High Performance Network Cybersecurity
DOE SBIR Commercialization

Some lessons learned and resource for peer small businesses

Richard Lethin, President, Reservoir Labs

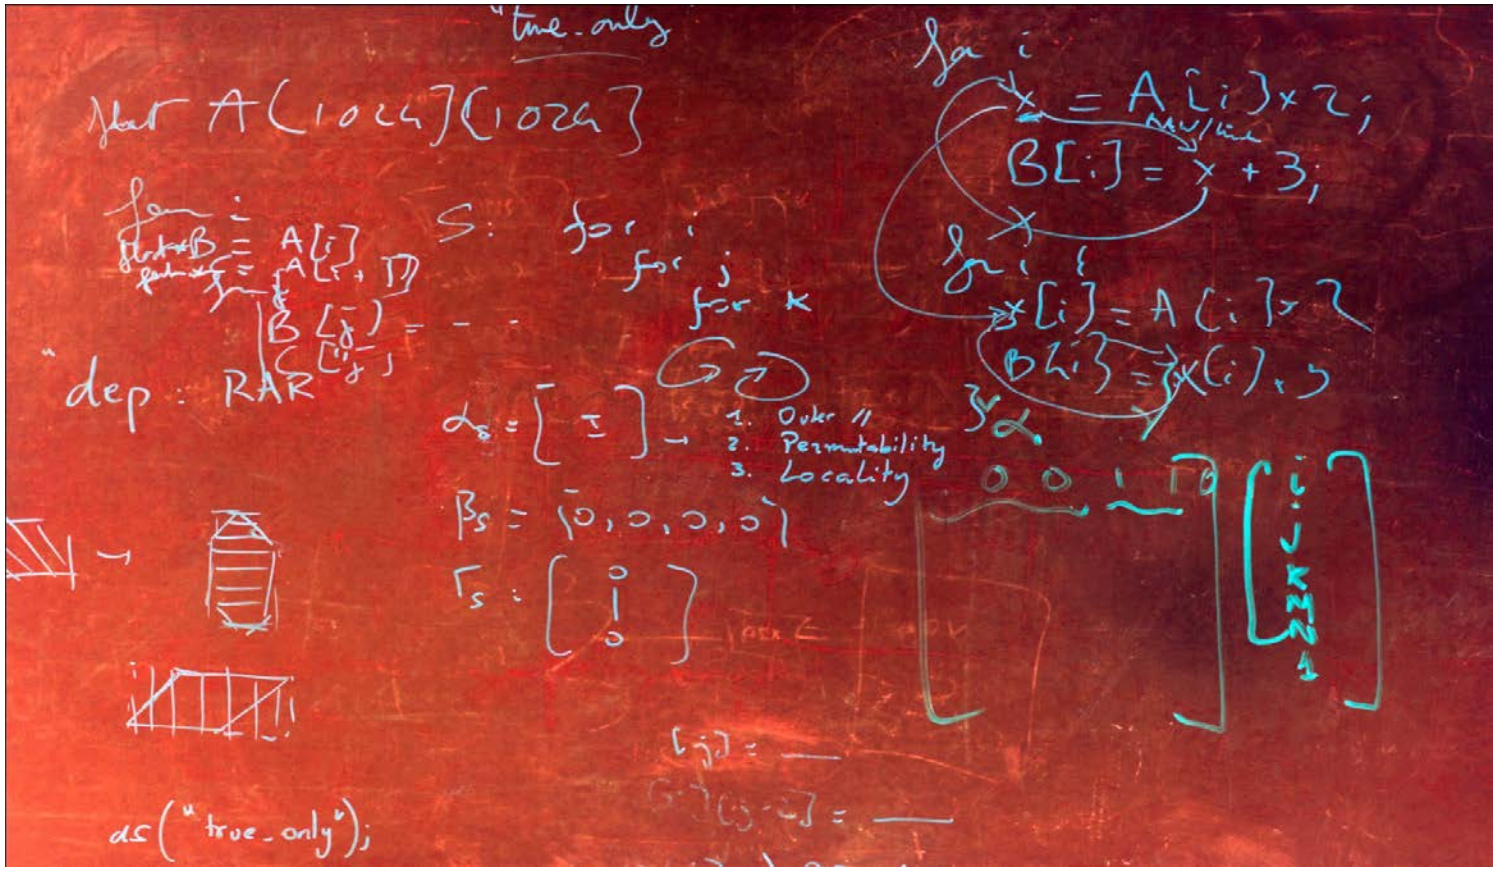

Representative Reservoir Technology and Expertise

Cybersecurity

R-Scope security appliance

Hypergraph data analytics of cyber sensors

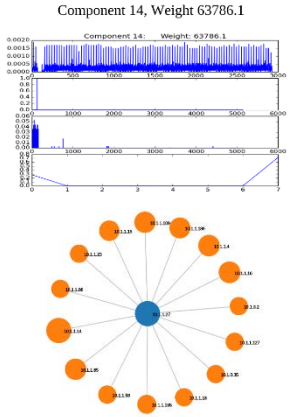

High Performance Computing

R-Stream automatic parallelizing compiler technology for ISR and exascale

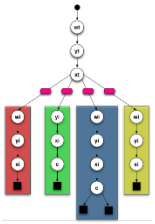

Software and programming tools for application-specific supercomputers

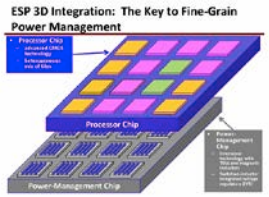

Analytics Verification

R-Check verification tools for software radio correctness and cybersecurity

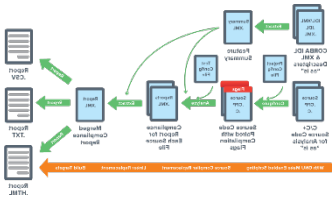

Data Analytics for DOD missions

Example 1: Incipient Alzheimer's data, CP-APR, 100 components

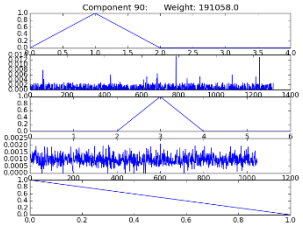

Algorithms

Spectral support graph algorithms

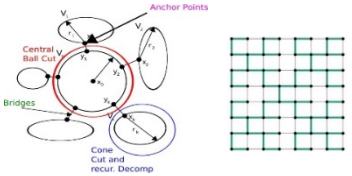

Fast Multipole Method (FMM)

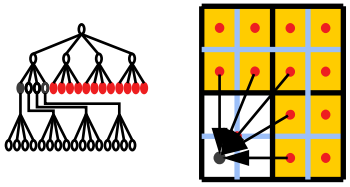

Compressive sensing

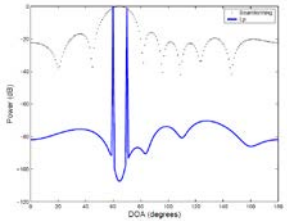

R-Scope® Network Sensor

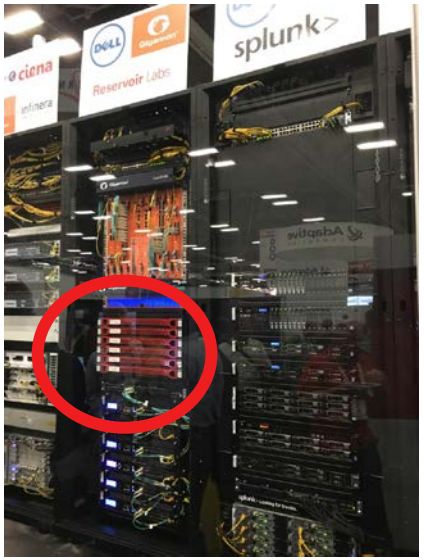

Product of DOE Advanced Scientific Computing Research (ASCR) Small Business Innovative Research (SBIR) funding

