Electronic Design in the Cloud

Dr. Raul Camposano, CEO, Nimbic
www.nimbic.com

Executive Summary

Computing is moving to the Cloud. So will design of electronic systems. The Cloud is shifting the computing paradigm once again. Simply put, large-scale commodity computing (millions of servers) delivered through the Internet are better than small compute centers, regarding virtually every criteria: Cost, scalability/performance, availability of service, utilization, security, provisioning, ease of use/maintenance, need for qualified IT personnel, etc. The cloud enables delivery of hosted software, platforms and infrastructure as a service.

The advantages of the cloud for designing integrated circuits, packages and boards are obvious. Imagine designing in a server farm, having access to a virtually unlimited number of computers, being able to run as many parallel jobs as necessary, provisioning new servers in minutes, exploiting parallelism in tools to get your jobs done with minimum turn around time. No need to purchase computers or tools upfront, installing them, maintaining them. All you need is an Internet connection. And you pay just for what you use, as you go.

This paper explores design in the cloud, in particular 3D electromagnetic field solving which happens to be a particularly good match for the cloud. We focus on the advantages of our solution that are difficult to implement in a traditional enterprise compute center setting. Design in the cloud is still largely a vision, but a new generation of design technology, partly written from the ground up, with the cloud and parallelism in mind, will emerge. Computing is moving to the Cloud. So will design of electronic systems.

Introduction

Much has and is being written about Cloud Computing, there is no need for yet another introduction (71,100,000 search matches). We share the view that over a period of time much (not all) computing will move to the cloud. As a paradigm shift it can be compared with client/server computing, where the clients are “small” computers or mobile devices and the server is the cloud (one or many clouds). A multitude of applications already shifted: Email, search, social media, gaming, web services (Google apps, file synchronization/backup…), smartphone apps (navigation, stargazing…) and many startups (such as ourselves) have moved their software development environment to the cloud. The speed of migration to the cloud will depend on the success stories and accidents on the way, but it will happen.

Migration to the cloud is certainly not being led by the electronic design automation industry. It can’t: Consumer apps have been one obvious driver for the early cloud; they are the perfect “clients”. Furthermore, Electronic Design Automation (EDA) is only a niche of about $5B in the over $200B large enterprise software segment. EDA users are mostly large companies with existing compute centers, not eager to rely on third parties for their compute needs and reluctant to let their designs move out of house.

This paper makes the case that, despite the above, design of electronic systems is ready to move to the cloud. It will naturally be adopted first by the long tail of small users, which benefit most right away. It is also likely that designers of printed circuit boards (PCB) and packages will be earlier adopters than Integrated Circuit (IC) designers, because there are fewer Intellectual Property (IP) issues and the overall magnitude of the problems in PCB/packages is smaller thus favoring smaller players.

In the first part of this paper we introduce nCloud, our cloud design environment, which we host in Amazon Web Services (AWS). We show how nCloud addresses the hurdles to adoption of the cloud for electronic design. In the second part we explore the advantages of using our 3D electromagnetic field solver nWave in the cloud.

Using the Cloud for Electronic Design

nCloud is Nimbic's cloud solution (Figure 1). It consists of a Web portal to access the services, infrastructure to host and deliver our services, and the software that constitutes our services (for the purposes of this paper nWave, our full 3D electromagnetic field solver).

