Using Cloud Computing to Reduce Simulation Turnaround Times and Increase Simulation Throughput

Shu Yang
IMMI
Westfield, IN 46074

Arend Dittmer
Penguin Computing
Fremont, CA 94538

Abstract

FEA simulation throughput directly impacts productivity in any engineering organization that uses computational simulations in their design workflows. In the typically iterative design process, higher simulation throughput and shorter turnaround times allow for exploring more design parameters which in turn results in product designs that are closer to the elusive 'optimal' solution. On the other hand, the limited availability of computational resources often limits simulation throughput and increases the turnaround time for simulations resulting in less optimal solutions or delayed schedules. Offering scalability and a pay-as-you go payment model cloud computing promises a way out of this dilemma. This paper provides an overview of Penguin Computing's public cloud infrastructure Penguin on Demand (POD) including a discussion of POD's security model. The paper then discusses how IMMI, a provider of advanced safety systems reduced job turnaround time and increased job throughput of LS-DYNA simulations, using a hybrid model of in-house compute resources and Penguin Computing's public cloud offering Penguin Computing on Demand (POD). Specific examples such as design of IMMI's FlexSeat, a three point belted seat for school buses, frontal crashing simulation of fire trucks as well as respective benchmark data will be provided.

Background

It is no secret that optimizations and design improvements applied in a product's virtual concept and design phase can be realized at much lower cost than fixes that are applied in later stages of a product's life cycle. For designs that are evaluated through FEA simulations a higher throughput of simulations means that more variations of a product design can be explored resulting in more optimized, higher quality products. Automated design optimization tools take the search for the optimal design to the next level allowing for the automation and aggregation of simulation results from multiple simulation tools. Regardless of the specific analysis tools that are used it is fair to state that engineering productivity is directly impacted by the availability of computational resources in an organization's infrastructure.

Hardware performance increases and supporting software improvements such as multi-threaded and distributed code bases allow for the simulation of increasingly larger, more complex designs and tighter tool integration. In light of increasingly complex product designs it is crucial for manufacturing organizations to use current High Performance Computing infrastructures to stay competitive.

Large manufacturing organizations with a critical mass of projects have been using HPC and leading edge digital manufacturing technologies for years. With many concurrent projects the demand for compute cycles typically ‘evens out’ allowing for the procurement of systems for an average workload. However, occasional simultaneous workloads peaks generated by multiple projects, unexpected natural events impacting the supply of power or high operational expenses may result in large organizations looking for easy access to external compute resources. Small- to medium-sized manufacturers often use comparatively basic computational capabilities and often still rely on physical prototyping. The Council on Competitiveness coined the term ‘missing middle’ to describe these small- to medium-sized manufacturers (SMMs) who are missing out on the benefits of advanced modeling, simulation and analysis using HPC systems. Smaller organizations typically face the following challenges:

- The capital expense associated with the purchase of a compute cluster is in many cases cost prohibitive
- Human resources for cluster deployment and management are expensive and HPC expertise is typically not available in-house
- Operational expenses such as space, power and cooling can be prohibitive
- Workloads are often project driven and bursty making it difficult to ‘right-size’ a cluster for an average workload; in other words clusters are typically either under-utilized or provide insufficient compute capacity if a workload peak occurs

Cloud Computing for HPC

The concept of using computational resources in the cloud on a pay-as-you-go basis seems to open a promising perspective for larger organizations as well as the ‘missing middle’:

- Only a minimal capital investment is required
- Resources can be scaled out and resource allocation can be adjusted according to the current workload
- There are no operational expenses beyond the cost of on-demand access

However, the concept of HPC simulations ‘in the cloud’ brings new challenges:

1. While computational resources are available in the cloud ‘on-demand’, application licenses are in most cases only available as ‘static’ perpetual licenses. Moreover, providing access to licenses managed by an in-house license server from a transient system instance ‘in the cloud’ can be challenging.
2. Even though they may be labeled HPC instances, many cloud computing resources do not deliver the performance for effectively turning around simulations. Distributed simulations with LS-DYNA for example greatly benefit from low-latency interconnects that are not available from most providers of cloud services
3. While the use of cloud resources saves capital and operational expenses significant expertise is required for effectively managing cloud computing resources. Computational resources need to be constantly allocated and de-allocated in order to use the additional compute resources effectively. With most cloud computing providers the consumer is left on his own with the burden of deploying the right software stack for running simulations effectively.
4. Depending on the data size, data movement from and to the cloud can be time consuming
5. Many organizations have security concerns about processing of simulations off-premise

