons are to leading order described by the single CGLE and within the weakly nonlinear theory they are always in a regime of bistability of stable oscillations and *defect chaos* [19]. Thus, rotating non-Boussinesq convection appears to be the first system that should allow a detailed comparison of one of the non-trivial regimes of the two-dimensional CGLE with precise experiments.

2 Weakly Nonlinear Description of Hexagons with Rotation

We are interested in a weakly nonlinear description of hexagon patterns in the presence of rotation. One of the best known systems exhibiting such patterns is non-Boussinesq convection, i.e. fluid convection in regimes in which the temperature dependence of the fluid parameters is of importance. In the weakly nonlinear regime the full fluid equations can be reduced to simpler equations for the overall amplitude of the convective pattern. We discuss two complementary approaches.

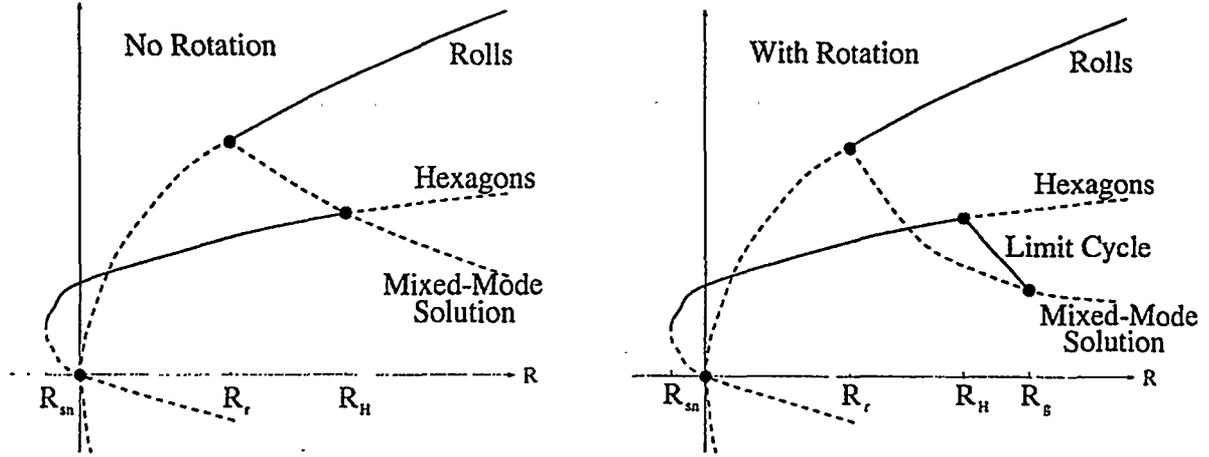


Figure 1: Bifurcation diagrams for weakly non-Boussinesq convection without and with rotation.

For small convection amplitudes the fluid equations can be reduced systematically to three coupled Ginzburg-Landau equations for the three Fourier modes making up the hexagonal pattern [17],

$$\partial_t A = RA + (\mathbf{n}_A \cdot \nabla)^2 A + \overline{BC} - A|A|^2 - (\nu + \bar{\nu})A|B|^2 - (\nu - \bar{\nu})A|C|^2. \quad (1)$$

In terms of the amplitudes A , B , and C , typical fluid quantities like the temperature in the midplane are given by

$$\theta = Ae^{iqx} + Be^{i(-qx/2 + \sqrt{3}qy/2)} + Ce^{i(-qx/2 - \sqrt{3}qy/2)} + c.c. + h.o.t. \quad (2)$$

The form of (1) can be derived using symmetry arguments. The bifurcation parameter in the amplitude equations is R . The breaking of the chiral symmetry by the rotation introduces a difference between the two cross-coupling coefficients as expressed by $\bar{\nu}$. In the absence of spatial gradients these equations yield typically the bifurcation diagrams shown in fig.1 [22, 23]. The main new feature introduced by the rotation is the branch of oscillating hexagons arising at R_H in a Hopf bifurcation and ending on the mixed-mode solution in a heteroclinic cycle at R_g .

To achieve a quantitative comparison with experiments we have also derived the coupled Ginzburg-Landau equations (1) for weakly non-Boussinesq convection directly from the fluid equations. A typical phase diagram obtained from that analysis is shown in fig.2. Since currently available experimental set-ups allow rotation rates up to 1Hz the regime of oscillating hexagons should be easily accessible in these systems.

It turns out that due to the rotation the steady hexagons can undergo an instability to a spatially disordered state with an isotropic Fourier spectrum. Since the coupled Ginzburg-Landau equations (1) only allow small deviations in the wavevector from the three preferred wavevectors they cannot capture such a state. We have therefore also made use of a phenomenological model of the Swift-Hohenberg type [20],

$$\partial_t \psi = R\psi - (\nabla^2 + 1)^2 \psi - \psi^3 + \alpha \psi^2 + \gamma \hat{e}_z \cdot (\nabla \psi \times \nabla(\nabla^2 \psi)), \quad (3)$$

which preserves the isotropy of the system. Here R is the control parameter, γ is a measure for the rotation, and ψ is the order parameter, which gives, e.g., the temperature of the layer in the mid-plane.

3 Instabilities of Steady Hexagons

The stability analysis of the steady hexagons within coupled Ginzburg-Landau equations as well as within Swift-Hohenberg-type equations reveals steady and oscillatory instabilities that can arise at long wavelengths or at finite wavelengths [20, 17]. A typical result from the Swift-Hohenberg-type equation (3) is shown in

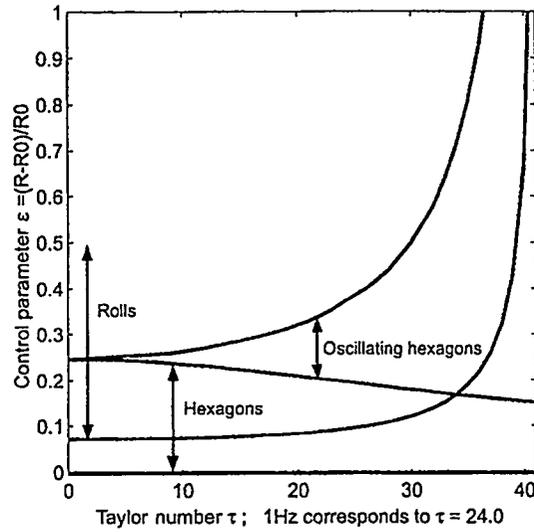


Figure 2: Phase diagram for non-Boussinesq convection in a layer of water of height $d = 0.106\text{cm}$ with mean temperature $T_0 = 47.5^\circ\text{C}$.

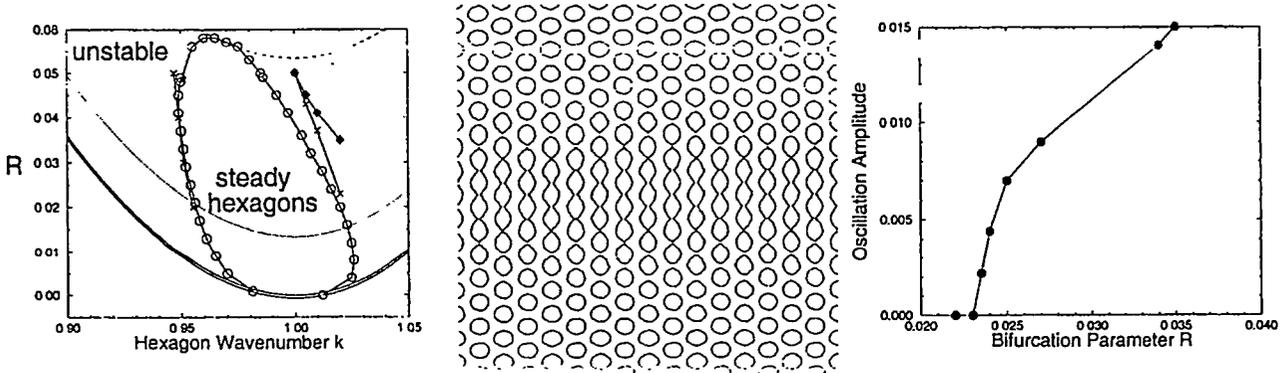


Figure 3: a) Stability region for steady hexagons in (3). Modulated hexagons arise between the lines marked by crosses and solid diamonds. Chaotic hexagons appear to the right of the solid diamonds (see fig.4). b) Contour lines of modulated hexagons. c) Bifurcation diagram for the oscillation amplitude of the modulated hexagons.

fig.3a. The stability limit for wavenumbers larger than the critical wavenumber is due to an oscillatory instability, which leads to spatially and temporally periodic *modulated hexagons*. A snapshot of this state is shown in fig.3b. The modulation is in the form of a standing wave: after half a period the contour lines will be disconnected in the center of the pattern and connected in the top and the bottom part. The bifurcation to the modulated hexagons is supercritical as shown in the bifurcation diagram fig.3c.

The modulated hexagons persist up to the line marked by solid diamonds. There, additional sideband modes come into play and destroy the order as illustrated in the time dependence of the dominant Fourier modes shown in fig.4a. The modes A , B , and C make up the hexagon pattern, while the mode P_{000} is responsible for the spatial modulation of the hexagons. Around $t = 4000$ additional modes become important and the modulated hexagons are supplanted by a spatio-temporally chaotic state. Two snapshots of that state are depicted in fig.4b,c. Strikingly, the hexagonal symmetry of the pattern does not break down completely. Thus, the Fourier spectrum still exhibits 6 peaks, which are, however, quite broad due to the

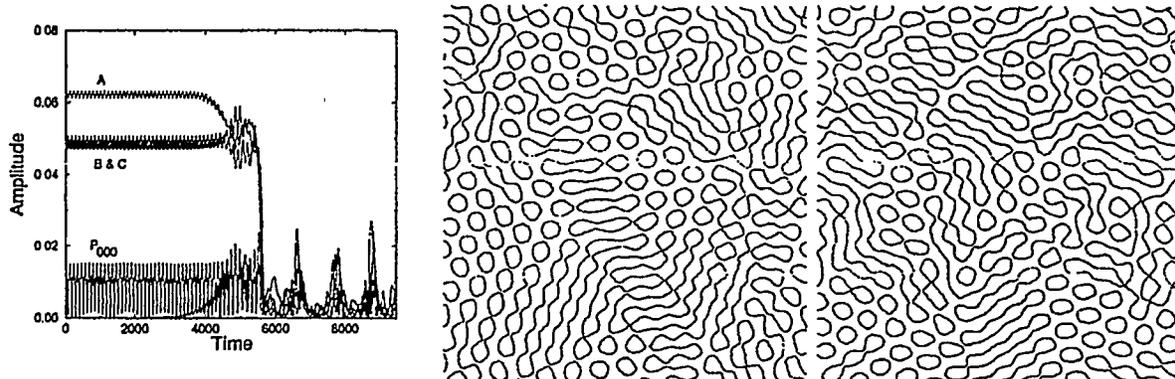


Figure 4: a) Crossing the solid diamonds in fig.4a the order of the modulated hexagons breaks down. Modes A , B , and C correspond to the modes making up the initial hexagon pattern. Mode P_{000} is the mode modulating the hexagons. b) and c) Successive snapshots indicating the rotation of the dominant orientation of the disordered patterns

disorder in the pattern. With time, the spectrum evolves and effectively rotates at a quite steady rate [20]. The rotation can be seen by comparing the two snapshots in fig.4b,c. In fig.4c the dominant orientation has rotated by about 30° compared to fig.4b. The rotation is also visible in the time-dependence of the dominant Fourier modes shown in fig.4a. Beyond $t = 6000$, the growth and decay of the modes A , B , and C representing the initial hexagon pattern shows the variation of the dominant orientation of the pattern relative to these modes.

4 Defect Chaos in Oscillating Hexagons

The oscillating hexagons arise in a secondary Hopf bifurcation off the steady hexagons. Near that bifurcation they can be described by a CGLE for the oscillation amplitude \mathcal{H} coupled to the phase vector $\vec{\phi} \equiv (\phi_x, \phi_y)$ that characterizes the two wave vectors of the hexagons [18],

$$\partial_T \mathcal{H} = \varepsilon \delta_1 \mathcal{H} + \xi \nabla^2 \mathcal{H} - \delta_2 \mathcal{H} \nabla \cdot \vec{\phi} - \rho \mathcal{H} |\mathcal{H}|^2, \quad (4)$$

$$\begin{aligned} \partial_T \vec{\phi} = & D_\perp \nabla^2 \vec{\phi} + D_\parallel \nabla (\nabla \cdot \vec{\phi}) + D_{x_1} (\hat{e}_z \times \nabla^2 \vec{\phi}) + D_{x_2} (\hat{e}_z \times \nabla) (\nabla \cdot \vec{\phi}) + \alpha \nabla |\mathcal{H}|^2 \\ & + \beta_1 (\hat{e}_z \times \nabla) |\mathcal{H}|^2 - i \beta_2 (\mathcal{H} \nabla \mathcal{H}^* - \mathcal{H}^* \nabla \mathcal{H}) + i \eta [\mathcal{H} (\hat{e}_z \times \nabla) \mathcal{H}^* - \mathcal{H}^* (\hat{e}_z \times \nabla) \mathcal{H}]. \end{aligned} \quad (5)$$

Two aspects of these equations are particularly noteworthy. First, to the cubic order considered in (1) the CGLE (4) decouples from the equation (5) for the phase vector if the hexagon wavenumber is at the bandcenter. Second, although the coefficients ξ and ρ in the CGLE depend on the values of ν and $\tilde{\nu}$ their variation turns out always to be restricted to a range in which there is bistability of ordered oscillations and defect chaos. Thus, if the non-Boussinesq effects in the fluid are sufficiently weak to allow a description of the oscillating hexagons by the coupled Ginzburg-Landau equations (1), then independent of the values of ν and $\tilde{\nu}$ the oscillating hexagons are predicted to exhibit defect chaos of the type studied theoretically in the CGLE [14]. Two snapshots of simulations of the coupled Ginzburg-Landau equations in that regime are shown in fig.5a,b. The disordered pattern shows domains in which the hexagons oscillate. Thus, they have a slight roll-like character with an orientation that 'rotates' in time. The time interval between fig.5a and fig.5b corresponds to half a period of the oscillation. In addition, there are also localized regions in which the hexagons are quite unperturbed and steady, e.g. in the bottom left corner of fig.5a,b. The latter positions correspond to locations of vanishing oscillation amplitude, i.e. dislocations (spirals) in the complex amplitude \mathcal{H} . This is made more clear in fig.5c in which the zero-lines of the real and of the imaginary part of \mathcal{H} are plotted for the same time as in fig.5a. The dislocations correspond to the points of intersection of these lines. Comparison with fig.5a,b shows that wherever two defects are close to each other the oscillation

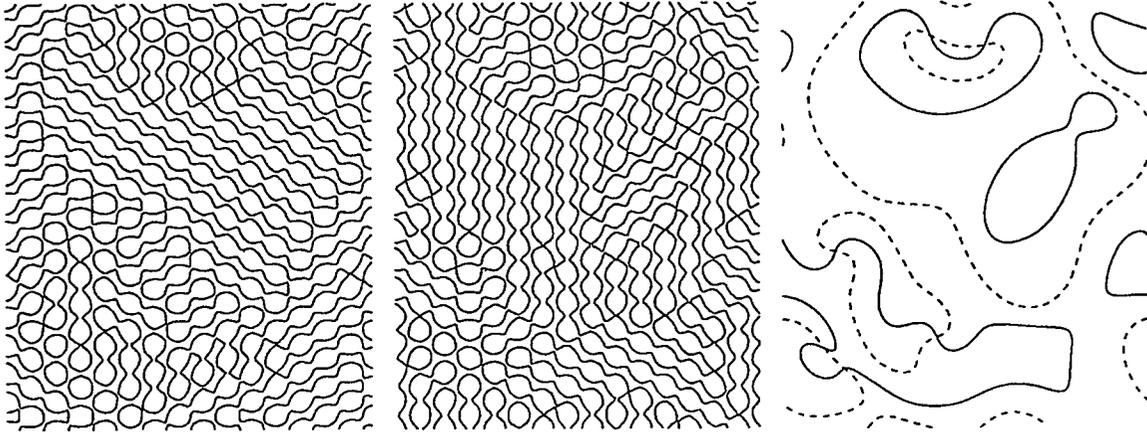


Figure 5: Snapshot of defect chaos in oscillating hexagons. a) and b) contour-lines of hexagon pattern half a oscillation period apart. c) zero-lines of the real and imaginary part of the oscillation amplitude. Intersections identify defects.

amplitude is strongly suppressed and the hexagon pattern is quite regular. For these parameters the spiral character of the dislocations is not very strong; the spirals do not persist long enough to develop into a clear spiral structure.

