

Seen As Stable Marriages

Hong Xu, Baochun Li

Department of Electrical and Computer Engineering
University of Toronto

Abstract—In this paper, we advocate the use of stable matching framework in solving networking problems, which are traditionally solved using utility-based optimization or game theory. Born in economics, stable matching efficiently resolves conflicts of interest among selfish agents in the market, with a simple and elegant procedure of deferred acceptance. We illustrate through one technical case study how it can be applied in practical scenarios where the impeding complexity of idiosyncratic factors makes defining a utility function difficult. Due to its use of generic preferences, stable matching has the potential to offer efficient and practical solutions to networking problems, while its mathematical structure and rich literature in economics provide many opportunities for theoretical studies. In closing, we discuss open questions when applying the stable matching framework.

I. INTRODUCTION

Matching is perhaps one of the most important functions of markets. The stable marriage problem, introduced by Gale and Shapley in their seminal work [1], is arguably one of the most interesting and successful abstractions of such markets. There are men and women, looking for partners, as shown in Fig. 1. Each has a ranking, or *preference*, over agents of the opposite gender. Given marriages that assign each man to a woman, the following is certainly not desirable with respect to individual rationality: there exists a pair of man and woman who both prefer each other to their assigned partners. Such a pair is *unstable* in the sense that they have a clear incentive to break up from the current marriage and marry each other instead. Therefore, a good marriage does not induce any such unstable pairs: it is *stable*. When the problem is extended beyond a one-to-one setting (such as in the college admissions problem [1]), it is more generally referred to as *stable matching*.

Matching problems also exist pervasively in networking, ranging from assigning channels to users and flows in wireless scheduling, to mapping video segments to servers in video-on-demand streaming systems. In many cases, network elements are controlled by some central entity, and a metric of utility may be easily defined. The problem can then be formed into an optimization problem that can be systematically solved in a centralized or distributed manner, without paying respect to individual rationality.

However, as technologies evolve and networks become complex, optimization becomes monolithic, sometimes even inept, to be used in practice. Its effectiveness hinges upon the availability of complete and accurate information, which may not be realistic in practical settings. Parameter values serving as inputs may be distorted by delayed, noisy, or inaccurate measurements, which may lead to a significant “drift” in a computed optimal solution from the actual optimum. Further, an optimization algorithm needs to enumerate combinations of

agents from both sides to find the optimal solution, which is computationally expensive.

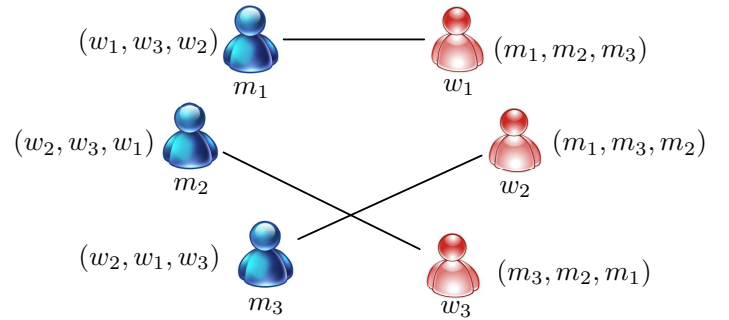


Fig. 1. A simple example of the marriage market. The matching shown is a stable matching.

Many large-scale networked systems, such as peer-to-peer networks, consist of numerous autonomous components. The interests of these selfish agents are not aligned, and individual rationality needs to be taken into account. While game theory is widely adopted in such scenarios, its existing use often relies on some rigorous definition of utility. Oftentimes it is difficult to quantize and unify the effects of various factors into a single utility function, not to mention the stringent requirement of complete information that is implicitly assumed. In practice, these techniques are not widely adopted due to these difficulties though theoretically sound.

In this paper, we advocate the use of stable matching as a general framework to tackle networking problems, where preferences are used to model each agent’s interest, and stability serves as the solution concept instead of optimality. As a case study, we present a problem related to content-eyeball ISP peering, and illustrate how it can be modelled and solved as a variant of stable matching problems. The use of the stable matching framework is a substantial and likely controversial shift from utility-based optimization or game-theoretic solution methods. Merits of the stable matching framework lie in the competitiveness of outcomes, generality of the preferences, efficiency and simplicity of its algorithmic implementations, and most importantly, its overall practicality. Stable matchings are in the core of the market that cannot be improved upon by a coalition of agents [2]. The generic preferences embrace many heterogeneous and complex considerations that different agents may have. The classical deferred acceptance algorithm can be applied in a centralized manner with little complexity. In addition, we propose a new decentralized mechanism that yields exactly the same outcome without centralized coordination at the helm.

