



DELL EMC POWER EDGE® R940 MAKES *DE NOVO* ASSEMBLY EASIER

Genome Assembly on Deep Sequencing data with SOAPdenovo2

ABSTRACT

De novo assemblies are memory intensive since the assembly algorithms need to compare every read with every other read imposing a time complexity of $O(n^2)$. The memory requirement has been growing as more sequence reads are being generated than ever before, so a big memory machine like the Dell EMC PowerEdge® R940 is required to assemble deep sequencing data. Here we introduce R940 with Intel® Xeon® Scalable Processor, which reduces the running time by 27% and consumes 39.5% less memory on SOAPdenovo2 to assemble 3.2 billion reads, in comparison to its predecessor, R930.

July, 2017

TABLE OF CONTENTS

EXECUTIVE SUMMARY	3
AUDIENCE.....	3
INTRODUCTION	4
<i>De novo</i> assembly?	4
Memory capacity limitations on a shared memory architecture	5
PowerEdge R930 vs R940	5
TEST CONFIGURATIONS AND DATA.....	6
PERFORMANCE WITH SOAPDENOV2	6
Running Time	6
Peak Memory Consumption	8
CONCLUSION	9
REFERENCES.....	10

EXECUTIVE SUMMARY

In the last decade, modern society has continued to improve the quality of life through better healthcare, producing and consuming sustainable food and energy, and protecting our environment. All of these societal advancements are tightly related to progress in the life science domain. Bioinformatics emerged from the massive amount of data that is now available in these fields and the advancement of computer technology takes a key role to finding solutions. In bioinformatics, sequence assembly refers to reconstructing the original genome sequence from a massive number of DNA fragments. Sequence assembly is used to generate a reference genome sequence and also enables comparative, functional and evolutionary analyses on a genome sequence. However, obtaining the right data (higher accuracy and low error rate) and resolving ambiguity are two main concerns in *de novo* assembly. These problems require *de novo* assembly to be performed multiple times with different algorithms; hence a big memory machine to run assembly algorithms efficiently is an important part of the solution. Dell EMC PowerEdge® R940 is a better server to perform *de novo* assembly faster and more efficiently than its predecessor.

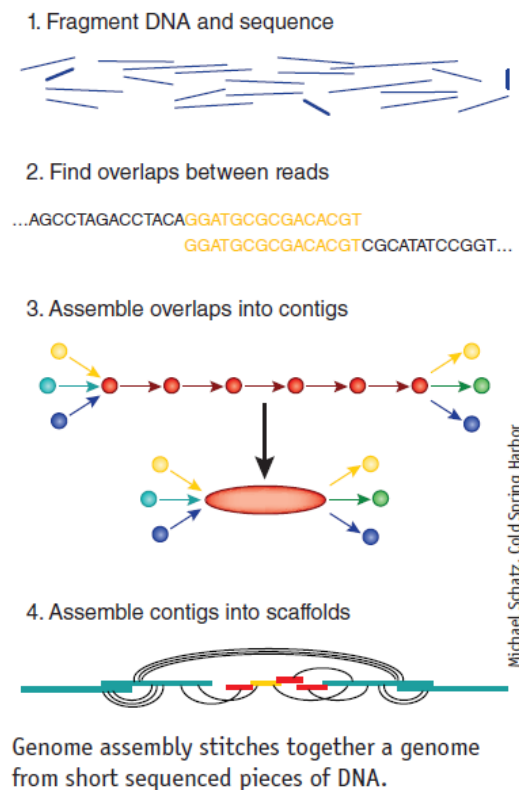
AUDIENCE

This white paper is intended for providing valuable insights to consider Dell EMC PowerEdge® R940 in the context of Life Science applications. System administrators, solution architects, researchers, and others who are involved in a decision making for a HPC system constitute the target audience.

INTRODUCTION

De novo assembly?

