

Scalable Approach to Tackle **Increasing Chip Complexity**

By **Anton Rozen**, Director of Backend Design, Mellanox Technologies, Tel Aviv, Israel

Increasing design complexity and multiphysics challenges hamper the productivity of system-on-chip (SoC) design teams. Engineers want electronic design automation tools that not only reduce runtime but also give them increased flexibility to critically examine and improve their designs. Mellanox engineers apply new solutions that leverage big data techniques and flexible computing resources to deliver this functionality.

High-speed networking is the backbone of connectivity in data centers. Extreme bandwidth and ultra-low-latency networking solutions are critical for the next era of data centers to efficiently process exponentially growing data from emerging AI, 5G and autonomous applications. Companies performing system-on-chip (SoC) designs for networking are challenged as chip size and complexity clash with ever-increasing time-to-market pressures. Grid complexity and the sheer number of gates increase dramatically each year, and network IC teams must design, analyze and tape out chips with dimensions of 400–500 mm or more.

Increased cross coupling of various multiphysics effects — including power and thermal reliability — pose significant challenges for FinFET design closure. Multiphysics analysis is critical to overcoming these challenges in order to design extremely large, complex and power-hungry chips, despite narrowing design margins and tighter project schedules.

Faced with this complexity, design teams must have software tools that deliver capacity, flexibility, speed and accuracy.

Mellanox, a leading supplier of end-to-end Ethernet and InfiniBand intelligent interconnect solutions and services for servers, storage and hyper-converged infrastructure, knows these challenges and trade-offs



ANSYS RedHawk-SC Introduction
[ansys.com/redhawk-intro](https://www.ansys.com/redhawk-intro)

Architecture (IR-drop methodology)	Design Complexity	IR-Drop Analysis Generation	Number of Machines	Max Node Count that Can Be Handled	Example
Monolithic	Up to 250 million nets	First 2003 onward	1	1B	ANSYS RedHawk and others
Distributed	Up to 350 million nets	Second 2013/14	32	4B	ANSYS RedHawk-DMP and others
Elastic	Greater than 350 million nets	Third 2016/18 onward	Scalable beyond 1,000 cores	Unlimited and scalable	ANSYS RedHawk-SC

As Mellanox has pushed designs into ultra-deep submicron nodes, the design complexity – and the need for more flexible and scalable design tool solutions – has increased.

firsthand. The design teams must manage and validate designs by making the most efficient use of computing resources and engineering time. To this end, the team relied on ANSYS RedHawk-SC software.

LOOKING FOR VISIBILITY

The Mellanox team needed fast turnaround time with pinpoint voltage drop accuracy to ensure power integrity and reliability for their highly complex network processors. But they also sought something that had eluded them in earlier years on other big, high-complexity designs: flexibility and speed of analysis. Because designs have evolved from a little more than 100 million nets at the 45nm node to nearly 350 million nets at 16nm, Mellanox estimates it will need to address nearly 450 million IC nets at 7nm.

This type of evolution requires tool capacity to match. A decade ago, in and around the 45nm process node, tool architectures were generally monolithic, and teams were restricted to a single machine that could handle up to 1 billion power and ground nodes at once. (A node is a connection point between any two elements in the power and ground network that are extracted. These elements could be parasitic resistance, inductance or capacitance

of the wire or device instance pin connected to the wire. Node count is a metric commonly used in power integrity analysis to predict the design size; it directly impacts the runtime and memory requirements for the analysis.)

In those days, tool capacity was an issue. When conducting multiple analyses for power integrity and reliability signoff, each run (in serial rather than parallel) might take more than 24 hours. This required large servers and considerable resource allocation to complete the analysis. Worse, the system occasionally had trouble managing the complexity and would crash. The analysis then would have to be restarted from scratch.

A second generation emerged to keep up with complexity. This generation leveraged distributed compute, could scale to up to 32 machines and could handle a maximum of 4 billion nodes. This was satisfactory until ICs became even more complex.

