

# EPIC: An Energy-Efficient, High-Performance GPGPU Computing Research Infrastructure

Magnus Sjalander, Magnus Jahre, Gunnar Tufte, and Nico Reissmann  
 Norwegian University of Science and Technology (NTNU)  
 firstname.lastname@ntnu.no

**Abstract**—The pursuit of many research questions requires massive computational resources. State-of-the-art research in physical processes using simulations, the training of neural networks for deep learning, or the analysis of big data are all dependent on the availability of sufficient and performant computational resources. For such research, access to a high-performance computing infrastructure is indispensable.

Many scientific workloads from such research domains are inherently parallel and can benefit from the data-parallel architecture of general purpose graphics processing units (GPGPUs). However, GPGPU resources are scarce at Norway’s national infrastructure.

EPIC is a GPGPU enabled computing research infrastructure at NTNU. It enables NTNU’s researchers to perform experiments that otherwise would be impossible, as time-to-solution would simply take too long.

## I. INTRODUCTION

The end of Dennard’s scaling left computing systems across all domains increasingly power constrained. Specialized hardware in the form of accelerators emerged as alternatives to perform computations more energy-efficient. Specifically, general-purpose, graphic-processing units (GPGPUs) became increasingly popular as a means to accelerate programs in high-performance computing (HPC) and artificial intelligence.

GPGPUs devote more compute resources to accelerate data-parallel applications by sacrificing resources that improve sequential program performance, rendering them more energy-efficient for data-parallel application domains. Nowadays, GPGPUs are significantly employed in High-Performance Computing (HPC) systems to meet performance demands while maintaining power constraints. For example, five of the ten most powerful supercomputers in the world rely on GPGPUs for their computational power [20]. Furthermore, eight of the top ten most energy-efficient supercomputers in the world rely on GPGPUs [4].

The EPIC research infrastructure is a project between the Department of Computer Science and the IT Division at the Norwegian University of Science and Technology (NTNU) that aims at providing a GPGPU compute platform. EPIC is a part of the NTNU Idun computing cluster [6], which provides a high-availability and professionally administrated compute platform for NTNU. Idun combines compute resources of individual shareholders to create a cluster for rapid testing and prototyping of HPC software. Currently, EPIC constitutes 48% of the total number of nodes in the IDUN cluster and 100% of the GPGPU resources.

EPIC is with its 90 GPGPUs one of Norway’s largest GPGPU enabled computational infrastructures. Norwegian national infrastructure has a very limited number of GPGPU resources, e.g., Saga [19] has only 32 NVIDIA Tesla P100 [15] and Abel [18] has 32 much older NVIDIA Tesla K20.

## II. THE IDUN CLUSTER

The Idun cluster is a Tier-2 [5] research cluster at NTNU meant as a stepping stone for the national infrastructure and serves as a platform for rapid testing and prototyping of HPC software, research into energy-efficient computing, and GPU-aided simulations and design-space exploration.

Currently, Idun consists of 73 nodes connected by two networks: one ethernet network and one high-throughput and low-latency infiniband (IB) network. The 1 Gb/s ethernet network serves as an administration and provisioning network, while the IB network is used for inter-node communication. The IB network is a mix of FDR (4x lanes each of 14 Gb/s) and EDR (4x lanes each of 25 Gb/s), as shown in Figure 1. Each node is connected with either FDR or EDR, resulting in 56 Gb/s or 100 Gb/s per node, respectively. The individual IB switches are connected in a tree structure with 3xFDR links between each switch, resulting in 168 Gb/s inter-switch connection speed.

Idun’s storage is provided by two storage arrays and a Lustre parallel distributed file system [12]. The storage arrays, one serves as Lustre metadata target (MDT) and one as Lustre object storage target (OST), are complemented with two Lustre metadata servers (MDS) and two object storage servers (OSS). The MDT and MDSs store the namespace data of the file system, such as filenames, directories, access permissions and file layouts, while the OST and OSSs store the file data. Together, the IB network and the Lustre file system, provide the means to efficiently transfer data to the compute resources, enabling an effortless scaling of the cluster in terms of nodes and/or GPUs.

## III. THE EPIC RESEARCH INFRASTRUCTURE

The EPIC research infrastructure consists of four distinct investments (see Table I), each with a distinct purpose:

The original **EPIC1** consists of eight nodes with two NVIDIA P100 GPUs and three N4L PPA 3560 power meters [1]. These nodes are used for energy-efficient computing research such as energy efficient resource management for latency-critical cloud services [14].