Substantial government and commercial sales

Value of resulting cyber security – priceless

Some Key Technology in R-Scope

- **New queuing algorithm** to reduce packet drops in hardware queues
- **Lockless bimodal producer-consumer queues** to eliminate multi-thread contention
- Algorithm to **dynamically shunt** traffic while maximizing **information entropy**.
- **Lockless hash tables** with low false negatives to eliminate memory contention overheads.
- **Multiresolution priority queues** to reduce the complexity of a priority queue down to O(1).

All patented or patent-pending

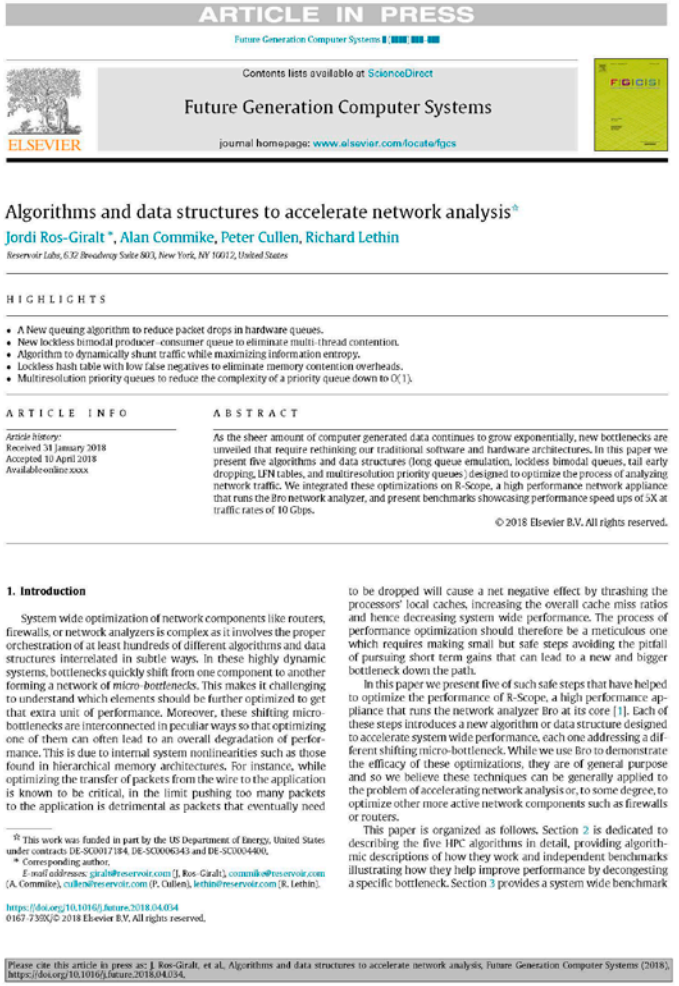

ARTICLE IN PRESS

Future Generation Computer Systems

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Algorithms and data structures to accelerate network analysis☆

Jordi Ros-Giralt *, Alan Commike, Peter Cullen, Richard Lethin

Reservoir Labs, 632 Broadway Suite 803, New York, NY 10012, United States

HIGHLIGHTS

- A New queuing algorithm to reduce packet drops in hardware queues.
- New lockless bimodal producer–consumer queue to eliminate multi-thread contention.
- Algorithm to dynamically shunt traffic while maximizing information entropy.
- Lockless hash table with low false negatives to eliminate memory contention overheads.
- Multiresolution priority queues to reduce the complexity of a priority queue down to O(1).

