Volume 4, Issue 5, May 2014 ISSN: 2277 128X
International Journal of Advanced Research in Computer Science and Software Engineering
Research Paper
Available online at: www.ijarcsse.com

ANCHOR: A Stable Matching Framework for Managing Cloud Resources

Anita Kadam*
Computer Science and Engineering,
VTU Regional Centre Gulbarga
India

Bhavana.S
Computer Science and Engineering,
VTU Regional Centre Gulbarga
India

Shailaja Shastri
Computer Science and Engineering,
PDA Engineering Collage, Gulbarga
India

Abstract— We present Anchor, a general resource management architecture that uses the stable matching framework to decouple policies from mechanisms when mapping virtual machines to physical servers. In Anchor, clients and operators are able to express a variety of distinct resource management policies as they deem fit, and these policies are captured as preferences in the stable matching framework. The highlight of Anchor is a new many-to-one stable matching theory that efficiently matches VMs with heterogeneous resource needs to servers, using both offline and online algorithms. Our theoretical analyses show the convergence and optimality of the algorithm. Our experiments with a prototype implementation on a 20-node server cluster, as well as large-scale simulations based on real-world workload traces, demonstrate that the architecture is able to realize a diverse set of policy objectives with good performance and practicality.

Keywords — Stable Matching, Cloud Computing, VM placement, Resource management

I. INTRODUCTION

Modern data centers heavily rely on virtualization to flexibly multiplex different applications onto physical servers, in order to efficiently utilize their resources. With virtualization, applications are packaged and run in the form of virtual machines (VM) that share the server infrastructure. Due to the multi-tenant nature of such virtualized data centers, resource management becomes a major challenge for cloud operators to achieve economies of scale. VMs impose extremely diverse resource requirements that need to be accommodated, as they run completely different applications owned by different clients. As such, they are entitled to distinct resource management policies depending on specific needs of their owners.

On the other hand, the infrastructure is managed as a whole by the cloud operator, who relies on a common resource management substrate, and has a wide variety of its own objectives to achieve, such as workload consolidation, cost minimization, and load balancing. Therefore, the resource management substrate must accommodate and orchestrate the needs and interests of both cloud operators and clients. However, current solutions provided by virtualization vendors are far from satisfactory: they are proprietary, hard-coded, and not easily customizable. There exists no interface for cloud clients to express resource management needs in their applications.

In this work, we present *Anchor*, a new architecture that decouples *policies* from *mechanisms* when it comes to managing resources in the cloud. Stakeholders in the cloud, including both cloud operators and their clients, are able to express and configure their high-level resource management policies, based on performance, cost, and network load, as they deem fit. These policies serve as input to guide mechanisms that manage cloud resources, so that conflicts of interest among stakeholders can be resolved. The output is a mapping between VMs and physical servers: *Anchor* allocates VMs to servers before they are run, and, if necessary, migrates running VMs away from their original hosts using live VM migration. *Anchor* is designed to be scalable to support hundreds of thousands of VMs and servers, to be expressive so that clients and operators can specify a wide variety of their policies and preferences with ease.

Our design of the *Anchor* architecture is based on a *stable matching* framework from economics theory, which elegantly and efficiently addresses common and conflicting interests of agents in a resource market. In our framework, the concept of *preferences* is used to express various high-level policies that stakeholders specify, and, rather than *optimality*, *stability* is used as the central solution concept of the matching mechanism. Merits of the stable matching framework lie in its competitiveness of outcomes, generality of preferences, efficiency and simplicity of its algorithmic implementations, and most importantly, its overall practicality.

The *Anchor* architecture consists of three main components: a resource monitor, a policy manager, as well as a matching engine, as shown below. The cloud operator configures resource management policies as input to the policy engine. When requests for allocating a new VM or migrating an existing VM arrive, *Anchor* assumes that detailed information about VM configurations and its own policy goals is available at the time. The policy manager then queries information required by both operator and client policies from the resource monitor, which maintains resource usage information through the management API, and translates these high-level objectives into preferences for the VM. It also configures server preferences in a similar manner. These preferences are fed into the matching engine that runs our stable matching mechanism. The highlight of our matching engine is a new multi-stage deferred-acceptance algorithm that matches multiple VMs of heterogeneous resource demands to a single physical server, corresponding to a many-to-one stable matching problem with size heterogeneity. The result determines the final matching between VMs and servers, and is executed through the management API.