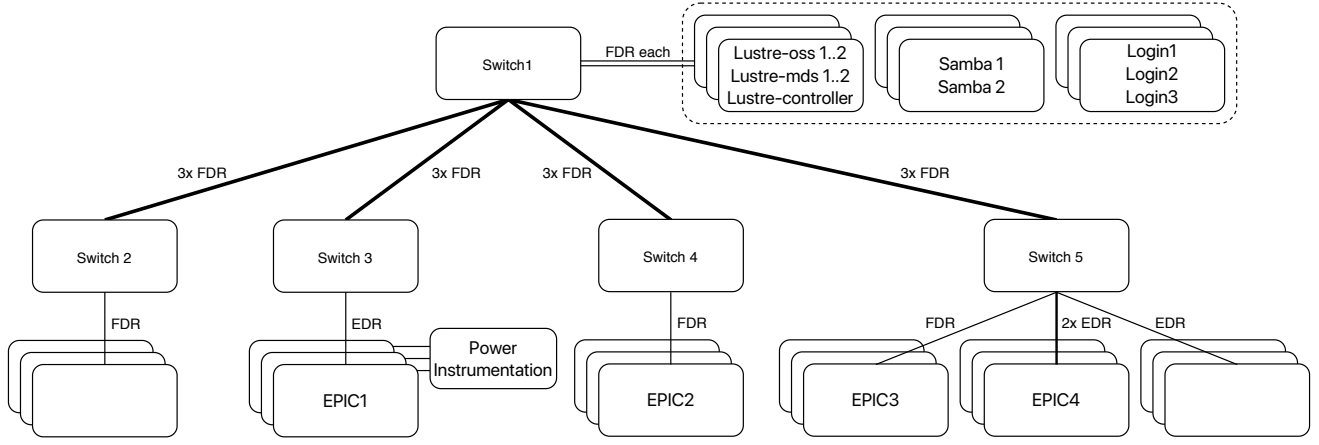


Fig. 1. The topology of the Idun with the EPIC research infrastructure.

TABLE I  
EPIC CONFIGURATION

Name	#Nodes	Machine Model	#CPUs	Processor model	#Cores	Memory	#GPUs	GPU model
EPIC1	8	Dell PE730	2	Intel Xeon E5-2695 v4 [8]	36	128 GiB	2	NVIDIA Tesla P100 16 GiB [15]
EPIC2	19	Dell PE730	2	Intel Xeon E5-2650 v4 [7]	24	128 GiB	2	NVIDIA Tesla P100 16 GiB [15]
EPIC3	5	Dell PE740	2	Intel Xeon Gold 6132 [9]	28	768 GiB	2	NVIDIA Tesla V100 16 GiB [16]
EPIC4	2	Dell DSS8440	2	Intel Xeon Gold 6148 [10]	20	768 GiB	8	NVIDIA Tesla V100 32 GiB [16]
	1	Dell DSS8440	2	Intel Xeon Gold 6148 [10]	20	768 GiB	10	NVIDIA Tesla V100 32 GiB [16]

**EPIC2** consists of 19 GPGPU nodes, each equipped with two NVIDIA P100 GPUs. These nodes complement EPIC1 and provide raw computational GPU power. These nodes are used for research in 3D object identification, physical simulations (e.g., nanomagnet ensemble dynamics modeled in MuMAX [22]), and deep learning.

**EPIC3** consists of five big-memory nodes, each equipped with two NVIDIA V100 GPUs. These nodes are meant for AI research that requires massive training sets, and therefore need more main memory.

**EPIC4** is an extension of EPIC2 providing another 26 GPUs for raw computational power. It consists of one node with ten V100 32 GiB GPUs and two nodes with eight V100 32 GiB GPUs. In addition, the big-memory GPUs (32 GiB instead of 16 GiB) enable larger working set sizes beneficial for 3D object identification and large AI models.

Even though the the purpose and configuration of EPIC1-4 differ, all 90 GPGPUs can be accessed as one distributed resource for massive GPGPU performance.

#### IV. RESEARCH OUTCOME

The EPIC cluster has been an indispensable resource for a wide range of research, e.g., efficient resource management, nanomagnetic modeling, 3D object identification, etc. Below is a non-exhaustive list of published articles that relied on EPIC to produce their results:

- Energy efficient resource management for latency-critical cloud services [14].
- Emergent computation on magnetic ensembles [11], [13].
- Bit-serial matrix multiplication acceleration [21].
- Intermediate representation (IR) for optimizing compilers [17].
- Nano-scale structures of aluminum alloys [2], [3].

#### V. CONCLUSION

EPIC is a multi-million investment by the Department of Computer Science in collaboration with the IT Division to provide GPGPU resources for NTNU’s researchers. The large number of GPGPUs enable research studies to be performed at a scale that otherwise would be impossible to conduct. Thus, EPIC’s computational resources help NTNU’s researchers to stay competitive and produce state-of-the-art results.

#### ACKNOWLEDGEMENT

The original EPIC1 was financed by an NTNU Advanced Research Equipment grant (AVIT) and EPIC3 was financed by the NTNU-Telenor AI Lab (now called the Norwegian Open AI Lab). EPIC2 and EPIC4 were financed by the Department of Computer Science. The network, storage, and maintenance of the cluster is provided by the HPC group of NTNU’s IT Division.

