

# Combinatorial Auction-Based Resource Allocation in the Fog

Lyndon Fawcett, Matthew Broadbent, and Nicholas Race

School of Computing and Communications, Lancaster University, UK  
{l.fawcett1, m.broadbent, n.race}@lancaster.ac.uk

**Abstract**—Network service composition is becoming increasingly flexible, thanks in part to advances in virtualisation and cloud technologies. As these penetrate further into networks, providers are often looking to leverage this infrastructure to improve their service delivery. This desire poses a number of obstacles, including a diversity in device capabilities and the need for a value exchange mechanism. In this demonstration, we present a platform that seeks to address a selection of these challenges.

**Index Terms**—Fog Computing, Combinatorial Auctions, Network Virtualisation, Provisioning, Orchestration

## I. INTRODUCTION

There is a growing trend towards highly configurable networks and services. Achieved through the movement of functionality into software [2], this progress enables network providers to dynamically adapt their provision in response to varying demands. Thus far, focus has been primarily on data centre and cloud environments. More recently, interest has shifted towards access and last-mile networks, encapsulating home and business environments in the process. This complete spectrum of networks and devices is called the *Fog* [1], [4].

When combined with Network Functions Virtualisation (NFV), the Fog presents a number of interesting opportunities for network operators: services can now be pushed even closer to the edge of networks, and in some cases, only a single hop away from the target devices. In conjunction with existing capabilities, this new enhancement allows for greater performance and efficiency, regardless of the type of service to be deployed.

The potential benefits are best described through use of an example. Content Delivery Networks (CDNs) store and host replica copies of content, primarily multimedia. These are then used to serve consumer requests, rather than retrieving the content from the origin server. In current networks, these content replicas are located in centralised strategic locations, such as Internet Exchange Points (IXPs).

Yet there are still a number of network hops to reach the final destination, which may ultimately be located in the user's home or place of work. CDNs exist primarily to improve the experience given to the user. Despite current deployment strategies, these final hops can nonetheless have an impact on the service delivered. In a Fog scenario, it would be possible to locate and serve content within the last-mile network, or even within the user's local network. This reduces the chance of network impairment impacting user experience.

## II. ORCHESTRATION

By making resources available at different and varied points within a network, a level of orchestration is required. This orchestration provides an abstraction for the delivery of services by handling the specifics of deploying such across heterogeneous hardware and software platforms. This is usually achieved using container technologies (such as Docker), which logically separate functionality into standalone objects.

Building on this, current NFV management and orchestration (MANO) platforms offer the ability to automatically deploy network services and scale them according to demand [5]. However, in the case of infrastructure management, they are generally restricted to existing cloud management systems. These platforms are designed for well managed, predictable and homogeneous infrastructures. Moreover, running a hypervisor on a resource-constrained Fog device may introduce noticeable overhead, or exclude some devices entirely.

These obstacles alone render current MANOs inappropriate for the diverse and fluid environment of the Fog, where devices need to be remotely managed, and alternative virtualisation techniques may be required. Furthermore, current MANOs provide rigid and limited orchestration logic, with no way of deciding where services should be placed. This is a key facet of the Fog, which cannot be exploited otherwise.

The unsuitability of current MANOs have led us to create our own orchestration system designed to be compatible with the uncertainty and heterogeneity of the Fog. This is the focus of the demonstration outlined in Section V.

## III. EVOLUTION OF THE FOG

Despite the identification of the aforementioned deficiencies, the rapid availability of related technologies shows the rate of progression towards fully realising such environments. We envision that as the field evolves, there will be three distinct generations of Fog deployment, each an evolution on the last:

- a) *1<sup>st</sup> Generation*: Infrastructure and service provider are the same entity. This is the current state-of-the-art.
- b) *2<sup>nd</sup> Generation*: Infrastructure provider and service provider are independent. Multiple service providers have a desire to use the same infrastructure (multi-tenancy).
- c) *3<sup>rd</sup> Generation*: Multiple independent infrastructure and service providers, including infrastructures owned by customers.

#### IV. MONETISATION

When Fog deployments reach the 2<sup>nd</sup> generation (or beyond), a clear value exchange must take place between the many service providers and the infrastructure provider. Market competition will be present, since there are likely to be multiple parties that wish to use the same set of resources.

To best handle this price competitive situation, and to ensure that the infrastructure provider is fairly recompensed for the use of their resources, we propose the use of an auction-based resource allocation function. Auctions are often adopted when a fair market value cannot be easily determined; buyers indicate the price they are willing to pay, and the winner is the buyer that bids the highest amount. Auction mechanisms are in widespread use today, including cloud resource provisioning. There are many ways of achieving this outcome, facilitated by a menagerie of usable methods.

