

High Fidelity Simulations of Large-Scale Wireless Networks Proposal No.: 16-0306

Investment Area: Defense Systems and Assessments
Project Intent: Create
Duration: 2 Years
Classified: No
Principal Investigator: Uzoma Onunkwo
Project Manager: Zachary Benz

1. Overview/Abstract: 310 word limit for section 1.1 and 1.2 combined

1.1 Project Purpose (250 words - *Define all acronyms.*)

The worldwide proliferation of wireless connected devices continues to accelerate. There are 10s of billions of wireless links across the planet with an additional explosion of new wireless usage anticipated as the Internet of Things develops. Wireless technologies do not only provide convenience for mobile applications, but are also extremely cost-effective to deploy. Thus, this trend towards wireless connectivity will only continue and Sandia must develop the necessary simulation technology to proactively analyze the associated emerging vulnerabilities.

Wireless networks are marked by mobility and proximity-based connectivity. The de facto standard for exploratory studies of wireless networks is discrete event simulations (DES). However, the simulation of large-scale wireless networks is extremely difficult due to prohibitively large turnaround time. A path forward is to expedite simulations with parallel discrete event simulation (PDES) techniques. The mobility and distance-based connectivity associated with wireless simulations, however, typically doom PDES and fail to scale (e.g., OPNET and ns-3 simulators).

We propose a PDES-based tool aimed at reducing the communication overhead between processors. The proposed solution will use light-weight processes to dynamically distribute computation workload while mitigating communication overhead associated with synchronizations. This work is vital to the analytics and validation capabilities of simulation and emulation at Sandia. We have years of experience in Sandia's simulation and emulation projects (e.g., MINIMEGA and FIREWHEEL). Sandia's current highly-regarded capabilities in large-scale emulations have focused on wired networks, where two assumptions prevent scalable wireless studies: (a) the connections between objects are mostly static and (b) the nodes have fixed locations.

1.2 Why is this proposed as an LDRD project? (60 words - *Define all acronyms.*)

There is great uncertainty in the ability to use traditional computational systems to simulate large-scale wireless networks. Mobility in wireless networks and the associated proximity-based communications makes such large-scale simulations extremely difficult. The underlying problems with simulations in this field – process synchronization and load balancing – are unsolved, but yet the need for such investigative studies is ever-increasing.

2. Proposed R&D

2.1 Technical Approach and Leading Edge Nature of Work:

The state of the art in simulating large-scale wireless networks relies on using significant compute resources (>10,000 processors) to achieve acceptable speedup (e.g., Rensselaer's Optimistic Simulation Systems or ROSS) and wall clock times. In 2013, the ROSS set the world record in a PHOLD benchmark test, simulating approximately 500 billion events per second on Lawrence Livermore National Lab's (LLNL's) Sequoia Blue Gene/Q supercomputer [1]. The study of DES scaling with ROSS was conducted by renowned PDES experts at LLNL and Rensselaer Polytechnic Institute. The ROSS effort



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



concretely illustrates that parallel discrete event simulations on multicore systems can achieve significant simulation speedups. However, the need for simulation platforms with very high processor count is not necessary; our target simulation platform is a standard server or cluster with tens to hundreds of processors. In addition, the nature of mobile wireless communications deviates significantly from the PHOLD benchmark due to large event count caused by proximity calculations and consequent dynamic communications, suggesting that ROSS scalability results may not hold for wireless network simulations.

Today, key computations surrounding wireless network simulations are naively executed using PDES engines. In particular, mobility and geo-based proximity calculations are generally executed using exhaustive distance calculations between the simulated data-transmitting object and all other simulated objects. These expensive calculations are performed when simulating transmission of broadcast messages because they are more intuitive for determining wireless connectivity between any pair of simulated objects. However, this usually leads to high communication overhead between processors and invariably yields poor parallel scaling.

Important pragmatic principles for achieving efficiency in parallel applications are: 1) computations are distributed evenly across processors and 2) communications between these processors are minimal. Under these circumstances, applications can often achieve linear speedups in the number of processors. Though these principles are well-known, they are still difficult to achieve in PDES and experts resort to naïve implementations that typically do not scale in systems with high mobility and proximity-based communications.

In most wireless network scenarios, node movement derives from human behavioral patterns resulting in several temporal clusters of wireless communication hubs (e.g., places of business). These communication hubs yield very large event generation rates in a discrete event simulation. Due to constant mobility in the system, statically assigning simulated nodes to processors incurs major pitfalls – uneven, highly varying computations on each processor and frequent synchronization communications between processors. Our differentiating approach in this area involves creating parallelized spatial indexing data structures for proximity-based computations and leveraging problem structure in light-weight dynamic load balancing for mitigating overhead due to synchronization communications. By creating parallelized spatial indexing algorithms, computational time for simulation of proximity-based communication events should be reduced by an order of magnitude yielding greater turnaround for what-if explorations of wireless networks [2].