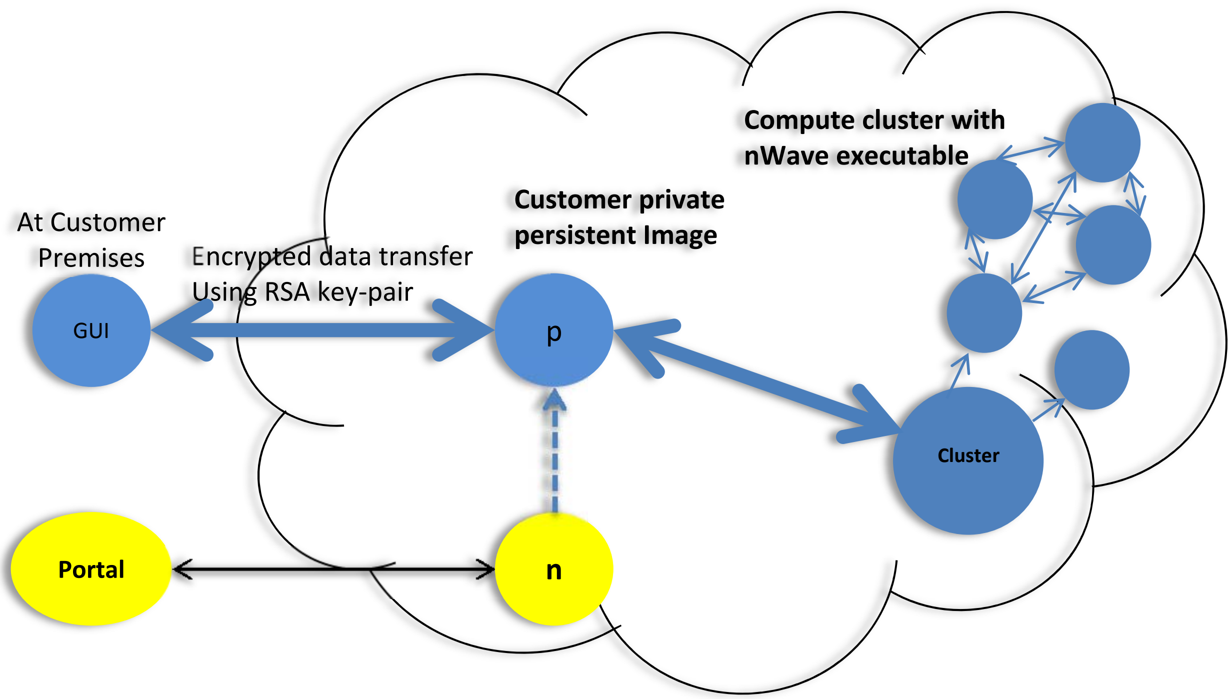

Figure 1. nCloud

To use nCloud, the customer logs in through the portal and establishes an account. The instance **n** (a virtual machine running in the cloud), which belongs to Nimbic, then instantiates a persistent virtual machine **p,** dedicated to the customer (of exclusive use by the customer, no one else has access to it). The customer uses the services, in our case runs nWave, by communicating with **p**. **p** instantiates as many additional virtual machines (the cluster) as necessary to execute the jobs submitted by the customer to **p**. Results are collected by **p** and sent to the customer. **p** is persistent for as long as the customer purchases services ("has an account"). In the sequel, we describe the main features of nCloud in more detail.

Licensing

A typical license of an EDA tool in 2011 consists of a one to three year subscription, often called "time-based license" (as opposed to "perpetual"). nCloud allows to "license" a service (using a tool) for any period of time. Strictly speaking, nCloud meters the use of a service. The minimal increment can be one hour, e.g. Amazon Web Services delivers infrastructure by the hour. Another way of looking at this is that you can purchase as little as a one hour license. We use buying a "license" and a "service" interchangeably in this paper. Designers are more used to purchasing tool licenses than services.

Provisioning

When a customer needs to submit a job, communication is established with the persistent instance **p,** which is running in the cloud. Communication happens at Internet speed, usually in tens to hundreds of milliseconds as determined by the latency and bandwidth of the connection. There is no new machine to provision. Once a job is submitted, **p** will provision the virtual machines necessary to execute the job. Typically this happens in minutes. Provisioning is straightforward and is done by the end user. If a project requires a spike of jobs to be run at a given time, additional licenses (services) can be purchased on a short-term basis and provisioned in minutes.

Parallelism and scaling

Using nCloud, the customer can run as many jobs in parallel as purchased. Running n jobs in parallel costs the same than running them sequentially (at the granularity of the service, in our case one day): Running 10 jobs during a day costs the same than running 1 job during 10 days. In addition, design tools such as nWave can exploit parallelism: by running many instances in parallel a job is accelerated. nCloud will allocate machines in the cluster to speed up a job according to how many copies a customer wishes to dedicate to a given job. In this way, a job that may take many days on one machine can be completed in hours by using dozens of machines in parallel. Jobs that take hours finish in a few minutes.

Transfer of data

To start a job, the input to a particular tool must be sent to the cloud. The customer does this by an encrypted data transfer to **p** using SSL (Secure Sockets Layer) encryption. Similarly, results of a job are retrieved transferring data from **p**. A common connection of 10Mb/sec allows transferring up to 75 Megabytes of data in one minute (theoretical maximum). The user has to decide if it is practical to transfer design data back and forth to the cloud for a given application. In the case of nWave for example, the design data and results are only a few Megabytes, allowing for a transfer of data to/from the cloud in seconds. Compared to the compute time required to solve a problem, often in the hours or days, this is negligible. For other applications which require more data an alternative is to leave the data in the cloud, using a service such as Amazon's S3.