ARTICLE INFO

Article history:
Received 31 January 2018
Accepted 10 April 2018
Available online xxxx

ABSTRACT

As the sheer amount of computer generated data continues to grow exponentially, new bottlenecks are unveiled that require rethinking our traditional software and hardware architectures. In this paper we present five algorithms and data structures (long queue emulation, lockless bimodal queues, tail early dropping, LFN tables, and multiresolution priority queues) designed to optimize the process of analyzing network traffic. We integrated these optimizations on R-Scope, a high performance network appliance that runs the Bro network analyzer, and present benchmarks showcasing performance speed ups of 5X at traffic rates of 10 Gbps.

1. Introduction

System wide optimization of network components like routers, firewalls, or network analyzers is complex as it involves the proper orchestration of at least hundreds of different algorithms and data structures interrelated in subtle ways. In these highly dynamic systems, bottlenecks quickly shift from one component to another forming a network of *micro-bottlenecks*. This makes it challenging to understand which elements should be further optimized to get that extra unit of performance. Moreover, these shifting micro-bottlenecks are interconnected in peculiar ways so that optimizing one of them can often lead to an overall degradation of performance. This is due to internal system nonlinearities such as those found in hierarchical memory architectures. For instance, while optimizing the transfer of packets from the wire to the application is known to be critical, in the limit pushing too many packets to the application is detrimental as packets that eventually need to be dropped will cause a net negative effect by thrashing the processors' local caches, increasing the overall cache miss ratios and hence decreasing system wide performance. The process of performance optimization should therefore be a meticulous one which requires making small but safe steps avoiding the pitfall of pursuing short term gains that can lead to a new and bigger bottleneck down the path.

In this paper we present five of such safe steps that have helped to optimize the performance of R-Scope, a high performance appliance that runs the network analyzer Bro at its core [1]. Each of these steps introduces a new algorithm or data structure designed to accelerate system wide performance, each one addressing a different shifting micro-bottleneck. While we use Bro to demonstrate the efficacy of these optimizations, they are of general purpose and so we believe these techniques can be generally applied to the problem of accelerating network analysis or, to some degree, to optimize other more active network components such as firewalls or routers.

This paper is organized as follows. Section 2 is dedicated to describing the five HPC algorithms in detail, providing algorithmic descriptions of how they work and independent benchmarks illustrating how they help improve performance by decongesting a specific bottleneck. Section 3 provides a system wide benchmark

☆ This work was funded in part by the US Department of Energy, United States under contracts DE-SC0017184, DE-SC0006343 and DE-SC0004400.

* Corresponding author.
E-mail addresses: jgiralt@reservoir.com (J. Ros-Giralt), commike@reservoir.com (A. Commike), cullen@reservoir.com (P. Cullen), lethin@reservoir.com (R. Lethin).

https://doi.org/10.1016/j.future.2018.04.034
0167-739X/

Please cite this article in press as: J. Ros-Giralt, et al., Algorithms and data structures to accelerate network analysis, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2018.04.034.

https://www.reservoir.com/publication/algorithms-data-structures-accelerate-network-analysis-

High Performance Network Topic from ASCR

FY2011 Topic 39d:

- "Other: In addition to the specific subtopics listed above, the Department invites grant applications in other areas relevant to this Topic. Contact: Richard Carlson, 301-903-9486, rcarlson@ascr.doe.gov"
- "The goal will be to achieve to terabit-scale end-to-end network infrastructures for the open science community."

Lesson:

Don't ignore the "Other" topic: program managers are looking for good proposals, and will recognize and support them.

Looking back at 2013:
Reservoir is offering a full cyber solution (sensor + Splunk analytics) for high speed NW cyber security based on Bro

R-Scope Network Security Appliance

Agile, Scalable, Advanced Network Monitoring
Analytics/Detection Engine
Advanced Analytics Provided
Turnkey solution

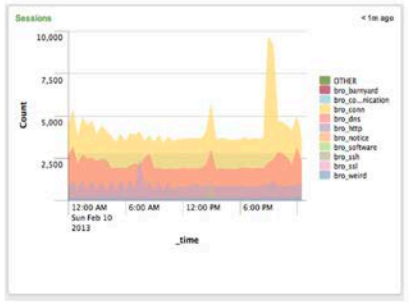

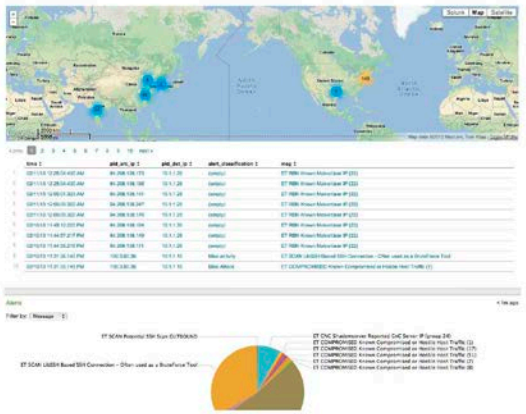

Reservoir Labs

Looking back at 2013:
This R-Scope solution is built working with excellent partner companies

R-Scope Network Security Appliance

R-Scope DOMINATE-T
R-Scope PACE-X
Bro open source software base
SIEM integration
Perimeter and insider protection
Configurable to 100 Gbps and beyond