Introducing Penguin Computing on Demand (POD)

Penguin offers dedicated physically distinct machine sets that are available exclusively to a single customer for a given period of time. This is of particular interest for larger organizations that want to cost effectively complement their existing in house resources. Penguin is for example accommodating multiple Japanese manufacturing organizations that are affected by the limited supply of power in the wake of the 2011 Tsunami.

For users that require a flexible way to allocate resources on a pay-as-you-go basis Penguin offers an on-demand HPC environment that comprises high-density compute nodes and direct attached high-speed storage. Compute jobs are executed directly on compute nodes avoiding a virtualization layer that negatively impacts performance. Compute nodes feature the current processor generations from Intel (Xeon 5600/E5-2600) and AMD (Opteron 6200) with clock speeds ranging from 2.2GHz to 2.9GHz, 24GB to 128GB RAM per server and up to 1TB local storage space per node. POD provides intra-node connectivity via a low latency QDR Infiniband infrastructure. Storage systems are accessed through a 10GbE connection. For demanding and specialized workloads POD offers high performance parallel files system options such as Lustre or Panasas that are attached via multiple 10GbE links or Infiniband.

Distributed LS-DYNA jobs greatly benefit from this infrastructure: The Infiniband interconnects provide a bandwidth of 40Gb per second and very low latencies in the microsecond range.

Through the localized network topology and topology aware scheduling distributed compute jobs are placed on nodes connected to the same switch for full bi-sectional bandwidth. Fig. 1 shows the scalability of a distributed LS-DYNA run based on a modified 'Neon' model on POD