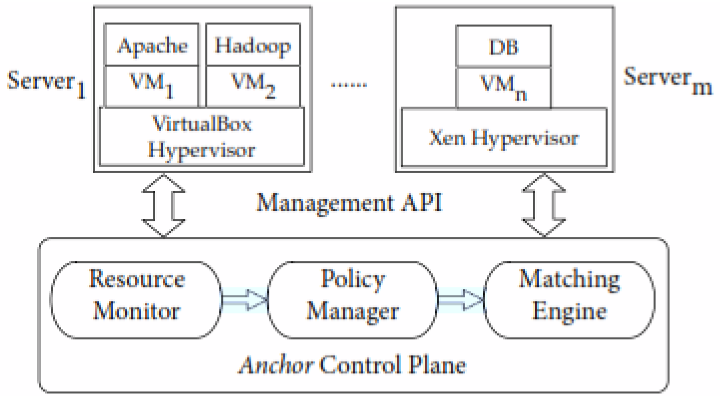

Fig. 1.ANCHOR architecture

II. BACK GROUND AND MODEL

A. A Primer on Stable Matching

We start by introducing the classical theory of stable matching in the basic one-to-one marriage model [19]. There are two disjoint sets of agents, $\mathbf{M} = \{m_1, m_2, \ldots, m_n\}$ and $\mathbf{W} = \{w_1, w_2, \ldots, w_p\}$, men and women. Each agent has a transitive preference over individuals on the other side, and the possibility of being unmatched [18]. Preferences can be represented as rank order lists of the form $p(m_1) = w_4, w_2, \ldots, w_i$, meaning that man m_1's first choice of partner is w_4, second choice is w_2 and so on, until at some point he prefers to be unmatched (i.e. matched to the empty set). We use $\succ_i$ to denote the ordering relationship of agent i (on either side of the market). If i prefer to remain unmatched instead of being matched to agent j, i.e. $\emptyset \succ_i j$, j is said to be *unacceptable* to i, and preferences can be represented just by the list of acceptable partners.

***Definition* 1:** An outcome is a *matching* μ: $\mathbf{M} \times \mathbf{W} \times \emptyset \rightarrow \mathbf{M} \times \mathbf{W} \times \emptyset$ such that $w = \mu(m)$ if and only if $\mu(w) = m$, and $\mu(m) \in \mathbf{W} \cup \emptyset$, $\mu(w) \in \mathbf{M} \cup \emptyset$, $\forall m, w$.

It is clear that we need further criteria to distill a "good" set of matchings from all the possible outcomes. The first obvious criterion is *individual rationality*.

***Definition* 2:** A matching is *individual rational* to all agents, if and only if there does not exist an agent i who prefers being unmatched to being matched with $\mu(i)$, i.e., $\emptyset \succ_i \mu(i)$.

This implies that for a matched agent, its assigned partner should rank higher than the empty set in its preference. Between a pair of matched agents, they are not unacceptable to each other.

The second natural criterion is that a *blocking set* should not occur in a good matching:

***Definition* 3:** A matching μ is *blocked* by a pair of agents (m, w) if they each prefer each other to the partner they receive at μ. That is, $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$. Such a pair is called a *blocking pair* in general.

When a blocking pair exists, the agents involved have a natural incentive to break up and form a new marriage. Therefore such an "unstable" matching is undesirable.

***Definition* 4:** A matching μ is *stable* if and only if it is individual rational, and not blocked by any pair of agents.
***Theorem* 1:** A stable matching exists for every marriage market.

This can be readily proved by the classic *deferred acceptance algorithm (DA)*, or the *Gale-Shapley algorithm* [19]. *It* works by having agents on one side of the market, say men, propose to the other side, in order of their preferences. As long as there exists a man who is free and has not yet proposed to every woman in his preference, he proposes to the most preferred woman who has not yet rejected him. The woman, if free, "holds" the proposal instead of directly accepting it. In case she already has a proposal at hand, she rejects the less preferred. This continues until no proposal can be made, at which point the algorithm stops and matches each woman to the man (if any) whose proposal she is holding. The woman-proposing version works in the same way by swapping the roles of man and woman. It can be readily seen that the order in which men propose is immaterial to the outcome.

B. Models and Assumptions

In a cloud, each VM is allocated a slice of resources from its hosting server. In this paper, we assume that the size of a slice is a multiple of an *atomic* VM. For instance, if the atomic VM has one CPU core equivalent to a 2007 Intel Xeon 1 GHz core, one memory unit equivalent to 512 MB PC-10600 DDR3 memory, and one storage unit equivalent to 10 GB 5400 RPM HDD, a VM of size 2 means it effectively has a 2 GHz 2007 Xeon CPU core, 1 GB PC-10600 DDR3 memory and 20 GB 5400 RPM hard disk. Note that the actual amount of resources is relative to the heterogeneous server hardware. Two VMs have the same size as long as performance is equivalent for all resources.

This may seem an oversimplification and raise concerns about its validity in reality. We comment that, in practice, such atomic sizing is common among large-scale public clouds to reduce the overhead of managing hundreds of thousands of VMs. It is also valid in production computer clusters [20], and widely adopted in related work [16], [21] to reduce the dimensionality of the problem. Appendix A provides more discussion on the validity of this assumption, especially with different job requirements.

We design *Anchor* for a setting where the workloads and resources demands of VMs are relatively stable. Resource management in the cloud can be naturally cast as a *stable matching* problem, where the overall pattern of common and conflicting interests between stakeholders can be re-solved by confining our attention to outcomes that are stable. Broadly, it can be modelled as a *many-to-one* problem [19] where one server can enroll multiple VMs but one VM can only be assigned to one server. Preferences are used as an abstraction of policies no matter how they are defined. In traditional many-to-one problems such as college ad-missions [19], each college has a quota of the number of students it can take. This cannot be directly applied to our scenario, as each VM has a different "size" corresponding to its demand for resources. We cannot simply define the quota of a server as the number of VMs it can take.