Bro

Reservoir Labs

Looking back at 2013:
The R-Scope solution is a complete system with DOMINATE and PACE models

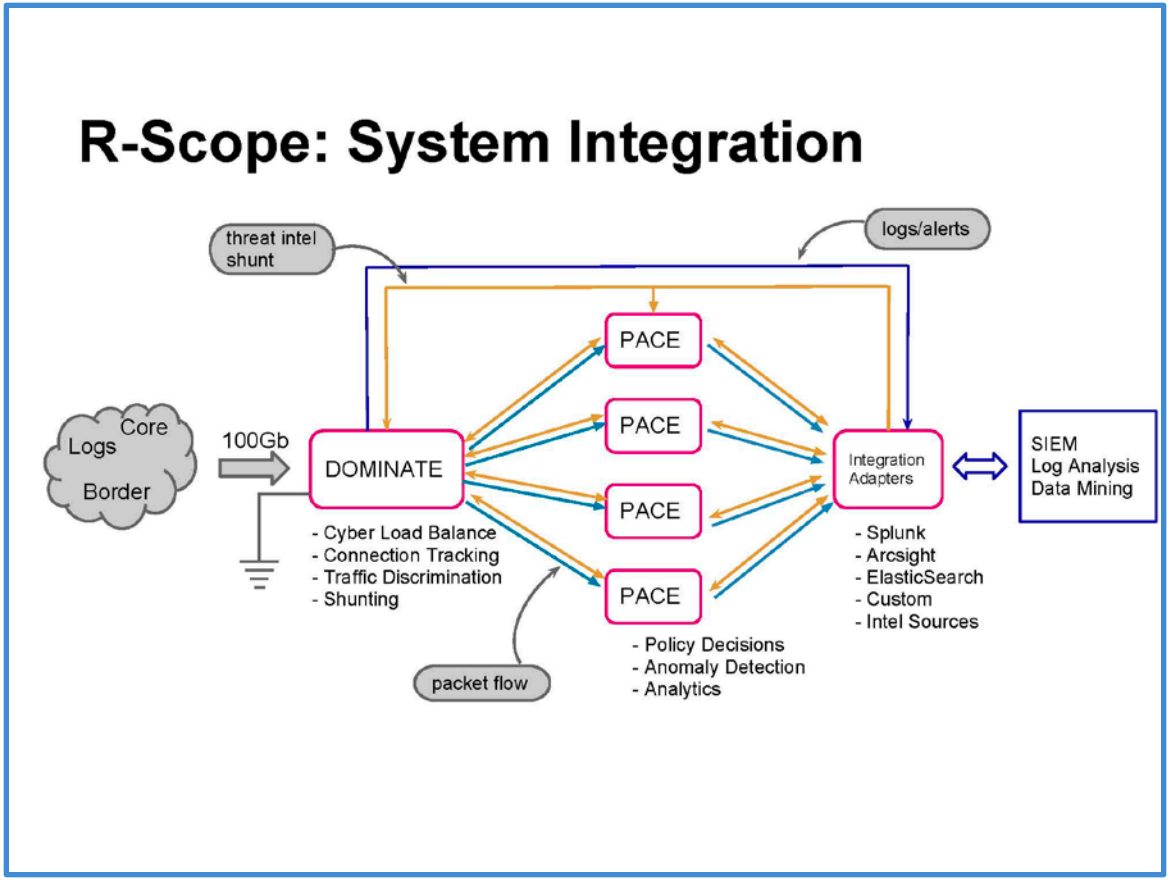

Looking back at 2013:
The high performance SBIR research is concentrated on the DOMINATE-T model