Availability of service

Even though large, distributed compute facilities have high availability of service, occasionally some cloud services may become unavailable. We have hosted our cloud in AWS, which offers several regions and zones for its infrastructure. If one particular zone goes down, nCloud allows restarting the services of affected customers in a different zone within minutes. Virtual machines are stored in persistent storage and can be reinstated anywhere. The jobs being run in the machines that are affected by an outage would of course be lost, just like in traditional compute centers.

Access

The services being purchased in the cloud are accessed through the Internet; in practical terms that means from anywhere. nCloud allows to restrict access to a fixed IP address to increase security. Any client that supports Internet access is supported, including desktops, laptops and mobile devices such as smartphones tablets (iPad™). Access is typically through a remote terminal or a browser. A particular tool may require an app to run on the client, for example an interactive graphic user interface (GUI).

Interactivity

Even though the Internet enables highly interactive apps such as on-line gaming, there are interactive applications, which look better when run locally, using a powerful graphics card, without the latency and bandwidth limitations of the Internet. A 3D object editor still looks better run locally. Our approach for cases with a large amount of computing to do in addition to an interactive GUI, is to split the GUI from the rest of the tool and install the GUI as a local app in the client. An example of a work session in this case is the sequence of local editing, transmitting data to the cloud, computing in the cloud, retrieving data from the cloud and local visualization.

Security

nCloud has been engineered with security in mind. The implemented security mechanisms can be divided into three groups: Isolation, security of data in transit and security of data at rest.

Each Customer account is allocated a dedicated private virtual machine **p** that is isolated from the world. After being instantiated by the Nimbic **n** instance, login is disabled and the customer generates a key pair to communicate with **p**. The only mode of communication is through data transfer via SSL. **p** autonomously governs all communications, autonomously allocates a private compute cluster for performing parallel tasks, submits the parallel jobs and collects results for the customer and finally autonomously shuts down the cluster when it is not needed. Additionally, **p** can be configured to only receive data from a fixed IP address. Summarizing, **p** is a virtual machine running somewhere in the Amazon cloud, with login disabled, only exchanging data encrypted with a customer owned key via SSL, autonomously executing a limited number of tasks. All Customer specific data, communication, and computations are completely isolated in dedicated virtual instances.

For data in transit, all data uploads and downloads to the Amazon cloud are done using secure transfer utilizing RSA key pairs. The customer locally generates and uploads public keys but the private keys stay known only to the customer. Inside the of Amazon cloud, all data transfers are among virtualized machines, so that it is practically impossible to packet sniff. In addition, all data exchange is done with secure protocols (SSH) using locally generated key-pairs known only on the customer dedicated compute instances.

The data at rest in the cloud is stored only on local “ephemeral” disks, which belong to an instance and are wiped once the instance is terminated. In addition the ephemeral disk is mounted as an encrypted drive that uses a locally generated random key at boot time that is not stored on disk. The format used does not signal that encrypted data is present. Data on the persistent **p** customer instance is not archival, it is deleted on request from the local ephemeral disk. The customer needs to retrieve and store this data with the option to encrypt and store it on persistent storage in the cloud.

Cost

Comparing the total cost of electronic circuit design using cloud computing with a traditional data center setup depends on too many factors to be able to make a meaningful comparison. In a survey of 155 companies that use cloud computing in August of 2010 (*State of Cloud Applications and Platforms: The Cloud Adopters' View,* B. Narasimhan and R. Nichols, Computer, March 2011), 67% stated that the total cost of ownership of cloud computing was lower than on premises. In the case of electronic circuit design, a large part of the cost lies in the EDA tools. In absence of a pricing model for design tools in the cloud (there is no commercial full design solution in the cloud) and without knowing the exact use model (which depends on the tool and the particular design) it is impossible to make cost comparisons. A business model for EDA in the cloud will need to be developed. In any case, we believe that the trend for compute infrastructure in the cloud is to become relatively cheaper over time as adoption increases, cloud compute facilities increasingly exploit economies of scale, the cloud matures and support cost go down and ease of use goes up. We are also firm believers that fair pricing and availability of licenses for periods as short as a day will lead to higher usage of design solutions in the Cloud.

Private Clouds

For those customers that have invested in large compute centers that satisfy their security needs and can operate them cost-efficiently, a private cloud may be of interest. The advantages of virtualizing a data center won't be discussed here; if additionally the particular interface of AWS is mimicked (such as several commercial offerings do), nCloud can be deployed on such a private cloud. Potential advantages include the easy spilling over into the public cloud when more capacity is needed, and the use of metering based licensing.