Our simulation framework design will advance parallel implementation of spatial indexing algorithms like the R*-tree, which is expected to greatly reduce communication overhead when simulating proximity-based communications that are dominant in wireless networks. Currently, most wireless network PDES simulators require frequent communications between *all* processes of the simulation in order to know who should receive the simulated event, but this severely dampens scalability performance. In our approach, communication is expected to be reduced to significantly fewer processes with simulated components because proximity-based simulations will rely on spatial indexing thereby reducing queries to smaller set of processes than exhaustive searches.

This work is critical to large-scale analytics in Sandia’s simulation and emulation capabilities. Previous attempts at large-scale simulation have either failed because the PDES engine did not perform to scale (e.g., OPNET and ns-3) or require unacceptable low-resolution simulation fidelity. Our team’s experience with Sandia’s simulation and emulation efforts ensure that this work will be synergistic and not duplicative.

2.2 Relationship to Prior and Other On-Going Work:

There are parallel discrete event simulators capable of simulating wireless networks with high fidelity. Well-known examples are the commercial OPNET, the open-source ns-3 and OMNET++ network simulators. While these simulators have proven useful for analytics and validation studies, they suffer



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



significant shortcomings in their PDES engines. In OPNET, previous studies have shown that in many cases, the PDES performance is actually worse than serial simulations for small-scale (~100) mobile wireless nodes in a network. In ns-3, the PDES is simply not supported for wireless simulations, but is supported for study of networks with wired point-to-point connections.

In this project, we aim to solve two critical problems that hinder good parallel scaling of discrete event simulations in general – mitigating communication bottlenecks in the PDES engine while preserving correct simulation results, and maximizing usage of available computational resources throughout the course of the simulation. Communication bottlenecks can be reduced in wireless simulations if synchronization events due to proximity-based data transfers are computed efficiently – using parallelized spatial indexing techniques. The mobility of simulated components in wireless networks causes atypical changes in process synchronization and computational resource usage. We anticipate that if we are successful in leveraging and extending current load-balancing techniques, we can determine an allocation pattern tailored specifically to wireless network simulations. We note that we are not advocating for the design of a new full-suite simulator - in replicating much of the work involved in the developing of simulators such as OPNET and ns-3. Rather, we are proposing to develop a novel wireless simulation framework that can (i) provide and demonstrate hooks that allow current simulators to use our engine and (ii) provide flexibility that allows users to design specific standalone wireless simulation studies like protocol reliability behavior. Designing a new full-suite of wireless network simulator toolbox would require multiyear effort and is unnecessary.

2.3 Key R&D Goals, Objectives, and Project Milestones:

2.3.1. Goal and Success Measure

Our primary goal is the design and implementation of an integrable and adaptable parallel discrete event simulation framework for enabling scalable, high fidelity simulations of large-scale wireless networks. This simulation software will follow a two-step design, where simulators (e.g., OPNET and ns-3) that have been verified and validated for wireless network simulations will seamlessly leverage our framework for very large wireless network simulations. Complementarily, this tool will also yield the ability to perform accurate standalone design and simulations that have more focused interest than provided by the full-suite simulators.

The success of this project hinges on the readiness of the developed solution to integrate seamlessly and scale efficiently to statewide and nationwide size of wireless network studies. This will, for the first time, enable high fidelity investigative studies of large wireless networks.

2.3.2. Key Objectives and Milestones

Objective/Milestone	Completion Date
Scalable DES simulation framework for wireless analysis	
Parallel spatial indexing for proximity communication	03/31/2016
Light-weight rebalancing for parallel processes	08/31/2016
Integration to verified and validated (V&V) OPNET/ns-3 simulators	
Integrate indexing and rebalancing and demonstrate valid simulation	01/31/2017
Develop a test case of wireless network simulation with 10,000 nodes	05/31/2017
Show vastly improved simulation rates over conventional	
V&V simulators on a multicore workstation	08/31/2017
Conference (Winter Simulation Conference) Paper	04/30/2017
Final SAND Report	09/01/2017



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



2.4 Technical Risk:

The known technical obstacles associated with this project, and their corresponding mitigation plans, are as follows:

- Fast determination of wireless nodes proximity for data communications in a large network without constantly iterating over all nodes.
 - Spatial indexing methods exist but need to be extended for parallel and fast execution, which will quickly enable wireless simulations at scale without expensive synchronization across all processors. This team has ties with members of Sandia's Cyber Engineering Research Laboratory (CERL) and Computing Science Research Institute (CSRI) to guide development of our algorithm approaches.
- Load balancing computations due to assignment of large simulated nodes to many-core computing clusters using light-weight processes.
 - The team is aware of load-balancing efforts in projects such as Sandia's ZOLTAN used in matrix multiplication and can exchange information with the team to drive a similar problem that deals with thin events. In this context, thin events refer to the small compute-to-communication parallel model as contrasted from the popular coarse events (coarsely parallel).
- Ability to interface *seamlessly* to verified and validated wireless network simulators.
 - The team is comprised of experts with years of experience with network simulators like OPNET and ns-3. This will give us leverage to avoid the long learning curve associated with beginner's experience on proper integration of our simulation framework to the aforementioned simulators.
 - Our design also allows for standalone but narrowly focused simulations that will parallel our envisioned two-step design that will integrate efficiently with existing V&V wireless network simulators.