We do not claim that we are the first to use stable matching in networking. There exist papers — though very few — that used the concept of stable matching, as we shall discuss soon. In contrast, the contribution of this paper lies in an explicit and general discussion of the stable matching theory and its application in networking with concrete case studies, and in provoking a broader understanding of its theoretic merit and practical value as another tool from economics. Broadly, stable matching can be applied in scenarios where coincidences and conflicts of interest among agents need to be resolved, but their considerations are difficult to be modelled quantitatively. It naturally fits into cases when agents’ individual rationality has to be respected. Alternatively, it also applies to cases with no hints of selfish agents, but optimization is clearly not possible or desirable.

II. BACKGROUND AND RELATED WORK

We start by introducing the basic theory of stable matching in the one-to-one marriage model, as shown in Fig. 1. There are two disjoint sets of agents, $\mathcal{M} = \{m_1, m_2, \dots, m_n\}$ and $\mathcal{W} = \{w_1, w_2, \dots, w_p\}$, men and women. Each agent has a complete and transitive preference over individuals on the other side, and the possibility of being unmatched [3]. Preferences can be represented as rank order lists of the form $p_{m_1} = w_4, w_2, \dots, \emptyset$, meaning that man m_1 ’s first choice of partner is w_4 , second w_2 and so on, until at some point he prefers to remain single (i.e. matched to the void set). We use $>_i$ to denote the ordering relationship of agent i (on either side of the market). If i prefers to remain single than being matched to j , i.e. $\emptyset >_i j$, then j is said to be *unacceptable* to i , and i ’s preference only include the acceptable partners. Preferences are *strict* if each agent is not indifferent between any two acceptable partners.

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

This implies that the outcome matches agents on one side to those on the other side, or to the empty set. Agents’ preferences over outcomes are determined solely by their preferences for their own partners in the matching.

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., $\emptyset >_i \mu(i)$.

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, $w >_m \mu(m)$ and $m >_w \mu(w)$. Such a pair is called a *blocking set* in general.

If there is a blocking set in the matching, the agents involved have an incentive to break up and form new marriages. Therefore such an “unstable” matching is not desirable.

Definition 4: A matching μ is *stable* if and only if it is individual rational, and is not blocked by any pair of agents.

The matching shown in Fig. 1 is a stable matching for the given preferences of agents.

Theorem 1: A stable matching exists for every marriage market.

This can be readily proved by the classic *deferred acceptance algorithm*, or the *Gale-Shapley algorithm* proposed in [1] (with men proposing). In the first round, each man proposes to his first choice if he has any acceptable ones. Each woman rejects any unacceptable proposals and, if more than one acceptable proposals are received, holds the most preferred and rejects all others. In each round that follows, any man rejected at the previous round makes a new proposal to his most preferred acceptable partner who has not yet rejected him, or makes no proposal if no acceptable choices remain. Each woman holds her most preferred offer *up to this round*, and rejects all the rest. When no further proposals are made, 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 men and women.

It is then observed that which side proposes in the algorithm has significant consequences. Specifically, the algorithm finds the two extremes among the set of stable matchings. The man-proposing version yields a man-optimal outcome that every man likes at least as well as any other stable matching, and the woman-proposing version a woman-optimal one. This is referred to as the polarization of stable matchings.

Following [1], the simple marriage model has been extended to other matching problems. Because of the richness of the literature, it is bold to even attempt a cursory survey of the existing results here. Instead, we choose to introduce some of the different models along with our case studies, and present relevant and important theoretical developments that can be applied to solve these problems, while using this section as a necessary background for a better understanding of the material.

From a practical perspective, due to the efficiency of stable matchings and the simplicity of implementation, the deferred acceptance algorithm has profound influence on market design. It has been adopted in a number of practical matching markets, prominent examples of which include the National Resident Matching Program of U.S. for medical school graduates, many medical labor markets in Canada and Britain, and recently school choice systems in Boston and New York City [3].