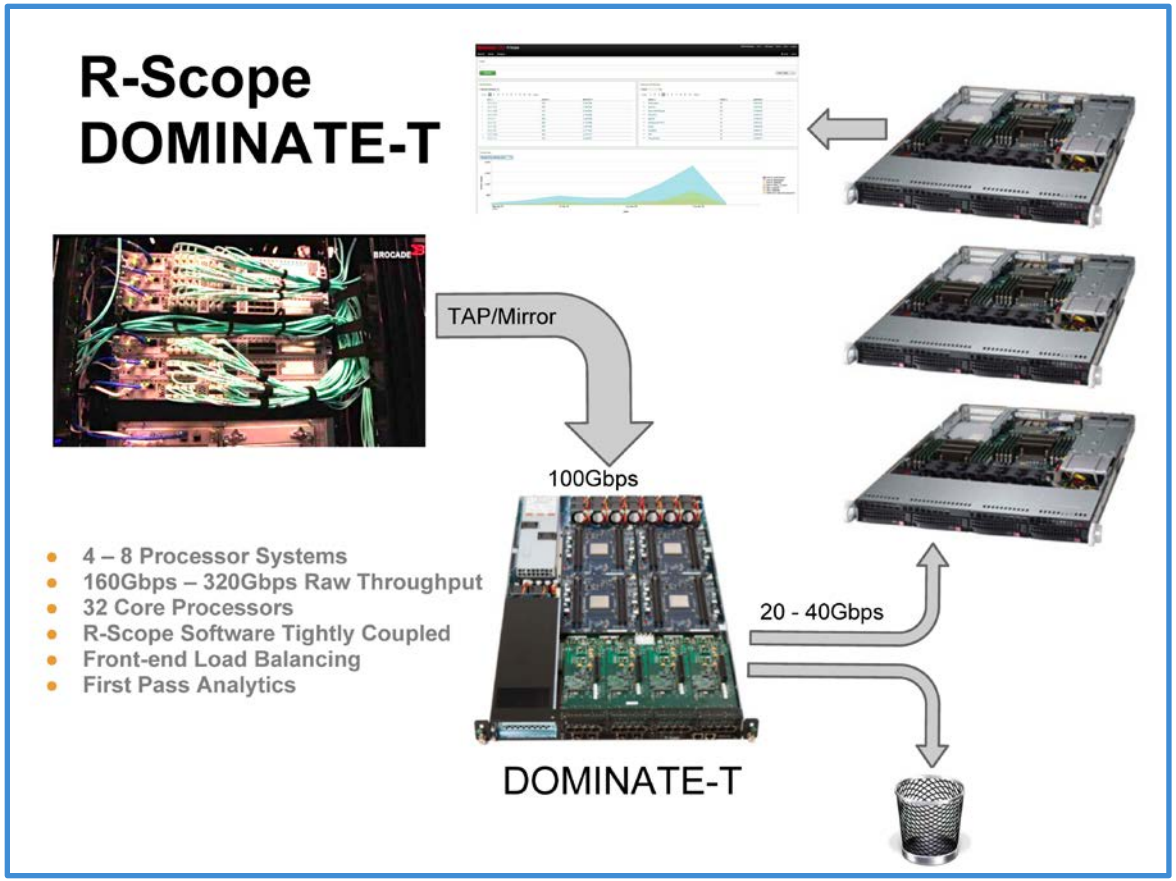

Looking back at 2013:
The DOMINATE-T model is based on Tilera chips and Tilera multi-daughtercard systems in a 1U form factor

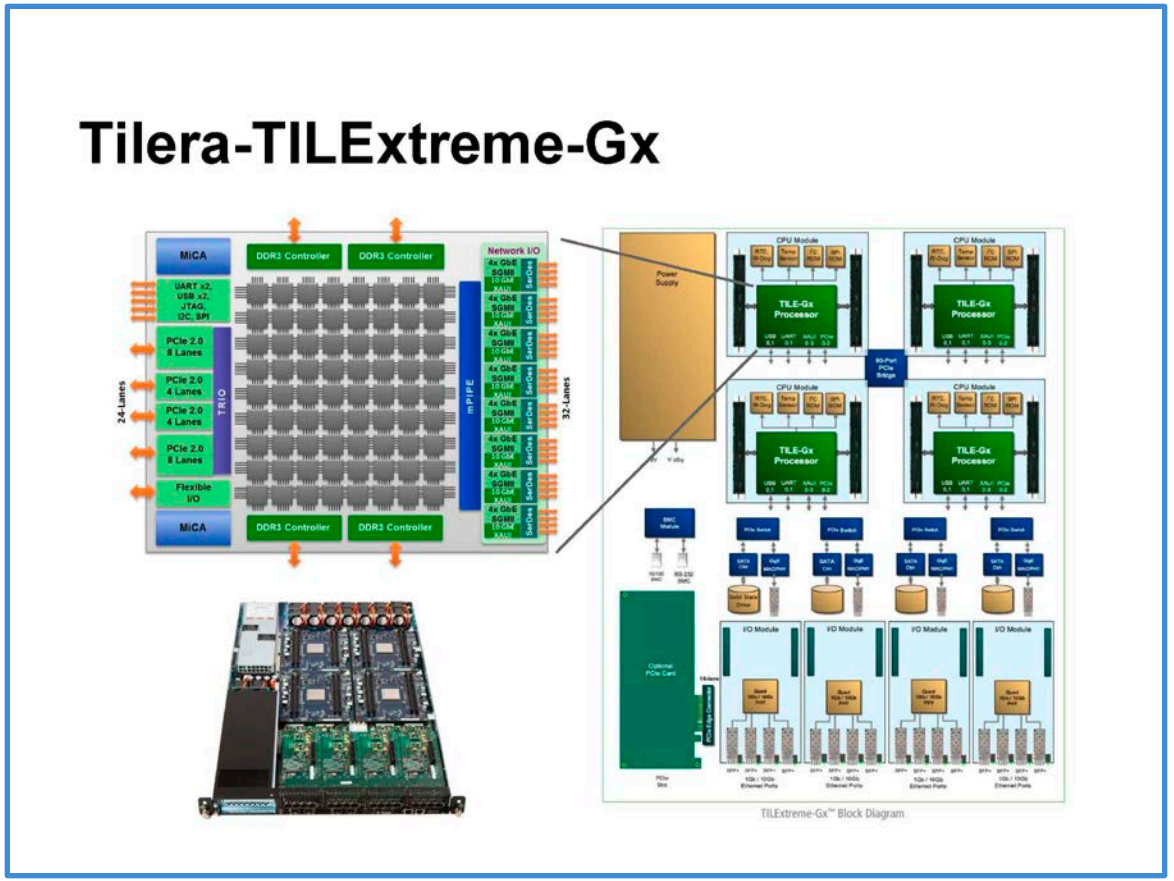

Looking back at 2013:
Reservoir innovates proprietary algorithms for flow regulation, load balancing, and many core chip management to utilize the Tilera hardware.