Genome assembly programs perform complex computations to join billions of fragmented reads of DNA. Short-read sequencing (Next Generation Sequencing (NGS)) technologies turn the haploid human genome into 2- to 3-billion pieces of jigsaw puzzle with 100 copies of each piece (1). This puzzle is complicated further by some pieces being missing or containing errors. These characteristics of sequencing instruments impose the algorithmic challenges in assembling large, complex genomes with large amount of repetitive regions. In order to compensate for these shortcomings in an easier way, much higher number of read depths¹ than 20 or more is intentionally generated. Deep and Ultra-deep generally refer over 100 reads per nucleotide to even 1000 reads. Although NGS technology reduced the cost of sequencing significantly, the read length becomes much shorter than the traditional Sanger sequencing, and shorter read length makes *de novo* assembly much harder (2) (3). Since every read has to be compared to every other read, the imposed time complexity is bounded to $O(n^2)$ where n represents the number of reads as shown in Figure 1 (3).



Genome assembly stitches together a genome from short sequenced pieces of DNA.

Figure 1 *De novo* assembly process: Baker, 2012

Besides nondeterministic polynomial time complexity², the memory requirement also grows exponentially with the number of reads. As sequence read data is represented as 'char' in a typical programming language, 3 billion reads with length of 100 base pairs (bp) can be translated to 300 TB of memory. With clever ways of represent these data in memory as well as avoiding redundancies, memory requirement can be minimized, for example, by using a Bruijn graph (4). However, due to the massive volume of data, *de novo* assembly still requires extremely large amounts of memory, nearly 2 TB of memory for 3 billion reads. Hence, it is logical to favor a shared memory architecture to assemble a genome.

¹ Coverage (or depth) in DNA sequencing is the number of reads that include a given nucleotide in the reconstructed sequence. Deep sequencing refers to the general concept of aiming for high number of replicate reads of each region of a sequence.

² Polynomial time is said to be an algorithm can be solved in polynomial time if the number of steps required to complete the algorithm for a given input is for some nonnegative integer, where is the complexity of the input. Nondeterministic polynomial time is a subset of polynomial time problems with some form of nondeterminism or "trial and error".

Memory capacity limitations on a shared memory architecture

Modern CPU micro-architecture supports multiple independent channels to the DIMMs through integrated memory controllers (IMCs) (5). Both Intel® Haswell and Broadwell support quadruple memory channels with two IMCs per processor (6) (7). These multiple independent channels enable concurrent access to multiple DIMMs and distributes the data amongst the DIMMs providing more bandwidth for accessing the data. However, when more DIMMs are added to the channel, the throughput per DIMM may decrease due to the increment of management overhead (8). The available bandwidth for read and write data may be reduced. While populating the channels with 2 DIMMs per channel (2 DPC) will not affect throughput dramatically—most systems allow to 2 DPC to run at native speed—populating 3 DIMMs per channel (3 DPC) or more will be an optimization challenge between obtaining the required capacity by using higher capacity DIMMs or taking maximum throughput hit by populating 3 DPC.

PowerEdge R930 vs R940

The R930 is designed to maximize memory capacity. This four socket server with 96 DIMM slots in 4U can support up to 12 TB of RAM with 128GB LRDIMM and Intel® Xeon E7 v4 or v3 CPU family. In order to increase memory capacity, each IMC is connected to 6 DIMMs via a Jordan Creek memory buffer. Two Jordan Creek memory buffers are implemented on a riser to support two memory channels. In this architecture, one IMC needs to control 12 DIMMs as shown in Figure 2 (9). Overall, each processor connects to 24 DIMMs, and it is considered as 3 DPC when DIMM slots are fully populated since each Jordan Creek enables two sub-memory channels. Although the design behaves as if 3 DPC is supported when actually 12 DIMMs assigned to one IMC, there is trade-off between memory capacity and memory performance. 16/32GB RDIMM, DDR4, dual rank which runs at 2400MHz natively slows down to 1600MHz with even 1 DPC configuration. It slows down even more to 1333MHz with 3 DPC configuration (9).