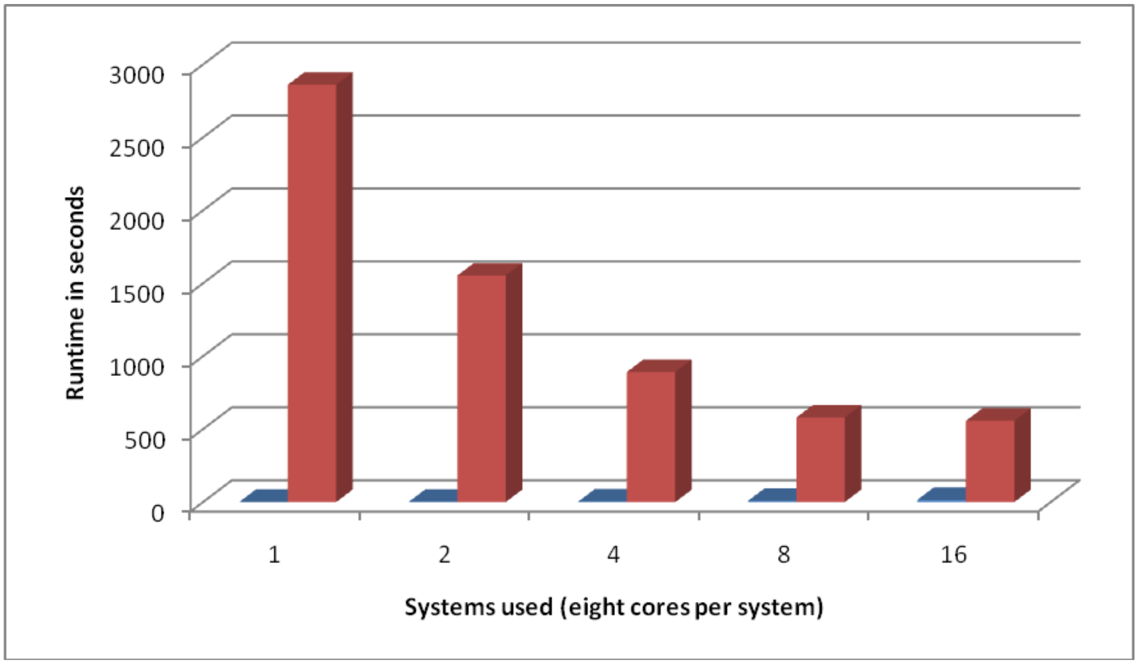

Fig. 1 Scalability of a distributed LS-DYNA run on POD

Users access the POD infrastructure through dedicated login nodes that provide a persistent and secure compute environment. Login nodes run as virtual machines and are securely accessed through an encrypted ssh connection from a dedicated IP address or range of IP addresses (see Fig. 2). From the login nodes users can access their home directory that is shared between the login node and the compute infrastructure. In the typical workflow users prepare compute jobs for execution on their login nodes and then submit the compute jobs through a scheduler. Other than through shared user directories and the scheduler's job submission commands the POD compute infrastructure is completely isolated from the login nodes. Users monitor job execution through the application's output files as well as the scheduler's command set. POD provides resource access through Sun Grid Engine as well as the Moab scheduler. For POD access from nodes that are not dedicated login nodes POD users can use 'podshell', a command line interface that supports file staging and submission of compute jobs from any system with internet access, using a dedicated web-service.

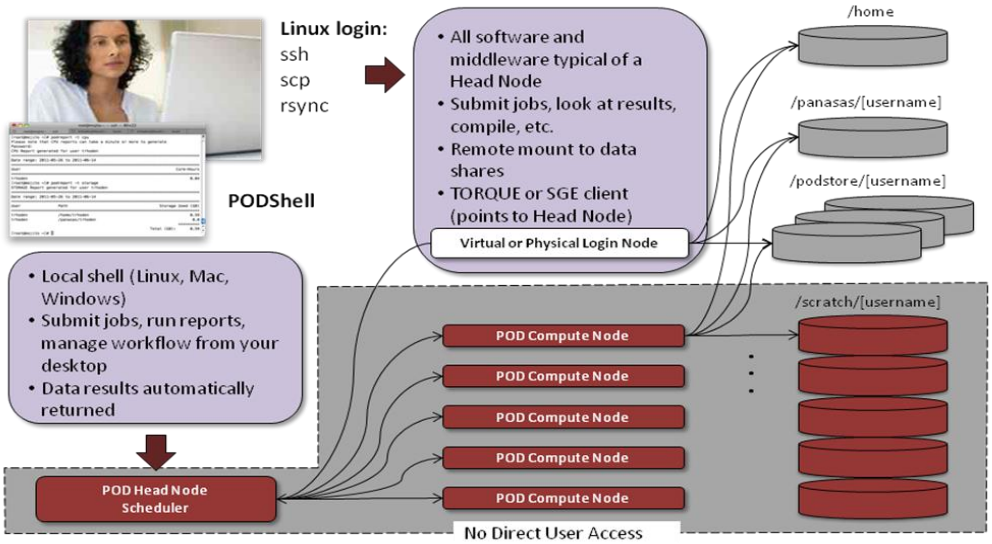

Fig. 2 Penguin Computing on Demand Architecture

POD also offers remote visualization options which reduces the pressure to copy data for local post-processing. On request Penguin's POD support team can also provide specialized compression algorithms such as FEMZIP for users who want to reduce the amount of data that needs to be transferred. For extremely large data sets POD offers a hard drive service.

Using LS-DYNA on POD

Through an agreement with LSTC, Penguin Computing offers LS-DYNA license access on-demand for US based organizations that have at least one in-house license. Using LS-DYNA on POD is completely transparent as the POD environment hosts an LS-DYNA license server with pre-allocated license tokens. LD-DYNA on-demand pricing is based on standard pricing for LS-DYNA and pro-rated to a minimum time window of 24-hours. Penguin Computing is the single point of payment for on-demand LS-DYNA licenses. This setup provides a high level of flexibility for on-demand users as it offers a pure pay-as-you-go payment model without a significant markup for licenses that are only used when needed.