In this work, we take an approach appropriate to the heterogeneous nature of a *Fog* environment: we propose the use of a combinatorial auction [3]. This allows bidders to define bids containing combinations of discrete sets of resources. For example, a bid could contain a *package* consisting of 128MB of RAM, 2 CPUs and 10GB of storage across 400 devices located specifically in local home networks. Importantly, the purpose of these bids is to reserve and allocate those resources for a fixed period of time; resources are auctioned again for other adjacent time slots. This system enables multiple providers to use the same set of resources at different times during the same day.

Returning to the previous CDN example, this system would allow such service providers to deploy content replication capabilities into networks previously considered out of reach, either due to associated maintenance costs or operational constraints. To maximise the impact of this content delivery, a CDN would need to successfully bid against other service providers, and possibly other rival CDNs, to reserve resources in time slots where requests for content are at their peak.

#### V. PLATFORM DESIGN AND DEMONSTRATION

To better illustrate how this platform can efficiently allocate resources in a way that satisfies both infrastructure and service providers, we will describe the intrinsic components and their interactions. This forms the basis for our proposed demonstration.

Figure 1 shows the various components of the architecture. The highest layer contains the *Service Providers*. These are the parties interested in using the infrastructure. They communicate directly with an *Auctioneer*, who provides details of the resources available and details of forthcoming auctions.

Using this information, *Service Providers* can bid on live auctions, submitting a number of alternative and/or complementary *packages* that they would like to reserve. Using a simple definition language, logic can be applied to these bids to ensure they receive the desired outcome. Once bidding is over, the *Auctioneer* will alert each *Service Provider* as to which of these reservations has been successful. The *Service Providers* will then supply details of the service they wish

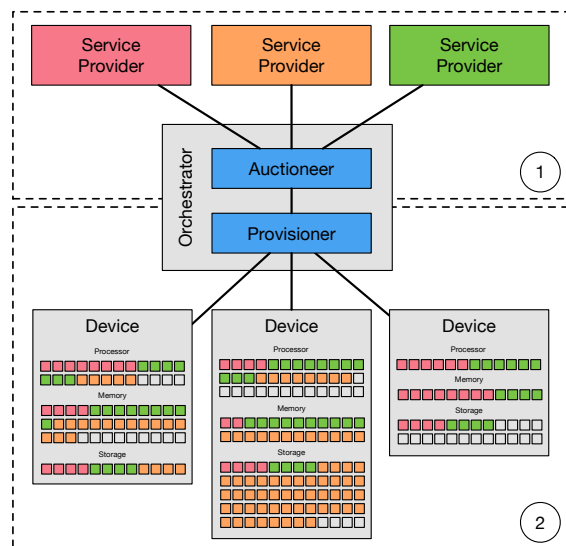


Fig. 1. Auction-Based Resource Allocation Platform

to run on their reserved infrastructure, using a supported templating language (such as Docker’s *Dockerfiles*). This is Phase 1 (as illustrated) completed.

Phase 2 begins the realisation of the aforementioned reservation. The *Auctioneer* informs the *Provisioner* of each successful reservation, which is then actioned on the constituent devices. This includes creating the relevant containers and services, and ensuring that they are kept within their particular resource constraints. The service is now deployed, and will remain until the reservation window expires. When this occurs, the *Provisioner* will clear the existing reservations, and make way for new services to be deployed.

The demonstration will show the auctioning and provisioning process in action, enabled through our orchestrator. It will utilise a variety of physical and virtual infrastructures, such as Raspberry Pi devices and OpenStack instances. This emulates the variety found in the *Fog* environment, and includes a competitive auction between a number of service providers. This will be illustrated through a live web-based interface, highlighting both phases, and demonstrating the provisioning of a service provider’s winning bid.

#### VI. ACKNOWLEDGEMENT

We are very grateful to Lancaster University for funding Lyndon Fawcett’s PhD studentship in association with the EPSRC funded TOUCAN project.

#### REFERENCES

- [1] F. Bonomi *et al.*, “Fog Computing and its Role in the Internet of Things,” in *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*. ACM, 2012, pp. 13–16.
- [2] M. Chiosi *et al.*, “NFV Introductory White Paper,” 2012. [Online]. Available: [http://portal.etsi.org/nfv/nfv\\_white\\_paper.pdf](http://portal.etsi.org/nfv/nfv_white_paper.pdf)
- [3] P. Cramton *et al.*, “Combinatorial auctions,” 2006.
- [4] A. V. Dastjerdi *et al.*, “Fog Computing: Principles, Architectures, and Applications,” *CoRR*, vol. abs/1601.02752, 2016. [Online]. Available: <http://arxiv.org/abs/1601.02752>
- [5] R. Mijumbi *et al.*, “Management and Orchestration Challenges in Network Functions Virtualization,” *IEEE Communications Magazine*, vol. 54, no. 1, pp. 98–105, 2016.