We formulate VM placement as a *job-machine stable matching* problem with size heterogeneous jobs. Each job has a size, and each machine has a *capacity*. A machine can host multiple jobs as long as the total job size does not exceed its capacity. Each job has a preference over all the acceptable machines that have sufficient capacities. Similarly, each machine has a preference over all the acceptable jobs whose size is smaller than its capacity. This is a more general many-to-one matching model in that the college admissions problem is a special case with unisize jobs (students).

III. THEORETICAL CHALLENGES OF JOB MACHINE STABLE MATCHING

We present theoretical challenges introduced by size heterogeneous jobs in this section.

Following convention, we can naturally define a *blocking pair* in job-machine stable matching based on the following intuition. In a matching μ, whenever a job j prefers a machine m to its assigned machine μ(j) (can be ∅ meaning it is unassigned), and m has vacant capacity to admit j, or when m does not have enough capacity, but by rejecting some or all of the accepted jobs that rank lower than j it
is able to admit j, then j and m have a strong incentive to deviate from μ and form a new matching. Therefore,

***Definition* 5:** A job-machine pair (j, m) is a *blocking pair* if any of the two conditions holds:

$$\text{(a):} \quad c(m) \geq s(j),\ j \succ_m \emptyset, \text{ and } m \succ_j \mu(j), \qquad (1)$$

$$\text{(b):} \quad c(m) < s(j),\ c(m) + \sum_{j'} s(j') \geq s(j),$$

$$\text{where } j' \prec_m j,\ j' \in \mu(m), \text{ and } m \succ_j \mu(j). \qquad (2)$$

c(m) denotes the capacity of machine m, and s(j) denotes the size of job j. Depending on whether a blocking pair satisfies condition (1) or (2), we say it is a *type-1* or *type-2* blocking pair. For example, in a setting shown in Fig. 2, the matching A − (a), B − ∅ contains two type-1 blocking pairs (b, B) and (c, B), and one type-2 blocking pair (c, A).

***Definition* 6:** A job-machine matching is *strongly stable* if it does not contain any blocking pair.

A. Non-existence of Strongly Stable Matchings

It is clear that both types of blocking pairs are undesirable, and we ought to find a strongly stable matching. However, such a matching may not exist in some cases. Fig. 2 shows one such example with three jobs and two machines. It can be verified that every possible matching contains either type-1 or type-2 blocking pairs.

p(j)	j	s(j)	c(m)	M	p(m)
A	a	2	2	A	c, a, b
A, B	b	1	1	B	b, c
B, A	c	1			

Fig.2. A simple example where there is no strongly stable matching. Recall that p() denotes the preference of an agent.

***Proposition* 1:** Strongly stable matching does not always exist.

Note that the definitions of type-1 and type-2 blocking pair coincide in classical problems with unisize jobs. The reason why they do not remain so in our model is the size complementariness among jobs. In our problem, the concept of capacity denotes the amount of resources a ma-chine can provide, which may be used to support different numbers of jobs. A machine's preferable job, which is more likely to be admitted in order to avoid type-2 blocking pairs, may consume less resources, and creates a higher likelihood for type-1 blocking pairs to happen on the same machine.

The non-existence result demonstrates the theoretical difficulty of the problem. We find that it is hard to even determine the necessary or sufficient conditions for the existence of strongly stable matchings in a given problem instance, albeit its definition seems natural. Therefore, for mathematical tractability, in the subsequent parts of the paper, we work with the following relaxed definition:

***Definition* 7:** A matching is *weakly stable* if it does not contain any type-2 blocking pair.

For example in Fig. 2, A−(c), B − (b) is a weakly but not strongly stable matching, because it has a type-1 blocking pair (b, A). Thus, weakly stable matchings are a superset that subsumes strongly stable matchings. A matching is thus called *unstable* if it is not weakly stable.

B. **Failure of the DA Algorithm**

With the new stability concept, the first theoretical challenge is how to find a weakly stable matching, and does it always exist? If we can devise an algorithm that produces a weakly stable solution for any given instance, then its existence is clear. One may think that the deferred acceptance (DA) algorithm can be applied for this purpose. Jobs propose to machines following the order in their preferences. We randomly pick any free job that has not proposed to every machine on its preference to propose to its favourite machine that has not yet rejected it. That machine accepts the most favourable offers made so far up to the capacity, and rejects the rest. Unfortunately, we show that this may fail to be effective. Appendix B.1 shows such an example.

Two problems arise when applying the classical DA algorithm here. First, the execution sequence is no longer immaterial to the outcome. Second, it may even produce an unstable matching. This creates considerable difficulties since we cannot determine which proposing sequence yields a weakly stable matching for an arbitrary problem.

Examined more closely, the DA algorithm fails precisely due to the size heterogeneity of jobs. Recall that a machine will reject offers only when its capacity is used up. In the traditional setting with jobs of the same size, this ensures that whenever an offer is rejected, it must be the case that the machine's capacity is used up, and thus any offer made from a less preferred job will never be accepted, i.e. the outcome is stable. However, rejection due to capacity is problematic in our case, since a machine's remaining capacity may be increased, and its previously rejected job may become favorable again.

