Research Statement

Roger Mailler

A great deal of my energy is spent on research. I love research and, like many faculty members, am driven by the excitement derived from making new discoveries. I take a very hands-on approach to conducting research and actively develop software and hardware, perform experiments, and write papers alongside my students. Acting as a PI, I have worked very hard to grow and promote my lab (a wet lab and computing lab) by obtaining funding ($2.7M since 2008), authoring papers (7 journal and 21 conference papers since 2008), and establishing myself as a top quality researcher (NSF Career award, 3 best paper awards, and 2 student research awards).

My research is guided by big questions and I find that addressing these questions requires persistence and the ability to see relationships between areas that seem quite unrelated. Historically, I have conducted research on **Distributed Problem Solving**. After finishing my undergraduate degree at Stony Brook University, I had the great pleasure to work with Prof. Victor Lesser. Victor, who many refer to as the grandfather of Multi-agent Systems, is also driven by big questions and as a side note in a conversation with me posed this question:

"When is it appropriate to use a distributed problem solving technique over a centralized one?"

This is a seemingly simple question and I believe that he was looking for a simple response, but the question intrigued me because so many factors go into making that decision. I have worked on that question for nearly 15 years and just within the past 3 years have found a solution that utilizes thermodynamic theory.

Interestingly, multi-agent systems was not the subject that attracted me to research. I have always been intrigued by the origins of intelligence. As an undergraduate, I took a number of courses like biological and cognitive psychology, as a graduate student, I took a course on neurobiology, and have studied neuroscience on my own for many years. A career, however, often drives our decisions and it wasn't until I joined the University of Tulsa in 2008 that I felt I could safely transition into the field of **Computational Neuroscience**. Currently, nearly all of my efforts go into this research area.

Distributed Problem Solving

One of my major research areas is the empirical study of distributed problem solving with a split focus on theoretical and application driven problems. Prior to joining TU in 2008, I largely worked on creating protocols for solving Distributed Constraint Optimization Problems (DCOP). The most influential of these is the Optimal Asynchronous Partial Overlay (OptAPO) algorithm [1]. OptAPO works by performing controlled, partial centralization in order to discover and take advantage of a problem's natural decomposability. This protocol represented an entirely new method for solving distributed problems and was the first known complete method based on partial centralization. According to Google scholar the original OptAPO paper has been cited 404 times since 2004.

Since joining TU, I have largely concentrated on improving the practical applicability of distributed problem solving methods by applying them to real-world problems and creating theoretical tools to evaluate their performance.

Traffic Mitigation

As part of her thesis work, my Ph.D. student (Melanie Smith) developed distributed techniques for cooperative traffic re-routing. The premise was that with the proliferation of GPS technologies, we could use DCOP techniques to allow the GPS units within cars to coordinate with one another when a significant traffic event occurs. This would have a marked advantage over the currently used methods, which involve either waiting or diverting to same secondary route that all of the other cars are utilizing. Using real-world data we obtained from Oakland County, Michigan (a suburb just north of Detroit), we developed a simulator and evaluated a number of protocols that used various amounts of information and coordination. In the end, we determined that the technology is viable assuming there are enough cars equipped with GPS devices that can coordinate [2]. Although Melanie's work made important strides, I still view this as a fruitful area for further investigation.

Tracking Objects in Low-earth Orbit

For the past six years, my lab has worked on a pressing real-world problem: tracking objects in Low-earth orbit (LEO). Currently, there are tens of thousands of objects circling the earth. These include objects like the International Space Station, the Hubble telescope, communication and weather satellites, and huge amounts of garbage. The US Air Force is tasked with tracking all of these objects and currently uses 29 tracking stations positioned all over the planet. Sadly, the number of space objects is growing at an exponential pace and our capacity to monitor them has already been exceeded.