Away from the band center the CGLE (4) becomes coupled to the phase equation (5). At the same time the values of the other coefficients in the CGLE change as well. The resulting dynamics are found to be very similar to that obtained in the CGLE alone. As the wavenumber of the underlying hexagons is increased the chaotic activity becomes weaker and the average number of defects decreases. A comparison of the density of dislocations in simulations of the coupled Ginzburg-Landau equations (1), of the phase-amplitude equations (4,5), and of the CGLE alone shows that even away from the band center the effect of the coupling to the phase is small. Specifically, we find a decrease in the chaotic activity and eventually a transition to a state like the frozen-vortex state found in the CGLE [16].

5 Conclusion

We have performed a weakly nonlinear analysis of hexagon patterns in systems with broken chiral symmetry. For the steady hexagons various regimes of interesting complex dynamics are possible, including the stable coexistence of steady hexagons, modulated hexagons, and spatio-temporally chaotic hexagons. The oscillating hexagons arising at slightly larger amplitudes are predicted to be described by the complex Ginzburg-Landau equation. The coefficients are found to be such that the oscillating hexagons should exhibit defect chaos that is bistable with the ordered oscillations. Calculations based on the fluid equations for weakly non-Boussinesq convection predict that this regime should be accessible in current experiments [24].

Acknowledgments

We have benefitted from discussions with M. Silber and gratefully acknowledge the support by the Engineering Research Program of the Office of Basic Energy Sciences at the Department of Energy (DE-FG02-92ER14303).

References

- [1] M. van Hecke. The building blocks of spatiotemporal intermittency. *Phys. Rev. Lett.*, 80:1896, 1998.
- [2] J. Burguete, H. Chaté, F. Daviaud, and N. Mukolobwiz. Bekki-Nozaki amplitude holes in hydrothermal nonlinear waves. *Phys. Rev. Lett.*, 82(16):3252–3255, April 1999.
- [3] S.W. Morris, E. Bodenschatz, D.S. Cannell, and G. Ahlers. Spiral defect chaos in large aspect ratio Rayleigh-Bénard convection. *Phys. Rev. Lett.*, 71:2026–2029, 1993.
- [4] D.A. Egolf. Dynamical dimension of defects in spatiotemporal chaos. *Phys. Rev. Lett.*, 81(19):4120–4123, November 1998.
- [5] A. La Porta and C. M. Surko. Phase defects and spatiotemporal disorder in traveling-wave convection patterns. *Phys. Rev. E*, 56(5):5351–5366, November 1997.
- [6] G.D. Granzow and H. Riecke. Ordered and disordered defect chaos. *Physica A*, 249:27, 1998.
- [7] M.C. Cross, D. Meiron, and Y. Tu. Chaotic domains: a numerical investigation. *Chaos*, 4:607, 1994.
- [8] Y. Hu, R.E. Ecke, and G. Ahlers. Convection under rotation for Prandtl numbers near one: part I. linear stability wavenumber selection and pattern dynamics. *Phys. Rev. E*, 55:6928, 1997.
- [9] G. Küppers and D. Lortz. Transition from laminar convection to thermal turbulence in a rotating fluid layer. *J. Fluid Mech.*, 35:609, 1969.
- [10] M. Dennin, G. Ahlers, and D.S. Cannell. Spatiotemporal chaos in electroconvection. *Science*, 272:388, 1996.
- [11] M. Treiber and L. Kramer. Coupled complex Ginzburg-Landau equations for the weak electrolyte model of electroconvection. *Phys. Rev. E*, 58:1973–1982, 1998.
- [12] W. Decker, W. Pesch, and A. Weber. Spiral defect chaos in Rayleigh-Bénard convection. *Phys. Rev. Lett.*, 73:648, 1994.
- [13] H.W. Xi, J.D. Gunton, and J. Vinals. Spiral defect chaos in a model of Rayleigh-Bénard convection. *Phys. Rev. Lett.*, 71:2030, 1993.
- [14] H. Chaté and P. Manneville. Phase diagram of the two-dimensional complex Ginzburg-Landau equation. *Physica A*, 224:348, 1996.
- [15] I. Aranson, L. Kramer, and A. Weber. Core instability and spatiotemporal intermittency of spiral waves in oscillatory media. *Phys. Rev. Lett.*, 72:2316, 1994.
- [16] G. Huber. Vortex solids and vortex liquids in a complex Ginzburg-Landau system. In P.E. Cladis and P. Palffy-Muhoray, editors, *Spatio-temporal patterns in nonequilibrium complex systems*, pages 51–62, Reading, 1995. Addison-Wesley.
- [17] B. Echebarria and H. Riecke. Instabilities of hexagonal patterns with broken chiral symmetry. *Physica D*, 139:97–108, 2000.
- [18] B. Echebarria and H. Riecke. Defect chaos of oscillating hexagons in rotating convection. *Phys. Rev. Lett.*, 84:4838, 2000.
- [19] B. Echebarria and H. Riecke. Stability of oscillating hexagons in rotating convection. *Physica D*, to appear.

- [20] F. Sain and H. Riecke. Instabilities and spatio-temporal chaos of hexagonal patterns with broken chiral symmetry. *Physica D*, accepted.
- [21] F. Sain. *Instabilities and Spatio-Temporal chaos of hexagons in the presence of rotation*. PhD thesis, Northwestern University, 2000.
- [22] J.W. Swift. Convection in a rotating fluid layer. In *Contemporary Mathematics Vol. 28*, page 435, Providence, 1984. American Mathematical Society.
- [23] A.M. Soward. Bifurcation and stability of finite amplitude convection in a rotating layer. *Physica D*, 14:227–241, 1985.
- [24] K.M.S. Bajaj, J. Liu, B. Naberhuis, and G. Ahlers. Square patterns in Rayleigh-Bénard convection with rotation about a vertical axis. *Phys. Rev. Lett.*, 81:806, 1998.

COMPLEX INTELLIGENT MACHINES

Herschel B. Smartt, Ph.D., Charles R. Tolle, Ph.D.,

Kevin L. Kenney

Idaho National Engineering and Environmental Laboratory

P.O. Box 1625

Idaho Falls, ID 83415-2210

ABSTRACT

The machine control problem is normally approached from the perspective of having a central body of intelligence (and control) in the machine [Albus, 1991]. However, we present a conceptual design of a machine using distributed learning and intelligence. This new design is loosely based on biological models of social insects. For example, in an ant colony each ant functions according to local rules of behavior [Hölldobler and Wilson, 1990, see chapters 8 and 9]. There is no "king" or "queen", although the latter name has been given to the reproducing ant. Following a similar approach, we present a modular machine architecture in which each machine element has local rules of behavior (and local learning) along with a global element that influences local behavior (but does not dictate actions). A prime goal is to develop methods of learning and behavior modification that ensure global stability and optimization of the total machine; we discuss the theoretical aspects of ensuring such optimal performance.

INTRODUCTION

James Albus [1991] at NIST has defined machine intelligence as "the ability of a system to act appropriately in an uncertain environment, where appropriate action is that which increases the probability of success, and success is the achievement of behavioral subgoals that supports the system's ultimate goal." Following Albus' intent, we can say that intelligent machines are those that either know or can learn everything they need to know to perform a process or task. Such machines may be able to perform a process or task autonomously (without operator intervention) or semi-autonomously (with operator intervention).

In this paper, we present a conceptual design of a machine using distributed learning and intelligence. Related work has been conducted, for example, by Dorigo and Colomi [1996] using ant-

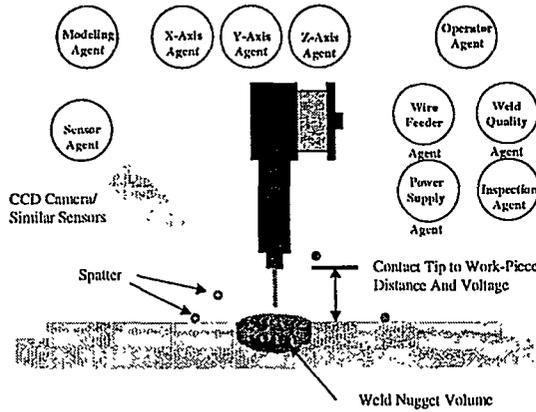


Figure 1. Arc spot welding machine with agents for the power supply, electrode wire feeder, positioner axes, sensor, and operator interface.

based local behavior of multiple agents to solve the Travelling Salesman Problem and other classical hard problems. Schatz et al. [1999] formulated a model for route learning in ants. Lambrinos et al. [2000] used a similar model for navigation of a mobile robot. Overgaard, Petersen, and Perram [1995, 1996] used local agent control of dynamic motion and path planning in multiple link robot arms.

Consider an intelligent machine in which various machine functions are carried out in a distributed manner. A schematic of such a machine for arc spot welding is shown in Figure 1. In addition to the machine hardware required (most of which is not shown) there are several “agents”. These agents have local control of various machine functions and are able to communicate with each other and with an operator agent, see Figure 2. The operator agent may be a human or may be an interface to a human (or even an interface to another machine). (Although it would be possible to focus on autonomous machines, we chose not to do so; our machines will interact with humans who have supervisory control authority.)

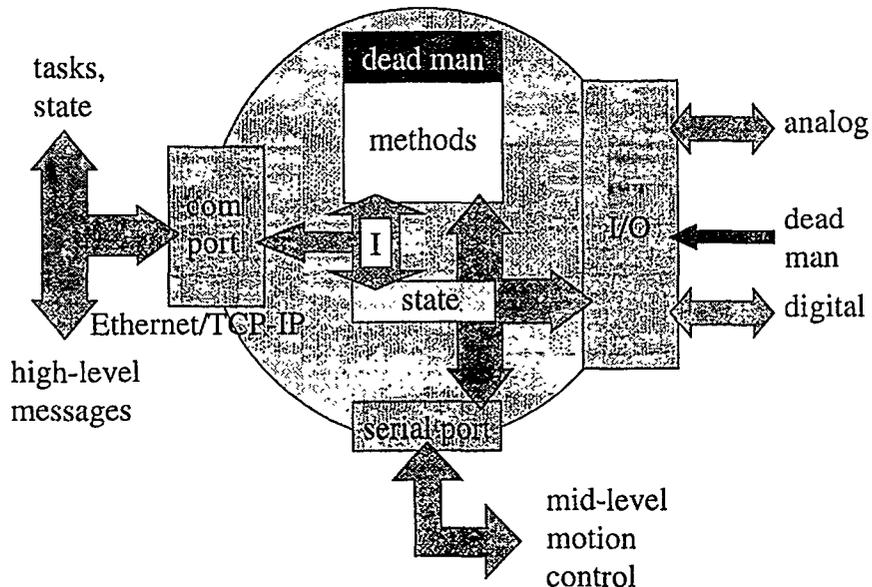


Figure 2. Agent block diagram.

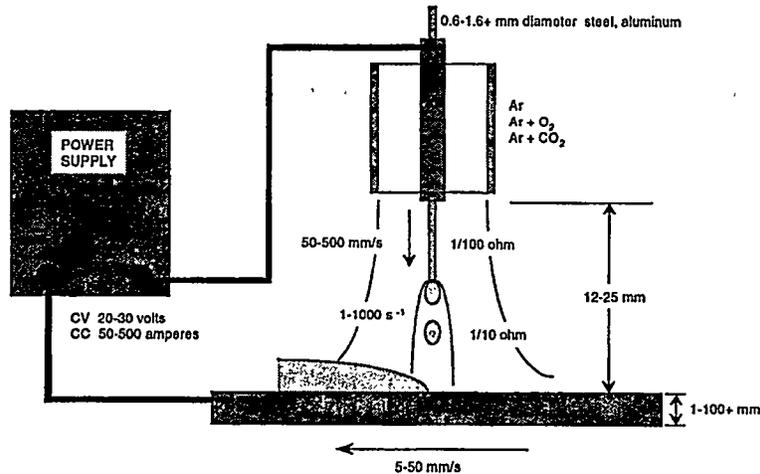


Figure 3. Schematic of gas metal arc welding process showing typical values of parameters.

The various agents will incorporate knowledge of how to perform tasks, the ability to learn from experience, and memory of past performance. The agents will also be able to optimize both their local behavior and the global behavior of the total machine.

To formulate such a machine, we need a variety of methods. In addition to distributed learning and control, we also chose to have our machines learn rules of behavior. This is distinct from learning control trajectories, a method frequently employed for machine learning. Our rules will be embodied using a variant of fuzzy logic [Johnson and Smartt, 1995] that allows the system to learn by back propagation [Rumelhart, 1986]. However, we discuss the application of iterative learning control^a to distributed intelligence. Iterative learning control is a recent set of methods for learning control trajectories that is well suited to iterative processes. However, iterative learning control methods may also be used to learn the weighting of rules for local optimization. We also discuss a new method of global optimization that uses artificial neural networks that learn the contribution of local behaviors to global cost.

SIMPLE INTELLIGENT WELDING MACHINE

Now consider the welding machine control problem. This is a much more complicated problem than two-dimensional motion control. First, there is a motion control problem involved. Even simple automated welding machines may have three degrees of motion. Consider welding in the flat position (e.g. joining two flat plates edge to edge with the plates in the horizontal plane). The welding torch must move along the weld joint. It also needs to be able to move at right angles to the weld joint (in the horizontal plane) to track misalignment of the joint with the axis of primary motion. Finally, it needs to move in a vertical direction to obtain changes in the contact tube to weldment distance. In addition, the weld torch may be mounted with a lead or lag angle relative to the weld joint. That is, the torch may be nominally vertical to the plates, but tilted backwards or forwards, respectively, with respect to the welding direction. For other weldment configurations, the torch may be leaned to one side or the other. Finally, the torch may be moved laterally with respect to the weld joint in a weaving pattern to effectively increase the width of the weld bead. Robotic welding systems may be even more complicated.

a. Uchiyama, 1978; Arimoto, Kawamura, and Miyazaki, 1984; Moore, 1993

Welding also involves selection of proper values of the process independent variables, see Figure 3. Disturbances to the process or uncertainties in the welding conditions may result in a need for the welding process independent variables to be changed during welding. Consequently, we need to consider that the trajectory we must obtain involves multiple degrees of motion via the robot as well as multiple degrees of motion through welding parameter space. What we seek is a set of generic rules that will ensure that the weld is made in some manner that will result in a structurally reliable weldment. Further, we want the welding machine to tune those rules to obtain a more robust process than would result from a fixed set of rules.

Consider a specific welding control problem. We desire to fabricate a steel structure using arc spot welding. Thus, steel sheet will be welded to an underlying structure by means of weld nuggets deposited into circular holes in the sheet. This geometry is shown in Figure 4.

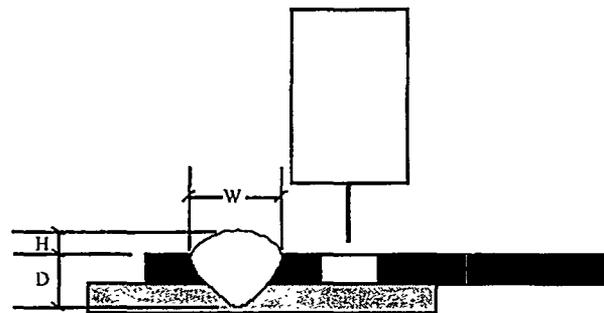


Figure 4. Cross section of gas metal arc spot weld showing a hole in the top sheet and a completed weld bead.