There only exists a very limited number of papers in the networking literature that used solutions designed to achieve stable matching [4]. In contrast, this paper seeks to present a systematic study on the feasibility and unique advantages of applying the stable matching framework to possibly a wider range of problems in networking, with an objective of drawing attention to and provoking discussions on this practical and effective tool from the field of economics.

III. CONTENT AND EYEBALL ISP PEERING AS MANY-TO-MANY STABLE MATCHING

In this section, we present a general many-to-many matching problem of interconnecting Internet Service Providers (ISPs), where stable matching is arguably the *only* practical solution.

The Internet consists of tens of thousands of ISPs, with profound heterogeneity. Roughly, they fall into three categories, *content*, *eyeball*, and *transit* [5]. Content ISPs (CPs) specialize in content delivery, such as Microsoft MSN and Google. Eyeball ISPs (EPs) are in the business of selling retail Internet access to end users, such as Verizon and Comcast. Traditionally it has been understood that CPs and EPs buy transit from tier-1 ISPs to deliver their traffic, and EPs establish settlement-free peering links among themselves, provided that their network sizes or traffic volumes are roughly equal in order to save costs.

This hierarchical picture is becoming increasingly invalid. Some CPs have built their own backbones for efficient delivery of content that constitutes the majority of Internet's traffic [6], and thus become an indispensable part of the network. EPs also possess increasing bargaining power because of their large customer base. To reduce the bulk of transit costs and improve latency, CPs are now *directly* peering with EPs, creating a much flatter Internet pyramid with tier-1 ISPs losing their grip on the ISP peering ecosystem [7].

This trend has been increasingly realized and independently observed among practitioners and researchers [6], [7]. Existing peering mechanism, however, may not work well for the emerging content-eyeball peering market. As of today, peering decision is usually made through bilateral negotiation. These decisions may look beneficial from a local perspective, but from a global perspective they are very unappealing [5]. The resulted market inefficiency is hard to rectify, since breaking the contract involves the risk of legal lawsuit and possibly anti-trust scrutiny, and incurs disruption to the Internet.

The difficulty in establishing a market mechanism is mainly the idiosyncratic factors involved in making peering decisions, including geographic coverage, traffic volume, routing requirements, marketing considerations, etc., alongside the common and conflicting interests of ISPs that will induce market failure if not resolved. Existing papers have studied novel settlement strategies of ISPs through economical analysis when an ISP interconnection topology is given [5], but few has touched the issue of how this topology is formed endogenously, especially among content and eyeball ISPs. Moreover, the market is inherently decentralized, aggravating the mechanism design problem.

The stable matching framework can be naturally applied here as a first attempt to understand and address the content-eyeball ISP peering problem, where individual rationality has to be respected. Content-eyeball peering can be seen as a many-to-many stable matching problem with quotas, and an extension of the deferred acceptance algorithm can serve as the centralized multilateral mechanism. Further, we provide a decentralized algorithm that preserves the privacy of preference information and the final matching outcome. Our theoretical results, therefore, offer novel perspectives by making use of the unconventional stable matching framework.

One may argue that, the usual game-theoretic approach, for example Nash Bargaining Solution based mechanisms, and the mechanism design approach, for example VCG auctions, can

also be applied here. The problem with game-theoretic approach is that some concrete utility function needs to be defined (similar to optimization), which is not possible considering the enormous complexity involved when making peering decisions. The mechanism design approach is also unlikely to succeed, due to collusions that are extremely common in real-world markets. Stable matching remedies both issues by adopting a general preference framework and a stability solution concept resistant to collusion of coalitions from both sides [2].

A. The Model

ISPs generally have different incentives to form peering relationships. As it improves latency which is critical to their revenue, CPs tend to exhibit strong incentives for peering, while some EPs may only be willing to enter a paid-peering contract, since it possesses bargaining power under the assumption that eyeball customers are less vulnerable to switching to another EP. Their interests can also be in conflict, especially when one ISP is favored by multiple ISPs.