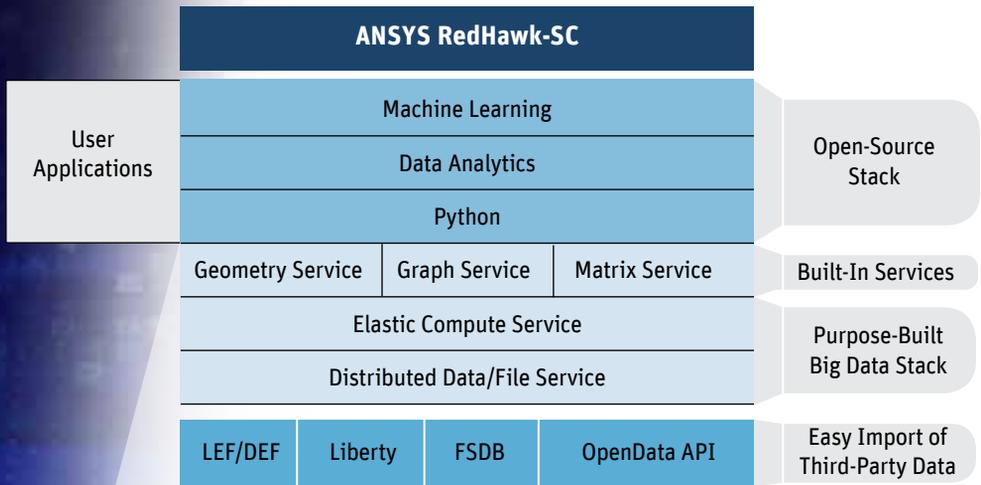
SCALING TO BIG DATA REQUIREMENTS

To deliver insights and enable the team to optimize its design, Mellanox needed a flexible, high-capacity solution that would scale for big data mining and analytics. Engineers began using ANSYS RedHawk-SC in 2018. RedHawk-SC is the latest SoC power integrity and reliability signoff platform built on ANSYS SeaScape – the world’s first custom-designed big data architecture for electronic system design and simulation. SeaScape provides per-core scalability, flexible design data access, instantaneous design bring-up and many other capabilities.

	Monolithic	Distributed	Full-Chip Elastic Computing
Technology	28nm	16nm	16nm
Chip size (die area, inst. count, no. of gates)	1/4" Full Chip 96M nets	Full Chip 225M nets	Full Chip 340M nets
CPU core usage/ Machines	1 machine of 1 TB	4 machines of 1.4 TB	150 works of 72 GB
Runtime	60 hours	72 hours	24 hours

SCALABILITY COMPARISON

◀ How the evolution of software has drastically reduced runtime for increasingly complex SoCs



ANSYS SeaScape big data elastic compute architecture

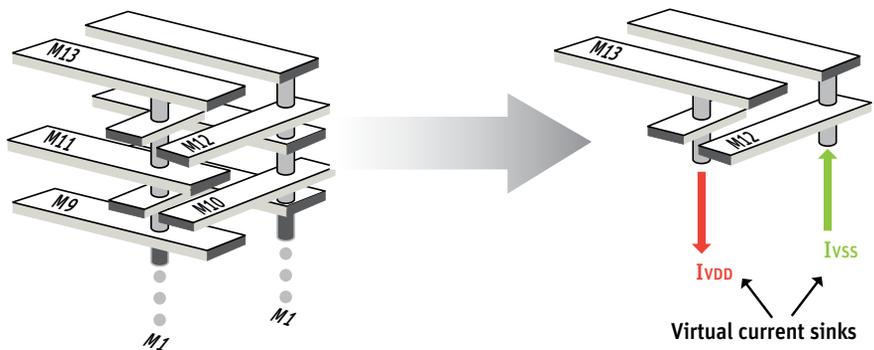
One of the keys to success lies in the elastic computing capabilities of RedHawk-SC. Elastic computing helps to process scenarios in parallel (or in serial), depending on the number of CPU cores available.