#### REFERENCES

- [1] *PPA3500Series 16 Phase Power Analyzer*, document ref: 527-001/1.
- [2] E. Christiansen, “Nanoscale characterisation of deformed aluminium alloys,” 2019.
- [3] E. Christiansen, C. D. Marioara, B. Holmedal, O. S. Hopperstad, and R. Holmestad, “Nano-scale characterisation of sheared  $\beta$ ” precipitates in a deformed al-mg-si alloy,” *Scientific reports*, vol. 9, no. 1, pp. 1–11, 2019.
- [4] “The GREEN 500,” <https://www.top500.org/green500/lists/2019/11/>, Nov. 2019.
- [5] M. Guest, “The Scientific Case for High Performance Computing in Europe 2012 - 2020,” Tech. Rep., 2012.
- [6] “NTNU Idun computing cluster,” <https://www.hpc.ntnu.no/display/hpc/Idun+Cluster>.
- [7] “Intel® Xeon® processor E5-2650 v4,” <https://ark.intel.com/content/www/us/en/ark/products/91767/intel-xeon-processor-e5-2650-v4-30m-cache-2-20-ghz.html>.

- [8] “Intel<sup>®</sup> Xeon<sup>®</sup> processor E5-2695 v4,”  
<https://ark.intel.com/content/www/us/en/ark/products/91316/intel-xeon-processor-e5-2695-v4-45m-cache-2-10-ghz.html>.
- [9] “Intel<sup>®</sup> Xeon<sup>®</sup> processor Gold 6132,”  
<https://ark.intel.com/content/www/us/en/ark/products/123541/intel-xeon-gold-6132-processor-19-25m-cache-2-60-ghz.html>.
- [10] “Intel<sup>®</sup> Xeon<sup>®</sup> processor Gold 6148,”  
<https://ark.intel.com/content/www/us/en/ark/products/120489/intel-xeon-gold-6148-processor-27-5m-cache-2-40-ghz.html>.
- [11] J. H. Jensen, E. Folven, and G. Tufte, “Computation in artificial spin ice,” in *Artificial Life Conference Proceedings*. MIT Press, 2018, pp. 15–22.
- [12] “Lustre<sup>\*</sup> software release 2.x, operations manual,”  
[http://doc.lustre.org/lustre\\_manual.pdf](http://doc.lustre.org/lustre_manual.pdf).
- [13] O. R. S. Lykkebø, J. H. Jensen, A. G. Penty, A. Strømberg, M. Sjölander, E. Folven, and G. Tufte, “Emergent computation on a magnetic lattice,” in *International Workshop on Theoretical and Experimental Material Computing*.
- [14] R. Nishtala, V. Petrucci, P. Carpenter, and M. Sjölander, “Twig: Multi-agent task management for colocated latency-critical cloud services,” in *Proceedings of the International Symposium High-Performance Computer Architecture*. New York, NY, USA: ACM, Feb. 2020.
- [15] “NVIDIA<sup>®</sup> TESLA<sup>®</sup> P100 GPU accelerator data sheet,”  
<https://images.nvidia.com/content/tesla/pdf/nvidia-tesla-p100-PCIe-datasheet.pdf>, Oct. 2016.
- [16] “NVIDIA<sup>®</sup> TESLA<sup>®</sup> V100 GPU accelerator data sheet,”  
<https://images.nvidia.com/content/technologies/volta/pdf/tesla-volta-v100-datasheet-letter-fnl-web.pdf>, Mar. 2018.
- [17] N. Reissmann, J. C. Meyer, H. Bahmann, and M. Sjölander, “RVSDG: An intermediate representation for optimizing compilers,” *arXiv preprint arXiv:1912.05036*, 2019.
- [18] “Abel GPU support,”  
<https://www.uio.no/english/services/it/research/hpc/abel/help/software/gpu-nvidia-cuda.html>.
- [19] “Saga,”  
<https://documentation.sigma2.no/quick/saga.html>.
- [20] “TOP 500 the list.”  
<https://www.top500.org/lists/2019/11/>, Nov. 2019.
- [21] Y. Umuroglu, D. Conficconi, L. Rasnayake, T. B. Preusser, and M. Sjölander, “Optimizing bit-serial matrix multiplication for reconfigurable computing,” *ACM Transactions on Reconfigurable Technology Systems*, vol. 12, no. 3, pp. 15:1–15:24, Aug. 2019. [Online]. Available: <http://doi.acm.org/10.1145/3337929>
- [22] A. Vansteenkiste, J. Leliaert, M. Dvornik, M. Helsen, F. Garcia-Sanchez, and B. Van Waeyenberge, “The design and verification of MuMax3,” *AIP advances*, vol. 4, no. 10, p. 107133, 2014.