In this situation, the weld torch may be moved to a suitable position over a weld site, using motion control as discussed earlier. The welding power supply contactor is activated, the power supply voltage is set, the shielding gas is turned on, and the electrode wire is fed downward. This will result in ignition of an arc with corresponding heat and mass transfer to the weldment. After a suitable time, the power supply contactor is deactivated and the electrode wire feed is stopped. A short time later the shielding gas is turned off. Although this is perhaps the simplest arc welding example we can consider, there are still important control decisions that ensure that the weld will meet its acceptance requirements.

A good weld in this example is one that is strong enough, does not excessively over or under fill the hole, has minimal spatter, and does not contain gross defects such as cracks or porosity that could lead to failure. To be strong enough, the weld bead must adequately penetrate the lower structure (but not excessively melt through that structure) while fusing into the upper sheet. For most applications, the cross-sectional area of the weld bead in the plane of the interface between the upper sheet and the lower structure needs to be equal to or greater than some critical amount.

To obtain a good weld in this example, the current must be high enough but not too high and the weld time (the time the arc is on) must be equal to or greater than some critical minimum. This will ensure that adequate heat and mass have been transferred to the weldment. It is also necessary for the voltage to be above some minimum (to reduce spatter) and below some maximum (to avoid melt through and burn back).

GLOBAL OPTIMIZATION

One of the key challenges to achieving intelligent distributed learning with global optimization within our welding machine is the interaction between the global optimization cost function and local agent optimization cost functions. Each independent optimization agent must be able to make changes to its locally controlled parameters while keeping in mind its effects on the global cost/process. Traditionally, industrial process optimization is broken down into sub-components then optimized locally. If there is time and resources, once the sub-components have been optimized, an engineering team is formed to globally optimizing the interactions between sub-components through manual trial and error adjustments away from the sub-components optima. Today there is a growing interest in global process self-optimization, or near optimization, through the use of "swarm intelligence" [Bonabeau et al., 2000]. One method in particular is based on how ants function in nature while searching for food [Dorigo et al., 1996]. Each ant acts as an independent agent making random decisions on where to search with each move. As traffic increases along a particular path, ants crossing that path will be biased with a greater probability to follow the already more traveled path by its higher pheromone level. This tendency is reinforced as ants travel back and forth along the path between food and the anthill faster than other ants on competing paths. This leads to a faster increase in the pheromone trail on the shorter paths with respect to the longer competing paths, which attracts additional ants to the short paths. This is a never-ending reinforcement of the global optimization cost function, i.e. increased food movement back to the anthill. However, as is pointed out by Bonabeau et al. [2000], there is also a kind of integral wind up effect in ant behavior. When a shorter path is introduced after the "best path" has been found, the system has a hard time finding it unless the dynamics are changed, e.g. the initial food source is used up. This limitation can be overcome by modeling the pheromone trail as evaporative [Bonabeau et al., 2000; Dorigo et al. 1996]. It is important to make clear the key ideas being presented within this optimization system: each ant adds its own piece of cost to the global cost function (food delivered), and they are able to communicate to each other about their success (pheromone trail).

A slightly different intelligent distributed learning system can be expressed in human terms as an everyday project team. Here each team member is an individual agent that contains a wide range of experience, talents, and education. Normally, such teams have a team/project leader whose role is key to their achieving intelligent distributed learning on a global optimization problem. The team leader, and his allowed interactions, differentiates this method of distributed learning from the ant's. Within this model of distributed learning, the global cost function is contained within the team leader and acts as an agent of its own. Tasks are distributed among team members as well as sub-groups of team members. These agents progress in solving their tasks, as well as developing localized communication paths between agents, i.e. real time reconfiguration of the sub-groups. More importantly, the key concept within this structure is that the team leader cannot dictate any particular action to a team member. In human terms, this is primarily due to the team leader's lack of technical details and/or conceptual understanding required to solve any particular subtask. Remember, the team leader is globally oriented. However, team leaders can attempt to influence a particular team member's actions in order to achieve the global optimal solution. For example, the team leader can inform a member/agent that by increasing the tolerances within their portion of the process, all of the remaining process can speed up dramatically, i.e. the other agents are waiting on that particular agent due to the extra time required to optimize his subtask, even though it will not add much to the global cost function. We propose that this interaction between the team leader and individual team members is the key to the successful development of an intelligent global distributed learning algorithm, as opposed to "swarm intelligence." Furthermore, we differentiate an intelligent global distributed learning algorithm from a centralized learning algorithm, such as traditional neural networks, by not allowing any agent to dictate to another agent its actions, i.e. the team leader is not allowed to force any agent into a particular action. In short, an intelligent global distributed learning algorithm must allow each agent its own localized cost function and the ability to solve its subtasks primarily by itself with non-dictated feedback from the other agents. Global optimization is actually

achieved through the local optimization procedures employed by each agent while taking into account the global effects of its choices.

Using the human project team model just discussed, consider the intelligent spot welding machine outlined in Figure 1. This machine has a team leader, the Weld Quality Agent. Its job is to evaluate the overall success of welds produced by the machine and supply feedback to the local agents—i.e., x-axis agent, y-axis agent, z-axis agent, power supply agent, wire feeder agent—on their effectiveness in producing quality welds. The key to designing a particular algorithm is how the team leader, the Weld Quality Agent, is allowed to interact with each of the local agents. A further complication is by what methods will the team leader pass localized global cost information to each agent.

With this in mind, we consider the following initial architecture and algorithm for study, as outlined in Figure 1. This algorithm is based on a global cost function maintained and calculated within the Weld Quality Agent using generic weld parameter information developed from machine's agents.

$$Weld_Quality = V_Q(V_D, V_M) + S_Q(S_I, V_I) + B_Q(S_I, V_I) + M_Q(P_I) + T_Q$$

where

V_Q represents the quality of the well nugget volume based on the desired volume, V_D , and the measured/calculated volume, V_M .

S_Q represents the quality cost of spatter produced during the welding process based on sensor inspection, S_I , and operator visual inspection, V_I .

B_Q represents the quality cost associated with burn back.

M_Q represents the quality of the mechanical joint produced by the machine based on the operator's physical inspection, P_I .

T_Q represents the cost to quality due to the time involved in producing the weld.

Now that the form of the global cost function has been chosen, the next step is to define the relationship and method for communicating global cost information to the local agents. We propose to accomplish this task by adding these effects of the localized global cost, C_{agent}^{global} , to the traditional local cost function, C_{agent}^{local} .

$$C_{agent} = \alpha C_{agent}^{local} + (1 - \alpha) C_{agent}^{global}$$

The α term is used to adjust the balance between local cost and global cost variations on the local optimization process. It is planned that for new welding setups, the machine may fix α to be 1 for an initial period of time, thereby allowing development of the initial relationships between the local agent costs and global costs before proceeding with the augmented local cost function above (i.e. $\alpha \neq 1$). Note that the effect due to the global cost is based on multiple terms within the weld quality cost function, e.g. $C_{agent}^{global}(V_Q, S_{QI})$. These mappings form the uniqueness and key difficulty of the proposed algorithm;

namely, how will the algorithm obtain a mapping from global cost effects to local cost effects, i.e. $Weld_Quality \rightarrow C_{agent}^{global}$? To simplify the initial algorithm and its solution, it will be assumed that the effect of the global cost on a particular agent is a linear combination of each of the $Weld_Quality$ sub-terms:

$$C_{agent}^{global}(V_Q, S_Q, B_Q, M_Q, T_Q) = a_1 V_Q + a_2 S_Q + a_3 B_Q + a_4 M_Q + a_5 T_Q$$

where the a_i 's are constants.

With this assumption in hand, it is planned to learn the forward direction map, from the local agent costs to the global cost function terms, via a neural network mapping:

$$(C_{WF}^{local}, C_{XA}^{local}, C_{YA}^{local}, C_{ZA}^{local}, C_{PS}^{local}) \xrightarrow{\text{neural_network}} M_Q$$

where C_{WF}^{local} represents the wire feeder agent's cost, C_{iA}^{local} represents the i^{th} -axis agent's cost, C_{PS}^{local} represents the power supply agent's cost, and as above the a_i 's are constants. This forward neural network mapping will then be used to produce the reverse gradient mapping by way of the back-propagation training method. In addition to training the neural network based on the traditional error prediction feedback of the forward network, the change in each of the global cost terms (ΔM_Q) due to changes in the local agents ($\Delta C_{WF}^{local}, \Delta C_{XA}^{local}, \Delta C_{YA}^{local}, \Delta C_{ZA}^{local}, \Delta C_{PS}^{local}$) will be back propagated through the network all the way to its inputs (this process is not used to update weights). By doing this, one is attempting to use the back-propagation training method to relate a change in global cost to the local agents by exploiting the gradient knowledge contained within the forward mapping of the neural network. We are not attempting to reverse map the input to outputs, instead we are only trying to obtain gradient information at the input of the neural network based on the change of the output of the network. In fact, the back-propagation algorithm is based on a gradient descent method, which back-propagates gradient information about the error in the outputs due to the inputs in a similar fashion.

It is planned to use radial basis neural networks within this part of the project because of their connection to Takagi-Sugeno fuzzy systems. This connection will be used to obtain a qualitative understanding of the mapping between the local agent costs and the global sub-cost. This is possible because radial basis neural networks and Takagi-Sugeno fuzzy systems have been shown to be mathematically the same, though developed from a different understanding [Spooner and Passino, 1996]. In essence, it is thought that one can develop a fuzzy model of the mapping process from local costs to global sub-costs by using a radial basis neural network [Passino, 1999]. This qualitative understanding of the relationships between local costs and global sub-costs can then be used in future model development for the welding process as well as in more traditional control systems for welding processes.

CONCLUSION

An approach to design of an intelligent machine has been presented based on distributed intelligence. Local agents are used to control individual machine functions and to process information needed by the machine functions. Examples of how this approach may be used to build a specific machine are presented for an arc spot welding application. A possible agent internal structure is presented that provides for local rules of behavior and safety considerations. An initial method for accomplishing distributed learning with global optimization has been presented. The learning method outline within this paper will form the basis for our continued research.

ACKNOWLEDGEMENT

This work is supported by the Basic Energy Sciences, Office of Science, U.S. Department of Energy, under DOE Idaho Operations Office Contract DE-AC07-99ID13727.

REFERENCES

- Albus, James. S., "Outline for a Theory of Intelligence," *IEEE Transactions on Systems, Man, and Cybernetics*, vol 21, no. 3, May/June 1991, pp. 473-509.
- Arimoto, S., S. Kawamura, and F. Miyazaki, "Bettering Operation of Robots by Learning," *Journal of Robotic Systems*, Vol. 1, pp. 123-140, March 1984.
- Bonabcau, Eric; Théraulaz, Guy. "Swarm Smarts," *Scientific American*, Mar. 2000, pp.72-79.
- Dorigo, Marco; Maniezzo, Vittorio; Colorni, Alberto. "Ant System: Optimization by a Colony of Cooperating Agents," *IEEE Transactions on Systems, Man, and Cybernetics-PartB: Cybernetics*, v 26, n 1, Fed. 1, 1996, pp. 29-41.
- Hölldobler, B. and Wilson, E. O., *the ANTS*, The Belknap Press of Harvard University Press, Cambridge, Mass.,1990.
- Johnson, J. A., and H. B. Smartt, "Advantages of an Alternative Form of Fuzzy Logic," *IEEE Transactions on Fuzzy Logic*, 3, pp. 149-157, 1995.
- Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., and Wehner, R., "A mobile robot employing insect strategies for navigation", *Robotics and Autonomous Systems*, 30 (2000) 39-64.
- Moore, K. L., *Iterative Learning Control for Deterministic Systems*, Series on Advances in Industrial Control, Springer-Verlag, 1993.
- Overgaard, L., Petersen, H. G., and Perram, J. W., "A General Algorithm for Dynamic Control of Multilink Robots", *The International Journal of Robotics Research*, Vol. 14, No. 3, June 1995, pp. 281-294.
- Overgaard, L., Petersen, H. G., and Perram, J. W., "Reactive Motion Planning: a Multiagent Approach", *Applied Artificial Intelligence*, 10:35-51, 1996.
- Passino, Kevin. "Intelligent Control," Workshop S-5, 38th IEEE Conference on Control, Phoenix, AZ, Dec. 7-10, 1999.
- Rumelhart, D. E., G. E. Hinton, and R. J. Williams, "Learning Internal Representations by Error Propagation," *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*, Vol. 1, ed. D. E. Rumelhart and J. L. McClelland, MIT Press, Cambridge, MA, pp. 318-362 (1986).
- Schatz, B., Chameron, S., Beugnon, G., and Collett, T. S., "The use of path integration to guide route learning in ants", *Nature*, Vol. 399, No., 24, June 1999, pp. 769-772.
- Spooner, Jeffrey; Passino, Kevin. "Stable Adaptive Control Using Fuzzy Systems and Neural Networks," *IEEE Transactions on Fuzzy Systems*, v 4, n 3, Aug. 1996, pp. 339-359.
- Uchiyama, M., "Formation of High Speed Motion Pattern of Mechanical Arm by Trial," *Transactions of the Society of Instrumentation and Control Engineering*, Vol. 19, pp. 706-12, May 1978.

LEARNING AND ADAPTATION IN MULTI-ROBOT TEAMS

Lynne E. Parker, Claude Touzet, and David Jung

Center for Engineering Science Advanced Research, Computer Science and Mathematics Division
Oak Ridge National Laboratory, Oak Ridge, TN 37831-6355

ABSTRACT

While considerable progress has been made in recent years toward the development of multi-robot teams, much work remains to be done before these teams are used widely in real-world applications. Of particular need are the development of mechanisms that enable robot teams to autonomously generate cooperative behaviors. This paper examines the issue of multi-robot learning and looks at various types of multi-robot learning. We briefly review various multi-robot learning approaches we have studied. The paper then presents the Cooperative Multi-robot Observation of Multiple Moving Targets (CMOMMT) application as a rich domain for studying the issues of multi-robot learning of new behaviors. We discuss the results of our hand-generated algorithm for CMOMMT and the potential for learning that was discovered from the hand-generated approach. We then describe our research in generating multi-robot learning techniques for the CMOMMT application and compare the results to the hand-generated solutions. Our results show that, while the learning approach performs better than random, naive approaches, much room still remains to match the results obtained from the hand-generated approach. The ultimate goal of this research is to develop techniques for multi-robot learning and adaptation that will generalize to cooperative robot applications in many domains, thus facilitating the practical use of multi-robot teams in a wide variety of real-world applications.

INTRODUCTION

Research in multi-robot cooperation has grown significantly in recent years. While the growth of this research is due in part to a pure scientific interest in teams of autonomous robots, much of the growth is due to the increasing realization by the user community that teams of robots may provide solutions to difficult problems that previously were untenable. Certainly, it has been shown (e.g., in [1] and elsewhere) that multi-robot teams can increase the reliability, flexibility, robustness, and efficiency of automated solutions by taking advantage of the redundancy and parallelism of multiple team members.

However, before multi-robot teams will ever become widely used in practice, we believe that advances must be made in the development of mechanisms that enable the robot teams to autonomously generate cooperative behaviors and techniques. With the current state of the art, the implementation of cooperative behaviors on physical robot teams requires expert behavior programming and experimentation, followed by extensive tuning and revision of the cooperative control algorithms. It is unlikely that a significant real-world impact of cooperative robot teams will occur as long as the current level of effort is required to implement these systems.