The USAF is considering the use of low-cost, remotely operated optical telescopes to help deal with the rapid increase in tracking tasks. The problem is that these telescopes need to be placed where light pollution is minimal. These areas are often remote creating significant communication difficulties. In addition, optical telescopes are notoriously unreliable due to mechanical issues and weather and can only be used for a fraction of the day during sunrise and sunset. These limitations create a complex, dynamic resource allocation problem where communications is slow and must occur over large distances because the satellites travel so quickly. Our approach to solving this problem has been to model it as a max-flow problem and develop distributed max-flow techniques that use geographically directed messages when a task reallocation must occur. We have found these techniques to be quite successful compared to a centralized technique, especially as communication delays begin to limit the flow of information [3, 4].

Although the USAF is no longer funding this work, I continue to work on this important problem as part of my NSF Career award. Currently, we have three prototype sensor platforms that we use to simulate remote optical telescopes and collect cloud cover data. These sensors provides us with real world data for our simulations and allow us to test our protocols under realistic conditions.

ALLOCATING INTELLIGENCE, SURVEILLANCE, AND RECONNAISSANCE (ISR) ASSETS

During the past four years, I have worked on two grants from the US Air Force that address the allocation of resources for intelligence gathering. As the pace of world events has increased, the need for timely intelligence has subsequently increased as well. Unfortunately, budgets have been decreasing and with that there has been an increasing desire to do more with existing surveillance assets. However, these assets are spread across the planet and are controlled by numerous organizations, so finding the right assets and getting them assigned to a task, while adhering to the policies and procedures of different organizations, is a time consuming process. Working with Dr. Rose Gamble, during our first grant, we developed policy-aware distributed resource allocation methods to solve this difficult, man-in-the-loop coordination problem.

Our approach is unique because it allows a policy engine to guide the coordination process. Unlike other techniques that have imbedded, hand-modeled, fixed utility functions, our system allows operators to dynamically modify these functions by adjusting usage policies for their assets [5]. So just like the telescope allocation problem this creates a complex dynamic, distributed resource allocation problem.

As a follow on to that grant, Dr. Gamble and I are currently working on dynamic mission planning for UAVs in adversarial environments. This project fuses together machine learning, path planning [6], and coordination to allow air assets to gather intelligence in unknown environments where the enemy can utilize both kinetic and electronic responses to thwart them. Our approach utilizes layered learning that attempts to model individual enemy assets, identify correlations in their activity, and then determine causal relationships with our actions. The results of the learning are a set of spatio-temporal constraints that are used for individual path planning and coordinated actions between multiple assets.

ANSWERING THE BIG QUESTION: A THEORY FOR DYNAMIC PROBLEM SOLVING

One of the greatest challenges facing the distributed problem solving community is the lack of a compelling theory that can determine when a distributed approach is warranted. It is my contention that distributed techniques are necessary when information is distributed, and communications combined with computation makes a centralized approach unresponsive to changes in the problem. However, the interaction between an algorithm and an environment that is in constant flux is complex and until recently the only way to compare techniques was to perform exhaustive testing.

In 2014, I discovered a mapping between Dynamic Constraint Satisfaction Problems (DynCSPs) and thermodynamic systems [7]. By considering the variables in the problem as particles and the constraints between the variables as interactions, I proved that DynCSPs follow Newton's three laws of thermodynamics. Alone, this is an impressive accomplishment, but I was also able to show that if a distributed protocol is operating to solve these every-changing problems that the solution quality forms a *stable* and *predictable* equilibrium. Using this mapping, I have, for the first time ever, come up with a method for evaluating the impact that dynamics has on a problem and characterizing the performance of an algorithm that is designed to solve it.

Since that time, my group has extended these results to include Dynamic, Distributed Optimization Problems (DynDCOPs) [8], has shown that the performance of some protocols can be derived without empirical studies [9], and, most recently, that we can directly characterize the impact of communication delays and information stagnancy [10]. By developing these analysis techniques, we believe that we can classify protocols that operate on dynamic problems in a similar manner to how we currently classify algorithms that operate on static ones. These tools would allow us to not only determine when a distributed protocol is warranted, but also which distributed technique will yield the best overall performance.

Computational Neuroscience

As mentioned previously, my original interest in computer science research stems from a desire to understand the basis for intelligence. Intelligence is hard to define. If you asked two different people to define intelligence, you are likely to get two completely different answers. However, using broad strokes, intelligence could be defined as the ability of an organism to rapidly and actively adapt to its environment in order to survive.

