

List of Publications with Abstracts

Orion Sky Lawlor
Department of Computer Science
University of Alaska at Fairbanks
olawlor@acm.org

- 2022 “Fluxed Melting for Rapid Regolith Digestion” [1]
Describes a way to pre-treat regolith, by fluxing it with sodium and melting it to glass, that allows it to be almost completely dissolved in hot hydrochloric acid. This experimentally validated regolith digestion approach allows the included metals and silica to be extracted using a solar powered closed-loop process.
- 2021 “Virtual reality vehicle simulator phase 1” [2]
An offroad vehicle simulator software package, designed for making people better at driving them. The vehicle simulator software (https://github.com/olawlor/AK_ATV/) is built in the Unity engine and is open source. This simulation approach has also proven useful for designing and sim-testing autonomous construction robots.
- 2021 “A Dual-Pin Tool Coupler for Robotic Excavation” [3]
Describes a coupler that lets robots pick up tools like buckets, hammers, logistics containers, and grinders. This coupler is lightweight (under 400 grams on the tool side), robust (able to sustain kilonewton forces), autonomous operation friendly, and still works covered in dust. Also see the poster version (<https://media1.virtualchair.net/spaces/lpic-2021-fall-meeting/posters/P019-poster.pdf>).
- 2020 “Forensic Analysis of Binary Structures of Video Files” [4]
This paper provides comprehensive details of a binary file forensic analysis technique for different media file containers, mostly focused on AVI and MP4/MOV container format. We also provide a considerable number of details to identify a forgery among video files. By analyzing the binary data structures and metadata, we can detect the use of editing tools, verify the purported source of a video file, and identify the true acquisition device model.
- 2020 “Case Study: An Automated Theft on Mars” [5]
Explores the implications of autonomous robotics for security on a Mars habitat. A winning entry in the Air Force sponsored MGMW-ERX Space Case Study Prize Challenge, for case studies in the space domain.
- 2017 “Between the Social and the Technical: Negotiation of Human-Centered Robotics Design in a Middle School Classroom” [6]

This paper presents a secondary school human-centered robotics (HCR) learning experience and the ways in which it supported students' orientation to technical and social aspects of Science, Technology, Engineering, and Mathematics (STEM).

- 2015 “AERO-Beam: An Open-architecture Test-Bed for Research and Education in Cyber-Physical Systems” [7]

As cyber-physical systems (CPS) increasingly become a part of daily life, a well-prepared workforce is needed to build and maintain, calling for new approaches and test-beds to introduce CPS concepts into classrooms at all levels. The literature has many examples of CPS relevant platforms, which are used successfully in research; however, especially the commercial ones are mostly closed-architecture and do not provide an in-depth engagement for the students in all aspects of the hardware and software. This paper introduces a novel platform, known as the AERO-Beam, which conveys many basics of CPS, while also introducing the dynamics and control fundamentals of quadcopters. Quadcopters are widely used in CPS projects, yet their limited batteries, rapid control response, and tendency to crash makes them difficult to use in CPS education. The AERO-Beam addresses the need in many ways, as demonstrated through a multi-level classroom implementation.

- 2014 “Exploiting Spatial Redundancy with Adaptive Pyramidal Rendering” [8]

Just as image data compression is designed to save space while preserving the essence of an image, we present an adaptive pyramidal rendering scheme designed to save rendering time while maintaining acceptable image qual-

ity. Our coarse-to-fine scheme predicts when and where it is safe to take less than one sample per output pixel, and exploits spatial redundancy to predict pixel colors in the resulting gaps, both of which can be performed at framerate in realtime on a modern GPU. As a lossy compression method, we present experimental data on the rendering time versus image quality tradeoff for several example renderers.

- 2014 “Algorithms: Rendering” [9]
- 2013 “Search and Rescue Robots for Integrated Research and Education in Cyber-Physical Systems” [10]

This paper presents a novel research-education integration approach that uses the interdisciplinary theme of Cyber-Physical Systems (CPS) to further advance teaching and learning in such systems and related areas, in electrical and computer engineering (ECE) and computer science (CS). The primary targeted learners in this emerging area are graduate and undergraduate students of ECE and CS, as well as teachers and students of STEM in high schools, who work together over a common CPS design project, with the graduate and undergraduate students teaching and mentoring the high school students. In the first year of this GK12 project, a networked ground and aerial vehicle based Search and Rescue Robotics system is selected as the CPS design project with an associated multi-school competition organized via the use of teleoperation over the internet. The first year results demonstrated great promise not only in terms of improved learning among students and teachers of multiple disciplines and levels, but also provided valuable hardware and software platforms for on-line experimentation and education in CPS, robotics, graphics and

animation, UGV and UAV control, communication and networking.