The SeaScape architecture is central to elastic computing. It rests on a distributed data/file service since data may be scattered around many locations. On top of this sits a distributed data analytics layer based on the MapReduce concept, which is fundamental to all big data analytics. This conceptually splits the data (mapping) into small chunks called shards and farms each shard for analysis. Processing can be distributed to servers as they become available, across as many servers as needed.

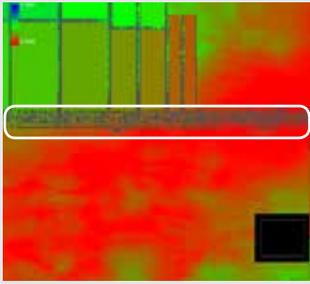
THE POWER PROBLEM

The challenge in these types of network processors is total power consumption and power dissipation. Unlike battery-powered designs, the types of designs that Mellanox works with can consume more than 200 W. So, engineers must achieve complete design analysis – accurate incremental power integrity and reliability analysis – while considering high power consumption without sacrificing accuracy or time to results.

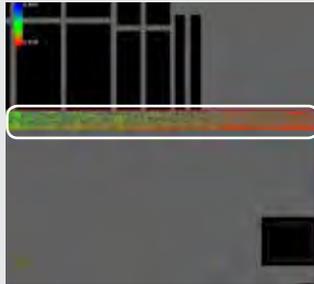
To speed up full-chip IR drop simulations, power grid roll-up methodology can be leveraged to abstract the low- and mid-level metals of the power and ground network. Such abstraction can be used in full-chip simulations. This allows teams to work at the unit level and then jump up to the top level for a comprehensive analysis of the full-chip design.



An example of a power integrity simulation using ANSYS RedHawk-SC roll-up methodology for abstracting low- and mid-level metal layers of the power grid for fast, incremental full-chip analysis

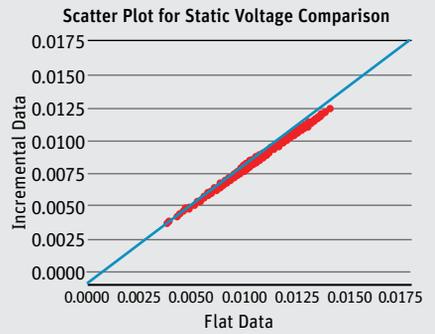


Top-level flat run

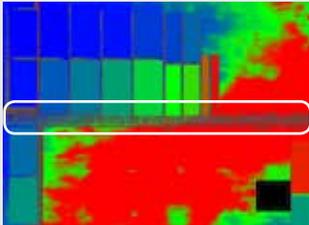


Top-level incremental run

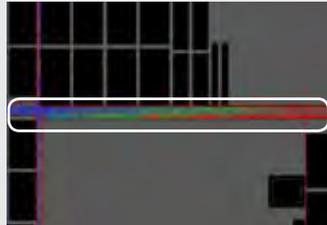
Static voltage drop at the interface regions showing very good correlation



Full chip flat vs. incremental static voltage drop analysis shows very good correlation.

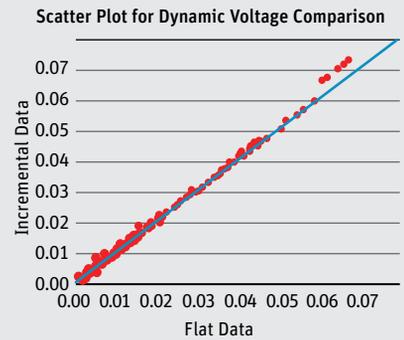


Top-level flat run

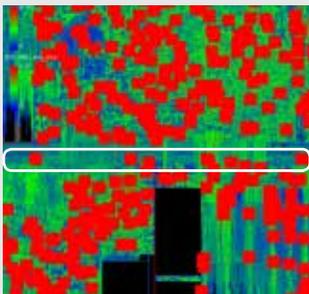


Top-level incremental run

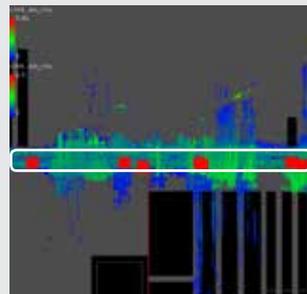
Dynamic voltage drop at the interface regions showing very good correlation



Full chip flat vs. incremental dynamic voltage drop analysis shows very good correlation.

