CLOUD PERSPECTIVE ON RECONFIGURABLE HARDWARE

Octavian Mihai MACHIDON, Titus Constantin BALAN, Radu CURPEN

"Transilvania" University of Brasov, Romania,

Abstract: *This paper focuses on the potential applications of reconfigurable hardware in cloud computing architectures, proposing a specific approach, identifying the sensitive areas that can benefit from hardware-favored implementations and also the challenges that reconfigurable computing brings along.*

Keywords: *Cloud computing, reconfigurable hardware, FPGA, services, virtual networks*

1. INTRODUCTION

Cloud computing, the next stage in the Internet evolution, is an architecture of computing concepts that aims to deliver all types of distributed resources (from computing power to computing infrastructure, applications, and business processes) to the customers as services. Cloud services include the delivery of software, infrastructure and storage over the Internet in a scalable way (according to the user demand), either as components or integrated platforms.

This new and revolutionary type of computing is changing the way companies use technology, re-shaping the consumer-provider relationship by reducing spending on technology infrastructure ("CapEx – Capital Expenditure"), streamlining processes, improving accessibility, minimizing the need for licensing new software and improving flexibility (including restructuring of the "OpEx – Operational Expenditure"). Basically, the cloud aims to eliminate many of the sensitive constraints that the traditional computing environment suffers from: space, time, power and cost.

The various types of cloud services are forming three distinct models or layers:

Infrastructure as a Service (IaaS) offers compute, storage and switching resources in a rather unstructured way (so, the OS – Operating Systems – are also "in the clouds").

Platform as a Service (PaaS) offers development configurations and environments, (more structured than IaaS, usually with local OS).

Software as a Service (SaaS) offers purpose-built business applications (assembled on a "canvas" from stand-alone consolidated software modules, remotely accessible, for instance, via APIs – Application Programming Interfaces).

As cloud computing is showing a rapid development, the attention turns to the hardware resources needed for providing cloud services.

Shifting more and more local applications to the cloud requires massive volumes of data on the servers to be delivered to users, while guaranteeing the security and availability of the data. However the growth of the datacenters brings on a new limitation: energy consumption, which needs to be surpassed by an efficient scaling of the resource use.

Due to these issues that need to be overcome, portable and embedded computing devices have been developed and are being integrated in the cloud systems [Romanca, Ogrutan, 2011].

Thus, hardware designers are facing new challenges with regard to the area and power requirements, but also design constraints like the time-to-market. In this context, a reconfigurable hardware, closer to the physical layer, is a promising asset that has the potential to help overcome many of the challenges mentioned above.

However, there are also a series of particularities specific to reconfigurable hardware that need to be addressed before a possible mainstream adoption of this technology in cloud systems.

2. RECONFIGURABLE HARDWARE AND THE CLOUD INFRASTRUCTURE

Reconfigurable hardware refers to devices that can change their physical circuitry on the fly. This offers a great deal of flexibility when implementing circuits, since the reconfigurable hardware resources are configurable (and usually re-configurable) after fabrication, making it possible for a single hardware device to implement various circuits.

Reconfigurable hardware is the building block for the newly emerged discipline of reconfigurable computing, which uses reconfigurable devices (like FPGAs – Field Programmable Gate Arrays) for computing purposes. These systems have impressive performances, with a series of strong points like:

-High speed;

-Reduced energy and power con-sumption (the circuitry being optimized for the application);

-Reduced size and component count.

These reconfigurable devices have a great potential due to their high adaptability and scalability reducing the need for dedicated circuitry, optimizing energy consumption and minimizing the hardware resources needed for a specific application.

Cloud computing, being new and evolving, provides some interest areas where reconfigurable hardware could help due to strong-points mentioned above.

The growths in the cloud computing market lead to the moving of more and more data from local storage to datacenters in the cloud for computation. This huge workload raises issues regarding efficiency that need to be addressed at the hardware level.

Energy consumption is one of the main concerns when dealing with growing data centers.

Reconfigurable hardware like FPGAs have the potential to offer a better scalability than one can obtain using traditional hardware resources.

Dynamic hardware resources are the latest step in the evolution of computing. Reconfigurable computing devices can be used as datacenter components that can easily scale and adapt to the fluctuations in load and demand that are a characteristic of cloud systems.

Another asset that reconfigurable computing brings is the potential of System-on-Chip (SoC) integrated structures with application-customized topology. This high configurability enables Network-on-Chip (NoC) architectures to be implemented with reconfigurable network topology and embedded protocols. Such architectures combine packet-switching with circuit-switching resources, thus gaining in flexibility while maintaining high energy-efficiency rates. These applications aim to ease the increased stress on network components (servers, routers) that the growth of cloud systems causes. [Stensgaard, Sparsø, 2008]

2.1 Content distribution networks

One of the applications where reconfigurable hardware could play an important role is in the implementation of dynamic content distribution networks (CDNs). These networks are a crucial part of today's Internet. A CDN is a distributed network of servers and file storage devices deployed to place content and applications in geographic proximity to users, that is reducing the load on the origin site infrastructure and the bandwidth.