- 2013 “In-memory Data Compression for Sparse Matrices” [11]

We present a high performance in-memory loss-less data compression scheme designed to save both memory storage and bandwidth for general sparse matrices. Because the storage hierarchy is increasingly becoming the limiting factor in overall delivered machine performance, this type of data structure compression will become increasingly important. Compared to conventional compressed sparse row (CSR) using 32-bit column indices, compressed column indices (CCI) can be over 90 percent smaller, yet still be decompressed at tens of gigabytes per second. We present time and space savings for 20 standard sparse matrices, on multicore CPUs and modern GPUs.

- 2013 “Crack Propagation Analysis with Automatic Load Balancing” [12]

This book chapter summarizes the Charm++ approach for finite element computations on distributed memory parallel computers. This includes automated load balancing, the ParFUM framework, and a case study involving cohesive-volumetric finite element crack simulation.

- 2013 “Remote Visualization of Cosmological Data Using Salsa” [13]

The experimental cosmology application ChaNGa simulates particle datasets consisting of tens of millions to billions of particles. Salsa, the visualization tool we describe in this chapter, is the simulated observatory intended to help extract useful information from the simulation results. Salsa is designed to visualize

these large multi-gigabyte multidimensional datasets by using parallel GPU rendering, and a technique known as volume impostors, to provide interactive analysis and visualization over a slow network link.

- 2011 “NetRun: Student Code in a Web Browser (Poster)” [14]

This poster presents a web integrated development environment called NetRun, which students can use to write, run, and debug code from inside any web browser. NetRun is a web server application that accepts requests from a student’s web browser, forwards the request to a variety of backend machines to be compiled and executed, and sends the output back to the user. Currently, NetRun supports code written in a variety of parallel languages, including pthreads, sockets, OpenMP, MPI, GLSL, NVIDIA CUDA, and our cudaMPI. Netrun can then run that code on a variety of machines, from an old Pentium II up to a modern cluster of graphics cards.

- 2011 “GPU-Accelerated Rendering of Unbounded Nonlinear Iterated Function System Fixed Points” [15]

Nonlinear functions, including nonlinear iterated function systems, have interesting fixed points. We present a non-Lipschitz theoretical approach to nonlinear function system fixed points which generalizes to non-contractive functions, compare several methods for evaluating such fixed points on modern graphics hardware, and present a nonlinear generalization of Barnsley’s Deterministic Iteration Algorithm. Unlike the many existing randomized rendering algorithms, this deterministic method avoids noncoherent branching and memory access, and takes advantage of programmable tex-

ture mapping hardware. Together with the performance potential of modern graphics hardware, this allows us to animate high quality and high definition fixed points in real time.

- 2011 “Embedding OpenCL in C++ for Expressive GPU Programming” [16]

We present a high performance GPU programming language, based on OpenCL, that is embedded in C++. Our embedding provides shared data structures, typesafe kernel invocation, and the ability to more naturally interleave CPU and GPU functions, similar to CUDA but with the portability of OpenCL. For expressivity, our language provides the FILL abstraction that releases control over data writes to the runtime system, which both improves expressivity and eliminates the chance of memory race conditions. We benchmark the new language on NVIDIA and AMD hardware for several small examples.

- 2010 “Interactive Volume Rendering Aurora on the GPU” [17]

We present a combination of techniques to render the aurora borealis in real time on a modern graphics processing unit (GPU). Unlike the general 3D volume rendering problem, an auroral display is emissive and can be factored into a height-dependent energy deposition function, and a 2D electron flux map. We also present a GPU-friendly atmosphere model, which includes an integrable analytic approximation of the atmosphere’s density along a ray. Together, these techniques enable a modern consumer graphics card to realistically render the aurora at 20–80fps, from any point of view either inside or outside the atmosphere.

- 2009 “Message Passing for GPGPU Clusters:

cudaMPI” [18]

We present and analyze two new communication libraries, cudaMPI and glmPI, that provide an MPI-like message passing interface to communicate data stored on the graphics cards of a distributed-memory parallel computer. These libraries can help applications that perform general purpose computations on these networked GPU clusters. We explore how to efficiently support both point-to-point and collective communication for either contiguous or noncontiguous data on modern graphics cards. Our software design is informed by a detailed analysis of the actual performance of modern graphics hardware, for which we develop and test a simple but useful performance model.