Thus, content-eyeball peering can be cast as a many-to-many stable matching problem with a set of CPs \mathcal{C} and a set of EPs \mathcal{E} . Each ISP has a quota, which is the maximum number of partners it allows for peering, due to technical overhead and personnel constraints of setting up the connections, and business considerations such as paid peering. Each ISP is assumed to have a strict preference ordering of its *acceptable* partners. This set of acceptable partners and their rankings could be produced by each ISP identifying and contacting their potential partners and evaluating the potential benefits.

Due to the variety and complexity of economic and business policy factors affecting this evaluation process, we can not and shall not attempt to define a generic preference framework or anything alike that each ISP adheres to, as in sharp contrast to most previous work [5]. The only assumption on the *structure* of the preferences we make is a *responsiveness* assumption in order to reduce the exponential complexity of expressing preferences over all subsets of acceptable ISPs. Specifically,

Definition 5: For any set of EPs \mathcal{E}_1 with $|\mathcal{E}_1| < q_c$, and any EP e and e' not in \mathcal{E}_1 , CP c prefers $\mathcal{E}_1 \cup e$ to $\mathcal{E}_1 \cup e'$ if and only if e is preferred to e' under c 's preference p_c , and prefers $\mathcal{E}_1 \cup e$ to \mathcal{E}_1 .

The same can be defined for EPs. Responsive preferences are a special case of preferences in which ISPs are substitutes rather than complements to the opposite side [3]. This essentially means that a CP always prefers adding an acceptable EP before reaching the quota and it always prefers replacing a EP with a better one when the quota is met. Clearly this is reasonable to assume, and it also establishes the sufficiency of preferences over individual ISPs to find a stable matching [8].

Definition 6: A matching $\mu \subset \mathcal{C} \times \mathcal{E}$ for a content-eyeball peering problem is such that, each CP $c \in \mathcal{C}$ appears once in at most q_c pairs where q_c denotes its quota, each EP $e \in \mathcal{E}$ appears once in at most q_e pairs, and each pair $(c, e) \in \mu$ is individually rational, *i.e.* mutually acceptable to c and e .

B. Solution Concepts and a Centralized Mechanism

The technical difficulty we encounter here is the choice of stability concepts. Many-to-many matching is a more general model than many-to-one matching. Several stability concepts can be defined. The most common ones in the literature are *pairwise stability*, *corewise stability*, and *setwise stability* [9].

Definition 7: A matching μ is *pairwise stable* if there are no ISPs c and e who are not partners in μ , but by becoming partners, possibly dissolving some of their partnerships given by μ to remain within quotas and keeping other ones, can both obtain a strictly preferred set of partners.

This is the usual stability concept we have seen in the one-to-one and many-to-one problems.

Definition 8: A matching μ is in the *core* (*corewise stable*) if there is no subset of ISPs who by forming all their partnerships only among themselves, can all obtain a strictly preferred set of partners.

Definition 9: A matching μ is *setwise stable* if there is no subset of ISPs who by forming new partnerships only among themselves, possibly dissolving some of their partnerships given by μ to remain within quotas and keeping other ones, can all obtain a strictly preferred set of partners.

Both corewise and setwise stability deal with coalitions of multiple ISPs. The key difference is that corewise stability requires all agents involved in a blocking set to be better off, while setwise stability only concerns the agents who form new partnerships in the blocking set. More detailed comparison and discussions of the concepts are available in [9].

Intuitively, setwise stability is the strongest definition and the other two are special cases of this concept. However, its properties and algorithmic implementations are less well understood compared to the usual pairwise stability. This dilemma does not exist for one-to-one and many-to-one cases where the three concepts are equivalent [9], which is not the case for many-to-many matching in general, simply because each agent can have multiple partners and form different coalitions [9].

Therefore, for empirical purposes, we restrict our attention to the case where the coalitions are formed between a single CP and a set of EPs, or a single EP and a set of CPs, but not between multiple CPs and EPs. This is a valid assumption in practice because a CP (EP) negotiates with groups of EPs (CPs) but not with other CPs (EPs). In this case, it is proved that the three concepts are equivalent when preferences are responsive [2]. Thus it is sufficient to work with pairwise stable matchings in our problem.

Theorem 2: Pairwise stable matchings always exist in the content-eyeball peering problem, when the ISPs have strict and responsive preference orderings [10].