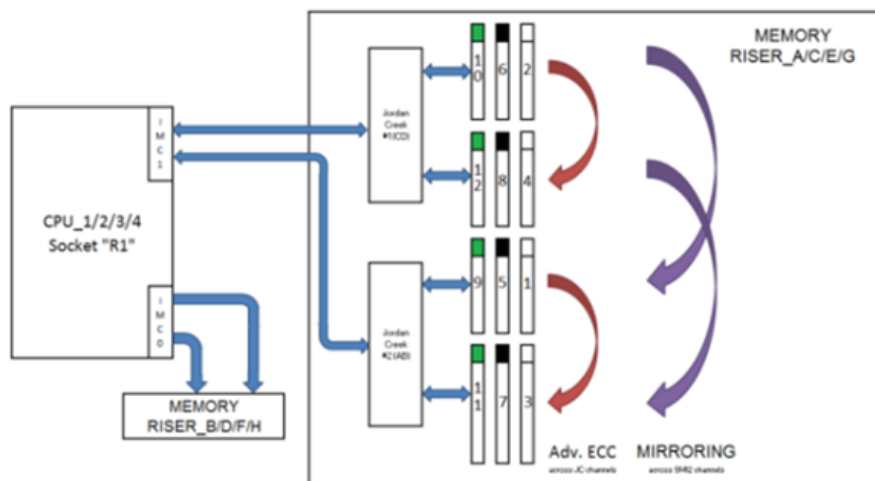


Figure 2 Memory configuration in PowerEdge R930

The Dell EMC 14th Generation Server, PowerEdge R940, is designed to support the new Intel® Xeon® Processor Scalable Family (micro-architecture also known as "Skylake-SP") providing three memory channels per IMC and a total of six channels per processor. This four socket server packed in 3U chassis supports up to 112 cores³, delivers 50% more QPI bandwidth compared to a two socket server. It supports 48 DIMMs with speed of up to 2667 MT/s, 6 memory channels per socket, and up to 6TB⁴ of memory without losing DIMM's native clock speed in contrast to PowerEdge R930. This is a big advantage for those CPU/memory intensive applications to handle memory consumption more efficiently.

Since sequencing technology changes so quickly from both NGS and computer technology, it is hard to assess assembly programs fairly. It is like tackling a new problem every time the data changes. Hence, the purpose of this study is to show PowerEdge® R940, 14th Generation server in comparison with its predecessor, it is not intended to report the performance of SOAPdenovo2 used in this study.

³ Up to four Intel® Xeon® Scalable processors, up to 28 cores

⁴ Supports 50% more NVMe drives than the R930, 48 DDR4 DIMM slots, Supports RDIMM /LRDIMM, Up to 12 NVDIMM, 384 GB Max, Supports registered ECC DDR4 DIMMs only

Test Configurations and Data

Since the main interest of this study is to make a reasonable observation for memory consumption behavior between two systems, all DIMM slots in both systems were fully populated to match the amount of memory to 1.5 TB. Although other components can influence the benchmark results like storage configuration, those influences should be marginal since the performance of *de novo* assembly applications is largely bounded by CPU/memory. The details of test configurations are listed in Table 1.

Table 1 Test configurations

	PowerEdge R940	PowerEdge R930
CPUs	4x Intel® Xeon® Platinum 8180M	4x Intel® Xeon® CPU E7-8890 v3
Base Frequency (GHz)	2.5	
# of Cores	28, total of 112	16, total of 64
TDP (W)	205	165
Memory	48x 32GB DDR4-2666MHz, total of 1.5 TB	96x 16GB DDR4-2400MHz, total of 1.5 TB
Storage	12x 1.9TB 12 Gps SSDs in RAID 0	Boot disk is 300GB RAID 1 with 1 spare (total of 3 disks), 3 TB RAID 0 based on 2x Intel SSDSC2BX01 SSD 1.5TB
Kernel	3.10.0-514.el7.x86_64	
Operating System	Red Hat Enterprise Linux Server release 7.3 (Maipo)	

The choice of test data is a whole human genome sequencing data set, [ERR318658](#) containing 3.2 billion reads (10). The number of reads in this whole genome sequencing data belongs to the range of deep sequencing. For the benchmark purpose, smaller data sets were generated from ERR318658 containing 1.6 billion reads, 160 million reads, and 80 million reads by truncate the original sequencing files.

Performance with SOAPdenovo2

SOAPdenovo2 is specially designed for human sized genomes and assembles Illumina Genome Analyzer short reads (11). The program creates new opportunities for building reference sequences and carrying out accurate analysis of unexplored genomes in a cost effective way.

Running Time

Figure 3 shows wall clock running times with various sizes of input data with 62 cores. In order to make a fair comparisons, two cores in R930 was unused to avoid unexpected performance degradation. PowerEdge R940 can process 3.2 billion read data 27% less time than PowerEdge R930 in core to core comparison.