- 2008 “Observation of volcanoes through webcams: Tools and techniques” [19]

This work explores techniques for deriving quantitative data from webcam observations. It illustrates the role that webcams can play in volcano monitoring, and shows our recently developed tools for the collation and dissemination of this data. Over the past 5 years, digital cameras have been installed at a number of volcanoes to allow the general public to see volcanic activity from the comfort of their own homes. In the last 3 years these webcam images have become part of the twice-daily volcano monitoring report by the remote sensing team of the Alaska Volcano Observatory (AVO). To allow comprehensive and systematic analysis, a database has been created containing all AVO webcam images as well as images from St. Helens and three KVERT webcams for Bezymyanny, Klyucheskoy and Shiveluch. In total, some 1.6 million images are currently held. The number increases daily as new images are ob-

tained and processed. The database holds additional information about each image such as both image-wide and localized-region statistics. Our tools have been developed to answer specific questions utilizing this data. Of the current 1.6 million images in the database, a very small percentage is considered interesting for volcano monitoring; the remainder can be ignored due to complete cloud cover or (for nocturnal images) lack of luminescence. We have developed a tool for automatically isolating uninteresting images (primarily based on image histograms.) Uninteresting images are tagged, which allows for them to be excluded from further processing. Our next tool is an automated system for isolating and measuring nocturnal luminescence. This tool has been developed using images of St. Helens and is being extended to work with other webcams where nighttime lava glow have been seen. The system works by first minimizing each camera's unique dark current and amplification noise signals and then establishes if any pixels fulfill a number of criteria that would indicate they are real "glow" pixels. A third tool integrates the webcam data into Google Earth, allowing the height of plumes to be easily estimated and also the location of features on the sides of the volcano to be accurately placed on to the ground/GoogleMap surface. This tool makes visualization of multiple webcams of a single volcano easy. The recent eruptive phase of St. Helens has prompted the development of a system for measuring 'dome growth'. A technique using stacks of images and the Irani-Peleg super-resolution algorithm to estimate subpixel region growth is being tested at the moment. The final tool presented (currently under development) is a simple volcanic plume finder. It works by analyzing the RGB histograms of key regions around the volcano vent and flag-

ging image statistics which indicate the possible presence of steam. It is postulated that when adjusted for local climatic conditions the amount of steaming may be indicative of change at a volcano. These techniques aim to move the analysis of webcam images into an operational framework on a sound scientific footing.

- 2007 "MPIglut: Powerwall Programming Made Easier" [20]

A powerwall is an array of separate screens that work together to provide a single unified display. Powerwalls are often driven by a small cluster, which requires parallel software to organize and synchronize the distributed rendering process. This paper describes MPIglut, our powerwall-friendly implementation of the popular sequential GLUT OpenGL 3D programming interface. MPIglut internally communicates using MPI to provide a single coherent display even across a distributed-memory parallel machine. Uniquely, MPIglut is source-code compatible with ordinary sequential GLUT code while providing high performance.

- 2006 "Interpolation-Friendly Soft Shadow Maps" [21]

We present Penumbra Limit Maps, a technique for extracting soft shadows from a modified shadow map. The shadow representation used by our method has excellent interpolation properties, allowing the shadow boundary to be rendered with sub shadowmap-pixel accuracy, which partially mitigates the resolution problems common to shadow map methods. Unlike similar shadow map methods, our method includes both inner and outer penumbræ, and is in fact physically correct for the simplest case of a straight object edge and infinitely distant extended light source. At object corners or where

multiple object edges overlap, our method is no longer physically exact, but still gives plausible results. Finally, we show the method can be implemented naturally and efficiently on programmable graphics hardware.

- 2006 “Multiple Flows of Control in Migratable Parallel Programs” [22]

Many important parallel applications require multiple flows of control to run on a single processor. In this paper, we present a study of four flow-of-control mechanisms: processes, kernel threads, user-level threads and event-driven objects. Through experiments, we demonstrate the practical performance and limitations of these techniques on a variety of platforms. We also examine migration of these flows-of-control with focus on thread migration, which is critical for application-independent dynamic load balancing in parallel computing applications. Thread migration, however, is challenging due to the complexity of both user and system state involved. In this paper, we present several techniques to support migratable threads and compare the performance of these techniques.