C. **Optimal Weakly Stable Matching**

There may be many weakly stable matchings for a problem instance. The next natural question to ask is then, which one should we choose to operate the system with? Based on the philosophy that a cloud exists for companies to ease the pain of IT investment and management, rather than the other way around, it is desirable if we can find a *job-optimal* weakly stable matching, in the sense that every job is assigned its best machine possible in all stable matchings.

The original DA algorithm is again not applicable in this regard, because it may produce type-1 blocking pairs even when the problem admits strongly stable matchings. Thus, our second challenge is to devise an algorithm that yields the job-optimal weakly stable matching. This quest is also theoretically important in its own right.

However, as we will show in Sec. 4.2, the complexity of solving this challenge is high, which may prevent its use in large-scale problems. Thus in many cases, a weakly stable matching is suitable for practical purposes.

IV. **A NEW THEORY OF JOB-MACHINE STABLE MATCHING**

In this section we present our new theory of job-machine stable matching that addresses the above challenges.

A. **A Revised DA Algorithm**

We first propose a revised DA algorithm, shown in Table 1, that is guaranteed to find a weakly stable matching for a given problem. The key idea is to ensure that, whenever a job is rejected, any less preferable jobs will not be accepted by a machine, even if it has enough capacity to do so.

TABLE 1

Revised DA

```
Input: c(m), p(m), ∀m ∈ M, s(j), p(j), ∀j ∈ J
Initialize all j ∈ J and m ∈ M to free
while ∃j who is free, and p(j) ≠ ∅ do
  m = j's highest ranked machine in p(j)
  if c(m) ≥ s(j) then
    j and m become matched, c(m) = c(m) − s(j)
  else
    Find all j′ matched to m so far such that j′ ≺m j
    repeat
    m sequentially rejects each j′ by setting it to free, in the order of p(m)
    c(m) = c(m) + s(j′), best rejected = j′
    until c(m) ≥ s(j) or all j′ are rejected
    if c(m) ≥ s(j) then
    j and m become matched, c(m) = c(m) − s(j)
    else
    j becomes free, best rejected = j
    for j″ ∈ p(m), j″ ≺m best rejected do
    Remove m from p(j″), j″ from p(m)
Return: the final matching, and remaining capacity c(m), ∀m ∈ M
```

The algorithm starts with a set of jobs **J** and a set of machines **M**. Each job and machine are initialized to be free. Then the algorithm enters a propose-reject procedure. Whenever there are free jobs that have machines to propose to, we randomly pick one, say j, to propose to its current favourite machine m in p(j), which contains all the machines that have not yet rejected it. If m has sufficient capacity, it holds the offer. Otherwise, it sequentially rejects offers from less preferable jobs j' until it can take the offer, in the order of its preference. If it still cannot do so even after rejecting all the j's, j is then rejected. Whenever a machine rejects a job, it updates the best rejected variable, and at the end *all* jobs ranked lower than best rejected are removed from its preference. The machine is also removed from preferences of all these jobs, as it will never accept their offers.

A pseudo-code implementation is shown in Table 1. We can see that the order in which jobs propose is immaterial, similar to the original DA algorithm. Moreover, we can prove that the algorithm guarantees that type-2 blocking pairs do not exist in the result.

***Theorem* 2:** The order in which jobs propose is of no consequence to the outcome in Revised DA.

***Theorem* 3:** Revised DA, in any execution order, produces a unique weakly stable matching.

Proof: The proof of uniqueness is essentially the same as that for the classical DA algorithm in the seminal paper [19]. We prove the weak stability of the outcome by contra-diction. Suppose that Revised DA produces a matching μ with a type-2 blocking pair (j, m), i.e. there is at least one job j' worse than j to m in μ(m). Since $m \succ_j \mu(j)$, j must have proposed to m and been rejected. When j was rejected, j' was either rejected before j, or was made unable to propose to m because m is removed from the preferences of all the jobs ranked lower than j. Thus $j' = \emptyset$, which contradicts with the assumption.

Theorem 3 also proves the existence of weakly stable matchings, as Revised DA terminates within $O(|\mathbf{J}|^2)$ in the worst case.

***Theorem* 4:** A weakly stable matchings exists for every job-machine matching problem.

The significance of Revised DA is multi-fold. It solves our first technical challenge in Sec. 3.2, and is appealing for practical use. The complexity is low compared to optimization algorithms. Further, it serves as a basic building block, upon which we develop an iterative algorithm to find the job-optimal weakly stable matching as we shall demonstrate soon. Lastly, it bears the desirable property of being insensitive to the order of proposing, which largely reduces the complexity of algorithm design.

Revised DA may still produce type-1 blocking pairs, and the result may not be job-optimal as defined in Sec. 3.3. In order to find the job-optimal matching, an intuitive idea is to run Revised DA multiple times, each time with type-1 blocking jobs proposing to machines that form blocking pairs with them. The intuition is that, type-1 blocking jobs can be possibly improved at no expense of others. However, simply doing so may make the matching unstable, because when a machine has both type-1 blocking jobs leaving from and proposing to it, it may have more capacity available to take jobs better than those it accepts according to its capacity before the jobs leaving. Readers may refer to Appendix B.2 for an example.