Researchers have recognized that an approach with more potential for the development of cooperative control mechanisms is autonomous learning. Hence, much current work is ongoing in the field of multi-agent learning (e.g., [2]). Brooks and Mataric [3] identify four types of learning in robotic systems:

- Learning numerical functions for calibration or parameter adjustment,
- Learning about the world,
- Learning to coordinate behaviors, and
- Learning new behaviors.

Our research has examined several of these learning areas. In the first area – learning for parameter adjustment – we have developed the L-ALLIANCE architecture [4], which enables robots to adapt their behavior over time in response to changing team capabilities, team composition, and mission environment. This architecture, which is an extension of our earlier work on ALLIANCE [5], is a distributed, behavior-based architecture aimed for use in applications consisting of a collection of independent tasks. The key issue addressed in L-ALLIANCE is the determination of which tasks robots should select to perform during their mission, even when multiple robots with heterogeneous, continually changing capabilities are present on the team. In this approach, robots monitor the performance of their teammates performing common tasks, and evaluate their performance based upon the time of task completion. Robots then use this information throughout the lifetime of their mission to automatically update their control parameters according to the L-ALLIANCE update mechanism. We note, however, that the parameter update strategy used by L-ALLIANCE is dependent upon the assumption of independent subtasks whose performance can be evaluated based upon the time of task completion. This assumption does not hold for the CMOMMT application domain that we describe in this paper.

We have also addressed approaches for learning in the second area – learning about the world. We have implemented a multi-robot system in which one robot learns to communicate symbolic information about the environment to another robot [6, 7]. In particular, we developed a two-robot team that was specifically designed to be heterogeneous, such that neither could successfully achieve the task alone. The robots were also endowed with a mechanism for learning to communicate task-specific symbolic information. To accomplish learning, a number of self-organized levels build upon each other. First, each of the robots learns a topological map of its environment. Once the topological map self-organizes sufficiently, the robots can learn successively higher-level relationships between locations, resulting in the emergence of a navigation capability. When the robots have a consistent concept of “location” as defined in terms of their sensory and behavioral suite, they will begin to have success in communicating task-specific information. Specifically, each robot is endowed with behavior that results in the robots communicating in order to ground specific symbols to particular locations. Having learned a shared grounding that maps a set of symbols to locations, the robots can communicate important task-specific locations to each other. The receiving robot can interpret the communicated information, drive to the communicated location, and perform mission-specific tasks. The robot soon learns that this reliably assists in the completion of its tasks, and is more efficient than performing the task without communication. The result is that the robots together learn to adapt their behavior toward more efficient mission completion by learning to represent and navigate around their environment and to communicate about it [8]. This approach has been successfully demonstrated in a laboratory cleaning task.

In the remainder of this paper, we discuss our research in the fourth topic area – that of learning new behaviors in multi-robot teams. The types of applications that are typically studied for this area of multi-robot learning vary considerably in their characteristics. Some of the applications include air fleet control [9], predator/prey [10, 11, 12], box pushing [13], foraging [14], and multi-robot soccer [15, 16]. Particularly challenging domains for multi-robot learning are those tasks that are *inherently* cooperative. By this, we mean that the utility of the action of one robot is dependent upon the current actions of the other team members. Inherently cooperative tasks cannot be decomposed into independent subtasks to be solved by a distributed robot team. Instead, the success of the team throughout its execution is measured by the combined actions of the robot team, rather than the individual robot actions. This type of task is a particular challenge in multi-robot learning, due to the difficulty of assigning credit for the individual actions of the robot team members.

Of these previous application domains that have been studied in the context of multi-robot learning, only the multi-robot soccer domain addresses inherently cooperative tasks with more than two robots while also addressing the real-world complexities of embodied robotics, such as noisy and inaccurate sensors and effectors in a dynamic environment that is poorly modeled. To add to the field of challenging application domains for multi-robot learning, we have defined and have been studying a new application domain – the Cooperative Multi-robot Observation of Multiple Moving Targets (CMOMMT) – that is not only an inherently cooperative task, but, unlike the multi-robot soccer domain, is also a domain that must deal with issues of scalability to large numbers of robots.

In the next section, we define the CMOMMT application. We then describe a hand-generated solution to this problem, along with the results we obtained with this approach. We then define a learning approach to enable robot teams to autonomously generate viable solutions to the CMOMMT application and compare the results to the hand-generated approach. The final section of the paper concludes with some summary remarks.

THE CMOMMT APPLICATION

The application domain that we are studying for use as a multi-robot learning testbed is the problem we entitle *Cooperative Multi-robot Observation of Multiple Moving Targets* (CMOMMT). This problem is defined as follows. Given:

- S : a two-dimensional, bounded, enclosed spatial region
- \mathcal{V} : a team of m robot vehicles, $v_i, i = 1, 2, \dots, m$, with 360° field of view observation sensors that are noisy and of limited range
- $\mathcal{O}(t)$: a set of n targets, $o_j(t), j = 1, 2, \dots, n$, such that target $o_j(t)$ is located within region S at time t

We say that a robot, v_i , is *observing* a target when the target is within v_i 's sensing range. Define an $m \times n$ matrix $B(t)$, as follows:

$$B(t) = [b_{ij}(t)]_{m \times n} \text{ such that } b_{ij}(t) = \begin{cases} 1 & \text{if robot } v_i \text{ is observing target } o_j(t) \text{ in } S \text{ at time } t \\ 0 & \text{otherwise} \end{cases}$$

Then, the goal is to develop an algorithm, which we call *A-CMOMMT*, that maximizes the following metric A :

$$A = \sum_{t=1}^T \sum_{j=1}^n \frac{g(B(t), j)}{T}$$

where:

$$g(B(t), j) = \begin{cases} 1 & \text{if there exists an } i \text{ such that } b_{ij}(t) = 1 \\ 0 & \text{otherwise} \end{cases}$$

In other words, the goal of the robots is to maximize the average number of targets in S that are being observed by at least one robot throughout the mission that is of length T time units. Additionally, we define *sensor_coverage*(v_i) as the region visible to robot v_i 's observation sensors, for $v_i \in \mathcal{V}$. Then we assume that, in general, the maximum region covered by the observation sensors of the robot team is much less than the total region to be observed. That is,

$$\bigcup_{v_i \in \mathcal{V}} \text{sensor_coverage}(v_i) \ll S.$$

This implies that fixed robot sensing locations or sensing paths will not be adequate in general, and instead, the robots must move dynamically as targets appear in order to maintain their target observations and to maximize the coverage.

The CMOMMT application is an excellent domain for embodied multi-robot learning and adaptation. CMOMMT offers a rich testbed for research in multi-robot cooperation, learning, and adaptation because it is an inherently cooperative task. In addition, many variations on the dynamic, distributed sensory coverage problem are possible, making the CMOMMT problem arbitrarily more difficult. For example, the relative numbers and speeds of the robots and the targets to be tracked can vary, the availability of inter-robot communication can vary, the robots can differ in their sensing and movement capabilities, the terrain may be either enclosed or have entrances that allow objects to enter and exit the area of interest, and so forth. Many other subproblems can also be addressed, including the physical tracking of targets (e.g. using vision, sonar, IR, or laser range), prediction of target movements, multi-sensor fusion, and so forth.

A HAND-GENERATED SOLUTION TO CMOMMT

We have developed a hand-generated solution [17, 18] to the CMOMMT problem that performs well when compared to various control approaches. This solution has been implemented on both physical and simulated robots to demonstrate its effectiveness. The hand-generated solution, which we call *A-CMOMMT*, is described briefly as follows. Robots use weighted local force vectors that attract them to nearby targets and repel them from nearby robots. The weights are computed in real time by a higher-level reasoning system in each robot, and are based on the relative locations of the nearby robots and targets. The weights are aimed at generating an improved collective behavior across robots when utilized by all robot team members.

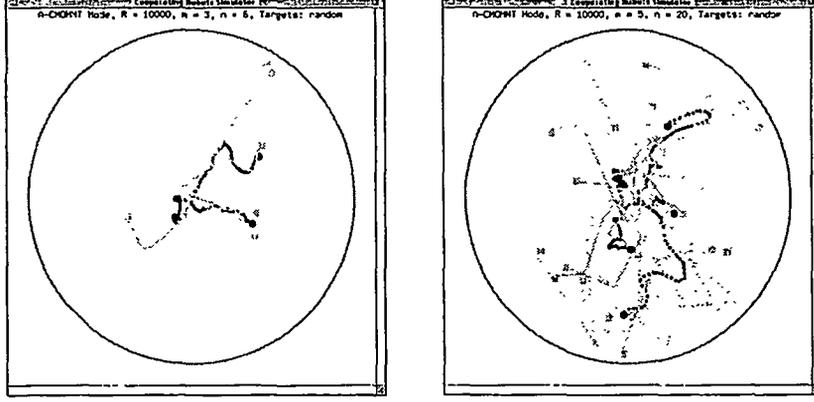


Figure 1: Simulation results of three robots and six targets (first image), and five robots and twenty targets (second image), with the robots using the hand-generated solution to CMOMMT, and the targets moving randomly.

The local force vectors are calculated as follows. The magnitude of the force vector attraction of robot v_l relative to target o_k , denoted $|\mathbf{f}_{lk}|$, for parameters $0 < do_1 < do_2 < do_3$, is:

$$|\mathbf{f}_{lk}| = \begin{cases} \frac{-1}{do_1} & \text{for } d(v_l, o_k) \leq do_1 \\ \frac{1}{do_2 - do_1} & \text{for } do_1 < d(v_l, o_k) \leq do_2 \\ \frac{-do_2}{do_3 - do_2} & \text{for } do_2 < d(v_l, o_k) \leq do_3 \\ 0 & \text{otherwise} \end{cases}$$

where $d(a, b)$ returns the distance between two entities (i.e., robots and/or targets). The magnitude of the force vector repulsion of robot v_l relative to robot v_i , denoted $|\mathbf{g}_{li}|$, for parameters $0 < dr_1 < dr_2$, is:

$$|\mathbf{g}_{li}| = \begin{cases} -1 & \text{for } d(v_l, v_i) \leq dr_1 \\ \frac{1}{dr_2 - dr_1} & \text{for } dr_1 < d(v_l, v_i) \leq dr_2 \\ 0 & \text{otherwise} \end{cases}$$

Determining the proper setting of the parameters do_1, do_2, do_3, dr_1 , and dr_2 is one approach to solving the CMOMMT multi-robot learning task using a *a priori* model-based technique.

Using only local force vectors for this problem neglects higher-level information that may be used to improve the team performance. Thus, the hand-generated approach enhances the control approach by including higher-level control to weight the contributions of each target's force field on the total computed field. This higher-level knowledge can express any information or heuristics that are known to result in more effective global control when used by each robot team member locally. The hand-generated approach expresses this higher-level knowledge in the form of a weight, w_{lk} , that reduces robot v_l 's attraction to a nearby target o_k if that target is within the field of view of another nearby robot. Using these weights helps reduce the overlap of robot sensory areas toward the goal of minimizing the likelihood of a target escaping detection.

The higher-level weight information is combined with the local force vectors to generate the commanded direction of robot movement. This direction of movement for robot v_l is given by:

$$\sum_{k=1}^n w_{lk} \mathbf{f}_{lk} + \sum_{i=1, i \neq l}^m \mathbf{g}_{li}$$

where \mathbf{f}_{lk} is the force vector attributed to target o_k by robot v_l and \mathbf{g}_{li} is the force vector attributed to robot v_i by robot v_l . To generate an (x, y) coordinate indicating the desired location of the robot corresponding to the resultant force vector, we scale the resultant force vector based upon the size of the robot. The robot's speed and steering commands are then computed to move the robot in the direction of that desired location.

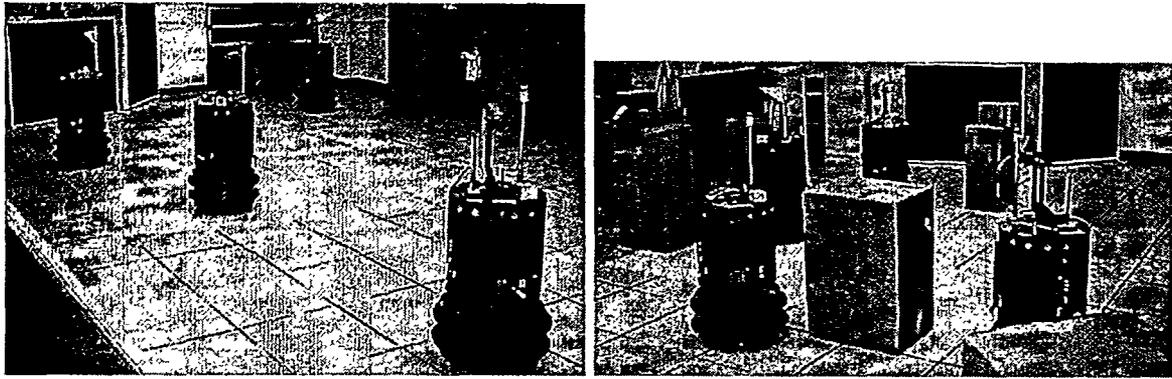


Figure 2: Robot team executing hand-generated solution to CMOMMT. The first photo shows robots operating in area with no obstacles. The second photo shows the robots amidst random distributed obstacles.

Results from Hand-Generated Solution

Figure 1 shows two of the simulation runs of the hand-generated algorithm (out of over 1,000,000 simulation test runs), in which (first image) three robots attempt to observe six targets, and (second image) five robots attempt to observe twenty targets. Figure 2 shows snapshots of two of the physical robot experiments (out of over 800) in which the robots perform the task either with no obstacles in the work area (first photo) or with randomly distributed obstacles (second photo).

The results of the hand-generated approach to CMOMMT vary depending upon a number of factors, including the relative numbers of robots and targets, the size of the work area, the motions of the targets (i.e., whether random or evasive), and the setting of the weights. In general, the *A-CMOMMT* algorithm performed best for a ratio of targets to robots greater than 1/2. We compared the hand-generated *A-CMOMMT* approach with a non-weighted local force vector approach, as well as two control cases in which robots either maintained fixed positions or are moved randomly. Figure 3 gives a typical set of these comparative results. Refer to [17] for more details on these results.

LEARNING IN THE CMOMMT APPLICATION

We have studied the CMOMMT problem from a learning perspective without the assumption of an *a priori* model [19]. This approach uses a combination of reinforcement learning, lazy learning, and a Pessimistic algorithm able to compute for each team member a lower bound on the utility of executing an action in a given situation. The challenges in this multi-robot learning problem include a very large search space, the need for communication or awareness of robot team members, and the difficulty of assigning credit in an inherently cooperative problem.

In this learning approach, lazy learning [20] is used to enable robot team members to build a memory of situation-action pairs through random exploration of the CMOMMT problem. A reinforcement function gives the utility of a given situation. The pessimistic algorithm for each robot then uses the utility values to select the action that maximizes the lower bound on utility. The resulting algorithm is able to perform considerably better than a random action policy, although it is still significantly inferior to the hand-generated algorithm described in the previous section. However, even with a performance less than that of the hand-generated solution, this approach makes an important contribution because it does not assume the existence of a model (as is the case in the Partially Observable Markov Decision Process (POMDP) domain), the existence of local indicators that help individual robots perform their tasks, nor the use of symbolic representations. The following subsections describe this approach and its results in more detail.