Figure 3 Running time comparisons with 62 cores: the total number of cores available in R930 is 64 while 112 cores are available in R940. The choice of 62 cores is to obtain a fair comparison between two systems. Two cores are kept for the system in R930 to avoid any unexpected performance degradation

In server to server comparison, R940 with 110 cores reduced running time by 37% for the tested configurations (Figure 4).

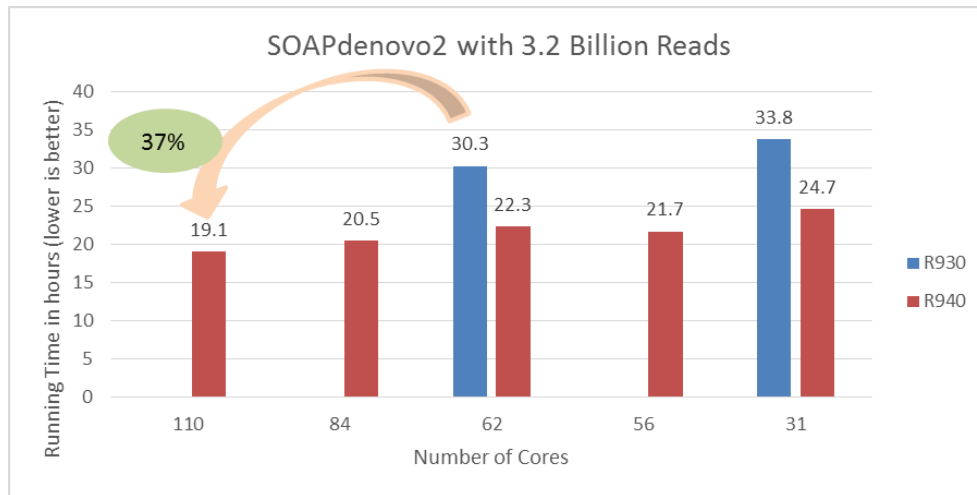


Figure 4 Running time comparisons between R930 and R940 as in *Table 1*: 2 free cores were reserved for both systems to avoid overstressing the systems.

Also, the scaling over various input sizes with a fixed number of cores shows nearly linear behavior on R940 as shown in Figure 5. Although scaling characteristic depends more on the quality of programs than hardware, this linear scalability over various data size demonstrates PowerEdge R940 is optimally built for CPU/memory intensive applications.

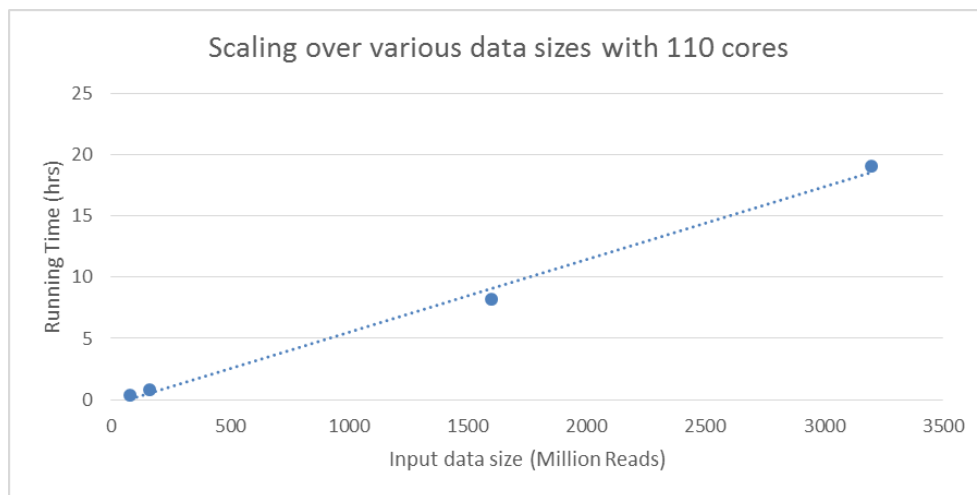


Figure 5 Scaling over various data size with a fixed number of cores for R940: regardless of the number of cores used, running time is linearly increased when the input data size grows.