- 2006 “ParFUM: A Parallel Framework for Unstructured Meshes for Scalable Dynamic Physics Applications” [23]

Unstructured meshes are used in many engineering applications with irregular domains, from elastic deformation problems to crack propagation to fluid flow. Because of their complexity and dynamic behavior, the development of scalable parallel software for these applications is challenging. The Charm++ Parallel Framework for Unstructured Meshes allows one to write parallel programs that operate on unstructured meshes with only

minimal knowledge of parallel computing, while making it possible to achieve excellent scalability even for complex applications. Charm++’s messagedriven model enables computation/communication overlap, while its runtime load balancing capabilities make it possible to react to the changes in computational load that occur in dynamic physics applications. The framework is highly flexible and has been enhanced with numerous capabilities for the manipulation of unstructured meshes, such as parallel mesh adaptivity and collision detection.

- 2006 “A system integration framework for coupled multiphysics simulations” [24]

Multiphysics simulations are playing an increasingly important role in computational science and engineering for applications ranging from aircraft design to medical treatments. These simulations require integration of techniques and tools from multiple disciplines, and in turn demand new advanced technologies to integrate independently developed physics solvers effectively. In this paper, we describe some numerical, geometrical, and system software components required by such integration, with a concrete case study of detailed, three dimensional, parallel rocket simulations involving system level interactions among fluid, solid, and combustion, as well as subsystem-level interactions. We package these components into a software framework that provides state of the art, common-refinement based methods for transferring data between potentially nonmatching meshes, novel and robust face-offsetting methods for tracking Lagrangian surface meshes, as well as integrated support for parallel mesh optimization, remeshing, algebraic manipulations, performance monitoring, and high-level data management and I/O.

From these general, reusable framework components we construct domain-specific building blocks to facilitate integration of parallel, multiphysics simulations from high level specifications that are easy to read and can also be visualized graphically. These reusable building blocks are integrated with independently developed physics codes to perform various multiphysics simulations.

- 2005 “Performance Degradation in the Presence of Subnormal Floating-Point Values” [25]

Operating system interference in parallel programs can cause tremendous performance degradation. This paper discusses the interference caused by the quiet generation of subnormal floating point values. We analyze the performance impact of subnormal values in a parallel simulation of a stress wave propagating through a three dimensional bar. The floating-point exception handling mechanisms of various parallel architectures and operating systems lead to widely differing performance for the same program. We show that a parallel program will exhibit greatly amplified performance degradation due to this interference. In addition we provide an simple example program that demonstrates underflow on a single processor. Finally we suggest a novel option for fixing these undesired slowdowns.

- 2005 “An Integration Framework for Simulations of Solid Rocket Motors” [26]

Simulation of solid rocket motors requires coupling physical models and software tools from multiple disciplines, and in turn demands advanced techniques to integrate independently developed physics solvers effectively. In this paper, we overview some computer science components required for such integration. We

package these components into a software framework that provides system support of high-level data management and performance monitoring, as well as computational services such as novel and robust algorithms for tracking Lagrangian surface meshes, parallel mesh optimization, and data transfer between nonmatching meshes. From these reusable framework components we construct domain-specific building blocks to facilitate integration of parallel, multiphysics simulations from high-level specifications. Through examples, we demonstrate the exibility of our framework and its components.

- 2004 “Impostors for Parallel Interactive Computer Graphics” [27]

We demonstrate an interactive parallel rendering system based on the impostors technique. Impostors increase the latency tolerance of an interactive rendering system, which allows us to use the power of a parallel machine even at high resolutions and framerates. Impostors also decrease the required rendering bandwidth, which makes possible the interactive use of a variety of advanced rendering techniques. These techniques are demonstrated by the interactive high-quality rendering of very large detailed models on large distributed memory parallel machines.

- 2004 “Debugging Support for Charm++” [28]

This paper describes a parallel debugger and the related debugging support implemented for CHARM++, a data driven parallel programming language. Because we build extensive debugging support into the parallel runtime system, applications can be debugged at a very high level.

- 2004 “Performance Modeling and Programming Environments for Petaflops Computers

and the Blue Gene Machine” [29]

We present a performance modeling and programming environment for petaflops computers and the Blue Gene machine. It consists of a parallel simulator, BigSim, for predicting performance of machines with a very large number of processors, and BigNetSim, an ongoing effort to incorporate a pluggable module of a detailed contention based network model. It provides the ability to make performance predictions for machines such as BlueGene/L. We also explore the programming environments for several planned applications on the machines including Finite Element Method (FEM) simulation.