Case Study: IMMI

Located in Westfield, Indiana, IMMI is a leading provider of advanced safety systems for commercial, industrial, military and emergency response vehicles as well as school buses, motor coaches and child restraints. Particularly in the last 15 years CAE has played an important role in the development of advanced safety systems. A primary example of a typical CAE driven development process at IMMI is the SafeGuard school bus seat, which is equipped with lap-shoulder seat belts. There are two key requirements to the bus seat structure design: Federal Motor Vehicle Safety Specification (FMVSS) 222 and 210. In FMVSS 222, the seat back's energy absorption must exceed given values when the seat back is pushed by force-application bars in the forward direction as well as the rearward direction (see Fig. 3-4). The seat back forcedeflection curves have to fall within the corridors specified in the standard. In addition, the displacement of the shoulder belt anchor points on the seat back cannot exceed certain limits when the shoulder belts are pulled forward at certain loads (Fig. 5).

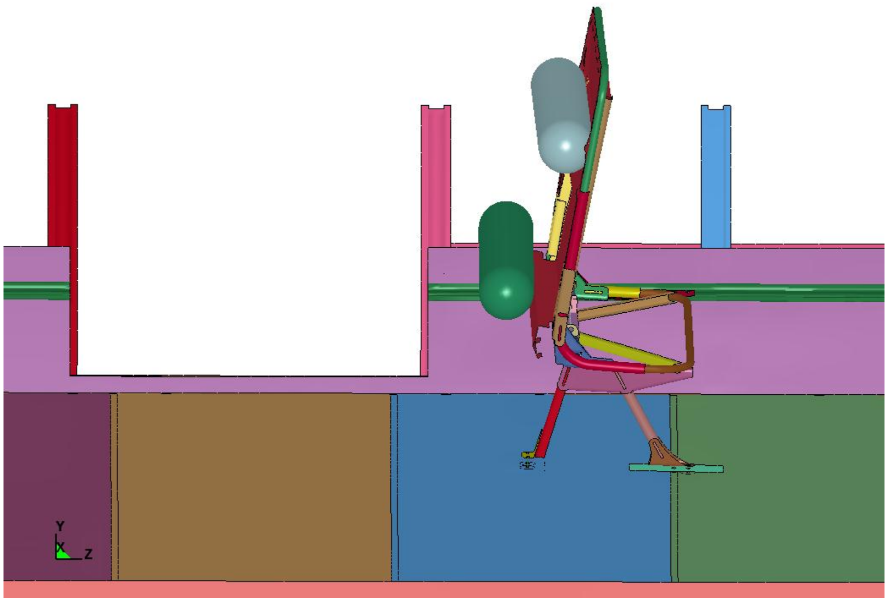

Figure 3. FMVSS 222 evaluation of the SafeGuard bus seat design in forward direction

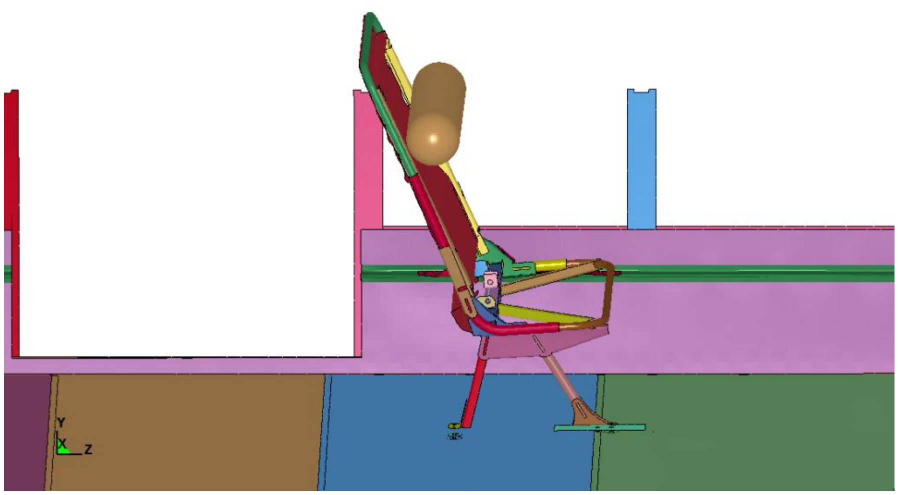

Figure 4. FMVSS 222 evaluation of the SafeGuard bus seat design in rearward direction

In FVMSS 210, the seat and seat belt anchorages cannot separate from the vehicle structure when certain loads are applied to the seat belts (see Fig. 6). Taking the two standards together, the seat needs to be strong enough to provide protection for belted occupants but not too strong or too soft so it provides compartmentalization for unbelted occupants.

