

LETTER

Open Access

Notes on Cloud computing principles

Thomas Sandholm^{1,2*} and Dongman Lee²

Abstract

This letter provides a review of fundamental distributed systems and economic Cloud computing principles. These principles are frequently deployed in their respective fields, but their interdependencies are often neglected. Given that Cloud Computing first and foremost is a new business model, a new model to sell computational resources, the understanding of these concepts is facilitated by treating them in unison. Here, we review some of the most important concepts and how they relate to each other.

Keywords: Resource allocation; Provisioning; IT-economics

Introduction

Imagine that you have to go on a trip to meet a friend in a different city. There are many modes of transportation available to you. You can drive there by car, take a taxi, share a ride in a van, take a bus or a train, or even fly there in an airplane. Your choice is determined by your general preference for these options. In particular, your choice depends on the economics and convenience of these alternatives given the characteristics of the trip, including distance to destination and time available. The cost of the choice you make in turn is related to how many other people are sharing the same mode of transportation, and how expensive it is to operate the transportation vehicle and infrastructure.

Now compare this choice to the choice of energy supplier that people faced in the early 20th century. You could buy your own electric generator, but it was not very cost efficient if your needs varied diurnally or seasonally. As it became apparent that electricity was as invaluable of a commodity as gas, water and the telephone, utility companies and national electrical grids that could aggregate and distribute electricity on demand replaced the privately owned generators.

Cloud computing [1] could be seen as an effort to commoditize computing, and distribute and operate it as efficiently as the electrical grid while still offering consumers the plethora of alternatives known from the transportation domain. The pre-cloud era could be compared to

everyone driving around in their own car and using their own generators. The cloud era allows computing to be used similarly to public transportation and makes it possible to tap into computing power with the same ease that you plug in your appliances to the electrical grid at home. To distinguish the Cloud from its predecessors it is often defined as a use of computing resources that are delivered as a service over a network. The way in which you provision these services holds the key to the innovation.

Cloud services need to be scalable, fault-tolerant, highly available, high-performance, reliable and easy to use, manage, monitor, and provision efficiently and economically. One early realization by Cloud computing pioneers was that meeting all these requirements for services handling massive amounts of data and huge numbers of concurrent users called for innovation in the software stack as opposed to the highly specialized hardware layer. The hardware is reduced to a commodity and the Quality of Services (QoS) are instead provided by a fully integrated and hardware agnostic software stack. Virtualization became the new silver bullet.

As the demand for computing power increased with more users coming on-line and more data being published on-line it became apparent that some drastic architectural changes had to be introduced to provision compute resources more efficiently. The most prominent enabler for efficient resource provisioning was data center consolidation. Instead of using spare cycles from arbitrary privately owned nodes in a network^a, it was more cost effective to provide high QoS by consolidating computing in highly streamlined data centers packed with low-cost dedicated compute and storage clusters in a highly reliable

*Correspondence: thomas.e.sandholm@hp.com

¹ HP Labs, 1501 Page Mill Rd, 94304 Palo Alto, CA, USA

² Department of Computer Science, KAIST, 291 Daehak-ro, 305-701 Daejeon, Korea

and fast network. These data centers were also frequently deployed in areas where energy and labor were cheap to further cut operational costs.

Data-center consolidation and more aggressive sharing of compute resources lead to the following key benefits of Cloud computing:

1. Lower cost of using compute resources
2. Lower cost of provisioning compute resources
3. Reduced time-to-market

The first benefit can be attributed to only paying for the resources when you use them. When you do not use them, the provider can allocate them to other users. Being able to host multiple users or tenants on the same infrastructure allows the provider to utilize the resources more efficiently and thereby increase the return on investment (ROI). This win-win relationship between users and providers is the reason why most companies switch to Cloud architectures. The growth and sudden popularity of Cloud computing was, however, not fueled by traditional, established companies. Start-ups were the pioneering users of Cloud technology as it reduced their time-to-market and provided them with less up-front risk to stand up a demo or beta version. If the users did not flock, not much harm was done, you just stopped paying for the resources. If there was an unexpected flash crowd of people bombarding the service, you would just pay for more resources. This type of usage is often referred to as the elasticity of the Cloud. The Cloud allows you to scale down as easily and as quickly as you scale up.

Below we will review some of the fundamental concepts of distributed computing at scale, and then relate these concepts to economic principles that help us understand the trade-offs governing their deployment. The main motivation for studying these economic principles is that solely maximizing systems metrics, such as, throughput, response time and utilization may not always be the most profitable strategy for a Cloud provider.

Before delving into these principles we will first take a look back at technologies that predated Cloud computing to see how the architecture of this new computing paradigm evolved into its current state.

Historical evolution

The vision of organizing compute resources as a utility grid materialized in the 1990s as an effort to solve grand challenges in scientific computing. The technology that was developed is referred to as Grid Computing [2], and in practice involved interconnecting high-performance computing facilities across universities in regional, national, and pan-continent Grids. Grid middleware was concerned with transferring huge amounts of data, executing computational tasks across administrative

domains, and allocating resources shared across projects fairly. Given that you did not pay for the resources you used, but were granted them based on your project membership, a lot of effort was spent on sophisticated security policy configuration and validation. The complex policy landscape that ensued hindered the uptake of Grid computing technology commercially. Compare this model to the pay-per-use model of Cloud computing and it then becomes easy to see what, in particular, smaller businesses preferred. Another important mantra of the Grid was that local system administrators should have the last say and full control of the allocation of their resources. No remote users should have full control or root access to the expensive super computer machines, but could declare what kind of software they required to run their jobs. Inherently in this architecture is the notion of batch jobs. Interactive usage or continuous usage where you installed, configured and ran your own software, such as a Web server was not possible on the Grid. Virtual machine technology [3] released the Cloud users from this constraint, but the fact that it was very clear who pays for the usage of a machine in the Cloud also played a big role. In summary, these restrictions stopped many of the Grid protocols from spreading beyond the scientific computing domain, and also eventually resulted in many scientific computing projects migrating to Cloud technology.