- 2003 “Bounding Iterated Function Systems using Convex Optimization” [30]

We present an algorithm to construct a tight bounding polyhedron for a recursive procedural model. We first use an iterated function system (IFS) to represent the extent of the procedural model. Then we present a novel algorithm that expresses the IFS-bounding problem as a set of linear constraints on a linear objective function, which can then be solved via standard techniques for linear convex optimization. As such, our algorithm is guaranteed to find the recursively optimal bounding polyhedron, if it exists. Finally, we demonstrate examples of this algorithm on two and three dimensional recursive procedural models.

- 2003 “Supporting dynamic parallel object arrays” [31]

We present efficient support for generalized arrays of parallel data driven objects. Array elements are regular C++ objects, and are scattered across the parallel machine. An individual element is addressed by its “index”, which

can be an arbitrary object rather than a simple integer. For example, an array index can be a series of numbers, supporting multidimensional sparse arrays; a bit vector, supporting collections of quadtree nodes; or a string. Methods can be invoked on any individual array element from any processor, and the elements can participate in reductions and broadcasts. Individual elements can be created or deleted dynamically at any time. Most importantly, the elements can migrate from processor to processor at any time. The paper discusses support for message delivery and collective operations in face of such dynamic behavior. The migration capabilities of array elements have proven extremely useful, for example, in implementing flexible load balancing strategies and for exploiting workstation clusters adaptively. We present the design, an implementation, and performance results.

- 2003 “Adaptive MPI” [32]

“Adaptive MPI”, or AMPI, implements virtual MPI processors, several of which may reside on a single physical processor. This virtualization allows MPI applications to use an automatic migration-based load balancer, automatically overlap computation and communication, and provides several other benefits. In this paper, we present the design of and recent work on AMPI, its low-level and application performance, and some of the advanced capabilities enabled by virtualization.

- 2002 “Parallel Objects: Virtualization and in-Process Components” []

We summarize the object-based virtualization model that we have been developing for the past decade, and demonstrate how it enables automatic optimizations, especially at runtime. The parallel programming paradigm represented by

the virtualization model has been implemented in the Charm++ and AMPI libraries. In this paradigm, the programmer species their parallel application as a collection of interacting entities, without any reference to processors. The runtime system is free to map these entities to processors, and migrate them at runtime as needed. This separation of concerns enables several runtime optimizations, involving message-driven execution, automatic load balancing and communication patterns. A recently developed component model is also shown to create new opportunities for runtime optimizations.

- 2002 “A Voxel-Based Parallel Collision Detection Algorithm” [33]

Two physical objects cannot occupy the same space at the same time. Simulated physical objects do not naturally obey this constraint. Instead, we must detect when two objects have collided—we must perform collision detection. This work presents a simple voxel-based collision detection algorithm, an efficient parallel implementation of the algorithm, and performance results.

- 2001 “A Grid-Based Parallel Collision Detection Algorithm” [34]

In this work, we present a scalable high-level parallel solution to a large subclass of collision detection problems. Our approach is to divide space into a sparse grid of regular axis-aligned voxels distributed across the parallel machine. Objects are then sent to all the voxels they intersect. Once all the objects have arrived, each voxel becomes a self-contained subproblem, which is then solved using standard serial collision detection approaches. This voxel-based approach efficiently and naturally separates many objects that cannot ever collide,

by placing them in separate voxels. Simultaneously, voxels bring together adjacent objects that may intersect.

- 2001 “Supporting Dynamic Parallel Object Arrays” [35]

We present efficient support for generalized arrays of parallel data driven objects. The “array elements” are scattered across a parallel machine. Each array element is an object that can be thought of as a virtual processor. The individual elements are addressed by their “index”, which can be an arbitrary object rather than a simple integer. For example, it can be a series of numbers, supporting multidimensional sparse arrays; a bit vector, supporting collections of quadtree nodes; or a string. Messages can be sent to any individual array element from any processor, and the elements can participate in reductions and broadcasts. Individual elements can be created or deleted dynamically at any time. Most importantly, the elements can migrate from processor to processor at any time. The paper discusses support for message delivery and collective operations in face of such dynamic behavior. The migration capabilities of array elements have proven extremely useful, for example, in implementing flexible load balancing strategies and for exploiting workstation clusters adaptively.

