

FIGURE 3: Scalability results for 200 iterations of parallel Lagrangian dual ascent on instances from QAPLIB.

they are part of many cutting-edge projects with high socio-economic impact. Using our methods, scientists and engineers will be able to solve Assignment Problems containing hundreds of thousands of vertices within a matter of minutes, leading to **transformative** discoveries. The parallel programming library developed during this research will provide a platform for solving large-scale Assignment Problems and comparing the performance of different algorithms, thus ensuring continued advancement of science and engineering.

**PUBLICATIONS AND DATA SETS**

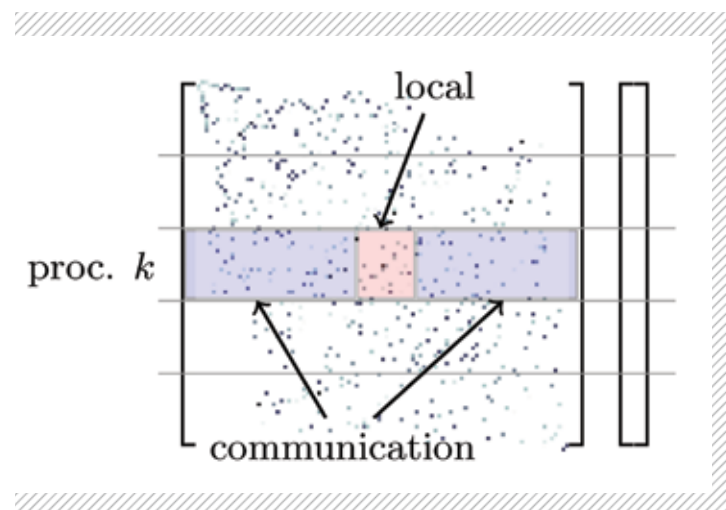
Date, K., R. Nagi, GPU-accelerated Hungarian algorithms for the Linear Assignment Problem. *Parallel Comput.*, (2016), <http://dx.doi.org/10.1016/j.parco.2016.05.012>

**THE NEXT GENERATION OF LARGE-SCALE SPARSE MATRIX COMPUTATIONS**

FIGURE 1: Communication in sparse matrix-vector multiplication

PI: Luke Olson<sup>1</sup>  
 Collaborators: Amanda Bienz<sup>1</sup>, Bill Gropp<sup>1</sup>, and Andrew Reisner<sup>1</sup>

<sup>1</sup>University of Illinois at Urbana-Champaign



**EXECUTIVE SUMMARY**

Sparse matrix computations are a common element in a range of high-performance computing applications and often dominate computation, particularly at scale. The goal of this project is to develop numerical schemes that limit parallel communication in sparse matrix methods, leading to efficient and scalable sparse operations without loss of accuracy.

Iterative methods for approximating the solution to sparse linear systems rely on sparse matrix-vector multiplication as a key computational kernel. Yet communication costs drive the complexity of these operations. The focus of this project is on algorithm design to alleviate communication costs at scale.




FIGURE 2: Effect of matrix sparsification on solver time.

This project has contributed two key developments: the ability to *sparsify* the matrix operations to limit communication and the method of redundant computation to localize computation, thus limiting communication distance and time.

operations throughout the simulation, leading to significant resource usage.

**INTRODUCTION**

Sparse matrix problems of the form  $Ax=b$  are common in many large-scale simulations. In this project, multigrid methods are considered as a method for the iterative approximation for this problem. Multigrid methods construct a series or hierarchy of smaller sparse matrices that are used to form an iterative refinement of the solution. This leads to fewer floating-point operations in the solution to  $Ax=b$ , but potentially very high communication costs and memory movement relative to the amount data in the computation.

A key kernel in this process is that of sparse matrix-vector multiplication  $w \leftarrow A v$ . Here, the sparsity pattern—or the location of non-zero entries in the matrix—governs the computational expense of the operation. A common bottleneck in the computation is the scenario where the sparse matrix resides on a high number of cores with a low number of matrix elements per core (Fig. 1).

Achieving efficiency in these computations is critical as many applications rely on these sparse

**METHODS & RESULTS**

Two methods were explored to address the communication demands in the problem. First, a method of sparsification whereby specific entries in the matrix are eliminated from the computation is explored. In multigrid, the matrices in the hierarchy are successively smaller, but relatively more dense, leading to wider communication patterns. A method is designed in [1] that avoids communicating work associated with *weak* or low-influence matrix entries in the problem. This leads to a notable speed-up of 2x, particularly at large core counts (Fig. 2). In addition, the *topology* of the system is also incorporated in the communication of sparse entries, preferring local on-node communication in the algorithm. This leads to a significant reduction in total message traffic, thereby reducing system contention and total simulation time (Fig. 3).