Utility computing [4] refers to efforts in the industry around the turn of the millennium to improve manageability and on-demand provisioning of compute clusters. At this time, companies were very skeptical to running their confidential workloads off premise and thus utility computing was often sold on a cluster-by-cluster basis and installed on a company-by-company or organization-by-organization basis. This deployment model made it very expensive to get up and running, which ironically had been one of the key claimed benefits of utility computing. Nevertheless, it started to become clear around this time that virtualization was the key to on-demand provisioning of compute resources. Web services and Service-Oriented Architectures [5] were touted as the solution to many of the problems seen in the earlier efforts of Utility and Grid computing. Providing a standard API would allow infrastructure to be allocated programmatically based on demand. The APIs and protocols were borne out of the evolution of the World Wide Web (WWW) that started to provide more dynamic and interactive content on Web pages leading to the phenomenon of mashups. Mashups in the early days essentially scraped HTML from various Web pages to dynamically create a value-adding service on a new Web page. As this was error prone, it was quickly realized that APIs were needed and the first Web services protocols, such as SOAP [6], were designed.

By the time Amazon launched their Elastic Compute Cloud (EC2) service in 2006, both Web service APIs

and virtualization technology (e.g. Xen[3]) were mature enough to form a compelling combination or a perfect storm to deliver the first real public utility computing service that had been envisioned a decade earlier.

In summary, the vision of the Grid combined with Virtual Machine technology and Web service APIs were the essential characteristics of the first Clouds. Next, we will review the fundamental distributed systems principles underlying today's Cloud systems.

Computational principles

Next, we provide a brief review of the most important computational systems principles related to the new Cloud computing era. These concepts have been presented in-depth in many pre-existing review articles, so here we only cover them at a level of detail that helps the reader appreciate the economic implications discussed in the second part of this letter. A reader already familiar with the systems-related principles of Cloud computing may skip forward to the section on Economic principles. However, to avoid confusion about taxonomy, this section may be revisited as a reference for technical definitions of these widely referred-to concepts.

Multi-tenancy

A tenant in the Cloud context is a user of Cloud infrastructure, i.e. Infrastructure-as-a-Service (IaaS) services [7]. A VM owner is an example of a tenant and if multiple VM owners are allocated on the same physical machine it is an example of multi-tenancy [8]. The difference between a multi-(end)-user service and a multi-tenant service is that a multi-user offering may benefit from having users know about each other and explicitly share social content to promote the network effect. A multi-tenant solution could internally benefit from shared physical resources but must give the impression of an exclusive offering to each of the tenants. As an example, hosting the Facebook service on a Web server in the Cloud would be an example of a multi-user service, but hosting both a Twitter Web server and a Facebook Web server in the same Cloud data center would be an example of multi-tenancy. From this definition, it is clear that the IaaS provider needs to provide mechanisms to isolate the tenants from each other.

Multiple tenants need to be isolated in terms of privacy, performance and failure:

- **Privacy Isolation.** Multiple tenants must not have access to each other's data. This may seem like an easy requirement to meet but in a typical file system there may be traces left after a file even after removing it, which would violate this property.
- **Performance Isolation.** Multiple tenants must not be effected by each other's load. If one tenant starts

running a CPU intensive task and other tenants see a drop in performance as a result, then this property is violated.

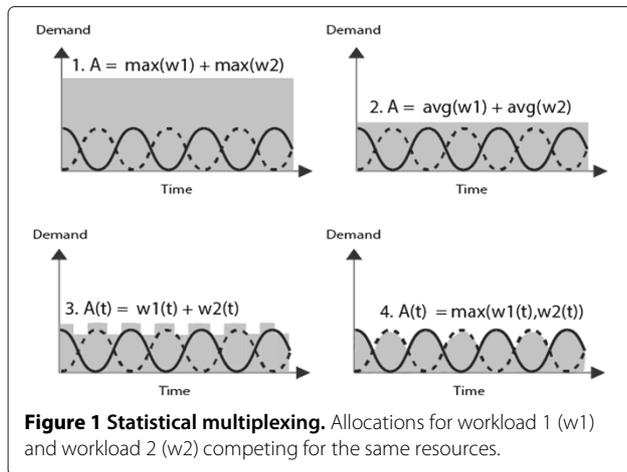
- **Failure Isolation.** If a tenant either inadvertently or maliciously manages to crash its compute environment, it should not effect the compute environment of other users. Imagine a Java VM hosting multiple applications such as a Tomcat Servlet engine. Now, if one servlet Web app crashes the VM, then the other apps in the same VM would also crash. This failure would in that case be a violation of the failure isolation property. Virtual machines offer a popular technique to ensure isolation, but in some cases the overhead of virtualization, of e.g. IO and network, is too high so a trade-off has to be made between isolation level and performance.

Ensuring these levels of isolation is closely related to the strategy used to allocate resources to tenants, which we will discuss next.

Statistical multiplexing

One major benefit related to data center consolidation that we discussed in the introduction is statistical-multiplexing [9]. The idea behind statistical multiplexing is that bursty workloads that are consolidated on the same Cloud infrastructure may in aggregate display a less bursty pattern. Figure 1 shows an example of statistical multiplexing with two workloads exhibiting complementing demand over time.