- 1999 “Room Capacity Analysis Using a Pair of Evacuation Models” [36]

We present two models for determining the amount of time it takes a given number of people to evacuate a given room. A room’s maximum capacity can be derived from this by imposing a maximum evacuation time. The maximum evacuation time must take into account factors such as the fire resistance of the room

and should be calculated, for example, by the Fire Marshall. We developed a graph-based network flow simulation. People are modeled as a compressible fluid which flows toward and out the exit. This model assumes people's interaction properties, based on industry research. We also developed a discrete particle simulation. In this model, people are modeled as disks that attempt to reach the exits. In this model, people's interaction properties emerge from local, person assumptions. In this paper, we develop and analyze both models. We then compare and evaluate the models' outputs, and finally analyze the capacity of a local dining hall, gymnasium, lecture hall, and swimming pool.

- 1999 “Automated DEM Production using ESA Tandem Mission Data for the Caribou-Poker Creek LTER Watershed, Alaska” [37]

The Alaska SAR Facility has developed an automated procedure for production of a digital elevation model (DEM) from ESA tandem mission data. This software processes from ASF Computer Compatible Signal Data (CCSD) or Level Zero (raw) products to a map-projected, ground-range 30m DEM. Several advanced techniques have been integrated into the procedures to improve accuracy and to allow full automation. This spring, a test site at the Caribou-Poker Long Term Ecological Research (LTER) watershed was analyzed in collaboration with Larry Hinzman (Principal Investigator) and the Arctic Region Supercomputing Center. An accuracy assessment of the results for four ESA tandem pairs showed average differences ranging from approximately 4 to 8 meters in elevation when compared with field measurements using differential GPS. The final mosaic had an average difference of 4.68 meters in elevation. Point target analysis of existing

3x6 and 2x3 arc-second USGS digital elevation data showed an average difference of 19.97 and 10.04 meters in elevation, respectively. These advances in accuracy are due to use of precision timing and orbital data in an interferometric SAR processor using an average Doppler, precise baseline refinement, and direct ground rectification. The Alaska SAR Facility continues to refine its interferometric SAR processor in support of NASA-approved users.

- 1998 “Automated Digital Elevation Model (DEM) Production Using ERS SAR Tandem Pairs” [38]
- 1998 “Repeat-Pass Satellite Interferometric Tools Available at the Alaska SAR Facility” [39]
- 1997 “Pursuit-Evasion Games in the Late Cretaceous” [40]

We have been asked by a group of paleontologists to model the hunting and evasion strategies of the predator *Velociraptor mongoliensis* as, alone or with a friend, he pursues the prey species *Thescelosaurus neglectus*; both extinct dinosaurs. The predator is 6/5ths faster than the prey, but has a 3-times larger turning radius at top speed. The chase continues for only 15 seconds, whereupon the predator must stop because of a lactic acid buildup in his muscles.

- 1997 “Towards Operational Application of Satellite SAR Images in Hydrological Studies” [41]
- 1997 “Generation of Fine Resolution DEM at Test Areas in Alaska Using ERS Tandem Pairs and Precise Orbital Data” [42]

The ASF Science division has released the world's first free end-to-end interferometric

digital elevation model (DEM) generation system. This software, which processes from raw signal data through to a map-projected, ground-range 20m DEM, is completely automated. Preliminary comparison with differential global positioning system (GPS) indicates that over a 100km swath, horizontal position errors are less than 120m, and comparison with the 2x3 arc second United States Geological Survey (USGS) DEM indicates an average vertical error of 7m, 25m RMS. This result was obtained from an ERS tandem pair over Delta Junction, Alaska. These advances in accuracy are due to use of precision timing and orbital data in an interferometric SAR processor using an average doppler, precise baseline refinement, and direct ground rectification. The computationally intensive nature of these algorithms was minimized through the creation of a parallel SAR processor and a linearized ground rectification procedure.