Lazy learning and Q-learning

Lazy learning [20] – also called instance-based learning – promotes the principle of delaying the use of the gathered information until the necessity arises (see Fig. 4). The same pool of information (i.e., memory) is used for different behavior syntheses. The lazy memory provides a good way of reducing the duration of any robotic learning application. In the context of reinforcement learning, lazy learning provides an instantaneous set of situation-action pairs (after the initial and unique sampling phase). Lazy learning samples the situation-action space according to a random action selection policy, storing the succession

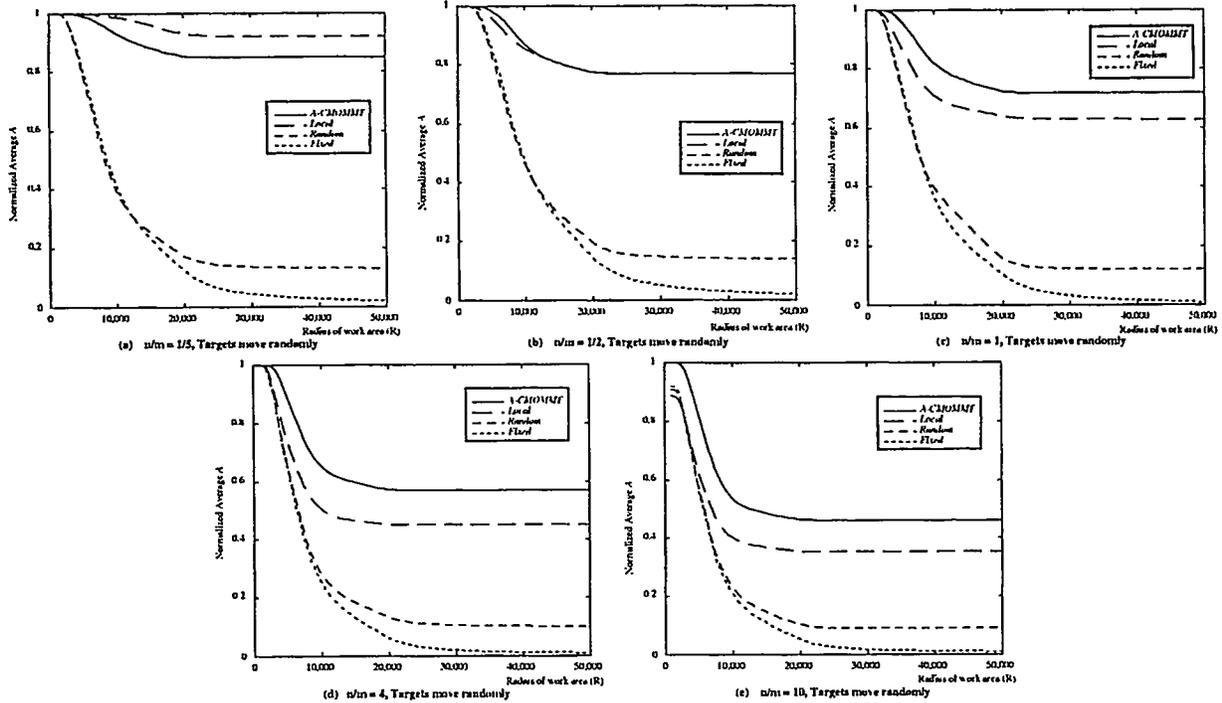


Figure 3: Simulation results of four distributed approaches to cooperative observation, for random/linear target movements, for various ratios of number of targets (n) to number of robots (m).

of events in memory and, when needed, probes the memory for the best action. The exploration phase is performed only once. By storing situation-action pairs, a lazy memory builds a model of the situation transition function.

In order to express a behavior, the memory must be probed. To do this probing, we use a modified version of the technique proposed in [21]. In [21] the objective is to provide a method for predicting the rewards for state-action pairs without explicitly generating them. For the current real world situation, a situation matcher locates all the states in the memory that are within a given distance. If the situation matcher has failed to find any nearby situations, the action comparator selects an action at random. Otherwise, the action comparator examines the expected rewards associated with each of these situations and selects the action with the highest expected reward. This action is then executed, resulting in a new situation. There is a fixed probability (0.3) of generating a random action regardless of the outcome of the situation matcher. New situation-action pairs are added to the memory, along with a Q-value computed in the classical way [22]. Among similar situation-action pairs in the memory, an update of the stored Q-values is made. However, there is a limit to the generality of this lazy memory because the Q-values associated with the situation-action pairs only apply for a particular behavior. With the desire of reducing the learning time as much as possible, as well as preserving the generality of the lazy memory, we modified the algorithm as follows: (1) the situation matcher always proposes the set of nearest situations – no maximum distance is involved, (2) there is no random selection of actions by the action comparator, and (3) the Q-values are not stored with the situation-action pairs, but computed dynamically as the need arises.

The Pessimistic Algorithm

We define a Pessimistic Algorithm for the selection of the best action to execute for a given robot in its current local situation as follows: find the lower bounds of the utility value associated with the various potential actions that may be conducted in the current situation, then choose the action with the greatest utility. A lower bound is defined as the lowest utility value associated with a set of similar situations.

The idea behind the Pessimistic Algorithm is that a local robot situation is an incomplete observation of the true state of the system. Thus, instead of trying to solve the observation problem by completing the observation (usual POMDP approach), we are only interested in ranking the utility of the actions. If we use a unique instance of the memory to obtain the utility of the situation, then chances are that the

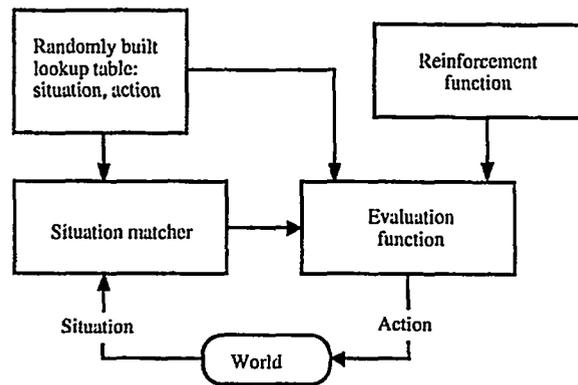


Figure 4: Lazy learning: randomly sampled situation-action pairs in the lookup table are used by the situation matcher to select the action to execute in the current situation. The reinforcement function qualifies the actions proposed, helping to select the best one.

utility attributed to this local situation is due in fact to other robot's actions. This probability decreases proportionally with the number of similar situations that are taken into account. If a large number of situations are considered, then there must be at least one for which the reward directly depends on the local situation. By taking the minimum utility value of the set of similar situations, we are guaranteed that, if the value is null, then the situation achieved does not imply losing target(s).

The Pessimistic Algorithm is then given as follows:

- Let M be the memory, a lookup table of situation-action pairs gathered during an exploration phase -- $M = [(s(1), a(1)), \dots, (s(t), a(t)), (s(t+1), a(t+1)), \dots]$.
- Let sit be the current situation.
- Find $S(sit)$, the set of n situations of M similar to sit .
- Let $S_{follow}(sit)$ be the set of the situations that directly follows each situation of $S(sit)$.
- Compute the lower bound (LB) of the utility value (U) associated with each situation $s(k) \in S_{follow}(sit)$:
 - $LB(s(k)) = \min(U(s(m)))$, for $s(m) \in S(s(k))$, the set of situations similar to $s(k)$.
- Execute the action that should take the robot to the new situation s^* : $s^* = \max(LB(s))$ and $s \in S_{follow}(sit)$.

The utility U associated with a given situation can be computed in many ways. It can be the exact value of the reinforcement function for this particular situation-action pair, or it can be a more elaborate variable. For example, in our experience we store the situation-action pairs, plus the number of targets under observation in the lookup table (M). However, the value that is used as utility is +1 if one or more targets have been acquired compared to the previous situation, -1 if one or more targets have been lost, or 0 otherwise. An exact Q value requires running the Q-learning algorithm with the samples stored in the memory.

Results of Learning Approach

We studied the efficiency of the Pessimistic Algorithm by comparing the performance of a team of robots with a purely random action selection policy, a user-defined non-cooperative policy and *A-CMOMMT*. In these experiments, each robot situation is a vector of two times 16 components. The first 16 components code the position and orientation of the targets. It simulates a ring of 16 sensors uniformly distributed around the robot body. Each sensor measures the distance to the nearest target. The sensor position around the body gives the orientation. The second ring of 16 components code in the same manner the position and orientation of neighboring robots. The maximum range for a target or a robot to be seen is 1, for an arena radius of 5. The actions of each robot are rotation and forward movement. The measure of performance is the mean observation time of all targets.

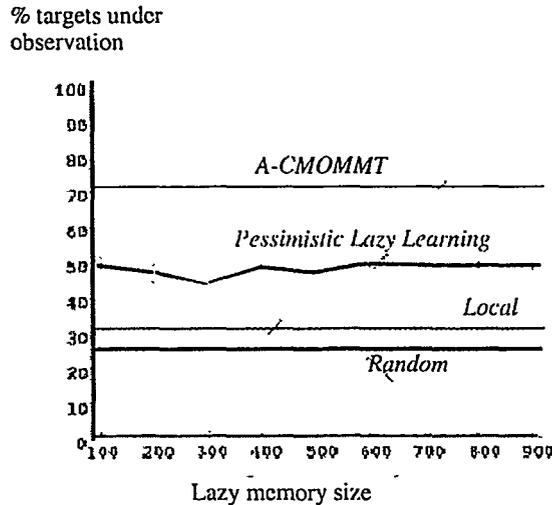


Figure 5: Performances of the Pessimistic lazy Q-learning approach compared to a random action selection policy, a user-defined non-cooperative policy and the hand-generated solution *A-CMOMMT*. The size of the lazy memory varies between 100 to 900 situation-action pairs. There are 10 robots and 10 randomly moving targets. The results are the mean of 10 different experiments per point for lazy learning policy, and 100 experiments for the other 3 policies. Each experiment duration is 1000 iterations.

Figure 5 shows the performance of a Pessimistic lazy Q-learning policy versus the size of the lazy memory, from 100 to 900 situation-action pairs. Each point is the average of 10 experiments. The standard deviation is also plotted on the graph. The lazy memories are obtained through an initial exploration involving from 15 to 25 targets and a single robot. During the sampling, the targets are fixed and the robot's policy is random action selection (with 5% chance of direction and orientation changes). The reinforcement function returns +1 if the total number of targets under observation increases, -1 if this number decreases, or 0 otherwise.

As we see there is an important performance gain associated with the Pessimistic lazy Q-learning over a purely random selection policy. This clearly demonstrates the importance of lazy Q-learning as a learning technique. Even more interestingly, lazy Q-learning performs much better than the user-defined non-cooperative policy (*Local*). It is important to note that neither policy is aware of the existence of the other robots. Both policies use the same sensory information – i.e., the distance and orientation of nearby targets. It is our opinion that the variation of performance is due to the fact that the lazy Q-learned behavior is somewhat less rigid than the user-defined policy. A lazy Q-learning guided robot will follow a target further than it could be, and, in doing so, will exhibit an erratic path, moving from one side of the target to another, back and forth without losing the target. In doing so, the surface under observation per unit of time is larger than the covered surface by the more rigid center-of-gravity-oriented robot. On the other hand, because it does not take into account the neighboring robots, it is easy to understand why the lazy Q-learned behavior performance cannot reach the level of the *A-CMOMMT* performance.

CONCLUSIONS

In this paper, we have proposed that the Cooperative Multi-robot Observation of Multiple Moving Targets (CMOMMT) application domain provides a rich testbed for learning and adaptation in multi-robot cooperative teams. We have described the need for learning and adaptation in multi-robot teams, and have defined the CMOMMT application, along with the characteristics that make it an interesting testbed for learning and adaptation. We reported on a hand-generated solution to the CMOMMT problem and discussed how the results from the implementation of this solution reveal the need for learning and adaptation in this domain. We discussed our work that uses the CMOMMT problem as a learning domain. The ultimate objective is to develop learning techniques using the CMOMMT domain that will generalize to other real-world domains, and will thus help realize the ultimate goal of enabling the widespread, practical use of multi-robot teams.

ACKNOWLEDGEMENTS

This research is sponsored by the Engineering Research Program of the Office of Basic Energy Sciences, U. S. Department of Energy. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U. S. Government purposes. Oak Ridge National Laboratory is managed by UT-Battelle, LLC for the U.S. Dept. of Energy under contract DE-AC05-00OR22725.

REFERENCES

- [1] L. E. Parker. On the design of behavior-based multi-robot teams. *Journal of Advanced Robotics*, 1996.
- [2] Gerhard Weiss and Sandip Sen, editors. *Adaption and Learning in Multi-Agent Systems*. Springer, 1996.
- [3] Rodney A. Brooks and Maja J. Mataric. Real robots, real learning problems. In Jonathan H. Connell and Sridhar Mahadevan, editors, *Robot Learning*. Kluwer Academic Publishers, 1993.
- [4] L. E. Parker. Lifelong adaptation in heterogeneous teams: Response to continual variation in individual robot performance. *Autonomous Robots*, 8(3), July 2000.
- [5] L. E. Parker. ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220-240, 1998.
- [6] David Jung, Gordon Cheng, and Alexander Zelinsky. Robot cleaning: An application of distributed planning and real-time vision. In Alexander Zelinsky, editor, *Field and Service Robotics*, pages 187-194. Springer, 1998.
- [7] David Jung and Alexander Zelinsky. An architecture for distributed cooperative planning in a behaviour-based multi-robot system. *Robotics and Autonomous Systems*, 26:149-174, 1999.
- [8] David Jung and Alexander Zelinsky. Grounded symbolic communication between heterogeneous cooperating robots. *Autonomous Robots*, 8(3), July 2000.
- [9] Randall Steeb, Stephanie Cammarata, Frederick Hayes-Roth, Perry Thorndyke, and Robert Wesson. Distributed intelligence for air fleet control. Technical Report R-2728-AFFA, Rand Corp., 1981.
- [10] M. Benda, V. Jagannathan, and R. Dodhiawalla. On optimal cooperation of knowledge sources. Technical Report BCS-G2010-28, Boeing AI Center, August 1985.
- [11] R. Korf. A simple solution to pursuit games. In *Working Papers of the 11th International Workshop on Distributed Artificial Intelligence*, pages 183-194, 1992.
- [12] Thomas Haynes and Sandip Sen. Evolving behavioral strategies in predators and prey. In Gerard Weiss and Sandip Sen, editors, *Adaptation and Learning in Multi-Agent Systems*, pages 113-126. Springer, 1986.
- [13] S. Mahadevan and J. Connell. Automatic programming of behavior-based robots using reinforcement learning. In *Proceedings of AAAI-91*, pages 8-14, 1991.
- [14] Maja Mataric. *Interaction and Intelligent Behavior*. PhD thesis, Massachusetts Institute of Technology, 1994.
- [15] P. Stone and M. Veloso. A layered approach to learning client behaviors in the robocup soccer server. *Applied Artificial Intelligence*, 12:165-188, 1998.
- [16] S. Marsella, J. Adibi, Y. Al-Onaizan, G. Kaminka, I. Muslea, and M. Tambe. On being a teammate: Experiences acquired in the design of robocup teams. In O. Etzioni, J. Muller, and J. Bradshaw, editors, *Proceedings of the Third Annual Conference on Autonomous Agents*, pages 221-227, 1999.
- [17] L. E. Parker. Distributed algorithms for multi-robot observation of multiple moving targets. *To appear in Autonomous Robots*, 2000.
- [18] L. E. Parker. Cooperative robotics for multi-target observation. *Intelligent Automation and Soft Computing, special issue on Robotics Research at Oak Ridge National Laboratory*, 5(1):5-19, 1999.
- [19] C. Touzet. Distributed lazy q-learning for cooperative mobile robots. *Submitted for publication*, 2000.
- [20] D. Aha, editor. *Lazy Learning*. Kluwer Academic Publishers, 1997.
- [21] J.W. Sheppard and S.L. Salzberg. A teaching strategy for memory-based control. In D. Aha, editor, *Lazy Learning*, pages 343-370. Kluwer Academic Publishers, 1997.
- [22] C. Watkins. *Learning from Delayed Rewards*. PhD thesis, King's College, Cambridge, 1989.