1. Without an elastic Cloud infrastructure, the most common way of provisioning resources to tenants is to allocate resources that meet the peak demand of each workload. Clearly, this leads to a major waste in resources for the majority of the time. Statistical multiplexing allows an allocation that is substantially lower than the sum of the peaks of the workloads.
2. Ideally if statistical multiplexing is applied on a large number of independent workloads, the aggregate will be stable, i.e. a straight line in the demand chart. If this is the case, it is enough to just allocate the sum of the averages of resource demand across all workloads.
3. Now assuming that we are in an elastic Cloud environment and we can slice resource allocations by time akin to how an OS time-shares CPU between processes. In this scenario, further reductions in resource allocations may be achieved by simply allocating the sum of resource demand across all workloads in each time slice.
4. Finally if each time slice only has a single workload active at any point in time, the allocation reduces to just the maximum demand across the workloads.



This model of perfect statistical multiplexing is hard to achieve in practice. The main reason for this is that workloads tend to be correlated. The effect is known as self-similarity. Self-similar workloads have the property that aggregating bursty instances will produce an equally bursty aggregate, something that is often observed in practice. However, there are many techniques to recreate the effects of statistical multiplexing without having to hope for it to occur organically. For instance, you could measure the correlation between workloads and then schedule workloads that are complementing on the same resources. These techniques are sometimes referred to as optimal packing of workloads or interference minimization [10]. Poor statistical multiplexing tends to lead to low utilization, or unmet demand, as we will discuss further when we review the economic principles governing under and over-provisioning.

Horizontal scalability

An application or algorithm that runs in the Cloud will not be able to scale up and down with the infrastructure unless it can run at least in part in parallel. Execution in the Cloud requires efficient scaling across machines, referred to as horizontal scalability. A local program running on a single machine on the other hand only needs to scale vertically, i.e. run faster as local resources such as CPU, memory, and disk are added. How well a program scales is thus related to the parallelizability of its algorithms. This effect is formalized in what is called Amdahl's Law [11]:

$$T(n) = T(1)(B + (1 - B)/n) \quad (1)$$

Amdahl's Law predicts the expected speed-up of a program or algorithm when run over multiple machines. $T(n)$ is the time taken to run on n machines. B is the fraction of the program that needs to run serially, i.e. that cannot be parallelized. Note that several disjoint sections in the

execution path may need to run serially to collect, distribute or synchronize parallel computations. It is clear that minimizing B maximizes the speedup. However, the most important consequence of Amdahl's Law is that it sets a theoretical cap on how many machines a program will benefit from running on, beyond which point adding new machines will not make the program faster. If B is close to negligible we can expect *linear scalability*. Adding x machines will make the program run x times faster. If the program speedup grows at a slower rate than the number of machines added, which is the common case due to various overheads of distribution, we refer to *sublinear scalability*. The program may also speedup at a faster rate than the machines being added, in which case the program is said to exhibit *superlinear scalability* (see Figure 2). This effect may happen if there is some common resource like a shared cache that benefits from more usage, e.g., more cache entries and fewer cache misses.

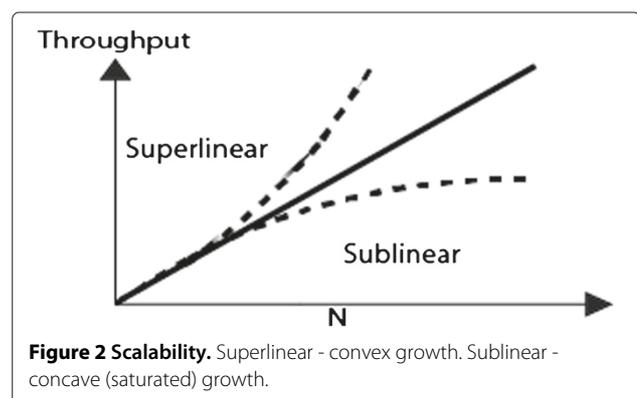
Many advances in the database community have emerged, and been popularized, to achieve horizontal scalability when processing massive amounts of data in Clouds. The most prominent concepts include: data partitioning and sharding [12]; consistent hashing and distributed hash tables (DHT) [13,14]; and eventual and quorum consistency [15,16].

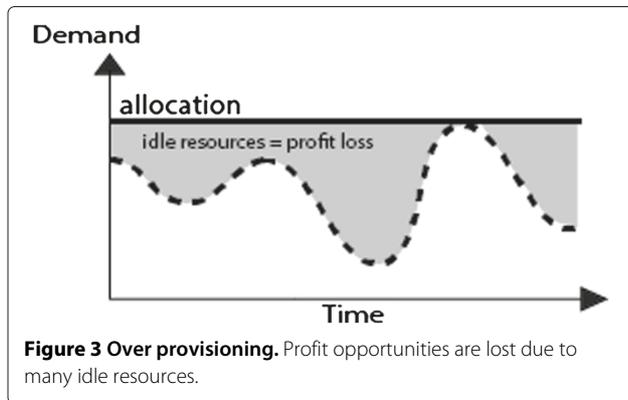
Economic principles

Now, we discuss the economic implications of the new Cloud systems principles to see how economic theories may help with various systems trade-offs that from a pure computational perspective seem insurmountable to overcome without fine-tuning heuristics and trial and error.

Over and under provisioning

As we alluded to in the section on statistical multiplexing, over-provisioning is a common strategy for allocating resources across tenants. Here we discuss the economic dilemma of over (Figure 3) versus under-provisioning (Figure 4) resources.