Cloud and CDNs work together in a holistic system that meets the demands from content delivery as well as economical computing power [Internap, 2013]. While cloud platforms are responsible for computing dynamic content, the large amounts of static content are being delivered by CDNs, with only the dynamic portion needing to come from the origin.

CDNs have the potential to act like catalysts for cloud acceptance by addressing several concerns regarding: security, availability of service, scalability and performance.

In order to cope with the increasing volumes of data that need to be delivered, considerable efforts are underway for improving CDN structure and web caching.

Research in this area has shown that besides different optimization strategies, an effective approach is the use of reconfigurable hardware in implementing the CDN servers [Cheng et al., 2012]. In this proposed implementation, most of the scheduling and downloading tasks are implemented in hardware, instead of the traditional software implementation. This highly customized approach makes the CDN web server reconfigurable, and with a higher throughput and lower power consumption [Cheng et al., 2012].

2.2 Service protocols

Another area in which reconfigurable hardware is showing great potential in assisting cloud computing is the hardware implementation of the web service protocol stack. These protocols are used in the definition, localization and implementation of web services, and also in the service interaction. Basically, the stack consists of the following four protocols:

-Service Transport Protocol;
-XML Messaging Protocol;
-Service Description Protocol;
-Service Discovery Protocol.

The hardware-based (FPGA) web servers that implement this protocol stack are proven to show an accelerated web processing rate, shortened web processing time and enhanced throughput. This is due to the pipelined hardware implementation of the protocol stack that is more efficient than the traditional software based implementations, and also the direct hardware execution – without an OS [Yu et al., 2011].

2.3 Virtual networking

Cloud computing's main motor is virtualization, which makes the rigid physical component available as a service, easy to use and software-managed. One of the benefits of virtualization is the sharing of network and storage devices, which increases the degree of utilization and makes application migration between hardware resources an easy task.

Network virtualization overcomes a series of challenges of the growing cloud, making it possible for multiple logical networks to be implemented over a single physical infrastructure, thus sharing network resources over multiple virtual networks. The traditional approach, software network virtualization, beside its known strong-points (high flexibility and support), suffers from the nature of the general purpose microprocessors which limit the potential performance. Configurable hardware solutions include ASIC (Application-Specific Integrated Circuits) and FPGA. With flexibility being a key asset when talking about virtual networks, FPGA emerges as a better choice, since it overcomes ASICs with regard to service customization, due to its re-configurability. The FPGA-based network virtualization model uses the chip to implement virtual routers that benefit from the platform's scalability, and can be easily customized thanks to reconfiguration.

Despite limitation in resources when having to cope to a large amount of virtual networks, integration of FPGA-based virtual routers with software- based virtual routers is a potential solution that overcomes all the above issues [Unnikrishnan et al., 2010].

3. APPLICATIONS IN CLOUD SECURITY

Along with the many advantages that cloud computing has, there are also some important problems that need to be addressed. Most of them are related to existing vulnerabilities in the cloud model of openness, which lead to an urgent need for data security and privacy protection issues to be solved as soon as possible. In a cloud system data is being shared between client and provider, and this raises challenges in the privacy protection of sensitive data. There are concerns about controlling what information can be revealed and who has access rights over the Internet to that piece of information. Our perspective has to go beyond traditional models of data-security and channel-security.

Securing the cloud infrastructure is a process relying on trusted computing and cryptography, and it needs the strict separation of sensitive data from non-sensitive data.

In cloud systems, the client does not have exclusive administrative control over the system where the data is placed. This could lead to a potential lack of trust, since data can protected by encryption only in transit to/from the datacenter and while in storage, but while in use, the data can be accessed by a malicious system administrator or other parties. Clients expect strong security guarantees to be provided as part of the Service Level Agreement (SLA) by the cloud operator. The traditional software-based systems offer some security, but are vulnerable to attacks from the inside. Conventional systems use a unified memory space for program and data, an architecture that offers little memory protection to attacks that can alter the program memory at runtime.

In our perspective, shifting the trust from software to hardware is a key improvement that reconfigurable FPGAs can bring to cloud data processing. By running sensitive applications or algorithms on the FPGA, the data is protected from malicious intruders that can affect the PC or the OS.

Research in this area has shown that reconfigurable hardware can solve most of the traditional system's flaws, due to its built-in bitstream protection, isolated memory spaces and computing parallelism.

The FPGA implementation of a cloud security mechanism provides a protected area within the untrusted environment for securely performing sensitive operations. Since data can be transferred safely to the device, here it can be manipulated without possible interference from outside factors (other system components or an administrator). Thus, the FPGA can play the role of a trusted computing device that decrypts the input cipher-text, performs the computational operations and re-encrypts the results. This potential implementation acts like an emulation of a homo-morphic encryption, which is a technique under development that allows computations to be carried out on the encrypted data, with the decrypted result matching the same result of operations performed on the original data [Eguro et al., 2012].