INFORMATION FUSION IN PHYSICAL SYSTEMS USING PHYSICAL LAWS

Nageswara S. V. Rao, David B. Reister, Jacob Barhen

Center for Engineering Science Advanced Research
Computer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831-6355

ABSTRACT

A physical system can be described by a set of parameters which are related to each other by certain physical laws. We consider that each parameter is either measured by sensors and/or estimated computationally. As a result, the estimated or measured values for a single parameter could be widely varying. We address the problem of fusing various measurements and/or estimates to improve the accuracy in estimating the parameter, when the error distributions of sensors and estimators are unknown. We propose a fusion method based on the least violation of physical laws that relate the parameters. Under the bounded variation condition of the physical law, we derive distribution-free performance bounds for a fusion rule computed using a finite sample. This result also implies the asymptotic convergence of the estimated fusion rule to the best possible rule which can be obtained under a complete knowledge of the error distributions.

INTRODUCTION

We consider a multiple sensor system that measures physical parameters of a system. Each parameter is either measured using an instrument or estimated using a computational method based on the measurements. There could be both systematic and random errors in the measurements as well as in the estimators. Furthermore, it may not be possible to know the actual parameter values, since all measurements and estimators (based on measurements) can introduce errors of different types. Consequently, there are a number of estimated and/or

measured values for each parameter. In general, very accurate sensor noise models could be derived from device properties. But such models are difficult to derive for estimators based on complicated computer codes. On the other hand, it is relatively easy to collect measurements using the sensors, and then compute the estimators based on measurements. Fusion rules based on measurements have been developed [6], and are shown to be very effective in practical engineering and robot systems. We consider the fusion of various measurements and estimators such that for each parameter the fused estimate is superior to the individual estimator or measurement. Since the actual parameter values are not known, the traditional pattern recognition or fusion solutions are not applicable here. The actual values, if available, could be used as the training data to design powerful fusers [6, 7, 5]. The lack of "traditional" training data motivated a new paradigm [8] that utilizes physical laws. In this paper, we extend the results of [8], which are valid for only Lipschitz physical laws, to include non-smooth physical laws.

The parameters of physical systems are related by physical laws, which are typically derived from first principles, and are verified by independent mechanisms. For example, for a simple mass sliding on a friction-less surface, we have $f = ma$, where f is force, m is mass and a is acceleration. If we choose a measurement or an estimator for each parameter, the accuracy of this set depends on how well the physical law is satisfied, and the "violation" of physical law is an indication of error. Thus, the set of estimators that achieves the least violation of the physical law is the most preferred. By fusing the measurements and estimators, one can achieve, in principle, performances superior to any set of estimators. The performance of the fuser, however, depends on the knowledge about the error distributions. If the sensor error distributions are known, the isolation fusers [5] can be designed to ensure fuser performance at least as good as best set of estimators. In the practical case, where we only have sensor measurements, we showed in [8] that (smooth) Lipschitz physical laws can be used to design the fuser. In particular, this result holds asymptotically (i. e. as sample size approaches infinity) and for finite samples, under Lipschitz properties of the physical law and fusion functions. These results are not valid if the physical law is discontinuous or the individual fusion functions are not smooth.

In this paper, we show that the bounded variation of the physical law is sufficient to obtain the finite sample as well as the asymptotic guarantees of the fusion procedure. This is achieved by employing fusers classes with the isolation property [5] and the bounded pseudo-dimension [1]; these conditions are satisfied by a number of fusers such as certain feedforward networks and linear combinations. The results of this paper enable us to utilize discontinuous physical laws and fusion rules to achieve performance superior to the best set of measurements. For finite sample sizes, we show distribution-free result that given large enough sample the fuser performs better than the best set of estimators within a specified precision and with a specified probability. This result also implies that the computed fuser asymptotically approaches the best fusion rule (computable under complete knowledge of the distributions) as the sample size increases.

In section 2, we describe the fusion problem originally formulated in [8]. We show how physical laws can be used to design a fuser in Section 3 under the above conditions. We briefly discuss fusion of data collected in the exploration of methane hydrates in Section 4.

PHYSICAL SYSTEMS AND LAWS

A physical system is specified by the parameters $P(z) = (p_1(z), p_2(z), \dots, p_n(z))$ with $p_i(z) \in \mathfrak{R}$, where z is one-dimensional variable such as time or position. Each parameter p_i is measured by a_i instruments and estimated by b_i estimators ($a_i \geq 0$, $b_i \geq 0$, and $a_i + b_i \geq 1$). The measurements corresponding to $p_i(z)$ are denoted by

$$m_i(z) = \{m_{i,1}(z), m_{i,2}(z), \dots, m_{i,a_i}(z)\}$$

and the corresponding estimators are denoted by

$$e_i(z) = \{e_{i,1}(z), e_{i,2}(z), \dots, e_{i,b_i}(z)\}.$$

Thus, there are $a_i + b_i$ competing values for each parameter, and in general we do not know which one is more accurate. The measurements are assumed to be noisy in that repeated measurements by a sensor of $p_i(z) = x$ for a fixed value are distributed independently according to the distribution $P_{m_{i,j}|x}$, which is denoted by $P_{m_{i,j}|p_i(z)}$. Thus, $m_{i,j}$ is a random variable. The estimator $e_{i,j}$ is a (deterministic) function of the measurements, and hence is also a random variable. The joint distribution of the measurements is denoted by $P_{m_1, m_2, \dots, m_n | p_1, p_2, \dots, p_n}$.

There is a physical law

$$L[p_1(z), p_2(z), \dots, p_n(z)] = 0$$

which relates the actual parameters corresponding to z . For the example of mass in the previous section, we have $L[f, m, a] = (f - ma)^2 = 0$. We assume that $L[\cdot]$ satisfies the reasonable *monotonicity* condition: for any y_1, y_2 , $|y_1| \leq |y_2|$, we have

$$|L[p_1(z), \dots, p_i(z) + y_1, \dots, p_n(z)]| \leq |L[p_1(z), \dots, p_i(z) + y_2, \dots, p_n(z)]|.$$

Monotonicity means that accurate parameter estimators yield no lesser "magnitude" of violation of the law compared to less accurate estimators.

Consider a single estimator or measurement \hat{p}_i for the parameter. The closeness of $L[\hat{p}_1(z), \hat{p}_2(z), \dots, \hat{p}_n(z)]$ to 0 determines how closely the law is satisfied. Let a *basic set*, denoted by S , be a set of measurements and estimators such that for each parameter we choose precisely one measurement or estimator (but not both). The total error due to S is given by

$$\hat{E}(S) = \sum_z L[\hat{p}_1(z), \hat{p}_2(z), \dots, \hat{p}_n(z)].$$

In all there are $\prod_{i=1}^n (a_i + b_i)$ possible basic sets, and \hat{S} be the one with least error such that $\hat{E}(\hat{S}) = \min_S \hat{E}(S)$. The expected error of S is denoted by

$$E(S) = \sum_z \int L[\hat{p}_1(z), \hat{p}_2(z), \dots, \hat{p}_n(z)] P_{m_1, \dots, m_n | p_1, \dots, p_n},$$

and let S^* be the one with the least expected error such that $E(S^*) = \min_S E(S)$. Note that S^* minimizes the expected error but \hat{S} in general does not. More detailed discussion of the physical laws can be found in [8].

DATA FUSION BASED ON PHYSICAL LAWS

A *fusion function* $f_i \in \mathcal{F}_i$ for parameter p_i combines the measurements and estimators such that $f_i(m_i(z), e_i(z))$ is an estimate of $p_i(z)$. Let $f = (f_1, \dots, f_n)$ denote the *fuser* for all parameters. The expected error due to the fused estimate is

$$E(f) = \sum_z \int L[f_1(m_1(z), e_1(z)), \dots, f_n(m_n(z), e_n(z))] dP_{m_1, \dots, m_n | p_1, \dots, p_n}$$

and let $f^* \in \mathcal{F}_1 \times \dots \times \mathcal{F}_n$ be the one with the least expected error. In general $E(f)$ cannot be computed if the error distributions are not known, and hence f^* is not computable. In stead, we compute \hat{f} that minimizes the empirical cost given by

$$\hat{E}(f) = \sum_z L[f_1(m_1(z), e_1(z)), \dots, f_n(m_n(z), e_n(z))],$$

based on a set of iid measurements (also called the sample)

$$\langle (m_1(z), e_1(z)), \dots, (m_n(z), e_n(z)) \rangle: z = 1, \dots, s \}.$$

Now we discuss methods that ensure $E(f^*) \leq E(S^*)$, and more importantly based on a computable \hat{f} that

$$E(\hat{f}) < E(S^*),$$

with a specified probability based entirely on the measurements and without any knowledge of the underlying distributions.

A fuser class $\mathcal{F}_i = \{f_i(y) : \mathfrak{R}^{a_i+b_i} \mapsto \mathfrak{R}\}$, for $y = (y_1, \dots, y_{(a_i+b_i)})$, has the *isolation property* [5] if it contains the function $\tau_j(y) = y_j$ for all $j = 1, 2, \dots, (a_i + b_i)$. If each \mathcal{F}_i satisfies the isolation property, then the following conditions are directly satisfied.

$$E(f^*) \leq E(S^*) \quad \text{and} \quad \hat{E}(\hat{f}) \leq \hat{E}(\hat{S}).$$

The first condition is useful only if f^* can be computed, which in turn requires the knowledge of the distributions. If the distributions are not known, then \hat{f} can be used as an approximation. In [8] we showed that with probability $1 - \delta$, we have

$$E(\hat{f}) - E(f^*) \leq \epsilon$$

given a sufficiently large sample, when the physical law and the fusers classes are Lipschitz. In general, however, physical laws may not be Lipschitz, especially if they involve discrete components or discontinuities. For example, consider the simple case of H_2O heated in a container, where p_1 denotes the temperature and $p_2 \in \{0, 1\}$ is the state, i. e. $p_2 = 0$ denotes liquid and $p_2 = 1$ denotes steam. Let T_0 denote the boiling temperature under this condition. Then, one of the physical laws is: $p_2 = 0$ if $p_1 < T_0$ and $p_2 = 1$ otherwise. This law can be represented as

$$L[p_1, p_2] = p_2 1_{\{p_1 < T_0\}} + (p_2 - 1) 1_{\{p_1 \geq T_0\}} = 0,$$

where the indicator function 1_C is 1 if condition C is true and is 0 otherwise. Here, $L[\cdot]$ is not Lipschitz. To address the cases typified by such $L[\cdot]$, we consider the class of functions

with bounded variation [2], which allow for discontinuities and discrete values, and include Lipschitz functions as a subclass.

Consider a function one-dimensional function $h : [-A, A] \mapsto \mathfrak{R}$. For $A < \infty$, a set of points $P = \{x_0, x_1, \dots, x_n\}$ such that $-A = x_0 < x_1 < \dots < x_n = A$ is called a *partition* of $[-A, A]$. The collection of all possible partitions of $[-A, A]$ is denoted by $\mathcal{P}[-A, A]$. A function $g : [-A, A] \mapsto \mathfrak{R}$ is of *bounded variation*, if there exists M such that for any partition $P = \{x_0, x_1, \dots, x_n\}$, we have $\sum(P) = \sum_{k=1}^n |f(x_k) - f(x_{k-1})| \leq M$. A multivariate function $g : [-A, A]^d \mapsto \mathfrak{R}$ is of bounded variation if it is so in each of its input variable for every value of the other input variables. The following are useful facts about the functions of bounded variation: (i) not all continuous functions are of bounded variation, e.g. $g(x) = x \cos(\pi/(2x))$ for $x \neq 0$ and $g(0) = 0$; (ii) differentiable functions on compact domains are of bounded variation; and (iii) absolutely continuous functions, which include Lipschitz functions, are of bounded variation.

We utilize the fuser classes with finite pseudo-dimension [1], which is described as follows. Let \mathcal{G} be a set of functions mapping from a domain X to \mathfrak{R} and suppose that $S = \{x_1, x_2, \dots, x_m\} \subseteq X$. Then S is *pseudo-shattered* by \mathcal{G} if there are real numbers r_1, r_2, \dots, r_m such that for each $b \in \{0, 1\}^m$ there is a function g_0 in \mathcal{G} with $\text{sgn}(f_b(x_i) - r_i) = b_i$ for $1 \leq i \leq m$. Then \mathcal{G} has the *pseudo-dimension* d if d is the maximum cardinality of a subset S of X that is pseudo-shattered by \mathcal{G} . If no such maximum exists, we say that \mathcal{G} has infinite pseudo-dimension. The pseudo-dimension of \mathcal{G} is denoted $\text{Pdim}(\mathcal{G})$. Pseudo-dimensions are known for several classes such as sigmoid neural networks, vector spaces, and linear combinations (see [1]).

Let \mathcal{G} be the class of functions from Z to into $[0, M]$, where $M > 0$, and let P be a probability measure on Z . Then $d_{L^1(P)}$ is the pseudo metric on \mathcal{G} defined by

$$d_{L^1(P)}(g_1, g_2) = E(|g_1 - g_2|) = \int_Z |g_1(z) - g_2(z)| dP(z)$$

for all $g_1, g_2 \in \mathcal{G}$. The *covering number* $\mathcal{N}(\epsilon, \mathcal{G}, d_{L^1(P)})$ of a function class \mathcal{G} is the smallest cardinality for a subclass $\mathcal{G}^* = \{g^*\}$ of \mathcal{G} such that $d_{L^1(P)}(g, g^*) \leq \epsilon$, for each $g \in \mathcal{G}$.

Theorem 1 Consider that the physical law is of bounded variation such that $|L(p)| \leq M_L$ for all p . Let parameters, estimators and measurements are bounded. Let each fuser class \mathcal{F}_i have finite pseudo-dimension d_i , and each fuser function g be bounded such that $|g(\cdot)| \leq M$ for all parameters. Let $d = \sum_{i=1}^n d_i$. Then given a sample of size

$$s = \frac{256M_L^2}{\epsilon^2} \left[4d \ln \left(\frac{128eM}{\epsilon} \right) + (n+1) \ln(4/\delta) \right],$$

we have

$$\mathbf{P} \left[E(\hat{f}) - E(f^*) > \epsilon \right] \leq \delta,$$

irrespective of the sensor distributions. Furthermore, $E(\hat{f}) \rightarrow E(f^*)$, as $s \rightarrow \infty$.

Proof: Consider the function class $\mathcal{L} = \{L(f_1, f_2, \dots, f_n) : f_1 \in \mathcal{F}_1, \dots, f_n \in \mathcal{F}_n\}$, where $L(f_1, f_2, \dots, f_n)$ is defined on a bounded domain. By combining Vapnik's argument (see [8] for details) with Theorem 3 of Hausser [4], we obtain

$$\mathbf{P} [E(\hat{f}) - E(f^*) > \epsilon] \leq 2E [\min (2\mathcal{N}(\epsilon/32, \mathcal{L}, d_{L^1}))] e^{-\frac{\epsilon^2 s}{256M^2}}. \quad (1)$$

We subsequently show that $\mathcal{N}(\epsilon, \mathcal{L}, d_{L^1(P)}) \leq 2^{2n} \left(\frac{4eM}{\epsilon} \ln \frac{4eM}{\epsilon}\right)^{2d}$, for any P . The sample size follows by using this cover bound in right hand side of Eq (1), and equating to δ and then solving for s .