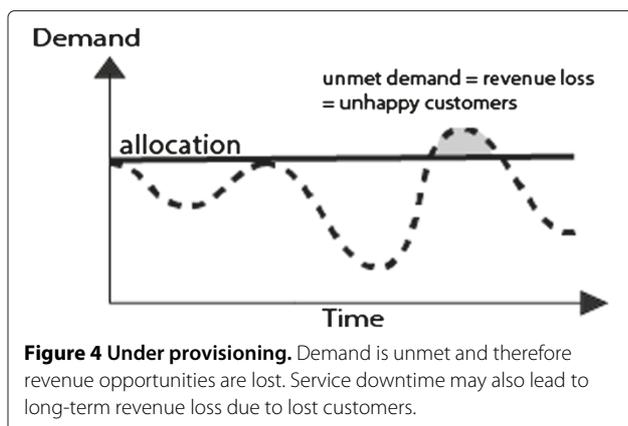


We can see that over-provisioning leads to a large area of idle resources over time. In financial terms this means high-operational cost, and lost opportunities to increase profit. To increase profit the IaaS provider may be tempted to lower the allocation to reduce the operational cost as seen in Figure 4. However, this leads to an even more severe drawback, unmet demand. Unmet demand means revenue loss, and can have long-term negative effects as customers who are denied access to a resource despite being willing to pay for it may not return. For this reason over-provisioning is more popular than under-provisioning. However, neither the IaaS provider nor the tenant may be able to perfectly predict the peaks, after all that is why they are running in the Cloud in the first place. In this case under-provisioning may occur inadvertently.

Hence, over-provisioning versus under-provisioning involves making a trade-off between profit and revenue loss.

Variable pricing

Given all the issues of allocating resources to bursty demand, it is natural to ask whether this burstiness can be suppressed somehow as opposed to being accommodated. That is exactly the idea behind variable pricing



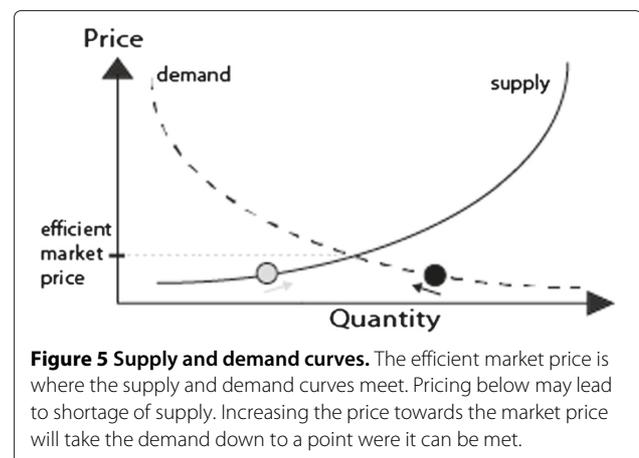
or demand-driven pricing. The idea is to even out the peaks and valleys with incentives. If the demand is high we increase the price. This leads to tenants who cannot afford the higher price to back-off and thereby demand is reduced. On the other hand, if the demand is low, a price drop may encourage tenants who would otherwise not have used some resources to increase their usage and thereby demand. The end result is a stable aggregate demand as in the statistical multiplexing scenario. The key benefits to IaaS providers include the ability to cash in on peak demand by charging premiums, and a mechanism to increase profit during idle times. Now, how can we ensure that the price is a good representation of demand? Here, microeconomic theory of supply and demand [17] helps.

If we plot the quantity of goods a supplier can afford to produce given a price for the good we get the supply curve. If we plot the quantity of goods requested by consumers given a price for the good we get the demand curve. The price at the point where the supply and demand curves meet is called the efficient market price as it is a stable price that a market converges towards (see Figure 5). To see why this is the case, consider the gray dot on the supply curve in Figure 5. In this case the supplier observes a demand that is higher than the current quantity of goods produced. Hence, there is an opportunity for the supplier to increase the price of the good to afford to produce more goods to meet this demand. Conversely, considering the black dot on the demand curve, we can see that the demand is higher than the volume of goods that the supplier can produce. In this case the demand will naturally go down and the consumers are likely to be willing to pay a higher price to get their goods.

In general, variable pricing allows a provider to allocate resources more efficiently.

Price setting

There are many ways to set prices for goods in a market. The most commonly known are various forms of auctions,



spot prices and reservations. In auctions, bidders put in offers to signal how much they are willing to pay for a good. In double actions, there are also sellers who put in asks denoting how much they are willing to sell the good for. The stock market is an example of a double auction. In computational markets, second price sealed bid auctions are popular since they are efficient in determining the price, i.e. reflect the demand, without too much communication. All bidders put in secret bids and the highest bidder gets the good for the price equalling the second highest bid.

In the case where there is not a completely open market price, and there is just a single provider selling off compute resources, spot pricing is a common way of setting demand based prices. The spot price is computed on a running basis depending on the current level of demand. There could for instance be a base pay that is discounted or hiked based on demand fluctuations. A spot market differs from a futures market in that goods are bought and consumed immediately. Futures markets such as options are less common in practical computational markets today.

Purchasing resources on a spot market involves a high risk of either having to pay more for the same allocation or being forced to reduce the allocation to stay within budget (see the section on Predictability below). A common way to reduce the risk for consumers is to offer a reservation market. A reservation market computes the expected spot demand for some time in the future and adds a premium for uncertainty to arrive at a reservation price. Essentially you have to pay for the provider's lost opportunity of selling the resources on the spot market. This way the risk is moved from the consumer of compute resources, the tenant, to the provider. I.e., the provider's actual cost or revenue when providing the resource may vary, whereas the cost for the tenant is fixed. If there is an unexpected hike in the demand and all resources have already been promised away in reservations there is no way for the provider to cash in on this demand, which constitutes a risk for the provider.