B. A Multi-stage DA Algorithm

We now design a *multi-stage* DA algorithm to iteratively find a better weakly stable matching with respect to jobs. The algorithm proceeds in stages. Whenever there is a type-1 blocking pair (j, m) in the result of previous stage μ_{t-1}, the algorithm enters the next stage where the blocking machine m will accept new offers. The blocking job j is removed from its previous machine $\mu_{t-1}(j)$, so that it can make new offers to machines that have rejected it before. (j)'s capacity is also updated accordingly. Moreover, to account for the effect of job removal, all jobs that can potentially form type-1 blocking pairs with $\mu_{t-1}(j)$ if j leaves (there may be other machines that j form type-1 blocking pairs with) are also removed from their machines and allowed to propose in the next stage (corresponding to the while loop in step 7). This ensures that the algorithm does not produce new type-2 blocking pairs during the course, as we shall prove soon. At each stage, we run Revised DA with the selected set of proposing jobs $\mathbf{J}'$, and the entire set of machines with updated capacity $c^{pre}_t(m)$. The entire procedure is shown in Table 2.

TABLE 2

Multi-stage DA

```
 1: Input: c(m), p(m), ∀m ∈ M, s(j), p(j), ∀j ∈ J .
 2: μ0 = ∅, t = 0, stop = false, J′ = ∅
 3: while stop == false do
 4:    t = t + 1, μ′ = μt−1
 5:    for m ∈ M do
 6:       c^pre_t(m) = c_t−1(m)
 7:    while Ω 6= ∅, where Ω is the set of jobs that form type-1 blocking pairs from μ′ with c^pre_t(m) do
 8:       for j ∈ Ω do
 9:          Add j′ to J′.
10:          if μ′(j) != ∅ then
11:             c^pre_t(μ′(j)) = c^pre_t(μ′(j)) + s(j).
12:             j′ is free and removed from the matching μ′.
13:    if J′ == ∅ then
14:       break
15:    (μt, ct(m)) = Revised DA(c^pre_t(m), p(m),
16:                   s(j), p(j), μ′, J′)
17:    if μt == μt−1 then
18:       stop = true
19: Return μt
```

We now prove important properties of Multi-stage DA, namely its correctness, convergence, and job-optimality.

a) Correctness

First we establish the correctness of Multi-stage DA.

***Theorem* 5:** There is no type-2 blocking pair in the matchings produced at any stage in Multi-stage DA.

Proof: This can be proved by induction. As the base case, we already proved that there is no type-2 blocking pair after the first stage in Theorem 3.

Given there is no type-2 blocking pair after stage t, we need to show that after stage t + 1, there is still no type-2 blocking pair. Suppose after t +1, there is a type-2 blocking pair (j, m), *i.e.*, $c_{t+1}(m) < s(j)$, $c_{t+1}(m) + \sum_{j'} s(j') \geq s(j)$, where $j' \prec_m j$, $j' \in \mu_t(m)$, $m \succ_j \mu_{t+1}(j)$. If $c^{pre}_{t+1}(m) \geq s(j)$, then j must have proposed to m and been rejected according to the algorithm. Thus it is impossible for m to accept any job j' less preferable than j in t + 1.

If $c^{pre}_{t+1}(m) < s(j)$, then j did not propose to m in t + 1. Since there is no type-2 blocking pairs after t, j' must be accepted in t + 1. Now since $c^{pre}_{t+1}(m) < s(j)$, the sum of the remaining capacity and total size of newly accepted jobs after t + 1 must be less than $c^{pre}_{t+1}(m)$, *i.e.* $c_{t+1}(m) + \sum_{j''} s(j'') \leq c^{pre}_{t+1}(m) < s(j)$, where j'' denotes the newly accepted jobs in t +1. This contradicts with the assumption that $c_{t+1}(m) + \sum_{j'} s(j') \geq s(j)$ since $\{j'\} \subseteq \{j''\}$.

If $c^{pre}_{t+1}(m) = 0$, then m only has jobs leaving from it. Since there is no type-2 blocking pair after t, clearly there cannot be any type-2 blocking pair in t + 1.

Therefore, type-2 blocking pairs do not exist at any stage of the algorithm. The uniqueness of the matching result at each stage is readily implied from Theorem 3.

b) Convergence

Next we prove the convergence of Multi-stage DA. The key observation is that it produces a weakly stable matching at least as good as that in the previous stage from the job's perspective.

***Lemma* 1:** At any consecutive stages t and t + 1 of Multi-stage DA, $\mu_{t+1}(j)\ _{j}\ \mu_t(j)$, $\forall j \in \mathbf{J}$.

Proof: Refer to Appendix C.1.