This can be readily proved by means of a centralized deferred acceptance mechanism, which is a generalized version of the one we discussed in Sec. II. Essentially, the mechanism works by letting agents on one side of the market propose to its most preferred subset of ISPs that have not rejected them, and agents on the other side accept their most preferred subset of proposals and reject the rest.

The analysis of this mechanism, including the polarization of the outcomes, is essentially the same as that in Sec. II with proper generalization. Due to space constraints, we omit the details in this paper. Note that since CPs have stronger incentives to peer, and EPs have the benefit of getting paid peering deals, it may be convincing to let CPs be the proposing side and the mechanism produces the CP-optimal outcome.

C. A Distributed Implementation

The content-eyeball peering market is inherently decentralized. Thus a distributed procedure with minimum centralized coordination is highly desirable. Moreover, in reality, an ISP would like to keep its preference and the final peering results *private* for obvious reasons, which cannot be achieved in the centralized mechanism. In this section, we present a new distributed mechanism that preserves information privacy while producing precisely the same stable matching given the same problem instance. Our mechanism is an extension of [11] in the marriage market. It consists of two procedures, *CP* and *EP* as presented in Algorithm 1 and 2, which are executed in each CP and EP, respectively. Note that the execution is *asynchronous*.

Procedure *CP*, after initialization, performs the following loop. If a CP c does not have enough partners confirmed, and its preference p_c is nonempty, it iterates by proposing to its first choice t , deleting t from p_c , and adding t to *list*, until it has q_c partners in *list*. It then waits on a message reply. If the reply is accept, it does nothing. If the reply is reject, it removes the sender from p_c . If the sender is in its *list*, it removes the sender from *list*, which means its proposal to the sender failed. Rejection is normally from EPs whom c has proposed to. It can also be from EPs that c did not propose to, which we will explain shortly in the *EP* procedure. This continues until a special stop message is received from a “registrar,” which we assume is a server that delivers all the encrypted messages exchanged between ISPs and thus can detect the quiescence.

Algorithm 1 Distributed CP Procedure

```

list  $\leftarrow \emptyset$ , end  $\leftarrow$  false;
while !end do
  while |list| <  $q_c$  and  $p_c \neq \emptyset$  do
     $t \leftarrow \text{pop}(p_c)$ ;
    sendMsg(propose,  $c$ ,  $t$ ); add(list,  $t$ );
  msg  $\leftarrow$  getMsg();
  switch msg.type
  accept: do nothing;
  reject: delete( $p_c$ , msg.sender);
         if msg.sender  $\in$  list then
           delete(list, msg.sender);
  stop: end  $\leftarrow$  true;

```

The *EP* procedure is more complicated. It keeps a loop that receives messages from CPs. Upon getting a proposal from c , it accepts if quota is not used up, and do an insertion sort to add c in order. A tricky issue here is that, if the resulted number of offers exceeds its quota, it sends messages to reject all CPs after the last element in *list*, simply because it will never accept

proposals from these CPs given the status quo. This reduces the possible number of proposals from those who will definitely be rejected, and speeds up the algorithm. This also ensures that when quota was met, EP e can accept a new proposal from c as long as c is currently in p_e . If this is the case, then again all CPs after the newly accepted one in p_e are also rejected and deleted from p_e . The EP procedure also terminates when it receives the stop message from the registrar.

The distributed and asynchronous stable matching mechanism eliminates the need of centralized coordination at the helm for the matching process. More importantly, it perfectly maintains the privacy of preference information and the final peering decisions. Each ISP only accesses its own preference, and at the end it has knowledge about only its own peering result. As the distributed mechanism shares the same deferred acceptance procedure, the final outcome is exactly the same CP-optimal stable matching produced by the centralized mechanism.