The research field of computational economies have tackled these problems as far back as the 1960s and 70s [18-20]. More recent computational market designs include [21-23]. Reviews of some of these designs can be found in [24,25].

In summary, reservation markets move the risk of uncertain prices from the tenant to the provider as uncertain demand.

The tragedy of the commons

The next principle we will discuss is a social dilemma referred to as the tragedy of the Commons [26]. The dilemma was introduced in a paper in 1968 by Garrett Hardin, where the following scenario was outlined.

Imagine a public, government-owned piece of land with grass, in the UK referred to as a Common. Now, a number of shepherds own sheep that they need to feed on this Common to keep alive. The shepherds will benefit economically from the sheep because they can, for instance, sell their wool. Each shepherd faces the financial decision whether it would be more profitable to purchase another sheep to feed on the Common and extract wool for, or provide more food to each sheep by sticking with the current herd. Given that it is free to feed the sheep on the Common and the reduction in available food is marginal, it turns out that it is always optimal for a selfish shepherd trying to optimize his profit to buy another sheep. This has the effect of driving the Common into a slump where eventually no more grass is available and all sheep die and all shepherds go bankrupt.

One could argue that less selfish shepherds who are wary of the benefits of the group of shepherds as a prosperous community will not let the situation end in tragedy. However, there are many examples of communities that have gone extinct this way. In general what these communities have in common is that there is a high degree of free-riders, i.e. community members who take more from the common resources of the community than they give back. Sometimes the effects are temporal and not as obvious since no one purposefully abuses the community. One example is the PlanetLab testbed [27] used by systems researchers in the US. The testbed is distributed across a large number of organizations to allow wide area and large-scale experiments. The weeks leading up to major systems conferences such as OSDI, NSDI, SOSP and SIGCOMM see extreme load across all machines in the testbed typically leading to all researchers failing to run their experiments.

The opposite of free-riding is referred to as altruism. Altruists care about the community and are the backbone of a sustainable and healthy community. A good example of this is the Wikipedia community with a small (compared to readers) but very dedicated group of editors maintaining the order and quality of the information provided. The opposite of the tragedy of the Commons is the network effect where more users lead to greater benefits to the community, e.g. by providing more content as in the Wikipedia case.

The balance between free-riders and altruists as well as the regulations and pricing of resource usage determines whether the tragedy of Commons or the network effect prevails.

This concept is closely related to what economists refer to as externality [28], individual actions impose an unforeseen positive or negative side-effect on the society. The archetypical example is factory pollution. Such side-effects are mainly addressed in the Cloud by various infrastructure isolation designs such as virtual machines,

or virtual private networks (see discussion in the section on Multi-tenancy above).

Incentive compatibility

One of the most frequently overlooked aspects of distributed systems is incentive compatibility [29]. Yet it is a property that all successful large-scale systems adhere to, the Cloud being no exception, and it is very often the main reason why proposed systems fail to take off. It is a concept borrowed from game-theory. In essence, an incentive compatible system is a system where it is in the interest of all rational users to tell the truth and to participate. In a systems context, not telling the truth typically means inserting incorrect or low quality content into the system to benefit your own interests. Incentive to participate is closely related to the notion of free-riding. If there is no incentive to contribute anything to a common pool of resources, the pool will eventually shrink or be overused to the point where the system as a whole becomes unusable. That is, the system has converged to a tragedy of the Commons. Ensuring that the system cannot be gamed is thus equivalent to ensuring that there is no free-riding and that all users contribute back to the community the same amount of valuable resources that they take out. A new, untested, system with a small user base also has to struggle with a lack of trust, and in that case it is particularly important to come out favorable in the individual cost-benefit analysis, otherwise the potential users will just pick another system. Tit-For-Tat (TFT) is an example of an incentive compatible algorithm to ensure a healthy and sustainable resource sharing system.

If Cloud resources are sold at market prices it ensures incentive compatibility, .i.e. ensuring that the price is following the demand (in the case of a spot market) or the expected demand (in the case of a reservation market) closely has the effect of providing an incentive for both suppliers and consumers to participate in the market. Earlier systems such as the Grid and P2P systems that did not have an economic mechanism to ensure incentive compatibility has historically had a much harder time of sustaining a high level of service over a long period of time due to frequent intentional and non-intentional free-riding abuses. Hence, demand-based pricing helps ensure incentive-compatibility.

Computational markets that have demand-driven pricing may however still not be incentive compatible. If it for instance is very cheap to reserve a block of resources ahead of time and then cancel it before use, it could lead to an artificial spike in demand that could dissuade potential customers from using the resource. This in turn would lead to the spot market price being lower, which could benefit the user who put in the original reservation maliciously. In economic terms, it is a classic example of someone not telling the truth (revealing

their true demand in this case) in order to benefit (getting cheaper spot market prices). Another classic example is an auction where the bidders may overpay or underpay for the resource, just to make sure competitors are dissuaded to participate or to falsely signal personal demand.

Efficiency

Shared resource clusters such as the Grid are commonly monitored and evaluated based on systems metrics such as utilization. A highly utilized system meant the resources typically funded by central organizations such as governments were being efficiently used. This type of efficiency is referred to as computational efficiency. It is a valuable metric to see whether there are opportunities to pack workloads better or to re-allocate resources to users who are able to stress the system more, .i.e. a potential profit opportunity (see the section above on Over and under provisioning). In a commercial system such as the Cloud it is also important to consider the value that the system brings to the users, because the more value the system brings to users the more they are willing to pay and the higher profit the Cloud provider is able to extract from a resource investment. This trade-off becomes apparent when considering a decision to allocate a resource to a user who is willing to pay \$0.1 an hour for some resource and utilize at close to 100% versus another user who is willing to use the same resource over the same period of time but at 90% utilization and paying \$0.5 an hour. There is likely more idle time and unused resources if the second user is accommodated but the overall profit will be higher ($0.5 - 0.1 = \$0.4/\text{hour}$).