DOMINATE-T leverages advanced Tilera hardware for throughput

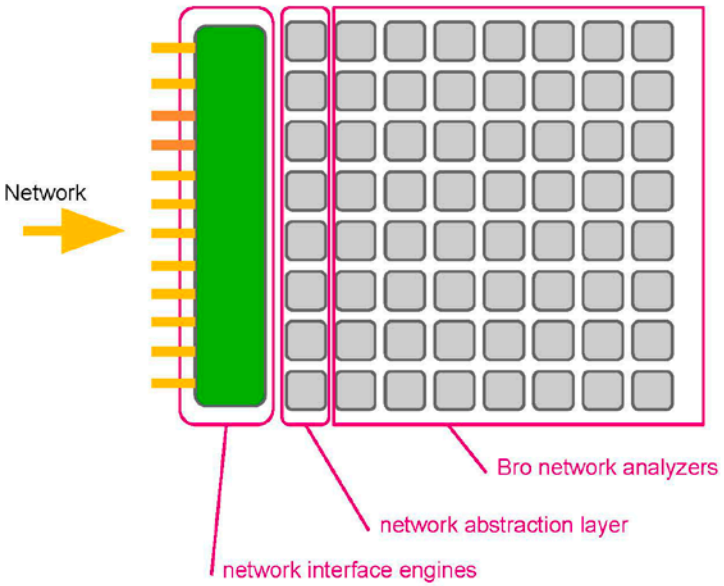

Looking back at 2013:
This leads to highly scalable network processing…

R-Scope DOMINATE-T events / sec

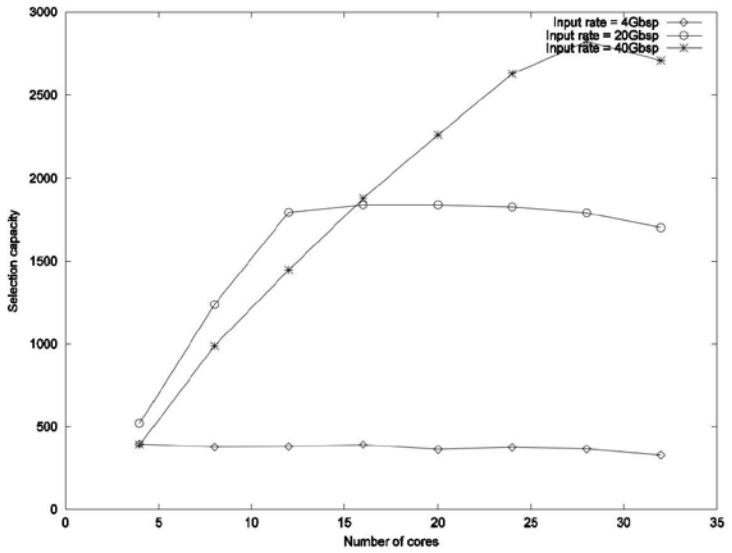

Single Tilera-Gx36
Synthetic traffic
Selection capacity events/sec

Indicates that a single Gx36 processor (36 cores) can handle 40 Gbps traffic at 2800 events/sec

Looking back at 2013:
The PACE-X model uses Reservoir packet processing innovations but on commodity Intel x86 processors with Myricom NIC, and hardens OS, adds Splunk connector

R-Scope-PACE-X

Security Appliance

- Validated HW/SW
- Signed updates
- CLI + GUI Management
- Splunk Integration
- Shunting Control

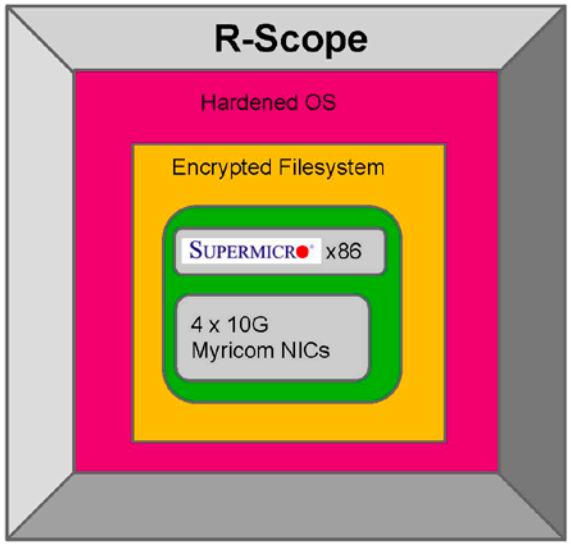

Looking back at 2013:
Leveraging National Lab resources - ESnet

DOMINATE-T at DOE

ESnet Research Testbeds

R-Scope DOMINATE-T
installed and in testing now at DOE Lab

100G Testbed

- High-speed protocol research
- Routed network from Argonne to NERSC (dedicated 100G)
- See recent HPCwire

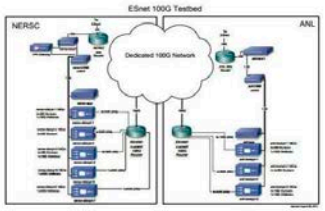

Dark Fiber Testbed

- Continental-scale fiber footprint for disruptive research

Lawrence Berkeley National Laboratory | U.S. Department of Energy | Office of Science

Looking back at 2013: Shipped the R-Scope DOMINATE-T to NERSC ESnet 100 Gbps testbed

