

# Rendering as a Service

Keith Bugeja

Department of Computer Science  
University of Malta

High-fidelity rendering requires a substantial amount of computational resources to accurately simulate lighting in virtual environments [1]. While desktop computing, boosted by modern graphics hardware has shown promise in delivering realistic rendering at interactive rates, rendering moderately complex scenes may still elude single machine systems [2]. Moreover, with the increasing adoption of mobile devices, which are incapable of achieving the same computational performance, there is certainly a need for access to further computational resources that would be able to guarantee a certain level of quality.



**Fig. 1.** Mnajdra site, generated using our high-fidelity renderer, Illumina PRT

Cloud computing is a distributed computing paradigm for service hosting and delivery. It is similar to grid computing, but leverages virtualisation at multiple levels to realise resource sharing and provision [3], providing resources on demand and charging customers on a per-use basis rather than at a flat rate. Cloud technologies have the potential of providing a large number of resources to dedicate to a single application at any given point in time [5]. This paradigm offers the possibility of interactive high-fidelity graphics for users that do not have access to expensive dedicated clusters and can be delivered on any system from mobile to tablet to a desktop machine.

The aim of our research is that of leveraging the cloud infrastructure to provide high-fidelity rendering and thus present a first attempt into rendering as a service, via a system that can provide parallel resources for rendering which could be either dedicated or adapt to the servers' workload. As such, there exist challenges across three broad aspects: efficient high-fidelity rendering, resource allocation and task distribution, and low-latency client-server communication.

The focus of this presentation is the task distribution framework for cluster computing [4] that is currently being researched and developed; whereas grid and cluster infrastructures favour job submissions for medium to long term executions, our task distribution framework is targetted at short term execution and interactivity, in the vein of Many-task Computing [6]. We also consider future directions that our research might take, specifically vis-a-vis generalisation and heterogenous computing.

## References

1. Vibhor Aggarwal, Kurt Debattista, Thomas Bashford-Rogers, Piotr Dubla, and Alan Chalmers. High-fidelity interactive rendering on desktop grids. *IEEE Computer Graphics and Applications, pre print*, (99):1–1, 2010.
2. Vibhor Aggarwal, Kurt Debattista, Piotr Dubla, Thomas Bashford-Rogers, and Alan Chalmers. Time-constrained high-fidelity rendering on local desktop grids. In *EGPGV*, pages 103–110, 2009.
3. M. Armbrust, A. Fox, R. Griffith, A.D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, et al. A view of cloud computing. *Communications of the ACM*, 53(4):50–58, 2010.
4. R. Buyya et al. High performance cluster computing: Architectures and systems (volume 1). *Prentice Hall, Upper SaddleRiver, NJ, USA*, 1:999, 1999.
5. P. Mell and T. Grance. The nist definition of cloud computing (draft). *NIST special publication*, 800:145, 2011.
6. I. Raicu, I.T. Foster, and Y. Zhao. Many-task computing for grids and supercomputers. In *Many-Task Computing on Grids and Supercomputers, 2008. MTAGS 2008. Workshop on*, pages 1–11. IEEE, 2008.