To evaluate the economic efficiency [30] one therefore often goes beyond pure system metrics. In economics, utility functions are used to capture the preferences or the willingness of a user to pay for a resource. Maximizing the overall utility across competing users is then a common principle to ensure an overall healthy and sustainable ecosystem. This sum of utilities across all users is referred to as the *social welfare* of the system. To compare two systems or two resource allocation mechanisms for the same system one typically normalizes the social welfare metric by comparing the value to an optimal social welfare value. The optimal social welfare value is the value obtained if all users (in the case of no contention) or the highest paying user receive all the resources that they desire. Economic efficiency is defined as the optimal social welfare over the social welfare obtained using an actual allocation strategy. A system with an economic efficiency of 90%, for instance have some opportunity, to allocate resource to higher paying users and thereby extract a higher profit.

In essence, ensuring economic efficiency involves optimizing social welfare.

There is however an argument to be made that always allocating to the highest paying user does not create a healthy sustainable ecosystem, which we will discuss next.

Fairness

Consider the case where some user constantly outbids a user by \$.0001 every hour in a competitive auction for resources. An economically efficient strategy would be to continuously allocate the resource to the highest bidder. The bidder who keeps getting outbid will however at some point give up and stop bidding. This brings demand down and the resource provider may lose out on long term revenue. It is hence also common practice to consider the fairness of a system. In economics, a fair system is defined in terms of envy between users competing for the same resource [31]. Envy is defined as the difference in utility that a user received for the actual allocation obtained compared to the maximum utility that could have been obtained across all allocations for the same resource to other users. The metric is referred to as envy-freeness and a fair system tries to maximize envy freeness (minimize envy). Having high fairness is important to maintain loyal customer, and it may in some cases be traded off against efficiency as seen in the example above. Fairness may not be efficient to obtain in every single allocation instance, but is commonly evaluated over a long period of time. For example a system could keep track of the fairness deficit of each user and try to balance it over time to allocate resources to a user that has the highest fairness deficit when resources become available.

In addition to fairness considerations, there could be other reasons why a resource seller may want to diverge from a pure efficiency-optimizing strategy. If information is imperfect and the seller needs to price goods based on the expected willingness to pay by consumers, it may be a better long-term strategy to set the price slightly lower to avoid the dire effects of losing trades by setting the price to high. Another reason may be that some consumers have less purchasing power than others, and giving them benefits, so they can stay in the market, improves the overall competitiveness (and liquidity, see below) of the market, which in turn forces the richer consumers to bid higher.

Liquidity

The central assumption in variable pricing models (see the section above on Variable pricing) is that the price is a proxy or a signal for demand. If this signal is very accurate, allocations can be efficient and incentives to use versus back off of resources are well aligned. If there are too few users competing for resources the prices may plummet and the few users left may get the resource virtually for free. It is therefore critical for a provider to have enough competing users and to have enough purchases of resources for all the market assumption to come into

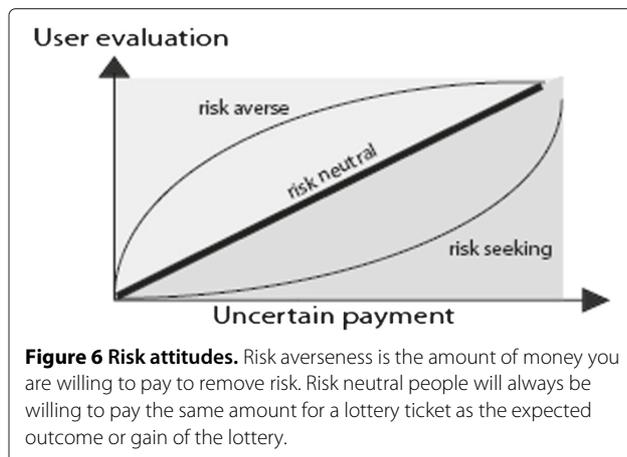
play. In particular, this means ensuring that the second part of incentive compatibility is met, i.e. users have an incentive to participate. Most providers fall back on fixed pricing if there is too little competition, but that may lead to all the inefficiency that variable pricing is designed to address. In economics, this volume of usage and competition on a market is referred to as liquidity [32]. Lack of liquidity is a very common reason for market failure, which is why many financial and economic markets have automated traders to ensure that there is a trade as long as there is a single bidder who sets a reasonable price. A provider may, for instance, put in a daemon bidder to ensure that resources are always sold at a profit.

Predictability

The biggest downside of variable pricing models is unpredictability. If the price spikes at some time in the future, the allocation may have to drop even though the demand is the same to avoid breaking the budget. Exactly how much budget to allocate to resources depends on the predictability of the prices, i.e. the demand. If the demand is flat over time, very little excess budget has to be put aside to cope with situations where resources are critically needed and demand and prices are high. On the other hand, if some application is not elastic enough to handle resource variation, e.g. nodes being de-allocated because the price is too high, a higher budget may need to be allocated to make sure the application runs at some minimal level of allocation.