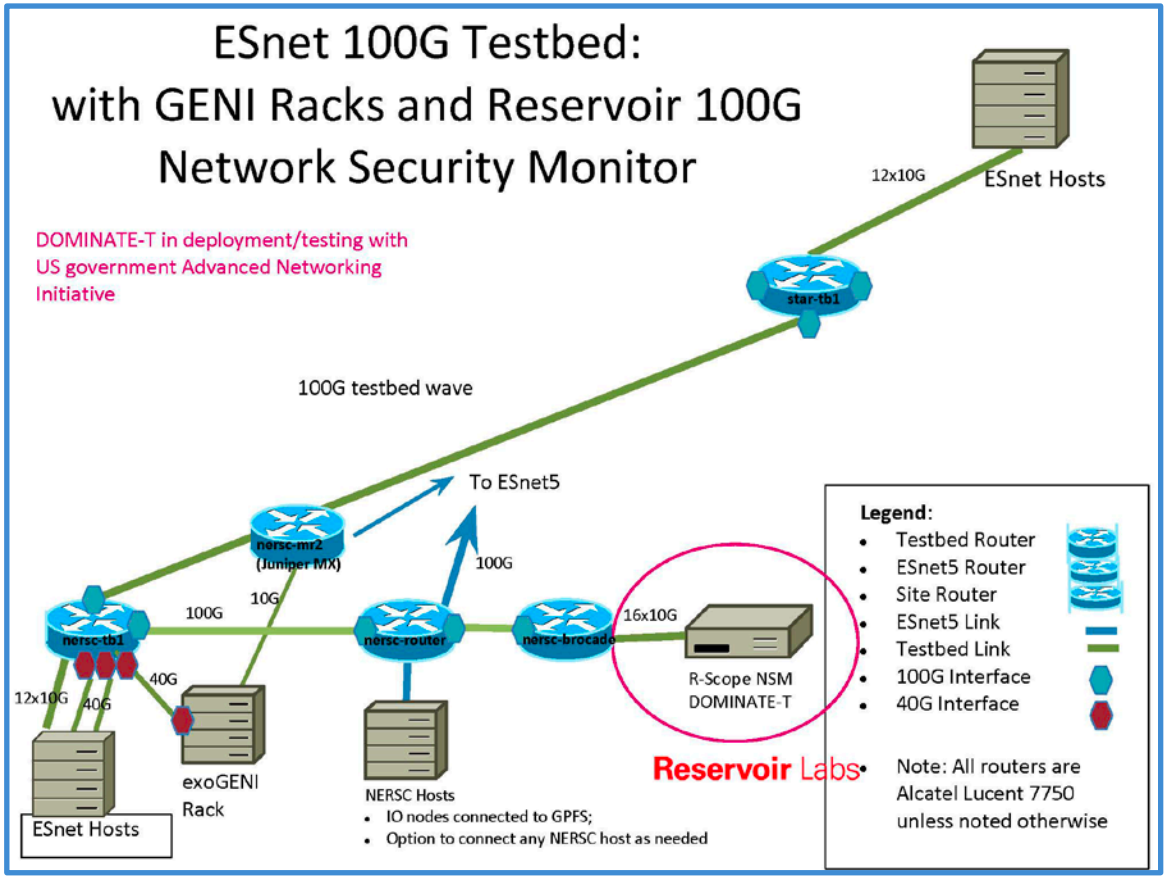

Selling *to* the National Labs

Even though we had DOMINATE inside ESnet, for testing, we didn't get any signals that ESnet was interested in buying.

It wasn't us. ESnet is in the national labs, and the labs aren't really interested in buying any SBIR technology.

To sell to the DOE requires

- Understanding and navigating complex procurement processes
- Leads through traditional selling (channel partners, integrators)

Lesson:

*The DOE SBIR program is oriented toward building business and transition **externally**. It is **not** for building technology to supply to DOE or its labs. This is a substantial difference from the DOD SBIR program which*

Supercomputing Conference Network: SCinet 2015-Present

ESnet ->

SCinet ->

Customers!

Lesson:

Leverage DOE resources, labs, PM relationships, peers, infrastructure. Network to gain access where needed. A DOE test environment is a stepping stone to customer environments and a possibly bright future.

SCinet

Supplier Risks

We experienced big supplier transitions

- Acquisitions, fissions, component cancellations

We found customer demand sweet spot is for PACE model

- Prospects value the resilience, hardening, integration
- A product is much more than performance

So, we focused on advancing PACE

- Increased our resilience further through move to new suppliers with deep support and deep resources, new features, …

Lesson:

Pay attention to supplier risk and engineer technology to be portable rapidly to new platforms. Be flexible (agile)!

When To Leave the SBIR Bubble

SBIR provides the opportunity to build technology without the pressures paying customers bring. This is a double-edged sword …

- Freedom to experiment and change
- Are you solving the right problem?

Once that first sale is made, there is no going back

- No time to fix foundational issues
- All hands on deck supporting customer needs

Lesson:

Take a hard look at readiness to determine the right time for initial sale. Think about "beta" to help set expectations with early customers.

Things Go Wrong

Lessons:

- *Assume your environment differs from your eventual customers' environments.*
- *Assume your test data is not representative of real customer data, and improve your tests constantly.*
- *Be prepared with tools to measure and characterize systems before engaging with potential customers.*
- *Be prepared with debugging and diagnostics in your system to determine what went wrong.*
- *Customers use your product in unanticipated ways and come to you with varying skill levels.*
- *Provide clear and pervasive mechanisms to determine how a customers environment/data differs from your research environment.*

The “Bedrock” Business Route

“Bedrock entrepreneurs describe 99.5 percent of all entrepreneurs who create more than 90 percent of all new wealth generated by entrepreneurs.”

Seasoned insight and reflection on the motivation, leadership skills, processes, and routes to success for bootstrapping a company.

Consider following the “Bedrock” route.