Therefore, the algorithm always tries to improve the weakly stable matching it found in the previous stage, whenever there is such a possibility suggested by the existence of type-1 blocking pairs. However, Lemma 1 also implies that a job's machine at t + 1 may remain the same as in the previous stage. In fact, it is possible that the entire matching is the same as the one in previous stage, i.e. $\mu_{t+1} = \mu_t$. This can be easily verified using the example of Fig. 2. After the first stage, the weakly stable matching is A − (c), B − (b). First b wishes to propose to A in the second stage. Then we assign b to ∅ and B has capacity of 1 again. c then wishes to propose to B too. After we remove c from A and update A's capacity, a now wishes to propose to A. Thus at the next stage, the same set of jobs a, b, c will propose to the same set of machines with same capacity, and the result will be the same matching as in the first stage. In this case, Multi-stage DA will terminate with the final matching that it cannot improve upon as its output (step 17-18 of Table 2). We thus have:

***Theorem* 6:** Multi-stage DA terminates in finite time. Note that in each stage, Multi-stage DA may result in new type-1 blocking pairs, and the number of type-1 blocking pairs is not monotonically decreasing. Thus its worst case run time complexity is difficult to analytically derive. In Sec. 6.3 we evaluate its average case complexity through large-scale simulations.

c) Job-Optimality

We now prove the most important result regarding Multi-stage DA:

***Theorem* 7:** Multi-stage DA always produces the *job-optimal* weakly stable matching when it terminates, in the sense that every job is at least as good in the weakly stable matching produced by the algorithm as it would be in any other weakly stable matching.

Proof: We provide a proof sketch here. A detailed proof can be found in Appendix C.2.

The algorithm terminates at stage t when either there is no type-1 blocking pair, or there is type-1 blocking pair(s) but $\mu_t = \mu_{t-1}$. For the former case, we show that our algorithm only permanently rejects jobs from machines that are impossible to accept them in all weakly stable matchings, when the jobs cannot participate any further. The outcome is therefore optimal. For the latter,

TABLE 3

Anchor's policy interface.

Functionality	*Anchor* API Call
create a policy group	g = create()
add/delete server	add/delete(g_o,s)
add/delete VMs	add/delete(g_c,v)
set ranking factors	conf(g,factor1,...)
set placement constraints	limit(g_c,servers)
colocation/anti-colocation	colocate(tenants,i,*g_c*)

case, we can also show that it is impossible for jobs that participated in t to obtain a better machine. Finally, we present another fact regarding the outcome of our algorithm.

***Theorem* 8:** Multi-stage DA produces a unique job-optimal *strongly* stable matching when it terminates with no job proposing.

The proof can be found in Appendix C.3.

d) An Online Algorithm

We have thus far assumed a static setting with a fixed set of jobs and machines. In practice, requests for job (VM) placement arrive dynamically, and we need to make decisions on the fly. It may not be feasible to re-run the matching algorithm from scratch every time when there is a new job. We further develop an online algorithm based on Revised DA that handles the dynamic case efficiently. Interested readers can find the detailed algorithm design and evaluation results in Appendix D and F.3, respectively.

V. SHOWCASES OF RESOURCE MANAGEMENT

A. *POLICIES WITH THE POLICY ENGINE*

We have presented the underlying mechanism of *Anchor* that produces a weakly stable matching between VMs of various sizes, as *jobs*, and physical servers, as *machines*. We now introduce *Anchor*'s policy engine which constructs preference lists according to various resource management policies. The cloud operator and clients interact with the policy engine through an API as shown in Table 3.

In order to reduce management overhead, we use *policy groups* that can be created with the create() call. Each policy group contains a set of servers or VMs that are entitled to a common policy. In fact some of the recent industry products have adopted similar ideas to help companies manage their servers in the cloud [22]. The policy is configured by the conf() call that informs the policy engine what factors to be considered for ranking the potential partners in a descending order of importance. The exact definition of ranking factors varies depending on the specific policy as we demonstrate in the following. With policy groups, only one common preference list is needed for all members of the group. Membership is maintained by add() and delete() calls. colocate() and limit() are used to set colocation/anti-colocation and placement constraints as we discuss in Appendix E.

It is also possible for the operator to configure policies on behalf of its clients if they do not explicitly specify any. This is done by enrolling them to the default policy group.

a) Cloud Operators

We begin our discussion from the operator's perspective.

Server consolidation/packing: Operators usually wish to consolidate the workload by packing VMs onto a small number of highly occupied servers, so that idle servers can be powered down to save operational costs. To realize this policy, servers can be configured to prefer a VM with a larger size. This can be done using conf(g_o,1/vm_size), where g_o is the operator's policy group. For VMs in the default policy group, their preference is ranked in the descending order of server load. One may use the total size of active VMs as the metric of load (conf(g_c, 1/server_load)), where g_c is the client's policy group. Alternatively, the number of active VMs can also serve as a heuristic metric (conf(g_c,1/num_vm)).

Notice that consolidation is closely related to packing, and the above configuration resembles the *first fit decreasing* heuristic widely used to solve packing problems by iteratively assigning the largest item to the first bin that fits.