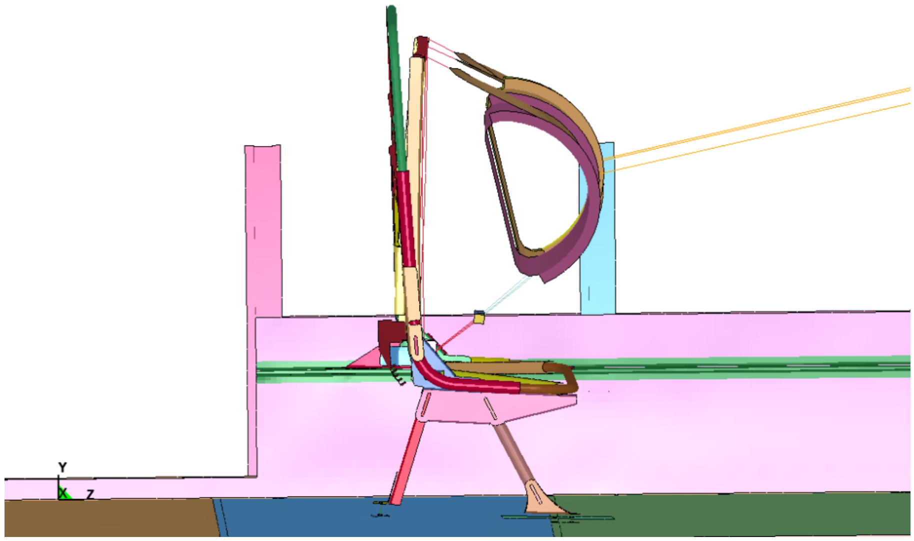

Figure 5. FMVSS 222 evaluation of the SafeGuard bus seat quasi-static test

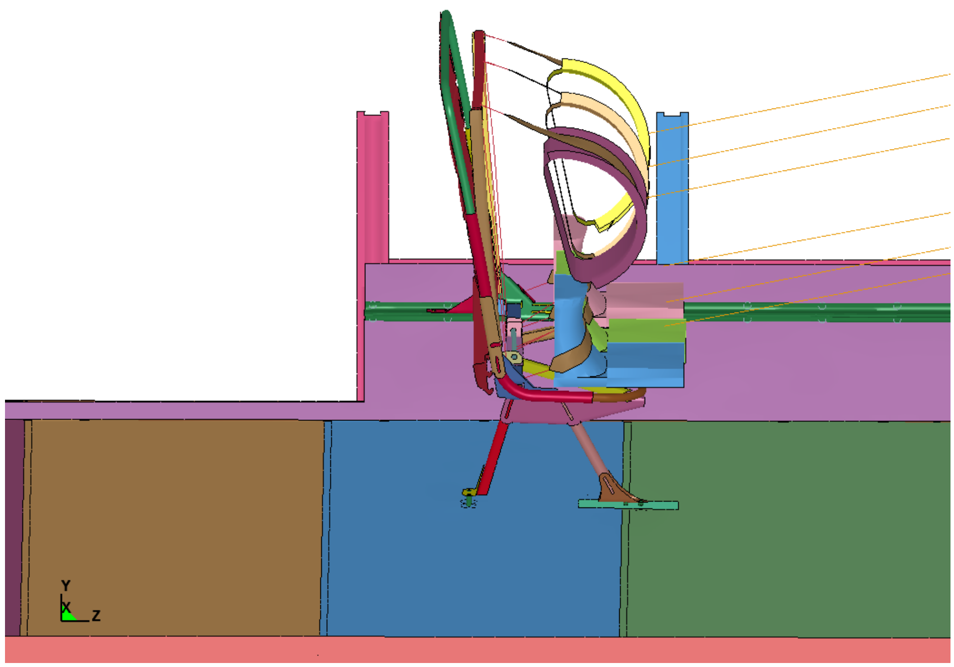

Figure 6. FMVSS 210 evaluation of the SafeGuard bus seat design

In IMMI's workflow design engineers first send their CAD designs to CAE engineers in the Applied Mechanics group. CAE engineers then build FE models and evaluate the design performance through simulations. If analysis results show that the design concept does not perform well, design engineers modify the design to address the concerns identified in the simulations. CAE engineers then re-evaluate the modified design. Typically there are over a dozen design-analysis iterations before physical prototypes are even built.

The Applied Mechanics group at IMMI supports internal product development and provides engineering analysis services to IMMI's customers. One example is frontal airbag sensor development and calibration for commercial vehicles. IMMI's Applied Mechanics Engineers work hand-in-hand with OEM Engineers to develop full vehicle models and perform crash analysis to assist sensor calibration and improve overall crashworthiness. Crash models are correlated with physical barrier crash test results and greatly expand the crash data available for calibration while reducing expensive full vehicle crash tests.

Since the Applied Mechanics group started in 1996, IMMI has invested in several servers. In 1996, IMMI purchased an SGI origin 2000 with 8 CPU's that was replaced with an SGI Altix 330 with 12 CPU's in 2005. At the time, these servers were sufficient to support internal product development efforts. It took less than an hour to perform component level stress analyses such as evaluating the strength of seat belt buckles. Larger, system level analyses such as the evaluation of a school bus seat design took about 8-12 hours to complete. Twelve hours of wall clock time for turning around large, complex models was considered satisfactory by CAE engineers because it allowed them to work on a model during the day, run the simulation overnight and obtain the results the next morning. However, crash analyses took significantly longer than system analyses. In 1999, it took over 40 hours on the Origin 2000 to simulate a vehicle crash model with around 200,000 elements. It still took between 20-30 hours on the Altix 330 to simulate a vehicle model with around 500,000 elements in 2006.