Lock-In

Finally, being locked into a particular cloud provider is often cited as a drawback. We plan to explore deploying nCloud on a second cloud in the future; several commercial offerings address this point.

3D Electromagnetic Field Solvers in the Cloud

We have ported Nimbic's 3D electromagnetic field solver nWave to the cloud making it the first service available on nCloud. nWave happens to be a particularly good application to run on the cloud. Figure 2 shows a snapshot of an early demo vehicle running in the cloud, and we'll use this example to illustrate the advantages of the cloud version.

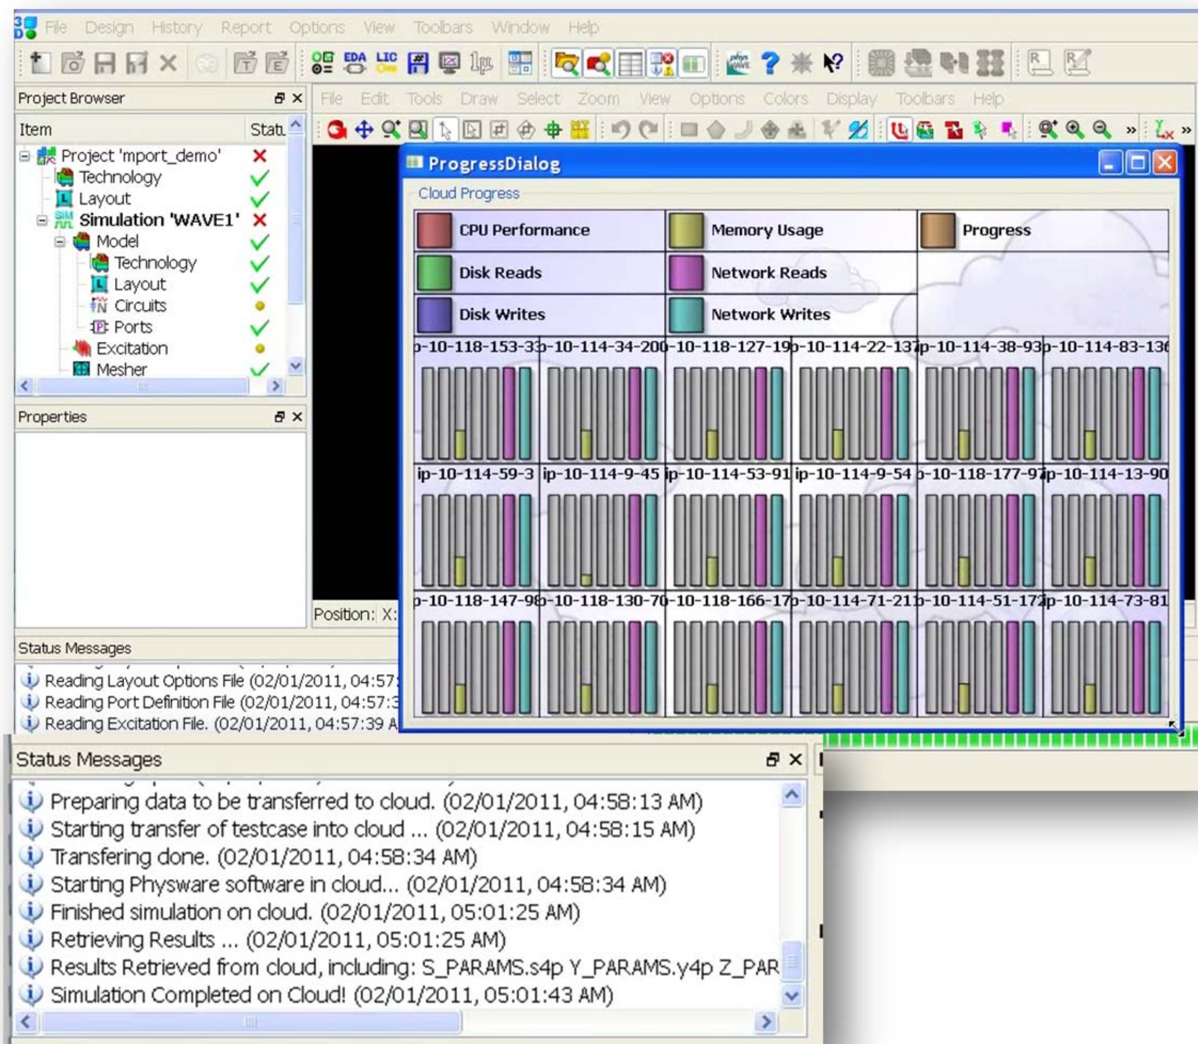

Figure 2. nWave running in the Cloud

Data I/O versus compute and memory required

nWave requires very little data I/O: just the design geometry of the package / board as input and S-,Y- and Z-parameters for a range of frequencies as an output. The data transfer to the cloud took only 19 seconds (dialog in figure 2) on a 1.5Mb/sec internet connection, retrieving the results is even faster. The simulation on 18 8-core large machines took 2 minutes and 51 seconds. Larger problems show an even higher ratio of compute to I/O. nWave on the cloud can be used keeping all data at customer premises and transferring it to the cloud only when starting a simulation. All data in the cloud is destroyed once a job is finished and results have been retrieved, unless the customer decides to encrypt them and store them in S3.

Parallelism – built from ground up

nWave was built with parallelism in mind from ground up. nWave is multithreaded, resulting in a speedup of approximately 6.5x on an 8 core machine. nWave can also partition a problem to run it on multiple machines that are connected by Ethernet and communicate using message passing (MPI), resulting in a speedup of approximately 3 times on 4 machines, and allowing to solve problems as big as the sum of the memory of the four machines. The third level of easy parallelism involves simply running different frequency points and parametric sweeps on different machines independently with almost 100% scalability. The example above used multi-threading (8 cores) and distribution of frequency points (18 machines). Compared to a run on a comparably fast four-core machine, which took 91 minutes, we obtained a speedup of approx. 32 times using 18*2=36 times more computing power. If we add the data transmission time to/from the cloud, we still get a speedup of 26 times.