Load balancing: Another popular resource management policy is load balancing, which distributes VMs across all servers to mitigate performance degradation due to application dynamics over time. This can be seen as the opposite of consolidation. In this case, we can configure the server preference with conf(g_o, vm_size), implying that servers prefer smaller VMs in size. VMs in the default policy group are configured with conf(g_c, server_load), such that they prefer less utilized servers. This can be seen as a *worst fit increasing* heuristic.

b) *Cloud Clients*

From the perspective of cloud clients, other than choosing to join the default policy group and follow the operator's configuration, they can also express their unique policies.

Resource hunting. Depending on the resource demand of applications, VMs can be CPU, memory, or bandwidth-bound, or even resource-intensive in terms of multiple resources. Though resources are sliced into fixed slivers, most modern hypervisors support dynamic resizing of VMs. For example, the hypervisor may allow a temporarily burst of CPU usage for a VM provided that doing so does not affect collected VMs. For memory, with a technique known as *memory ballooning*, the hypervisor is able to dynamically reduce the memory footprints of idle VMs, so that memory allocation of heavily loaded VMs can be increased.

Thus, clients may configure their policies according to the resource usage pattern of their VMs, which is un-known to the operator. CPU-bound VMs can be added to a *CPU-bound* policy group, which is configured with a call to conf(g_c,1/server_freecpu). Their preferences are then ranked in the descending order of server's time average available CPU cycles. Similarly, memory-bound and bandwidth-bound policy groups may be configured with the call conf(g c, 1/server_freemem) and conf(g_c, 1/server_freebw), respectively.

Anchor supports additional policies besides what we list here, including colocation/anti-colocation, tiered services, etc. Due to space limit, readers are directed to Appendix E for more details.

VI. RELATED WORK

This work is related to research in the following fields.

Stable Matching. A large body of research in economics has examined variants of stable matching [19] (see [18], [27] and references therein). Algorithmic aspects of stable matching have also been studied in computer science [28], [29]. However, the use of stable matching in networking is fairly limited. [16] uses the DA algorithm to solve the coupled placement of VMs in datacenters. Our recent work [30], [31] advocate stable matching as a general framework to solve networking problems. To our knowledge, all these works assume a traditional unisize job model, while we study a more general size-heterogeneous model.

VM Placement. VM placement on a shared infrastructure has been extensively studied. Current industry solutions from virtualization vendors such as VMware vSphere [3] and Eucalyptus [4], and open-source efforts such as Nimbus [5] and CloudStack [6], only provide a limited set of pre-defined placement policies. Existing papers develop specifically crafted algorithms and systems for specific scenarios, such as consolidation based on CPU usage [7]–[9], energy consumption [10]–[12], bandwidth multiplexing [13]–[15], and storage dependence [16]. They are thus complementary to *Anchor*, as the insights and strategies can be incorporated as policies to serve different purposes without the need to design new algorithms from the ground up.

OpenNebula [32], a resource management system for virtualized infrastructures, is the only related work to our knowledge that also decouples management policies with mechanisms. It uses a simple first fit algorithm based a configurable ranking scheme to place VMs, while we use the stable matching framework that addresses the conflict of interest between the operator and clients.

There is a small literature on online VM placement. [33]–[35] develop systems to predict the dynamic resource demand of VMs and guide the placement process. [15] Considers minimizing the long-term routing cost between VMs. These works consider various aspects to refine the placement process and are orthogonal to our work that ad-dresses the fundamental problem of VM size heterogeneity.

VII. CONCLUSION

We presented *Anchor* as a unifying fabric for resource management in the cloud, where *policies* are decoupled from the management *mechanisms* by the stable matching framework. We developed a new theory of job-machine stable matching with size heterogeneous jobs as the under-lying mechanism to resolve conflict of interests between the operator and clients. We then showcased the versatility of the *preference* abstraction for a wide spectrum of resource management policies for VM placement with a simple API. Finally, the efficiency and scalability of *Anchor* are demonstrated using a prototype implementation and large-scale trace-driven simulations.

Many other problems can be cast into our model. For instance, job scheduling in distributed computing platforms such as Map Reduce, where jobs have different sizes and share a common infrastructure. Our theoretical results are thus potentially applicable to scenarios beyond those described in this paper. As future work, we plan to extend *Anchor* for the case where resource demands vary, and VMs may require to be replaced, where specific considerations for VM live migration [21] are needed.

Result: Anchor enables efficient resource utilization of the infrastructure and improves performance of its VMs, by allowing individual clients to express policies specific to its resource needs.

Revised DA is effective and practical for large-scale problems with thousands of VMs, and offers very close-to-optimal performance for VMs.

REFERENCES

[1] Zenoss Inc., “Virtualization and cloud computing survey,” 2010.

[2] R. Campbell, I. Gupta, M. Heath, S. Y. Ko, M. Kozuch, M. Kunze, T. Kwan, K. Lai, H. Y. Lee, M. Lyons, D. Milojiciˇc,´ D. O’Hallaron, and Y. C. Soh, “Open Cirrus™ cloud computing testbed: Federated data centers for open source systems and services research,” in *Proc. USENIX HotCloud*, 2009.

[3] VMware vSphere™, http://www.vmware.com/products/drs/ overview.html.

[4] Eucalyptus, http://www.eucalyptus.com.

[5] Nimbus, http://www.nimbusproject.org.