Just before leaving SRI International, I decided to revisit this interest, which led me to the second big question that guides my research:

"Can we build a biologically accurate simulation of the entire nervous system of a living organism?"

The scientific impact of achieving this goal is profound as it would help biologist understand how a nervous system produces behavior and may usher in a new era in the field of neural networks.

The organism I chose for this project is the nematode *Caenorhabditis. elegans*. I chose this organism because it is a model organism in the fields of neurobiology, developmental biology, and genetics and as such there is an extensive amount of literature available.

C. elegans has a simple nervous system with only 302 neurons. In 1986 the nervous systems was nearly completely mapped [11] and since then many of the roles of the neurons have been discovered. It appeared that it should be reasonable to do a complete neuron-level simulation. However, just as Victor thought he was asking a simple question with a simple answer, this question is far from easy to answer.

The first step in approaching this problem was to develop a simulator that accurately models the physical environment and body of *C. elegans* to allow the nervous system to be embodied. This simulator, which is called ALIVE, uses a high-fidelity physics engine to accurately reproduce the forces involved in the movement of this organism. It faithfully reproduces the weight, size, shape, and musculature of the organism along with the surface tension and viscosity that it experiences. The simulation was cross-validated against live worms using video recordings that were converted to quantifiable data using in house image processing software [12]. This resulted in a best paper award at BIOSYSCOM 2010 [13].

C. ELEGANS LOCOMOTION

After building the physics simulation, we began to explore the properties of the neurons it uses to produce locomotion. Although quite a bit is known about the neurons, a number of unresolved questions still exist about how the worm produces its characteristic sinusoidal pattern of movement.

Working from the outside in, we began by developing a conductance-based model of the muscle cells. *C. elegans* muscle cells are quite different from mammalian cells. The major difference is the lack of voltage gated Na^+ channels that are used by mammalian cells to fire action potentials. Instead, *C. elegans* uses an L-type voltage gated, Ca^{2+} inactivated Ca^{2+} channel, called EGL-19. In addition, the muscle cells have Kv1 voltage gated K^+ channels, called SHK-1, and BK-type Ca^{2+} activated K^+ channels called SLO-2. It was long believed that the lack of sodium channels meant that the muscle cells used graded potentials and computational models have been developed based on this belief.

Recent findings by Liu et al. [14] show that these cells produce action potentials. Using the Hodgkin and Huxley (HH) model as a basis, we developed a conductance based model of the muscle cells using published electrophysiological data. Our model produced robust action potentials that verified their findings [15]. However, our model did not exactly fit the data despite making a number of adjustments. Working with the Wang lab at the University of Connecticut Health Center we found secondary dynamics associated with the inactivation of the SLO-2 channels. This finding led to the discovery that SLO-2 channels can be inhibited by the subunit BKIP-1 [16]. Additional, our work with the Wang lab has led to the discovery of rectifying gap junctions between the command interneurons and motor neurons in *C. elegans* [17].

We have also developed a model of the locomotion circuit using a leaky integrate-and-fire (IF) current based model. IF models are easy to work with because each neuron can be represented using a single Ordinary Differential Equation (ODE) versus the four or more needed for an HH based neuron. However, IF models assume that the duration and magnitude of the AP is not information bearing. In *C. elegans*, which have muscle cells with long APs (50 ms or longer) and neurons that have graded or biphasic potentials, this assumption does not holds.

To account for this, we developed a new leaky IF model that uses state information to control the timing and target potential for the cell. For example, a muscle cell can be in one of three states; resting, rising, or recovering. In the resting phase, the leaky component of the IF model attempts to keep the cell at the resting potential of

-26mV using a fairly small leakage conductance of 22 S/F. However, once external current causes the membrane potential to reach the threshold voltage of -18mV, EGL-19 channels begin to open causing the cell to change to the rising state. During rising, the muscle cell drives toward the peak of the AP, which occurs at +22mV. This occurs rapidly as the conductance goes as high as 290 S/F.

This new model, which we call the state-based IF (SBIF) model, is similar to the HH model because it can faithfully reproduce the timing and magnitudes of action, graded, and biphasic potentials seen in the neurons. However, like the IF model, it only requires a single ODE to represent a cell, making it computationally efficient enough to run real time simulations.