Decoupled GUI

Our GUI allows to manipulate a design at the geometric level, for example to merge a package and a board, add a ball grid array (BGA) and identify the parts of the geometry that are to be simulated. This involves zooming, panning and rotating 3D objects in real time (as the cursor is moved by the designer). A native solution running on a PC looks better than doing this in the cloud, so we separated the GUI from the actual field solving and offer the GUI as an app that runs the client on customer premises. The snapshot in Figure 2 depicts that GUI. Alternatively, nWave can be run totally in the cloud with some shortcomings in the display of 3D objects being manipulated, but fully functional otherwise. We plan to develop a web GUI that will fully function in a browser so that it can also be used in mobile devices such as tablets and allow some geometry editing, start/stop jobs, monitor progress and visualize results.

Advanced Features

nCloud's ability to start a large number of jobs in parallel opens exciting new applications which were not possible before the cloud. The most obvious is very fast turn around times which allow to include field solving in the design loop. Parametric designs based on parametric sweeps lead the way to multi-objective optimization. Figure 3 shows the computation of the tradeoff between cross talk and return loss by doing hundreds of simulations and computing the Pareto front. All these simulations can be done in parallel in the cloud, obtaining results almost instantly. A similar scenario applies to yield optimization using Monte Carlo techniques and optimization under statistical variability.

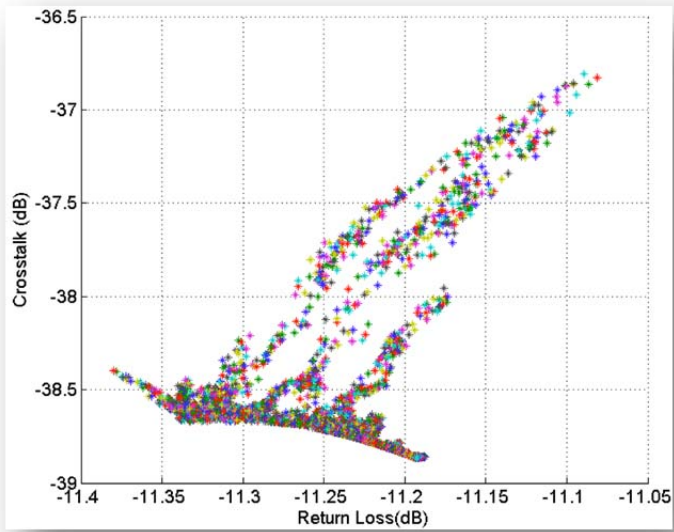

Figure 3 – Pareto front for Multi-dimensional Optimization
(Courtesy of V. Jandhyala, ACE Lab, UWEE)

We hope that we have been able to convey the many advantages of electronic design in the Cloud, in particular of 3D electromagnetic field solvers, and that we have motivated the reader to venture into the cloud and try nCloud Wave. As cloud technology matures, the value proposition will only get better.

To learn more about how we are enabling electronic design in the cloud, visit www.nimbic.com