In the rest of the proof we establish the bound on $\mathcal{N}(\cdot)$. Since $L(\cdot)$ is of bounded variance, it can be represented as a sum of two monotone functions $L = L_1 + L_2$. For $i = 1, 2$, let

$$\mathcal{L}_i = \{L_i(f_1, f_2, \dots, f_n) : f_1 \in \mathcal{F}_1, \dots, f_n \in \mathcal{F}_n\}.$$

Then let $\mathcal{L}_{ij} = \{L_i(p_1, \dots, p_{j-1}, f_j, p_{j+1}, \dots, p_n) : f_j \in \mathcal{F}_j\}$, which is a class of function obtained by composing a monotone function with functions from \mathcal{F}_i with bounded pseudo dimension. By Theorem 11.3 of [1], we have $\text{Pdim}(\mathcal{L}_{ij}) \leq \text{Pdim}(\mathcal{F}_i)$. Then by using Theorem 6 of [4] we have

$$\mathcal{N}(\epsilon, \mathcal{L}_{ij}, d_{L^1(P)}) \leq 2 \left(\frac{2eM}{\epsilon} \ln \frac{2eM}{\epsilon}\right)^{d_j}$$

for any measure P . By applying this cover bound for every component of L_i , we obtain

$$\mathcal{N}(\epsilon, \mathcal{L}_i, d_{L^1(P)}) \leq 2 \prod_{j=1}^n \left(\frac{2eM}{\epsilon} \ln \frac{2eM}{\epsilon}\right)^{d_j} = 2^n \left(\frac{2eM}{\epsilon} \ln \frac{2eM}{\epsilon}\right)^d$$

by the product rule. Since $L = L_1 + L_2$ we obtain

$$\mathcal{N}(\epsilon, \mathcal{L}, d_{L^1(P)}) \leq \mathcal{N}(\epsilon/2, \mathcal{L}_1, d_{L^1(P)}) \mathcal{N}(\epsilon/2, \mathcal{L}_2, d_{L^1(P)}) \leq 2^{2n} \left(\frac{4eM}{\epsilon} \ln \frac{4eM}{\epsilon}\right)^{2d}.$$

By noting that this bound is independent of P , we obtain

$$2E [\min (2\mathcal{N}(\epsilon/32, \mathcal{L}, d_{L^1}))] \leq 4N \left(\epsilon/32, \mathcal{L}, d_{L^1(P)}\right) \leq 2^{2n} \left(\frac{128eM}{\epsilon} \ln \frac{128eM}{\epsilon}\right)^{2d},$$

which yields the sample size as shown above. The asymptotic convergence follows from the Borel-Cantelli Lemma by showing

$$\sum_{l=1}^{\infty} \mathcal{N}(\epsilon, \mathcal{L}, d_{L^1}) \leq \sum_{l=1}^{\infty} 2^{2n} \left(\frac{128eM}{\epsilon} \ln \frac{128eM}{\epsilon}\right)^{2d} e^{-\frac{\epsilon^2 s}{256M^2}} < \infty$$

for every $\epsilon > 0$ in a manner identical to that in [8]. \square

The following corollary is a weaker version of Theorem 1 since $E(f^*) \leq E(S^*) \leq E(\hat{S})$.

Corollary 1 *Let \mathcal{F}_i satisfy the isolation property for all $i = 1, 2, \dots, n$. Under the same conditions as Theorem 1, we have following conditions satisfied.*

$$\mathbf{P} [E(\hat{f}) - E(S^*) > \epsilon] \leq \delta \quad \text{and} \quad \mathbf{P} [E(\hat{f}) - E(\hat{S}) > \epsilon] \leq \delta.$$

Informally speaking, this corollary shows that the error of the computed fuser \hat{f} is not likely to be much higher than that of the best basic set, and could be much smaller. Theorem 1 states that \hat{f} will be closer to f^* which can have much smaller error than S^* .

METHANE HYDRATES WELL LOGS

Gas hydrates are crystalline substances composed of water and gas, in which gas molecules are contained in cage-like lattices formed by solid water. One of the challenging problems is to predict the presence of hydrates using measurements collected at wells located in certain locations such as off the US coast in mid-Atlantic and Mackenzie Delta in Northwest Canada. At each well, a number of measurements are collected using a suite of sensors. These measurements include density, neutron porosity, acoustic transit-time, and electric resistivity, collected at various depths in the well [3]. Our focus is on the estimation of the *porosity* at various depths. Our data consists of 3045 sets of measurements each collected at different depths in a single well. There are a variety of methods to estimate porosity based on different principles and utilizing different measurements. We employed six known methods for estimating the porosity based on neutron measurements ($\hat{\phi}_1$), density measurements ($\hat{\phi}_2$), fluid velocity equation ($\hat{\phi}_3$), acoustic travel time based on S-wave ($\hat{\phi}_4$), time-average equation based on P-wave ($\hat{\phi}_5$), and Wood's equation ($\hat{\phi}_6$).

One of the well-established physical laws relates the parameters of porosity (ϕ), density (ρ), and hydrate concentration (ψ), as follows

$$L[\phi, \psi, \rho] = (\phi[\rho_m - (1 - \psi)\rho_w + \psi\rho_h] - \rho + \rho_m)^2 = 0,$$

where ρ_m , ρ_w , and ρ_h are known constants. In this equation, we use the only one measurement for density $\hat{\rho}$ and a single estimator $\hat{\psi}$ for the hydrate concentration using the Archie's equation. We consider a fuser based on the linear combination of the estimators

$$\hat{\phi}_F = w_7 + \sum_{i=1}^6 w_i \hat{\phi}_i,$$

where $(w_1, \dots, w_7) \in \mathbb{R}^7$ is the weight vector that minimizes the error based on measurements. The error achieved by $\hat{\phi}_F$ is about 20 times better than that of the best estimator $\hat{\phi}_4$ (details can be found in [8]). Note that $L[\cdot]$ and the fusers employed here satisfy the conditions of Corollary 1. Incidentally, they also satisfy the smoothness conditions of [8].

CONCLUSIONS

We presented an information fusion method that applies to physical systems wherein accurate measurements of physical parameters are not possible. We presented a method that combines various measurements and estimators to achieve performance at least as good

as the best set of measurements. We showed that a close approximation to the this optimal fuser can be computed such that with a high probability the solution performs at least as good as the best set of measurements, given large enough sample size. This work is an advance our earlier work [8] which is applicable to only Lipschitz laws and fusers. The study of projective fusers and meta-fusers [7] for the proposed formulation will be of future interest. It is also of interest to see if the boundedness of pseudo-dimension in Theorem 1 can be replaced by that of fat-shattering index [1], which would result in a weaker condition.

ACKNOWLEDGEMENTS

This research is sponsored by the Engineering Research Program of the Office of Basic Energy Sciences, U.S. Department of Energy, under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC, the Office of Naval Research under order No. N00014-96-F-0415, and by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory.

REFERENCES

- [1] M. Anthony and P. L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.
- [2] T. M. Apostol. *Mathematical Analysis*. Addison-Wesley Pub. Co., 1974.
- [3] S. R. Dallimore, T. Uchida, and T. S. Collett, editors. *Scientific Results from JAPEX/JNOC/GSC Mallik 2L-38 Gas Hydrate Research Well, Mackenzie Delta: Geological Survey of Canada Bulletin 544*. Geological Survey of Canada, Bulletin 544, 1999.
- [4] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100:78–150, 1992.
- [5] N. S. V. Rao. A fusion method that performs better than best sensor. In *First International Conference on Multisource-Multisensor Data Fusion*. 1998. 19-26.
- [6] N. S. V. Rao. Multiple sensor fusion under unknown distributions. *Journal of Franklin Institute*, 336(2):285–299, 1999.
- [7] N. S. V. Rao. On optimal projective fusers for function estimators. In *IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 1–6. 1999.
- [8] N. S. V. Rao, D. B. Reister, and J. Barhen. Fusion method for physical systems based on physical laws. In *Proceedings of 3rd International Conference on Information Fusion*, 2000.

INTEGRATION OF NANOSENSORS IN MICROSTRUCTURES

Li Shi, Guanghua Wu, Arun Majumdar
Department of Mechanical Engineering
University of California, Berkeley, CA 94720

ABSTRACT

This paper reports two examples where nanosensors have been integrated with microfabricated structures to provide new functional devices. The first is the integration of nanoscale temperature sensors on probe tips that have been used to study thermophysics of low-dimensional nanostructures such as carbon nanotubes. Thermal images at 50 nm spatial resolution are revealing the dissipation mechanisms in multiwall and single-wall carbon nanotubes. We report the direct observation of defect scattering on phonon transport in such nanotubes. The second example involves the generation of nanomechanical motion of a cantilever beam using specific biological reactions such as DNA hybridization and protein-ligand binding. We report here some new observations as well as the thermodynamic principles of how motion is created at nanoscales.

INTRODUCTION

A major engineering challenge that is common to several areas of nanotechnology is the integration of nanostructures into well-defined micro-patterns. For example, microelectronics and MEMS currently use optical lithography for generating such patterns. Optical lithography, however, is limited to a spatial resolution of about 100 nm whereas many of the exciting new nanoscale phenomena occur in the range of 1-10 nm. It is critical to bridge this 1-100 nm length scale gap such that nanostructures could be interfaced with MEMS and microelectronics systems and thereby interface with the macroscopic or "human" length scales. Although electron beam lithography can be used to fabricate structures in the 10-20 nm range, it is not truly scalable because it is performed sequentially. Lithography is often called the "top-down" approach to

making nanostructures since they are used to etch out a pattern. Self assembly of nanostructures, on the other hand, is a "bottom-up" approach where the process of integration or aggregation is thermodynamically driven. The combination of the top-down and bottom-up process will provide a means to integrate nanostructures in microstructures, thereby enabling engineering systems containing nanoscale components to be built.

In this paper, we report two examples where we have integrated nanoscale sensors onto microfabricated structures. By doing so, we have developed new functional devices that are now resulting in new scientific discoveries. The examples that we report here span both the physical and biological sciences, exemplifying the wide application of integrating systems across length scales.

SCANNING THERMAL MICROSCOPY OF NANOSTRUCTURES

In recent years, a number of low dimensional materials with length scale smaller than 100 nm have been developed. Carbon nanotubes, silicon nanowires, and semiconductor/metal nanocrystals are examples of such synthesized nanostructures. There is a great interest to experimentally investigate electron and phonon transport and heat dissipation phenomena in these materials. Such transport and dissipation phenomena are also important in ultra large-scale integration (ULSI) devices, whose minimum feature sizes are scaling down to sub-100 nm. Traditional measurement techniques, however, cannot resolve thermal features below 100 nm. For example, the spatial resolution of far-field optical thermometry techniques based on infrared and laser reflectance is diffraction limited to be on the order of wavelength, which is much larger than the length scale of sub-100 nm nanostructures currently of interest.

Scanning thermal microscopy (SThM) is capable of thermally investigating nanostructures and ULSI devices with spatial resolution in the sub-100 nm regime. The SThM maps surface temperature distribution by raster scanning a sharp temperature-sensing tip across the surface [1]. The tip is mounted on a micro cantilever beam such that a constant tip-sample contact force is maintained by the force feedback of an atomic force microscope (AFM). Tip-sample heat transfer changes the tip temperature, which is measured and used to determine the sample temperature.

The key element of SThM is the thermal probe. Figure 1 shows the schematic diagram of a SThM probe, which contains a thermocouple junction at the tip end. The tip-sample heat transfer mechanisms include solid-solid conduction through the contact, liquid conduction through a liquid film bridging the tip and sample, and air conduction. The thermal design of the cantilever probes is extremely important for SThM performance. The thermal resistance network in Fig. 1 suggests that for given ambient, T_a , and sample, T_s , temperatures, the tip temperature T_t , can be written as $T_t = T_s + (T_a - T_s)/(1 + \phi)$, where $\phi = R_c/R_{ts}$ is the ratio of the cantilever, R_c , and tip-sample, R_{ts} , thermal resistances. Hence, changes in sample temperature can be related to changes in the tip temperature as $\Delta T_t/\Delta T_s = \phi/(1 + \phi)$. This relation suggests that the accuracy and sensitivity of sample temperature measurement by the tip depends on ϕ , which must be large for

better SThM performance. Hence, the thermal design of the cantilever resistance, R_c , is extremely important. The spatial resolution, Δx of SThM measurements can be expressed as $\Delta x = \Delta T_n / (dT_s/dx)$ where ΔT_n is the noise in the temperature measurement and dT_s/dx is the measured temperature gradient. Because the tip and sample temperatures are related through ϕ , the spatial resolution can be expressed as

$$\Delta x = \frac{\Delta T_n}{(dT_s/dx)} \left(\frac{1 + \phi}{\phi} \right) \quad (1)$$

Equation (1) clearly suggests that small values of ϕ lead to poor spatial resolution of SThM.

Previous experiments have reported that $R_{ts} \approx 10^5$ K/W [2]. Therefore, the thermal design of the cantilever must require $R_c > 10^5$ K/W. In the past, SThM probes were made of a high thermal conductivity material such as metal or silicon, with no attention to paid to thermal design. This often led to probes with $R_c \ll R_{ts}$ which led to inaccuracies, loss of resolution, and artifacts. In addition, they were usually fabricated individually, making the process very time consuming and irreproducible. Recently, several groups have attempted to batch fabricate probes for scanning thermal microscopy. Two groups fabricated thermal probes using only optical lithography and wafer-stage processing steps. However, the probes were made of silicon, leading to the aforementioned inaccuracies and artifacts [3,4]. Silicon nitride thermal probes were also fabricated with a thermocouple junction defined at the tip by electron beam lithography [5]. The low throughput of electron beam lithography prohibits the process for being used for large-volume fabrication.

To address these issues, we have thermally designed and fully batch fabricated cantilever probes for SThM [6]. Based on heat transfer modeling, we chose SiO_2 and SiN_x as the tip and cantilever materials, respectively, in order to increase R_c as much as possible. It was shown by the modeling results that compared to silicon probes with similar geometry parameters, our current design could largely improve the thermal isolation of the sensor from ambient. In addition, Pt and Cr were chosen for the thermocouple materials for their high thermopower difference and low thermal conductivity. Finally, we optimized the geometrical parameters of the probe for increasing thermal resistance of the cantilever.

The thermal probes were fully batch-fabricated using wafer-stage process steps, with more than 300 probes fabricated on one single wafer. Figure 2 shows two micrographs of a finished cantilever probe with the tip containing a Pt-Cr thermocouple junction, with the Pt and Cr lines patterned along each cantilever arm. The tip region containing

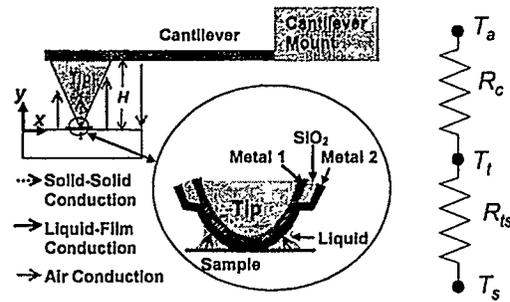


Fig. 1 Schematic diagram of a cantilever probe used for scanning thermal microscopy. The heat transfer mechanisms between the tip, cantilever, and the sample are also indicated, as is the thermal resistance network.

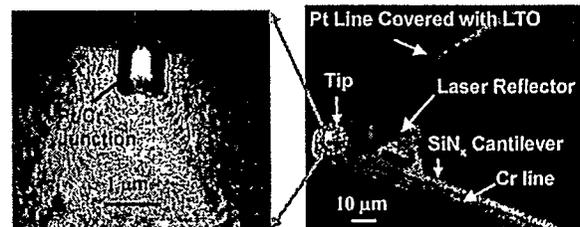


Fig. 2 Electron micrograph of the scanning thermal microscope probe. The picture on the right shows the cantilever beams and the tip whereas the one on the left shows the tip structure with a Pt-Cr thermocouple junction on the tip.

the overlap of Pt and Cr thin films was 0.5 μm tall and had a tip radius of about 50 nm. The height of this region controlled in the fabrication process to be in the range of 0.1-0.5 μm . Because the thermal resistance of these probes was very high ($\approx 10^6$ K/W), a low-power laser beam (≈ 1 mW) directed at the tip would increase its tip temperature by 80-90 $^\circ\text{C}$. Hence, to optically measure the cantilever deflections for atomic force microscopy, a thermally isolated laser reflector was fabricated.

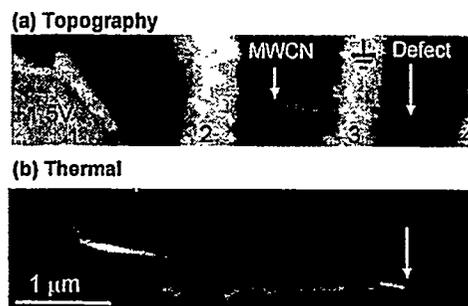


Fig. 3(a) Topographic and (b) thermal images of a multiwall carbon nanotube circuit under dc current of 27 mA for applied dc voltage of 1.5 V between contacts 1 and 3. Contacts 2 and 4 are floating.