Algorithm 2 *Distributed EP Procedure*

```

list  $\leftarrow \emptyset$ , end  $\leftarrow$  false;
while !end do
  msg  $\leftarrow$  getMsg();
  switch msg.type
    propose:  $t \leftarrow$  msg.sender;
      if |list| <  $q_e$  then
        sendMsg(accept,  $e$ ,  $t$ );
        insertionSort(list,  $t$ );
      if |list| =  $q_e$  then
        for each  $i$  after  $t$  in  $p_e$  do
          sendMsg(reject,  $e$ ,  $i$ );
           $p_e \leftarrow p_e - i$ ;
      else if  $t \in p_e$  then
        sendMsg(reject,  $e$ , last(list));
        sendMsg(accept,  $e$ ,  $t$ );
        insertionSort(list,  $t$ );
        for each  $i$  after  $t$  in  $p_e$  do
          sendMsg(reject,  $e$ ,  $i$ );
           $p_e \leftarrow p_e - i$ ;
      else then
        sendMsg(reject,  $e$ ,  $t$ );
  stop: end  $\leftarrow$  true;

```

IV. LESSONS LEARNED AND CONCLUDING REMARKS

In this paper, we articulated the stable matching framework, which advocates a novel perspective on solving networking problems practically. Instead of striving to achieve any concrete notion of optimality, our framework pursues the unconventional *stability* as the central solution concept using only the ordering information in two-sided matching problems. Through the case study, stable matching has been evidently demonstrated as a simple, efficient, and practical solution framework applicable to complex problems where the traditional utility-based approach is inapplicable or impractical.

In trying to be impartial, we wish to critically examine what is less satisfactory from the stable matching framework. In retrospect, one of the most compelling concerns is the polarization

of stable outcomes that favors the proposing side of the market. It is desirable in many cases to find a “fair” stable matching that does not favor either side as an alternative operating point of the system. There are efforts in the economics literature that pursue this direction, such as the *egalitarian* stable matching that minimizes the total rank sum of the outcome in the marriage model [12]. This is a feasible extension to our framework.

Strategy-proofness is also crucial to the effectiveness of stable matching mechanisms. Is it the dominant strategy for agents to *truthfully* report, or act according to their preferences and quotas? Can they benefit from manipulating this information? In this regard, some negative results have been proven on the theoretical impossibility of achieving truthfulness for both sides of the market in some models [2]. On the other hand, many real-life implementation experiences suggest that the impact of misbehavior is extremely limited, especially when each agent has a small number of acceptable partners in a thick market [3]. Given this mixed picture, it is interesting to explore the issue in a networking context, and understand the extent to which this impacts the suitability of the new framework.

Our attempt of advocating stable matching in favor of conventional optimization or game-theoretic approaches here may seem ambitious. Yet, we believe that stable matching, due to its rich theoretical foundation and practical implementation experiences, has unique merits in solving problems in the networking domain. Towards this vision, this paper is meant to serve as a first step towards applying stable matching to design practical networking solutions. With unique advantages that preferences are more flexible to express practical considerations, and that its solutions are efficient and may be decentralized, stable matching has the potential to become a general framework to solve a wide range of networking problems.

REFERENCES

- [1] D. Gale and L. S. Shapley, “College Admissions and the Stability of Marriage,” *Amer. Math. Mon.*, vol. 69, no. 1, pp. 9–14, 1962.
- [2] 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.
- [3] A. E. Roth, “Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions,” *Int. J. Game Theory*, vol. 36, pp. 537–569, 2008.
- [4] F. Fazel and D. R. Brown III, “On the Endogenous Formation of Energy Efficient Cooperative Wireless Networks,” in *Proc. Allerton*, 2009.
- [5] R. T. Ma, D. Chiu, J. C. Lui, V. Mistry, and D. Rubenstein, “On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers,” in *Proc. CoNEXT*, 2008.
- [6] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, “Internet Inter-Domain Traffic,” in *Proc. ACM SIGCOMM*, 2010.
- [7] G. Goth, “New Internet Economics Might Not Make It to the Edge,” *IEEE Internet Computing*, vol. 14, no. 1, pp. 7–9, 2010.
- [8] A. E. Roth, “The College Admissions Problem Is Not Equivalent to the Marriage Problem,” *J. Econ. Theory*, vol. 36, pp. 277–288, 1985.
- [9] F. Echenique and J. Oviedo, “A Theory of Stability in Many-to-Many Matching Markets,” *Theoretical Economics*, vol. 1, pp. 233–273, 2006.
- [10] A. E. Roth, “Stability and Polarization of Interest in Job Matching,” *Econometrica*, vol. 53, pp. 47–57, 1984.
- [11] I. Brito and P. Meseguer, “Distributed Stable Matching Problems,” in *Proc. Constraint Programming*, 2005, pp. 152–166.
- [12] 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.