Essentially users as well as applications have different sensitivity to risk of losing resource allocations or resources being more expensive. In economics the attitude towards risk is described in the risk-aversion or risk attitude property of a user. There are three types of users that differ in how much they are willing to spend to get rid of risk (variation) [33]. Risk-averse users will spend more money than the expected uncertain price (i.e. hedge for future spikes c.f. the discussion on over-provisioning and under-provisioning) [34]. Risk-neutral users will spend exactly the expected price. Finally, risk-seekers will put in a lower budget than the expected price to meet their allocation needs (see Figure 6). An application that is perfectly elastic and that may scale down or up over time as long as the long term performance is guaranteed may choose a risk neutral strategy. Risk seekers are less common in computational markets, but they may be bettering on demand going down in the future. Risk-averse users are the most common group, and the premium they pay above the expected price is a good indicator for how much a resource provider can charge for reservations, which essentially eliminates this uncertainty.

In summary, the elasticity of a Cloud application is highly related to the risk-aversion of the resource purchase, i.e. how much to pay to hedge uncertainty.



Discussion

Even though many of the economic approaches to computational resource allocation have been known since the 1960s, the adoption has been very slow. One of the reasons may be the assumption of instant, low-latency scaling and friction-less allocation, a.k.a. elasticity assumed by the economic models. Another, may be the limited opportunities to large-scale sharing and co-location of workloads, as many private firms are very sensitive to share their computational resources with others. The success of Public Clouds, such as Amazon EC2, has brought many of these economic concepts back into mainstream usage again. One example is the now fully operational Amazon Spot Market. There have been a multitude of attempts in the past to deploy such markets, but they do not start providing tangible benefits to consumers and providers until the markets reach a certain level of maturity and liquidity. It is informative to study when Amazon thinks Spot instances should be used [35]. The key stated reasons include, scaling out large, low-risk computations whenever either the market demand is low, and thus prices are low, or when large computations need to run unexpectedly and the planned capacity is not sufficient.

Another recent example of a more sophisticated computational market is the Deutsche Börse Cloud Exchange [36]. This exchange allows IaaS providers to sell resources at a centralized exchange to avoid vendor lock-in and to spur competition for more efficient pricing of commodity Cloud resources. The main argument to sellers is that their sales volume would increase, and to consumers that the prices would be lower. Trust is also an important factor that would allow smaller providers to sell through a well-known stock exchange. As these markets and others mature, the Economic principles discussed here will start having a bigger impact on how we provide, consume, and design Cloud resource infrastructure in the future.

Summary

We have discussed some computational principles underlying the efficient design of Cloud computing infrastructure provisioning. We have also seen how economic principles play a big role in guiding the design of sustainable, profitable, and scalable systems. As Cloud computing becomes more commonplace and more providers enter the market, the economic principles are likely to play a bigger role. The sophistication of the market designs depends very much on the level of competition and usage, a.k.a. as the liquidity of a market.

The key to a successful market design is to align the incentives of the buyers and sellers with those of the system as a whole. This will ensure participation and liquidity. Most computational principles in the Cloud are governed by the notion that large scale distributed systems see failures so frequently that failover and recoverability must be an integral part of the software design. In order to failover successfully one needs to have full programmatic control from hardware to end-user application. An ongoing trend has been to develop platforms and cloud operating systems that offer this level of software control of hardware to automate administration, management, and deployment dynamically based on demand.

Endnote

^adone in many P2P networks at the time.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

This text is based on a distributed systems class co-taught and co-developed by the authors at KAIST during the fall of 2013. TS contributed mostly to the economics section, and DL contributed mostly to the systems section. Both authors read and approved the final manuscript.

Acknowledgements

This work was partially supported by the IT R&D program of MSIP/KEIT. [10045459, Development of Social Storyboard Technology for Highly Satisfactory Cultural and Tourist Contents based on Unstructured Value Data Spidering]. We thank the students in the class, which this text is based on, for their feedback. We also thank Bernardo Huberman for influencing many of the ideas presented in the sections on Economic principles. Finally, we would also like to thank Filippo Balestrieri for reviewing early drafts of this letter.

Received: 20 August 2014 Accepted: 25 November 2014

Published online: 05 December 2014

References

1. Armbrust M, Fox A, Griffith R, Joseph AD, Katz R, Konwinski A, Lee G, Patterson D, Rabkin A, Stoica I, Zaharia M (2010) A view of cloud computing. *Commun ACM* 53(4):50–58. doi:10.1145/1721654.1721672
2. Kesselman C, Foster I (1998) The grid: Blueprint for a new computing infrastructure. Morgan Kaufmann Publishers, Inc, San Francisco, CA, USA
3. Barham P, Dragovic B, Fraser K, Hand S, Harris T, Ho A, Neugebauer R, Pratt I, Warfield A (2003) Xen and the art of virtualization. In: Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles. SOSP '03, ACM, New York, NY, USA, pp 164–177. doi:10.1145/945445.945462. <http://doi.acm.org/10.1145/945445.945462>
4. Rolia J, Friedrich R, Patel C (2002) Service centric computing - next generation internet computing. In: Calzarossa M, Tucci S (eds).