[6] CloudStack, http://www.cloudstack.org.

[7] B. Urgaonkar, P. Shenoy, and T. Roscoe, “Resource overbooking and application profiling in shared hosting platforms,” *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 239–254, December 2002.

[8] X. Meng, C. Isci, J. Kephart, L. Zhang, E. Bouillet, and D. Pen-darakis, “Efficient resource provisioning in compute clouds via vm multiplexing,” in *Proc. Intl. Conf. Autonomic Computing (ICAC)*, 2010.

[9] M. Chen, H. Zhang, Y.-Y. Su, X. Wang, G. Jiang, and K. Yoshi- hira, “Effective VM sizing in virtualized data centers,” in *Proc. IFIP/IEEE Intl. Symp. Inte. Netw. Manag. (IM)*, 2011.

[10] R. Nathuji and K. Schwan, “Virtualpower: Coordinated power management in virtualized enterprise systems,” in *Proc. SOSP*, 2007.

[11] A. Verma, P. Ahuja, and A. Neogi, “pMapper: Power and migra-tion cost aware application placement in virtualized systems,” in *Proc. ACM Middleware*, 2008.

[12] M. Cardosa, M. R. Korupolu, and A. Singh, “Shares and utilities based power consolidation in virtualized server environments,” in *Proc. IFIP/IEEE Intl. Symp. Inte. Netw. Manag. (IM)*, 2009.

[13] X. Meng, V. Pappas, and L. Zhang, “Improving the scalability of data center networks with traffic-aware virtual machine place-ment,” in *Proc. IEEE INFOCOM*, 2010.

[14] D. Breitgand and A. Epstein, “Improving consolidation of virtual machines with risk-aware bandwidth oversubscription in com-pute clouds,” in *Proc. IEEE INFOCOM*, 2012.

[15] J. W. Jiang, T. Lan, S. Ha, M. Chen, and M. Chiang, “Joint VM placement and routing for data center traffic engineering,” in *Proc. IEEE INFOCOM*, 2012.

[16] M. Korupolu, A. Singh, and B. Bamba, “Coupled placement in modern data centers,” in *Proc. IPTPS*, 2009.

[17] Y.Rekhter, T. Li, and S. Hares, “A border gateway protocol 4 (BGP-4),” RFC 4271, http://datatracker.ietf.org/doc/rfc4271/, January 2006.

[18] A. E. Roth, “Deferred acceptance algorithms: History, theory, practice, and open questions,” *Int. J. Game Theory*,

vol. 36, pp. 537–569, 2008.

[19] D. Gale and L. S. Shapley, "College admissions and the stability of marriage," *Amer. Math. Mon.*, vol. 69, no. 1, pp. 9–15, 1962.

[20] The RICC Log, http://www.cs.huji.ac.il/labs/parallel/ workload/l ricc/index.html.

[21] T. Wood, P. Shenoy, A. Venkataramani, and M. Yousif, "Black-box and gray-box strategies for virtual machine migration," in *Proc. NSDI*, 2007.

[22] http://gigaom.com/cloud/tier-3-spiffs-up-cloud-management/.

[23] Oracle VirtualBox, http://www.virtualbox.org.

[24] Apache Mahout, http://mahout.apache.org.

[25] Olio, http://incubator.apache.org/olio.

[26] Faban, http://www.opensparc.net/sunsource/faban/www.

[27] A. E. Roth and M. Sotomayor, *Two-sided Matching: A Study in Game Theoretic Modeling and Analysis*, ser. Econometric Society Monograph. Cambridge University Press, 1990, no. 18.

[28] R. Irving, P. Leather, and D. Gusfield, "An efficient algorithm for the "optimal" stable marriage," *J. ACM*, vol. 34, no. 3, pp. 532–543, 1987.

[29] D. F. Manlove, R. W. Irving, K. Iwama, S. Miyazaki, and Y. Morita, "Hard variants of stable marriage," *Elsevier Theoretical Computer Science*, vol. 276, pp. 261–279, 2002.

[30] H. Xu and B. Li, "Seen as stable marriages," in *Proc. IEEE INFOCOM*, 2011.

[31] ——, "Egalitarian stable matching for VM migration in cloud computing," in *Proc. IEEE INFOCOM Workshop on Cloud Comput-ing*, 2011.

[32] B. Sotomayor, R. Montero, I. Llorente, and I. Foster, "Virtual infrastructure management in private and hybrid clouds," *IEEE Internet Comput.*, vol. 13, no. 5, pp. 14–22, September 2009.

[33] Z. Gong, X. Gu, and J. Wilkes, "PRESS: PRedictive Elastic Re-Source Scaling for cloud systems," in *Proc. IEEE CNSM*, 2010.

[34] A. Kalbasi, D. Krishnamurthy, J. Rolia, and M. Richter, "MODE: Mix driven on-line resource demand estimation," in *Proc. IEEE CNSM*, 2011.

[35] Z. Shen, S. Subbiah, X. Gu, and J. Wilkes, "CloudScale: Elastic resource scaling for multi-tenant cloud systems," in *Proc. ACM SoCC*, 2011.