Peak Memory Consumption

The improvement in peak memory consumption is unusual given that an application typically uses a fixed amount of memory according to input data size and/or the number of threads allocated. However, SOAPdenovo2 consumes 39.5% less memory on PowerEdge R940 than PowerEdge R930 for 3.2 billion read data. Figure 6 shows the peak memory consumption over various data sizes with 62 cores on both systems. The reduction is relatively consistent over different data sizes as well as different number of cores used as shown in Figure 7. The plot shows the peak memory consumptions with 31/62 cores on R930 and 31/62/110 cores on R940. The linear trend lines are well fit onto the data points from the both systems.

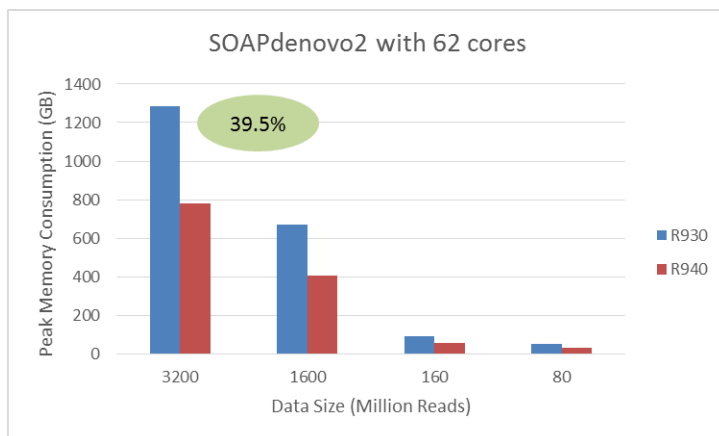


Figure 6 Peak memory consumption with 62 cores. Peak memory consumptions grow over the input data size.

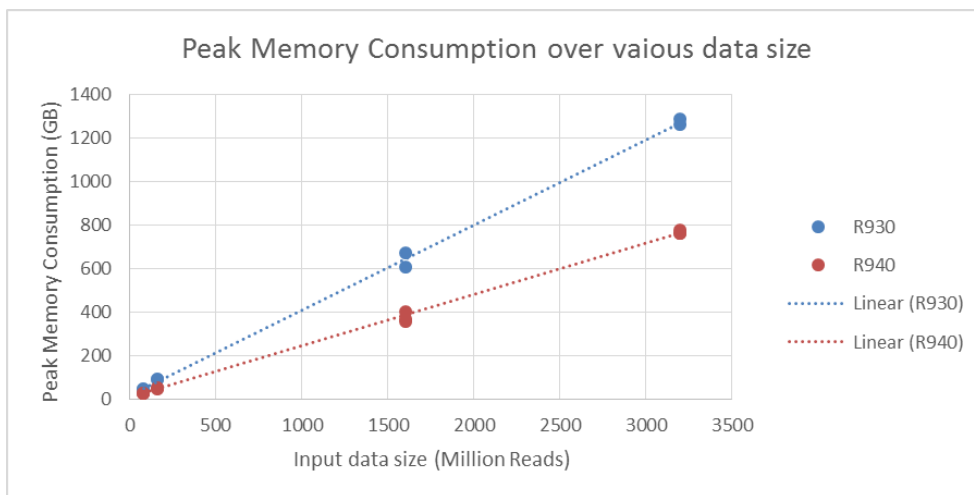


Figure 7 Peak memory consumptions over various input data sizes and the number of cores: the data points for R930 were generated by using 31/62 cores with 80/160/1600/3200 million reads data. For R940, the results with 31/62/110 cores for the same inputs were plotted.

One possible explanation for the peak memory reduction is larger memory bandwidth on R940. Although 2400MHz DIMMs are populated in R930, the memory clock speed must be downgraded to 1333MHz due to the 3 DPC configuration which is theoretically 50% of the memory clock speed on R940. Actual memory bandwidth tests show more than 50% lower results for read/write on larger data sets with loop unrolling codes as shown in Figure 8 and Figure 9. The bandwidth tests were performed with parallel memory bandwidth (pmbw) benchmark/measurement (12).