There is a genuine risk that we could fail as our projected goals do not accomplish the results we anticipate. However, our team and clear objectives will enable us to quickly stipulate go/no-go areas. The technical risks are moderate to high, but manageable with the specified mitigation measures. The knowledge gained will be very useful to Sandia DS&A mission and enable us answer questions involving ever-growing cyber concerns over the behavior of and vulnerabilities associated with global scale wireless networks.

3. Resources

3.1 Key Research Team Members:

Name	Org	Role
Uzoma Onunkwo	9336	PI
Robert Cole	5624	Simulations & Academic Collaborations
Anand Ganti	9336	Analytics & Validation
Brian Van Leeuwen	5624	Wireless Simulations
Jean-Paul Watson	1464	Algorithm Development

3.2 Qualifications of the Team to Perform This Work:

Members of this LDRD team have many years of experience developing and studying network models, performing discrete event simulations, and developing high performance algorithms.

Uzoma Onunkwo has years of vast experience in network simulations, high performance computing, and information processing. He successfully led the completion of a recent LDRD project in 2013 that



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



resulted in a conference publication (“Harnessing many-core processors for scalable, highly efficient, and adaptable firewall solutions”) that earned a Best Paper award at the 2013 IEEE International Conference on Computing, Networking and Communications.

Anand Ganti has experience with algorithms, information systems, and mathematical modeling.

Brian Van Leeuwen is a member of the Critical Systems Security Department where he performs systems analysis of secure networked information systems. The work has included development of methods for performing cyber analysis with simulation and also using live, virtual, constructive (LVC) approaches with an emphasis on secure network communications. He has experience in modeling and simulation of secure wireless networked systems.

3.3.2 Subcontracts, Unusual Budget Items, and Capital Purchases:

None

4. Strategic Alignment and Potential Benefit

4.1 Anticipated Outputs:

This LDRD will yield a tool that enables scalable simulation of large wireless networks, and enables a core capability to understand new vulnerabilities in military and commercial domains due to the exploding growth in the wireless sector. Today, it is virtually impossible to answer questions similar to the following:

- a) In a brigade or division size (from about 1K to 20K devices) of wirelessly connected devices (e.g., battlefield studies), how fast or reliable can protocol X push critical commands to device-holding individuals versus protocol Y?
- b) When a critical infrastructure in a smart grid is exposed to requests from thousands of networked devices (mostly, wireless), how do the rest of the grid respond? How about if we change our designs, what are the cost-risk benefits?

If we are successful in the proposed milestones, we believe that our solution can appropriately answer these questions and many similar ones, which cannot be addressed today in current simulation technologies.

4.2 Leveraging Results:

We will communicate directly with the emulytics team (FIREWHEEL, MINIMEGA) to determine the usefulness of our implementation in emulating wireless devices. If our framework is sufficiently scalable, we will publish our results in open journals and conferences. Finally, we will provide the SAND report that clearly documents the results of this project.

4.3 Tie to Investment Area Call:

Sandia National Laboratories has a vital and unmet need for valid simulations of large-scale wireless networks. The proliferation of numerous wireless systems, along with the expected booming growth in wireless networks with the proliferation of Internet of Things (IoT), raises national security questions about possible disruptions and threats to our nation’s cyber capabilities. High fidelity solution for investigative analysis of large wireless networks enables us to appropriately address these major concerns for the first time at scales beyond thousands of wireless devices.

4.4 Relevance to DOE and National Security Missions: (100 word limit - *Define all acronyms.*)

We aim to expand our nation’s capabilities in vulnerability assessment and tactical studies via large-scale wireless simulations (e.g., military battlefield, smart grid, emergency response, and IoT studies). This project fits well with the Department of Defense strategic goal to protect critical infrastructure. This



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



project will address our cyber security needs of reducing vulnerabilities that can damage our Nation's critical cyber infrastructures exposed to wireless communications.



U.S. DEPARTMENT OF
ENERGY



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



**Sandia
National
Laboratories**

REFERENCES

1. Barnes Jr, P. D., Carothers, C. D., Jefferson, D. R., & LaPre, J. M. (2013, May). Warp speed: executing time warp on 1,966,080 cores. In *Proceedings of the 2013 ACM SIGSIM conference on Principles of advanced discrete simulation* (pp. 327-336). ACM.
2. Agarwal, P. K., de Berg, M., Gudmundsson, J., Hammar, M., & Haverkort, H.J., (2001).. Box-trees and R-trees with near-optimal query time. In *Proceedings of the seventeenth annual symposium on Computational geometry* (pp. 124-133), ACM,



U.S. DEPARTMENT OF
ENERGY



National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2011-9283P. SCG.MV.



**Sandia
National
Laboratories**