Ultimately, sparse matrices in the multigrid hierarchy are too small to reside on the full set of processing elements. Therefore a redistribution of data is needed to map matrix data to a subset of processors. This comes at a cost, which was explored in this project. By performing a redistribution of data based on a performance model, nearly scalable results can be achieved. As an example, for  $p$  processors,




FIGURE 3: Effect of localized message communication on MPI messages.

it was observed that the standard approach to redistribution of data yielded a growth of  $p^{1.1}$  in simulation time, while a redistribution designed to limit communication yields around  $p^{0.17}$ . This is a significant saving, particularly as  $p$  approaches one million cores.

### WHY BLUE WATERS

Blue Waters was key to this work as it provided access to the fast Cray Gemini interconnect as well as to the large core count needed for accurate scalability studies. Sparse matrix algorithms need to take advantage of multiple aspects of the compute architecture, and Blue Waters allowed for the development of both on-node features and long-distance communication decisions in the multigrid algorithm. As multigrid methods are expected to be a key solver technology on future systems, access to Blue Waters has been instrumental in advancing these methods.

### NEXT GENERATION WORK

Increased concurrency will continue to challenge the scalability of sparse matrix computations. One goal is to capture the topology features of these systems more directly in the algorithms. This project provided several steps toward that goal. A future element of this work would be to add heterogeneous compute elements to the multigrid algorithm and to enhance the underlying performance models to effectively capture their use.

## EVALUATING DATA-DRIVEN DETECTORS OF ELECTRICITY THEFT IN SMART GRIDS

**Allocation:** Illinois/50.0 knh  
**PI:** William H. Sanders<sup>1</sup>  
**Co-PI:** Varun Badrinath Krishna<sup>1</sup>  
**Collaborator:** Juran Kirihara<sup>1</sup>

<sup>1</sup>University of Illinois at Urbana-Champaign

### EXECUTIVE SUMMARY

Electricity theft is a billion-dollar problem faced by utilities around the world and current measures are ineffective against sophisticated theft attacks that compromise the integrity of smart meter communications. We are devising algorithms that detect such theft attacks, and that are based on mathematical techniques in statistics and machine

learning. The goal is to detect and mitigate theft by identifying anomalies in consumption patterns of electricity consumers. We used Blue Waters to evaluate the effectiveness of the autoregressive integrated moving average (ARIMA)-based approach to detect simulated anomalies in smart meter data. The best parameters for these algorithms need to be found using scanning techniques, and they need

to account for a wide range of attack parameters that produce anomalies. Our evaluation is based on a large dataset obtained from a real smart meter deployment.

### INTRODUCTION

Bloomberg News reported that electricity theft in India contributes to blackouts and costs \$17 billion in lost revenue annually. According to the World Bank, electricity theft contributes to a loss in electricity delivery of over 25% of generated supply in India, 16% in Brazil, 6% in China and the U.S., and 5% in Australia. Theft in these countries is almost always achieved by tapping into electric distribution lines. To detect these thefts, utility companies such as BC Hydro have been convincing consumers to install smart meters. However, there has been some push-back as consumers have begun to realize that smart meters are vulnerable to cyber intrusions. In 2010, the Cyber Intelligence Section of the FBI reported that smart meter consumptions were being under-reported in Puerto Rico, leading to annual losses for the utility estimated at \$400 million. In 2014, BBC News reported that smart meters in Spain were hacked to cut power bills. Given that smart meters can be compromised, the roll-out efforts of utilities such as BC Hydro may only increase the attack surface for cyber intrusion-based theft methods.

We identified seven classes of electricity attacks, some of which distribute the monetary loss to consumers, at no loss to the utility. Therefore, this problem is not only important to utilities, but also consumers around the world.

### METHODS & RESULTS

The methods in this project are detailed in our earlier work [2], where we simulated electricity theft attacks on 500 consumers and tested our detector's false positive and false negative rates on them. The detector fitted an ARIMA model to the consumption data time series and then flagged outliers using a confidence interval created from the model.

We used methods in [3] to fit the ARIMA time series but learned from the larger simulation on Blue Waters that these methods do not scale well and are very sensitive to outliers. Also, we found unexpected results that used Python packages built by third parties. Upon further investigation, we identified errors in the algorithms coded in those Python

packages (specifically the statsmodels.tsa.arima\_model.ARIMA package). The algorithms simply do not implement ARIMA models correctly and use the differencing order term in the ARIMA model in a manner that is inconsistent with the theory.

### WHY BLUE WATERS

Our earlier work published in [2] was performed at a much smaller scale (500 consumers) on a regular server rack, consuming inordinate amounts of processing time. We wanted to perform evaluation studies of our detector at a larger scale (2900 consumers). Also, we wanted to try out many parameters for our detector at that scale. Without Blue Waters, it would have taken years to complete these tasks.

### NEXT GENERATION WORK

Our experience with Blue Waters helped us identify problems with third party software packages. We are redeveloping those packages ourselves and hope to complete them in time for the September 2016 call for proposals. With the correct algorithms, we hope to discover insights on our electricity theft detector from our use of Blue Waters.