Figure 8 ScanRead tests by pmbw suit

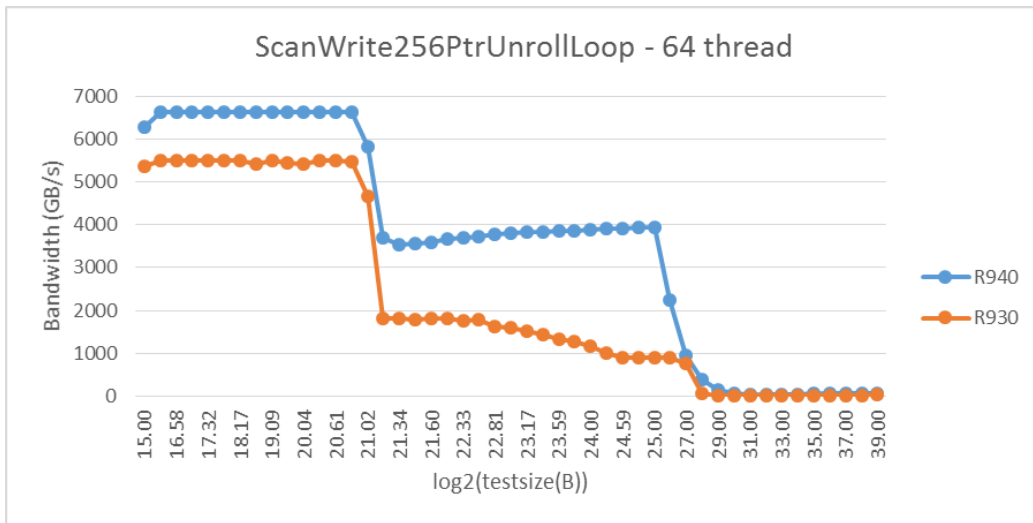


Figure 9 ScanWrite tests by pmbw suit

Conclusion

The Dell EMC 14th Generation server, PowerEdge R940, is superior to its predecessor in several ways. By taking advantage of more memory channels supported with Skylake-SP and faster memory speed, R940 demonstrates speed gain as well as more efficient memory usage. SOAPdenovo2 actually consumes less memory since it process more data faster on R940. This means that even 6 billion reads data can be processed on R940 with 1.5TB memory which was not possible to run on R930 (based on estimations from linear curve fittings of peak memory consumptions from Figure 7). Overall Dell EMC PowerEdge R940 is strongly recommended for *de novo* assembly applications over R930 for the speed and the efficiency.

References

1. What is Next-Generation DNA Sequencing. [Online] <https://www.ebi.ac.uk/training/online/course/ebi-next-generation-sequencing-practical-course/what-you-will-learn/what-next-generation-dna->.
2. Sanger Sequencing. [Online] https://en.wikipedia.org/wiki/Sanger_sequencing.
3. *De novo* genome assembly: what every biologist should know. Baker, Monya. 2012, NATURE METHODS, pp. 333–337.
4. De Bruijn graph. [Online] https://en.wikipedia.org/wiki/De_Bruijn_graph.
5. Multi-channel memory architecture. [Online] https://en.wikipedia.org/wiki/Multi-channel_memory_architecture.
6. Products formerly Haswell. [Online] <https://ark.intel.com/products/codename/42174/Haswell#@server>.
7. Products formerly Broadwell. [Online] <http://ark.intel.com/products/codename/38530/Broadwell#@server>.
8. 13G PowerEdge Server Performance Sensitivity to Memory Configuration. [Online] <http://en.community.dell.com/techcenter/b/techcenter/archive/2016/09/23/13g-poweredge-server-performance-sensitivity-to-memory-configuration>.
9. PowerEdge R930 Source Book. [Online] https://dell.gosavo.com/ViewResource.aspx?ResourceId=39121622&ResourceType=DocumentHtmlPreview&ts=636008149826626561#_Toc452652715.
10. ERR318658. [Online] <http://www.ebi.ac.uk/ena/data/view/ERR318658>.
11. SOAPdenovo2: an empirically improved memory-efficient short-read *de novo* assembler, Luo et. al., 1, 2012, GigaScience, Vol. 1.
12. pmbw - Parallel Memory Bandwidth Benchmark / Measurement. [Online] <https://panthema.net/2013/pmbw/>.