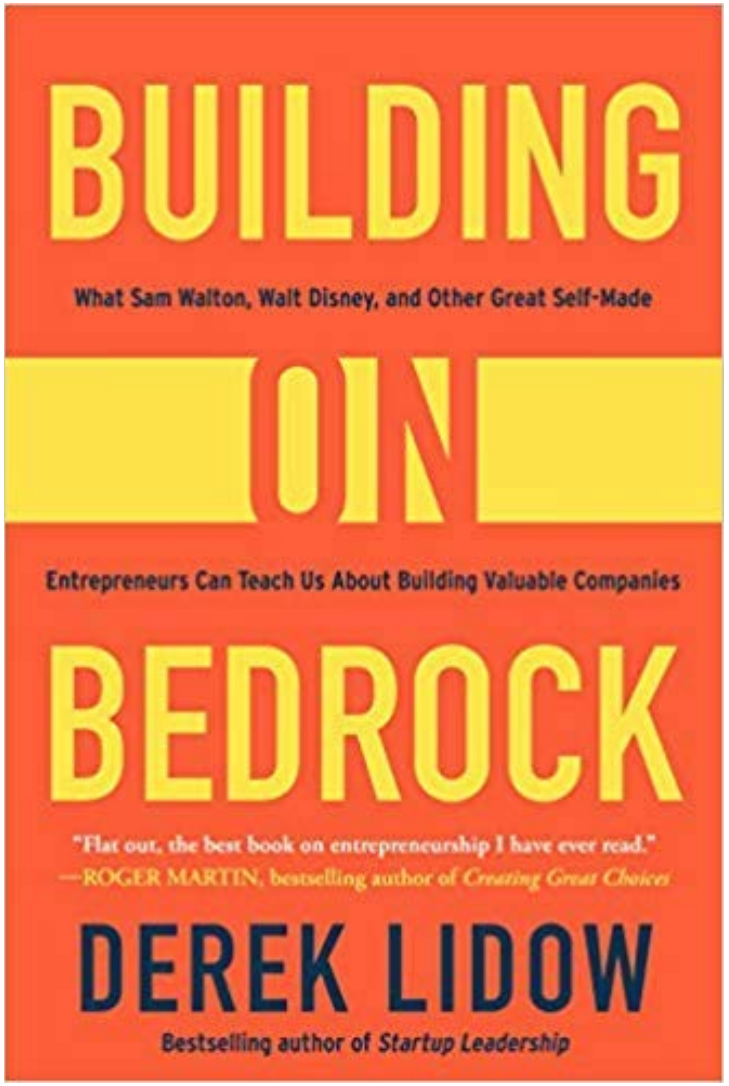

Understand the SBIR Program and Advocate

Get insights into the origins of the SBIR program in the 1980's, and the politics and pitfalls around it.

Understand "SBIR Data Rights" and how to use them.

Maintain communication with your congressional delegation to advocate for Small Business and the SBIR program.

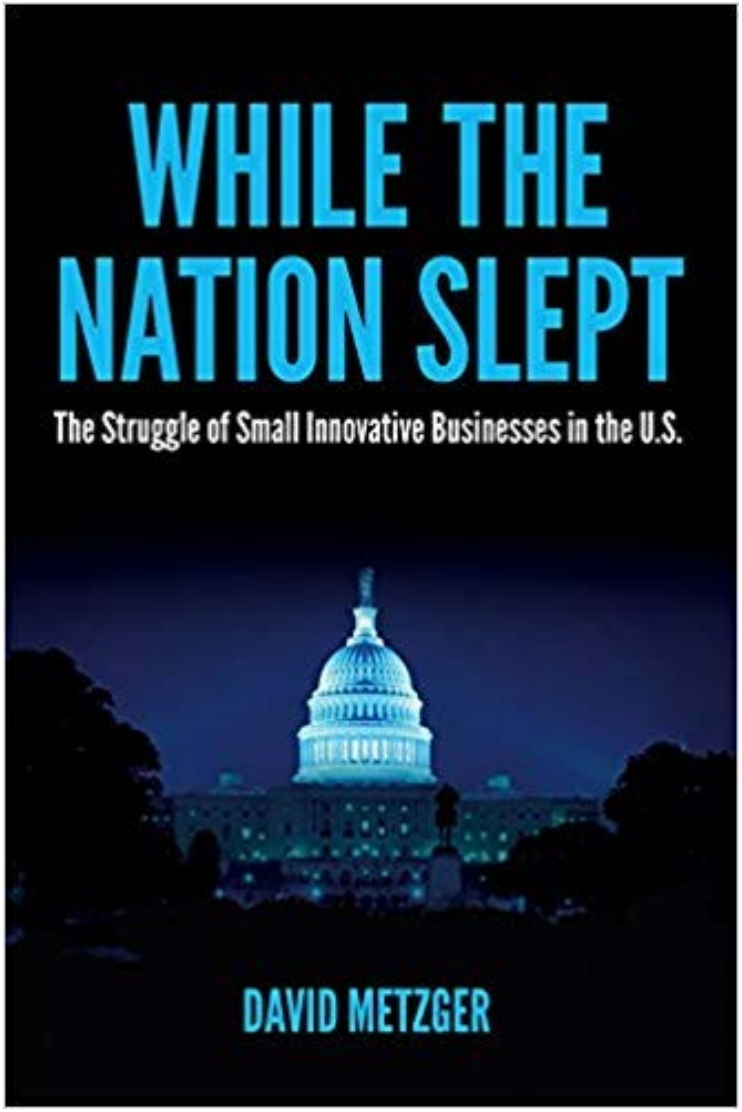

Bring Your DOE SBIR Project to Success Too

Be aware of and use the resources (lab facilties, commercialization assistance) that DOE brings, beyond funding.

Network intensely.

Be flexible and agile.

Exit the SBIR bubble gracefully.

Customers! Be ready.

Follow “Bedrock” principles.

Understand how “the system” works and help shape it.