We are currently working on a publication for this new model, but verifying the model has been difficult even though we have an accurate physics-based simulator. Many researchers have assumed that the tension of the muscle cells can be estimated based on the posture of the worm's body. However, posture is the result of a number of factors, so a worm could be in a bent posture, but have no muscle activity at all. This indicates that even if we reproduce the movement of the worm, the exact muscle activation pattern may not have been replicated.

Recent advances in optogenetics now make it possible to directly record the activity of the muscle cells. Currently, we are using a worm that expresses the Genetically Encoded Calcium Indicator (GECI) GCaMP2 in its muscle cells (Pmyo-3::GCamp2) to record muscle activity while it freely moves on an agar surface. To do this, we use a microscope that we recently acquired using a Defense University Research Instrumentation Program (DURIP) grant. This scope is quite unique as it allows us to capture 4 megapixel images at up to 100 images per second. Additionally, using a sequencer board that I developed, we can rapidly change the lighting conditions of the image from brightfield, which allows us to track to worm and measure its body posture, to the excitation wavelength for GCaMP2, which allows us to measure the intensity of the fluorescence in the muscle cells.

The amount of data generated from these experiments in considerable. The video alone weighs in at about 14 GB per minute and once processed creates one record per image with 120 data fields. Handling this amount of data has been challenging and has necessitated the creation of a custom video compression format along with new software for controlling the microscope, capturing, compressing, editing, and analyzing the video, and finally, processing the resulting data.

Our preliminary analysis of the data we have collected has revealed some interesting results that run contrary to conventional wisdom. First, it has been commonly believed that worms propagate waves down the entire length of their body. Largely this is not true. Second, waves travel down the dorsal and ventral sides simultaneously instead of down only one side at a time. Third, as we suspected, body posture and muscle activation are not entirely correlated.

These tools represent enabling technology for our lab because we can we use them to verify our model and, in combination with other types of genetic manipulation and laser ablation, to further unravel the intricacies of the nervous system.

ANSWERING THE BIG QUESTION: C. ELEGANS PLANNING

In fact, my most recent proposal to AFOSR builds on these models and technologies to move beyond locomotion into planning. Like all living organisms, the *C. elegans* nervous system receives myriad information from both internal and external sources. This information is integrated by the nervous system and eventually drives movement. With our new hardware and software, we are now able to directly measure the activity of neurons in the anterior ganglion, which is referred to as the nerve ring. Additionally, we can construct new worm strains that allow us to directly stimulate or inhibit neurons in this region using light. This will allow my lab to uncover the roles and relationships between neurons that up till this point could only be speculated about. The results of these experiments will be directly added to the locomotion model.

GRANTS AND FUNDING

To fund my research lab, I am very active in writing grant proposals. During my 8 years at TU, I have authored 25 proposals generating over $2.7M in research funding. I currently have two research proposals under review. The first was submitted to AFOSR to continue my work on the nervous system of *C. elegans* for $725,581. The second is as a member of the AIS team for the DARPA Agile Computing Technology (ACT2) program. This program is an IDIQ worth approximately $960M. Below is a table listing the proposals I have won since joining TU.

Organization	**Title**	**Date**	**Amount**	**PI or Co-PI**
DOD/SRI	MABLE: Modular Architecture for Bootstrap Learning Experiments	Aug 2008	$294,041	PI
TU	Mechanisms of Locomotion in *Caenorhabditis elegans*	Oct 2008	$1,000	PI

NIH INBRE	*C. elegans* Mutant Construction for Locomotion Circuit Functional Analysis	Feb 2010	$10,000	PI
AFRL	Distributed Constraint Optimization for Telescope Scheduling	Jul 2010	$121,694	PI
TU	3D Reconstruction of *Caenorhabditis elegans* using Scanning Electron Microscopy	Oct 2010	$1,000	PI
AFRL	Dynamic, Policy-Aware Distributed Coordination	Nov 2012	$490,741	PI
NSF	Career: Problem Solving in Dynamic, Distributed Environments	Jul 2013	$452,137	PI
AFOSR	Circuit Models for Robust, Adaptive Neural Control	May 2015	$582,726	PI
AFOSR	Acquisition of a Microscope for Studying the Locomotion Circuitry of Caenorhabditis elegans.	Aug 2015	$232,163	PI
AFRL	Information Driven, Adaptive Distributed Planning	Sep 2016	$519,065	PI