This type of encryption, if implemented efficiently, would have great applications in the outsourcing of private computations, which is the case with cloud computing.

Although an effective and fully operational homo-morphic encryption algorithm has not been developed to this date, in our perspective the use of reconfigurable hardware like FPGA should enable the behavior of such an encryption to be emulated due to the computational closed environment that this device offers.

Using FPGAs for cryptography in the cloud also has important social implications since it could lower the cost of these services, this way contributing to a more widespread adoption of anonymity [Madhavapeddy et al., 2011].

4. RECONFIGURABLE HARDWARE AS SERVICES

Access to the reconfigurable resources through web services makes it possible to integrate reconfigurable computing and distributed applications in a synergy that brings a series of benefits for both technologies.

This could be the chance for reconfigurable computing to make the jump from a niche activity with low accessibility to a mainstream computing technology used together with other heterogeneous elements like GPU (Graphical Processing Units) and DSP (Digital Signal Processors) in cloud systems.

In our perspective, by applying SOA – Service-Oriented Architecture concepts to reconfigurable computing platforms, a minimizing complexity of the redesign and improved flexibility in IP (Intellectual Property) Core extensions support can be achieved.

This could potentially solve the problems that FPGA researchers still have with regard to issues like high design complexity and limited programmability. Providing IP Cores and hardware processors as function units available through web services, with a high degree of parallelism, this interaction can ease achieving high levels of performance with improved energy consumption management [Wang et al., 2012].

The major side-step in integrating reconfigurable platforms like FPGAs in cloud computing systems is the engineering effort of developing hardware applications that cloud customers have to deal with.

HDL – Hardware Description Languages might prove to be unattractive for cloud customers to use in their application development. In our perspective, the engineering of "hardware description services" is needed as new programming models in order to offer a development environment with high-level language compilers. Integrating such an environment in the cloud could make reconfigurable hardware more attractive for application deployment [Sergiyenko et al., 2013].

4. CONCLUSIONS

Cloud computing presents unique benefits for the IT industry but brings along also specific challenges. Given the complexity and rapid growth of the cloud computing infrastructure, and also the issues regarding the integrity and security of the data, the focus is turning more and more towards the underlying hardware resources.

In this paper we have identified the areas where implementations using reconfigurable hardware like FPGAs have the potential of bringing important benefits to cloud computing and we suggested specific approaches in a service-oriented perspective.

Reconfigurable computing implementations have great adaptability and scalability potential, which are important assets in cloud systems due to the permanent variations in load and demand. Another area where this approach represents an asset is the cloud security, where shifting from software to hardware is a major improvement to the traditional security models, the computational closed environment that devices like FPGAs could unlock encryption possibilities that are unachievable with traditional software implementations.

To fully benefit from the huge potential of reconfigurable hardware, and also to make "hardware description services" widely available for developers, this paper proposes focusing on the development of integrated development environments with high-level language compilers. Nevertheless, in our perspective, methods of cloud configurability could transcend from macro- to micro-structures, being „scalability-independant".

REFERENCES

1. Cheng, G., Zhu, Y., Rong, G., Qiu, M. (2012). *Prototyping High Efficiency Cloud Computing Architecture: Implementation of a Content Delivery Network Server on FPGA*. 7th International Conference on Computing and Convergence Technology (ICCCT)

2. Eguro, K.,Venkatesan, R. (2012). *FPGAs for Trusted Cloud Computing*. 22nd International Conference on Field Programmable Logic and Applications

3. Madhavapeddy, A., Singh, S. (2011) *Reconfigurable Data Processing for Clouds*. IEEE 19th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM), 2011

4. Romanca M., Ogrutan P.L. (2011), *Systems with embedded computers,* "Transilvania" Publishing House, Brasov

5. Sergiyenko, A., Simonenko, V., Vinogradov, Y., Gluchenko, K. (2013). *IP Core Synthesis in a Cloud.* Third International Conference "High Performance Computing" HPC-UA 2013

6. Stensgaard, M., Sparsø, J., (2008). *ReNoC: A Network-on-Chip Architecture with Reconfigurable Topology.* Second ACM/IEEE International Symposium on Networks-on-Chip

7. Unnikrishnan, D., Vadlamani, R., Liao, Y., Dwaraki, A., Crenne, J., Gao, L., Tessier, R. (2010) *Scalable Network Virtualization Using FPGAs.*

8. Wang, C., Li, X., Chen, P., Zhang, J., Feng, X., Zhou, X. (2012). *Regarding Processors and Reconfigurable IP Cores as Services*. IEEE Ninth International Conference on Services Computing

9. Yu, J., Zhu, Y., Xia, L.,Qiu, M., Fu, Y., Rong, G. (2011). *Grounding High Efficiency Cloud Computing Architecture: HW-SW Co-Design and Implementation of a Stand-alone Web Server on FPGA.* Fourth International Conference on the Applications of Digital Information and Web Technologies (ICADIWT), 2011

10. *** (2013) *CDN: A Cloud Accelerant, Whitepaper,* Internap Network Services Corporation.