References

- [1] Orion Sky Lawlor, Samuel Lucas, and Joren Bowling. Fluxed melting for rapid regolith digestion. In *Space Resources Roundtable (SRR)*, 2022.
- [2] Orion Sky Lawlor and Ben Talotta. Virtual reality vehicle simulator phase 1. In *PacTrans / USDOT accession 01723934*, 2021.
- [3] Orion Sky Lawlor. A dual-pin tool coupler for robotic excavation. In *Lunar Surface Innovation Consortium (LSIC) fall 2021 meeting*, 2021.
- [4] Abir Hasan, Nusrat Jahan, and Orion Lawlor. Forensic analysis of binary structures of video files. *Proceedings of CSP 2021*, 2021.
- [5] Orion Lawlor. Case study: An automated theft on mars. In *MGMWerx Space Challenge*, 2020.
- [6] Andrea Gomoll, Selma Šabanović, Erin Tolar, Cindy E. Hmelo-Silver, Matthew Francisco, and Orion Lawlor. Between the social and the technical: Negotiation of human-centered robotics design in a middle school classroom. *International Journal of Social Robotics*, 10:309–324, June 2018.
- [7] Orion Lawlor, Seta Bogosyan, Yasin Vural, Isaac Thompson, Michael Moss, and Metin Gokasan. Aero-beam: An open-architecture test-bed for research and education in cyber-physical systems. *Proceedings of IECON2015*, pages TS–171, 2015.
- [8] Orion Sky Lawlor and Jon Genetti. Exploiting spatial redundancy with adaptive pyramidal rendering. *Proceedings of WSCG*, page J37, 2014.
- [9] Orion Sky Lawlor. Algorithms: Rendering. In Phillip A. Laplante, editor, *Encyclopedia of Computer Science and Technology*. CRC Press, 2017. <https://www.taylorfrancis.com/books/9781351652490>.
- [10] Orion Lawlor, Mike Moss, Steven Kibler, Charles Carlson, Shaun Bond, and Seta Bogosyan. Search and rescue robots for integrated research and education in cyber-physical systems. In *Proceedings of IECON/ICELIE13 Vienna*, 2013. <http://cyberalaska.cs.uaf.edu/>.
- [11] Orion Sky Lawlor. In-memory data compression for sparse matrices. In *Proceedings of IA3, SC13*. ACM, 2013.
- [12] Philippe H. Geubelle Orion Sky Lawlor, M. Scot Breitenfeld and Gengbin Zheng. Crack

- propagation analysis with automatic load balancing. In Laxmikant V. Kale and Abhinav Bhatele, editors, *Parallel Science and Engineering Applications: The Charm++ Approach*, pages 187–210. CRC Press, 2013.
- [13] Orion Sky Lawlor and Thomas R. Quinn. Remote visualization of cosmological data using Salsa. In Laxmikant V. Kale and Abhinav Bhatele, editors, *Parallel Science and Engineering Applications: The Charm++ Approach*, pages 137–145. CRC Press, 2013.
- [14] Orion Sky Lawlor. Netrun: Student code in a web browser (poster). In *Workshop on Parallel and Distributed Computing Education*. NSF/TCPP, May 2011.
- [15] Orion Sky Lawlor. GPU-accelerated rendering of unbounded nonlinear iterated function system fixed points. *ISRN Computer Graphics*, December 2011.
- [16] Orion Sky Lawlor. Embedding OpenCL in C++ for expressive GPU programming. In *IPDPS WOLFHPC*, May 2011. <http://www.cs.uaf.edu/sw/EPGPU>.
- [17] Orion Sky Lawlor and Jon Genetti. Interactive volume rendering aurora on the GPU. *Journal of WSCG*, page H41, 2011.
- [18] Orion Sky Lawlor. Message passing for GPGPU clusters: cudaMPI. In *IEEE Cluster PPAC Workshop*, 2009. <http://www.cluster2009.org/paperw102.php>.
- [19] Joe Lovick, Orion Lawlor, Ken Dean, and John Dehn. Observation of volcanoes through webcams: Tools and techniques. In *AGU Fall Meeting Abstracts*. American Geophysical Union, December 2008.
- [20] Orion Sky Lawlor, Matthew Page, and Jon Genetti. MPIglut: Powerwall programming made easier. *Journal of WSCG*, pages 130–137, 2008.
- [21] Orion Sky Lawlor. Interpolation-friendly soft shadow maps. In *Proceedings of Computer Graphics and Virtual Reality '06*, pages 111–117. CSREA Press, 2006.
- [22] Gengbin Zheng, Orion Sky Lawlor, and Laxmikant V. Kale. Multiple flows of control in migratable parallel programs. In *Proceedings of 8th Workshop on High Performance Scientific and Engineering Computing (HPSEC-06)*. IEEE Press, August 2006.
- [23] Orion Sky Lawlor, Sayantan Chakravorty, Terry L. Wilmarth, Nilesh Choudhury, Isaac Dooley, Gengbin Zheng, and Laxmikant V. Kale. ParFUM: A parallel framework for unstructured meshes for scalable dynamic physics applications. *Engineering With Computers*, 22(3):215–235, 2006.
- [24] Xiangmin Jiao, Gengbin Zheng, Phillip A. Alexander, Michael T. Campbell, Orion S. Lawlor, John Norris, Andreas Haselbacher, and Michael T. Heath. A system integration framework for coupled multiphysics simulations. *Engineering with Computers*, 22(3–4):293–309, 2006.
- [25] Orion Sky Lawlor, Hari Govind, Isaac Dooley, Michael Breitenfeld, and Laxmikant Kale. Performance degradation in the presence of sub-normal floating-point values. In *Proceedings of the International Workshop on Operating System Interference in High Performance Applications*, September 2005.