Works Cited

[1] R. Mailler and V. Lesser, "Solving Distributed Constraint Optimization Problems Using Cooperative Mediation," in *Proceeding of the Third International Joint Conference on Autonomous Agent and MultiAgent Systems (AAMAS 2004)*, 2004.

[2] M. Smith and R. Mailler, "The Effect of Congestion on Coordinated Routing," in *Proceedings of the 14th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA-2011)*, 2011.

[3] F. Hafizoglu and R. Mailler, "Telecope Management for Satellite Tracking: A Decentralized Approach," in *Proceeding of the Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, 2013.

[4] F. Hafizolglu and R. Mailler, "Decentralized Telescope Management for Satellite Tracking," in *Proceedings of the 2013 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2013)*, 2013.

[5] A. Marshall, S. Alqahtani, A. Ridgway, C. Walter, R. Gamble and R. Mailler, "Combining coordination with usage policies to improve mission scheduling resilience," in *Resilience Week*, 2015.

[6] S. Alqahtani, I. Riley, S. Taylor, R. Gamble and R. Mailler, "MTL Robustness for Path Planning with A*," in *Proceedings of the Seventeenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, 2018 (Under Review).

[7] R. Mailler and H. Zheng, "A New Analysis Method for Dynamic, Distributed Constraint Satisfaction," in *Proceeding of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, 2014.

[8] R. Mailler, H. Zheng and A. Ridway, "Dynamic, distributed constraint solving and thermodynamic theory," *Journal of Autonomous Agents and Multi-Agent Systems,* pp. 1-30, 2017.

[9] A. Ridgway and R. Mailler, "Dynamic Theoretical Analysis of the Distributed Stochastic and Distributed Breakout Algorithms," in *Proceedings of the Fourtheenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, 2015.

[10] S. Samadidana and R. Mailler, "Analyzing the Effect of Information Stagnancy on the Distributed Stochastic Algorithm," in *Proceedings of the Seventeenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, 2018 (Under Review).

[11] J. G. White, E. SouthGate, J. N. Thomas and S. Brenner, "The Structure of the Nervous System of the Nematode Caenorhabditis elegans," *Philosphical Transations of the Royal Society of London,* Vols. Series B, Biological Sciences, pp. 1-340, 1986.

[12] J. Graves and R. Mailler, "Quantifying the accuracy of C. elegans image analysis," in *The Third International Conference on Bioinformatics, Biocomputational Systems and Biotechnologies*, 2011.

[13] R. Mailler, J. Graves, N. Willy and T. Sarratt, "A Biologically Accurate Simulation of the Locomotion of Caenorhabditis elegans," *International Jounral on Advances on Life Sciences,* vol. 2, no. 3 & 4, 2011.

[14] P. Liu, Q. Ge, B. Chen, L. Salkoff, M. Kotlikoff and Z.-W. Wang, "Genetic dissection of ion currents underlying all-or-none action potentials in C. elegans body-wall muscle cells," *Journal of Physiology,* vol. 589.1, pp. 101-107, 2011.

[15] C. Johnson and R. Mailler, "Modelling Action Potentials of the Body Wall Muscles in C. elegans: A Biologically Founded Computational Approach," in *Proceedings of the 7th International Conference on Bioinformatics and Computational Biology*, 2015.

[16] L.-G. Niu, P. Liu , Y. Shui, R. Mailler and Z.-W. Wang, "BKIP-1, an auxiliary subunit critical to SLO-1 function, inhibits SLO-2 potassium channel in vivo," *Scientific Reports,* 2017 (Accepted).

[17] P. Liu, B. Chen, R. Mailler and Z.-W. Wang, "Antidromic-rectifying gap junctions amplify chemical transmission at functionally mixed electrical-chemical synapses," *Nature Communications,* 2017.