The batch-fabricated thermal probes have been used for quantitative temperature measurement of VLSI via structures and for studying dissipation in multi-wall (MW) and single-wall (SW) carbon nanotube (CN) circuits. Figure 3a shows the AFM topography of one MWCN circuit that was imaged using the SThM probe. The sample contained a 14 nm diameter MWCN and four 30 nm thick gold contacts on an oxidized silicon wafer. Resistance measurement found that the tube was broken between contacts 3 and 4. The defect could be located in a high resolution AFM image. Figure 3b shows the thermal image of the sample obtained for a DC current of 27 μA flowing in the segment between contacts 1 and 3. The full width half maximum (FWHM) of the temperature profile across the CN was on the order of 50 nm, indicating the spatial resolution of the thermal imaging technique. This was approximately equal to the tip diameter. The temperature rise in the nanotube between contacts 1 and 3 can be clearly observed. To verify that the image was not due to topography-induced artifact, the thermal image was taken at different applied voltages and the thermal signals were found to increase with the voltage. Another possible artifact in the thermal image could be caused by current flow from the sample into the probe because of the difference in electrostatic potentials of the tip and sample. To rule out this possibility, we measured the tip-sample contact electric resistance when the tip was on top of the CN and the contacts, and found the resistance to be larger than the 1 G Ω measurement range of an Ohmmeter for the low contact force used in thermal imaging. The lack of electric contact was due to a chrome oxide layer formed at the tip during probe fabrication. This fact suggested that the thermal images were not due to electron flow at the tip-sample junction. In addition, the thermal probe was connected to a voltage amplifier with floating ground, such that no current would have flowed into the probe even if electric contact had been established at the tip-sample junction. We further confirmed the absence of electrostatic potential-induced artifacts in the thermal images by raising the electrostatic potential of the entire CN circuit without passing current through it. As we did so, no noticeable thermoelectric signal could be measured using the thermal probe scanned on the circuit. Therefore, the electrostatic potential did not introduce artifacts in the thermal images, and the thermal images were indeed due to phonon coupling instead of electron coupling at the junction. Furthermore, what the probe measured was the phonon temperature of the sample, which might or might not be at equilibrium with the electron temperature.

It is interesting to note that although no current flowed in the segment between contacts 3 and 4, the temperature of the left part of this segment was higher than that of the segment between contacts 2 and 3. This leads us to speculate that heat might be dissipated at the contacts and not in the bulk of the CN. Spatially uniform bulk dissipation would have led to parabolic temperature profiles instead. Note that switching the polarity of the applied voltage did not change the temperature distribution of the tube, indicating no thermoelectric effects at the contacts. The fact that the CN section between contacts 3 and 4 did not involve any electron transport and yet appeared hot suggests that phonon transport is very efficient in carbon nanotubes. In addition, it also demonstrates that SThM measured the phonon temperature and not electron temperature. The temperature between contacts 3 and 4 dropped rapidly at a point indicated by the arrow. We suspect that this was due to scattering by a defect in the CN, presumably the same defect that also blocked electron transport.

This study raises several questions, namely: Does the dissipation indeed take place at the contacts or along the nanotube? How efficient is phonon transport in the nanotube? What are the mean free paths for elastic and inelastic electron and phonon scattering? What are the roles of defects in nanotube transport phenomena? Some of these questions are currently being addressed experimentally using the SThM and will be presented later.

In conclusion, we have carefully designed and batch-fabricated probes for scanning thermal microscopy. The probes have been used to obtain thermal images of electrically heated carbon nanotube circuits. Our experimental results demonstrate that SThM measured the phonon temperature of the sample, and the tip-sample thermal coupling was dominated by heat conduction through a liquid bridge. Consequently, the spatial resolution of SThM was limited by the tip radius and was found to be 50 nm in this study. With this resolution, SThM offers the promising prospects of studying electron-phonon interaction and phonon transport in some low dimensional materials such as carbon nanotubes.

FROM BIOLOGY TO MOTION

Understanding the mechanisms of how biological reactions produce motion is fundamental to several physiological processes [7]. While most of the past effort has focused on studying single molecular motors [8], recent experiments [9,10] using microcantilever beams have reported the collective effect of multiple DNA hybridization and antigen-antibody reactions to produce nanomechanical motion. While this offers the promising prospects of interfacing molecular biology with micro and nanomechanical systems, an understanding of how this motion is produced has, however, remained elusive. Here we show that cantilever motion is created due to the interplay between changes in conformational entropy and intermolecular energetics induced by specific biomolecular reactions. The entropy contribution can be critical since it determines the direction of motion. Using these thermodynamic principles in conjunction with DNA hybridization experiments, we demonstrate that both the direction and the magnitude of cantilever motion can be controlled and optimized. This thermodynamic framework is also used to explain the nanomechanical motion created by protein-ligand binding.

Figure 5 illustrates the experiment that we used for studying nanomechanical motion created by multiple specific biomolecular reactions. The cantilevers used in our study were made of silicon nitride (SiN_x), with a gold coating on one surface. The experiment started by first placing a Au/ SiN_x cantilever in a fluid cell and then injecting a solution of sodium phosphate buffer (PB) at pH ~ 7.0 into the cell. The next step was to immobilize the probe molecule on the cantilever surface, which was followed by injection of a solution containing target molecules.

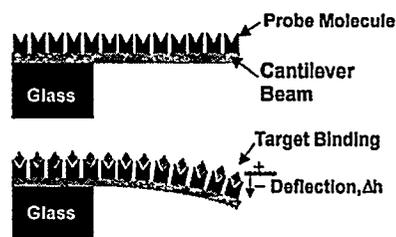


Fig. 5 Specific biomolecular interactions between target and probe molecules alters the intermolecular energy interactions within a self-assembled monolayer on one side of a cantilever beam. This produces sufficiently large force to bend the cantilever beam and generate motion. The origin of this nanomechanical motion lies in the interplay between changes in conformational entropy and the intermolecular energetics.

The cantilever motion was optically monitored at both the immobilization and probe-target binding steps. To form a self-assembled monolayer of probe ssDNA on the Au-coated cantilever surface, the ssDNA was modified with thiol groups attached to either the 5' end. Figure 6A shows the cantilever deflection as a function of time for a 50nt long probe ssDNA. Here, negative deflection represents the downward

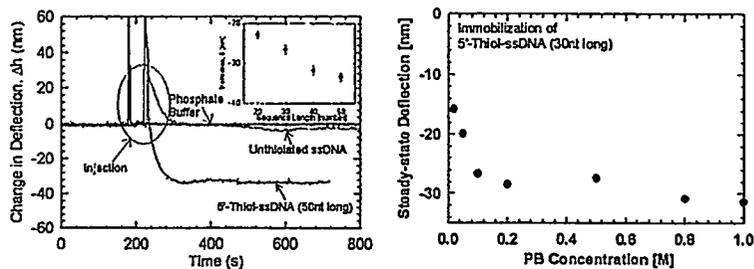


Figure 6 (A) Change in Au/ SiN_x cantilever deflection as a function of time for three different experiments: (i) exposure to 0.1 M phosphate buffer (PB); (ii) exposure to unthiolated probe ssDNA; (iii) exposure to probe ssDNA thiolated at the 5' end. Concentrations of unthiolated and single-end thiolated ssDNAs were all 50 ng/ μl or approximately 3.2 μM . Unthiolated ssDNA and pure PB solutions did not produce any significant deflection. The inset shows the steady-state cantilever deflection as a function of the length of the probe ssDNA thiolated. The results indicate that immobilization of probe ssDNA produces compressive stress bending the cantilever down. (B) Steady-state cantilever deflections caused by immobilization of ssDNA at different PB concentrations.

bending of the cantilever with the probe molecules on the top surface. Also shown are the deflection plots after the injection of unthiolated ssDNA and only PB. The inset in Fig. 6A shows the steady-state cantilever deflection as a function of the length of the probe ssDNA. Figure 6 B shows the cantilever deflection for immobilizing 30nt-long probe ssDNA as a function of PB concentration. It is clear from these experiments that regardless of the length of ssDNA or the ionic strength, the repulsive interactions between immobilized ssDNA created a compressive stress to bend the cantilever downwards.

After immobilizing the probe ssDNA, the complementary target ssDNA was injected into the fluid cell at the same PB concentration that was used to immobilize the probe ssDNA. Figure 7A shows the deflection plots for the hybridization reactions where the probe ssDNA was 20nt long and the target ssDNA were of four different lengths (20nt, 15nt, 10nt and 9nt) and distally complementary. The nanomechanical signal was sufficiently sensitive to detect single nucleotide length differences. The observation that the cantilever bent upwards in all cases suggests that hybridization relieved the compressive stress created during immobilization of thiolated probe ssDNA. To confirm that the signals were due to hybridization, a solution of a non-complementary target ssDNA was used and was found to produce no deflection signal. Figure 7B plots the steady-state deflection signal for the hybridization reaction under different PB

concentrations. An optimum PB concentration of 0.2-0.4 M was seen to produce the maximum deflection.

The fact that the cantilever deflections for both the immobilization and hybridization steps were influenced by the PB concentration suggests that electrostatic repulsive forces between neighbouring DNA molecules must have produced the compressive stress that bent the cantilever down. Because each nucleotide carries a negative charge due to the presence of a phosphate group, one would expect that hybridization would cause even more repulsion due to the presence of additional negative charge. However, the data in Fig. 7 clearly indicates that regardless of the PB concentration in the range of 0.05-1 M, hybridization always relieved the stress and produced upward cantilever motion. Therefore, electrostatic or steric repulsion alone cannot explain the behavior.

It is worth noting that the persistence length of ssDNA is about 0.75 nm [11] whereas that of double-stranded DNA (dsDNA) is about 50-80 nm [12]. We propose that the conformational entropy of the ssDNA provides the driving force to bend the cantilever down. Upon hybridization, the conformational entropy contribution is significantly reduced which then relieves the compressive stress resulting in upwards cantilever motion (see Fig. 8).

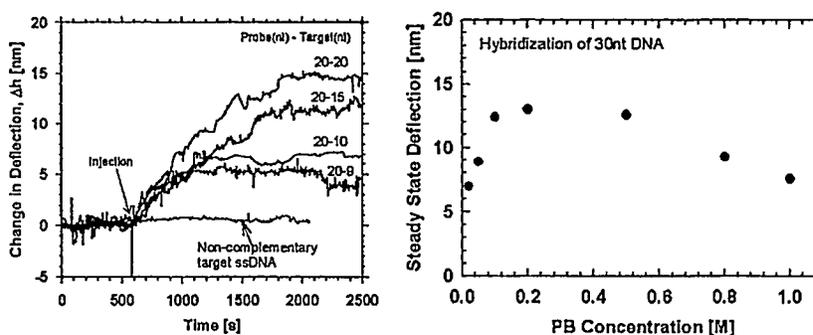


Figure 7 (A) Changes in Au-Si cantilever deflection due to hybridization of a probe ssDNA (50 ng/μl or 8 μM concentration) in the distal end with complementary target ssDNA of different lengths — 20nt, 15nt, 10nt, and 9nt (40 ng/μl or 3-6 μM concentration). Also shown is the absence of cantilever deflection for a non-complementary target ssDNA. The data clearly suggests that differences in nanomechanical motion due to one nucleotide difference in length can be observed. (B) Steady-state changes in cantilever deflection for hybridization of 30-nt long ssDNA at different PB concentrations. Note that immobilization of probe ssDNA were at the same PB concentration as the hybridization reaction.

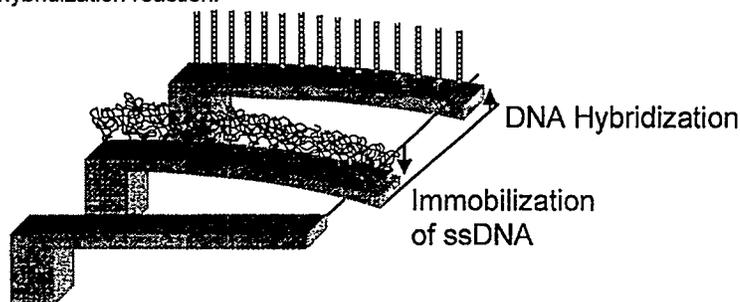


Figure 8 Schematic diagram illustrating the mechanism of motion generation due to DNA immobilization and hybridization. Immobilization of ssDNA on the top surface bends the cantilever down. The persistence length of ssDNA is 7.5 Å which leads to higher conformational entropy resulting in compressive stress. Hybridization increases the persistence length to about 50 nm, which significantly reduces the conformational entropic driving force, thereby relieving the compressive stress and producing an upward cantilever motion.

ACKNOWLEDGEMENT

This work was performed under the support of the Engineering Program of DOE Basic Energy Sciences. Our many thanks to Paul McEuen, Sergei Plyasunov and Adrian Bachtold of UCB-Physics for the carbon nanotube experiments, and to Thomas Thundat, Karolyn Hansen, and

Haifeng Ji of Oak Ridge National Lab, Ram Datar and Richard Cote of USC-Pathology, and Arup Chakraborty of UCB-Chemical Engineering, for their immense help and support for the DNA experiments.

REFERENCES

1. A. MAJUMDAR, "Scanning Thermal Microscopy," *Annu. Rev. Mater. Sci.* **29**, 505 (1999).
2. K. LUO, Z. SHI, J. VARESI, A. MAJUMDAR, "Sensor Nanofabrication, Performance, and Conduction Mechanisms in Scanning Thermal Microscopy," *J. Vac.Sci. Tech. B* **15**, 349 (1997).
3. T. LEINHOS, M. STOPKA, E. OESTERSCHULZE, "Micromachined fabrication of Si cantilevers with Schottky diodes integrated in the tip," *Appl. Phys. A* **66**, 65 (1998).
4. Y. ZHANG, Y. ZHANG, J. BLASER, T. S. SRIRAM, A. ENVER, R. B. MARCUS, "A thermal microprobe fabricated with wafer-stage processing," *Rev. Sci. Inst.* **69**, 2081 (1998).
5. G. MILLS, H. ZHOU, A. MIDHA, L. DONALDSON, J. M. R. WEAVER, "Scanning thermal microscopy using batch fabricated thermocouple probes," *Appl. Phys. Lett.* **72**, 2900 (1998).
6. L. SHI, O. KWON, G. WU, A. MAJUMDAR, ASME Int. Mech. Eng. Congress & Exposition, MEMS **1**, 93 (1999); The details of the modeling, design, and fabrication process of the thermal probes will be published in a separate paper.
7. R. D. VALE & R. A. MILLIGAN, "The way things move: Looking under the hood of molecular motor proteins," *Science* **288**, 88 (2000).
8. D. KELLER & C. BUSTAMANTE, "The mechanochemistry of molecular motors," *Biophys. J.* **78**, 541 (2000).
9. R. RAITERI, G. NELLES, H.-J. BUTT, W. KNOLL, P. SKLADAL, "Sensing of biological substances based on the bending of microfabricated cantilevers," *Sensors and Actuators B* **61**, 213 (1999).
10. J. FRITZ, et al., "Translating biomolecular recognition into nanomechanics," *Science* **288**, 316 (2000).
11. S. B. SMITH, Y. CUI, C. BUSTAMANTE, "Overstretching B-DNA: The elastic response of individual double-stranded and single-stranded DNA molecules," *Science* **271**, 795 (1996).
12. C. G. BAUMANN, S. B. SMITH, V. A. BLOOMFIELD, C. BUSTAMANTE, "Ionic effects on the elasticity of single DNA molecules," *Proc. Natl. Acad. Sci.* **94**, 6185 (1997).