As demand for better accuracy on vehicle crash analyses arose, vehicle models included more details, finer meshes and an increasing numbers of the elements. It took an increasing amount of time to run crash models on the Altix system. In addition, sometimes IMMI design engineers needed a quick decision on design changes due to manufacturing capability or material availability. CAE engineers at IMMI struggled to achieve a quick turnaround time with in-house computing resources. Another challenge the Applied Mechanics group encountered was the limited numbers of jobs that could run at the same time.

About 90% of all compute cycles at IMMI are used for LS-DYNA. In early 2010 it became clear that additional compute resources were required and IMMI had to decide whether to deploy additional in-house resources or to start using HPC resources in the cloud. After studying the Cloud Computing marketplace and comparing the cost of cloud offerings with expenses that would be incurred by deploying internal resources, IMMI decided to pursue the cloud computing option. POD was chosen as IMMI's cloud computing provider for LS-DYNA simulations because it provides secure, easy to access and on-demand services on a 'pay-as-you-go' basis. The cloud computing solution has met IMMI's computing needs very well. The demands on the computing resource fluctuate depending on the projects that CAE engineers are working on. For many non-vehicle-crash projects, the in-house Altix machine is sufficient. Only large jobs such as vehicle crash models or jobs that require faster turnaround are sent to the cloud. For a vehicle crash model with over 1 million elements, the wall clock time plus the time for transferring data is less than 12 hours. CAE engineers that submit compute jobs before they leave work in the evening get their analysis results when they are back in the office the next morning. On the school bus seat analysis, the typical job can be completed within an hour. Using POD resources CAE engineers are able to provide timely evaluation results to design engineers, speeding up product development decisions.

In the past, due to resource constraints, few design optimizations have been conducted at IMMI. With cloud computing on POD, design optimizations are becoming a standard practice in the development of new products.

References

[1] Digital Manufacturing Report, June 2011, "Hope for the Missing Middle", John Kirley
http://www.digitalmanufacturingreport.com/dmr/2011-06-06/hope_for_the_missing_middle.html