- Performance Evaluation of Complex Systems: Techniques and Tools. Lecture Notes in Computer Science, vol. 2459. Springer-Verlag Berlin Heidelberg, Germany, pp 463–479. doi:10.1007/3-540-45798-4_19. http://dx.doi.org/10.1007/3-540-45798-4_19
5. Gottschalk K, Graham S, Kregger H, Snell J (2002) Introduction to web services architecture. *IBM Syst J* 41(2):170–177. doi:10.1147/sj.412.0170
 6. Box D, Ehnebuske D, Kakivaya G, Layman A, Mendelsohn N, Nielsen HF, Thatte S, Winer D (2000) Simple Object Access Protocol (SOAP) 1.1. W3C note, World Wide Web Consortium. See <http://www.w3.org/TR/SOAP/> Accessed 1/12/2014
 7. Lenk A, Klems M, Nimis J, Tai S, Sandholm T (2009) What's inside the cloud? an architectural map of the cloud landscape. In: Proceedings of the 2009 ICSE Workshop on Software Engineering Challenges of Cloud Computing. CLOUD '09. IEEE Computer Society, Washington, DC, USA, pp 23–31. doi:10.1109/CLOUD.2009.5071529. http://dx.doi.org/10.1109/CLOUD.2009.5071529
 8. Wilder B (2012) Cloud Architecture Patterns: Using Microsoft Azure. O'Reilly Media, Beijing Cambridge Farnham Köln Sebastopol Tokyo
 9. Knightly EW, Shroff NB (1999) Admission control for statistical qos: theory and practice. *Netw IEEE* 13(2):20–29. doi:10.1109/65.768485
 10. Delimitrou C, Kozyrakis C (2014) Quasar: resource-efficient and qos-aware cluster management. In: Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems, ACM, New York, NY, USA, pp 127–144
 11. Amdahl GM (1967) Validity of the single processor approach to achieving large scale computing capabilities. In: Proceedings of the April 18-20, 1967, Spring Joint Computer Conference. ACM, New York, NY, USA, pp 483–485
 12. Stonebraker M (1986) The case for shared nothing. *IEEE Database Eng Bull* 9(1):4–9
 13. Karger D, Lehman E, Leighton T, Panigrahy R, Levine M, Lewin D (1997) Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the world wide web. In: Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing. ACM, New York, NY, USA, pp 654–663
 14. Stoica I, Morris R, Karger D, Kaashoek MF, Balakrishnan H (2001) Chord: A scalable peer-to-peer lookup service for internet applications. In: ACM SIGCOMM Computer Communication Review, vol. 31. ACM, New York, NY, USA, pp 149–160
 15. Vogels W (2009) Eventually consistent. *Commun ACM* 52(1):40–44
 16. Lamport L (2001) Paxos made simple. *ACM Sigact News* 32(4):18–25
 17. Mas-Colell A, Whinston MD, Green J (1995) *Microeconomic theory*. Oxford university press, Oxford
 18. Sutherland IE (1968) A futures market in computer time. *Commun ACM* 11(6):449–451. doi:10.1145/363347.363396
 19. Nielsen NR (1970) The allocation of computer resources—is pricing the answer? *Commun ACM* 13(8):467–474
 20. Ellison CM (1975) The utah tenex scheduler. *Proc IEEE* 63(6):940–945
 21. Waldspurger CA, Hogg T, Huberman BA, Kephart JO, Stornetta WS (1992) Spawn: a distributed computational economy. *Softw Eng IEEE Trans* 18(2):103–117
 22. Chun BN, Culler DE (2000) Rexec: A decentralized, secure remote execution environment for clusters. In: Network-Based Parallel Computing. Communication, Architecture, and Applications. Springer-Verlag Berlin Heidelberg, Germany, pp 1–14
 23. Rogers O, Cliff D (2012) A financial brokerage model for cloud computing. *J Cloud Comput* 1(1):1–12
 24. Sandholm T (2008) Statistical Methods for Computational Markets. Doctoral Thesis ISRN SU-KTH/DSV/R-08/6–SE, Royal Institute of Technology, Stockholm
 25. Yeo CS, Buyya R (2006) A taxonomy of market-based resource management systems for utility-driven cluster computing. *Soft: Pract Exp* 36(13):1381–1419
 26. Hardin G (1968) The tragedy of the commons. *Science* 162(3859):1243–1248. doi:10.1126/science.162.3859.1243. <http://www.sciencemag.org/content/162/3859/1243.full.pdf>
 27. Chun B, Culler D, Roscoe T, Bavier A, Peterson L, Wawrzoniak M, Bowman M (2003) Planetlab: an overlay testbed for broad-coverage services. *ACM SIGCOMM Comput Commun Rev* 33(3):3–12
 28. Buchanan JM, Stubblebine WC (1962) Externality. *Economica* 29(116):371–384
 29. Arrow KJ (1977) 1921– Studies in resource allocation processes. edited by Kenneth J. Arrow, Leonid Hurwicz. Cambridge; New York: Cambridge University Press, xiv, 482 p.ill. 24 cm. HB135 .A79 ISBN: 0521215226
 30. Papadimitriou C (2001) Algorithms, games, and the internet. In: Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing. ACM, New York, NY, USA, pp 749–753
 31. Varian HR (1974) Equity, envy, and efficiency. *J Econ Theory* 9(1):63–91
 32. Keynes JM (1965) *A Treatise on Money, Vol. 1: The Pure Theory of Money*. Macmillan, London, UK
 33. Pratt JW (1964) Risk aversion in the small and in the large. *Econometrica* 32(1/2). doi:10.2307/1913738
 34. Arrow KJ (1965) Aspects of the Theory of Risk-bearing. Yrjö, Helsinki, Finland
 35. Amazon EC2 Spot Instances. <http://aws.amazon.com/ec2/purchasing-options/spot-instances/>. (Accessed 17/11/2014)
 36. Jaeger H (2014) A new era in sourcing. Technical report, Deutsche Börse Cloud Exchange http://www.dbcloudexchange.com/wp-content/uploads/2014/08/A-new-era-in-sourcing_290813.pdf (Accessed 1/12/2014)

doi:10.1186/s13677-014-0021-5

Cite this article as: Sandholm and Lee: Notes on Cloud computing principles. *Journal of Cloud Computing: Advances, Systems and Applications* 2014 **3**:21.

Submit your manuscript to a SpringerOpen® journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com