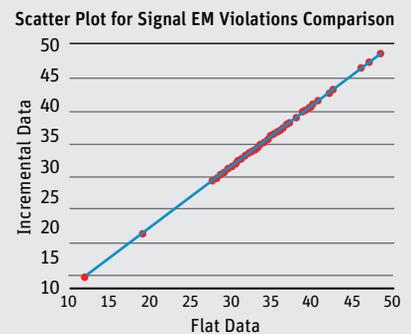


Top-level flat run



Top-level incremental run

Signal EM violations at the interface regions showing very good correlation



Full chip flat vs. incremental signal EM analysis shows very good correlation for all top-level signal nets.

Mellanox's experience reinforced the compelling features of ANSYS RedHawk-SC, including:

- 1. CAPACITY** – Ability to run large designs and be able to query and visualize them in the GUI smoothly. With RedHawk-SC's elastic scalability, full-chip analysis was completed within 24 hours.
- 2. FLEXIBILITY** – RedHawk-SC manages computing resources flexibly, setting a new level of effectiveness for EDA tool resources in the industry.
- 3. SPEED AND ACCURACY** – Big data analytics techniques enable faster top-level runs, shorter ECO loops and faster ECO fixes with accuracy. With RedHawk-SC and actionable analytics, the productivity increased significantly by a factor of three because teams can parallelize runs and understand the impact of voltage drop in the blocks in the full chip context.

Doing a full-chip flat run is resource-intensive and time-consuming. By performing incremental analysis enabled by techniques using big data analytics, designers can create a detailed view of a specific block and abstract everything else. This enables them to perform faster analysis and conduct quicker engineering change order (ECO) fixes more easily with visibility.

ANSYS RedHawk-SC, with its elastic computing capabilities and big data-enabled analytics, gave engineers the visibility they needed to overcome some previous challenges. The team particularly appreciated RedHawk-SC's self-sustaining stability to monitor its own jobs and to renew the job if it fails.

The team also leveraged RedHawk-SC's elastic computing and its MapReduce-enabled analytics to gain key insights. MapReduce gives designers a bird's-eye view and zeros in on hotspots very smoothly. It provides powerful capabilities such as bringing up the GUI to view a full chip database in less than two minutes and navigating different areas easily, like Google Maps' functionality.

Additionally, it enables vastly more powerful compute flexibility. With RedHawk-SC's elastic scalability, large chip areas that once required huge computing resources can be broken into very small pieces for analysis. The nature of the architecture lets those elements be distributed through a company's computing resources. In this way, it maximizes hardware resource utilization and optimizes cost.

TACKLING CHIP COMPLEXITY

With the soaring complexity of networking IC designs, a new approach to full-chip power integrity and reliability signoff is required. This means leveraging highly parallel computing concepts to analyze large chunks of data to drive improvements in visibility, time to results, and overall design productivity and efficiency.

ANSYS RedHawk-SC's elastic compute scalability and big data analytics techniques for full chip power integrity analysis allowed Mellanox to run large designs with production-proven accuracy in less than 24 hours with pinpoint accuracy. Combining incremental power integrity/reliability and signal line electromigration analysis helped increase productivity by a factor of three. 📌

RESOURCE

Cohen, R.; Rozen, A.; Abhijith, M.V.; Agarwal, R.; Ramachandran, S.; Johnson, S. "Fast and Accurate Incremental Power and Signal Integrity Analysis." www2.dac.com/56th/proceedings/posters/125_3.pdf (08/01/2019)