- [26] Xiangmin Jiao, Gengbin Zheng, Orion Lawlor, Phil Alexander, Mike Campbell, Michael Heath, and Robert Fiedler. An integration framework for simulations of solid rocket motors. In *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, Tucson, Arizona, July 2005.
- [27] Orion Sky Lawlor. *Impostors for Parallel Interactive Computer Graphics*. PhD thesis, University of Illinois at Urbana-Champaign, December 2004.
- [28] Rashmi Jyothi, Orion Sky Lawlor, and L. V. Kale. Debugging support for Charm++. In *PADTAD Workshop for IPDPS 2004*, page 294. IEEE Press, April 2004.
- [29] Gengbin Zheng, Terry Wilmarth, Orion Sky Lawlor, Laxmikant V. Kalé, Sarita Adve, David Padua, and Philippe Geubelle. Performance modeling and programming environments for petaflops computers and the Blue Gene machine. In *NSF Next Generation Systems Program Workshop, 18th International Parallel and Distributed Processing Symposium (IPDPS)*, Santa Fe, New Mexico, April 2004. IEEE Press.
- [30] Orion Sky Lawlor and John Hart. Bounding iterated function systems using convex optimization. In *Proceedings of Pacific Graphics 2003*, pages 283–292, October 2003.
- [31] Orion Sky Lawlor and L. V. Kalé. Supporting dynamic parallel object arrays. *Concurrency and Computation: Practice and Experience*, 15:371–393, 2003.
- [32] Chao Huang, Orion Sky Lawlor, and L. V. Kalé. Adaptive MPI. In *Proceedings of the 16th International Workshop on Languages and Compilers for Parallel Computing (LCPC 03)*, College Station, Texas, October 2003.
- [33] Orion Sky Lawlor and L. V. Kalé. A voxel-based parallel collision detection algorithm. In *Proceedings of the International Conference in Supercomputing*, pages 285–293. ACM Press, June 2002.
- [34] Orion Sky Lawlor. A grid-based parallel collision detection algorithm. Master’s thesis, Dept. of Computer Science, University of Illinois, May 2001.
- [35] Orion Sky Lawlor and L. V. Kalé. Supporting dynamic parallel object arrays. In *Proceedings of ACM 2001 Java Grande/ISCOPE Conference*, pages 21–29, Stanford, CA, June 2001.
- [36] Gregg A. Christopher, Orion Lawlor, and Jason Tedor. Room capacity analysis using a pair of evacuation models. *UMAP*, 20.3, 1999.
- [37] Rick Guritz, Mike Ayers, Tom Logan, Shusun Li, and Orion Lawlor. Automated DEM production using ESA tandem mission data for the Caribou-Poker Creek LTER watershed, Alaska. In *Fringe ’99 Meeting*, November 1999.
- [38] Rick Guritz, Orion Lawlor, Tom Logan, Robert Fatland, Joanne Groves, Shusun Li, and Verne Kaupp. Automated Digital Elevation Model (DEM) production using ERS SAR tandem pairs. In *Proceedings of IGARSS ’98*, 1998.
- [39] Rick Guritz, Orion Lawlor, Tom Logan, Shusun Li, , and Verne Kaupp. Repeat-pass satellite interferometric tools available at the Alaska SAR Facility. In *American Geophysical Union 1998 Fall Meeting Transactions, EOS*, volume 79.45, pages 21–22, 1998.

- [40] Gordon Bower, Orion Lawlor, and James Long. Pursuit-evasion games in the late Cretaceous. *UMAP*, 18.3, 1997.
- [41] Shusun Li, Rick Guritz, Tom Logan, Z. Wang, and Orion Lawlor. Towards operational application of satellite SAR images in hydrological studies. In *American Geophysical Union EOS*, volume 78.45, 1997.
- [42] Orion Lawlor, Tom Logan, Rick Guritz, Robert Fatland, Shusun Li, Z. Wang, and Coert Olmsted. Generation of fine resolution DEM at test areas in Alaska using ERS tandem pairs and precise orbital data. In *12th International Conference on Applied Geologic Remote Sensing*, November 1997.