

Existence of Market Equilibria for Grid Computing

Gireesh Shrimali¹, Giri Kumar Tayi²

¹Indian School of Business, Hyderabad, India

²School of Business, Albany, USA

Email: gireesh_shrimali@isb.edu, gtayi@albany.edu

Received July 15, 2012; revised August 27, 2012; accepted September 5, 2012

ABSTRACT

Grid computing has emerged as an effective mechanism for allocating globally available surplus computational capacity to applications whose requirements exceed local capacity. It is often viewed as a commodity exchange with additional grid computing specific constraints that may arise due to requirements on multiple resources (e.g., disk space) in addition to computing power. These constraints are related to complementarity and substitution effects among resources, and significantly alter the assumptions typically used for demonstrating the existence of market equilibrium. However, prior work in grid computing has simply assumed that market equilibria exist. Our work fills this gap by studying the existence of market equilibrium under the grid computing environment. To do so, we first establish an economic framework that incorporates the grid computing specific constraints into a commodity market. We next derive some intuitive necessary conditions based on the computing requirements of individual agents. We finally establish the existence of regular markets as a competitive equilibrium, given that these necessary conditions are met and that the agents' utility functions satisfy some minimal requirements. In the process, we also show existence of competitive equilibrium for the special case of grid computing as a pure exchange economy¹.

Keywords: Grid Computing; Competitive Markets; Exchange Economy; Market Equilibrium

1. Introduction

1.1. Motivation

As the number of computers connected to the Internet grows, the number of idle ones at any instance also grows. This phenomenon results in the availability of a substantial amount of Internet-connected surplus computing resources (e.g., microprocessors, memory devices, storage disks, etc.) at any point in time. On the other hand, at the same time, the proportion of Internet-connected devices that may require the use of this surplus capacity also grows. As an example of the latter, consider the enormous growth of ubiquitous computing that relies on mobile and embedded devices [1,2]. These devices, through their interaction with human beings, produce a huge amount of data that need to be processed efficiently anytime anywhere. However, such devices often have limited resources in terms of CPU, storage, battery power and communication bandwidth. Thus, there is a need to transfer ubiquitous computing application services to more powerful computational resources available elsewhere.

Grid computing has emerged as an effective mecha-

nism for allocating this surplus capacity to computers (and devices) that run applications whose requirements exceed locally available resources [3]. It is essentially a virtualization of computing in which computing resources can be combined using the Internet to fulfill the needs of computation hungry applications. Grid computing may appear in the literature in many different forms. For example, it may refer to volunteer computing [4], peer-to-peer (P2P) computing [5], or cloud computing [6], to name a few. At a high level, all of these allow exchange of computing resources in some form, and therefore, for the rest of the paper we refer to all of them collectively as grid computing. In essence, a computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities [7].

Grid computing originated in academia and research labs in order to solve CPU-intensive research problems, such as SETI@home [4]. Over time, it has been making its way into large enterprises due to its potential to create business value [8]. For owners of surplus computing resources, grid computing creates more revenue-generation opportunities—for example, consider Amazon's S3 platform [9]. Further, for many firms it has the potential to lower IT costs through outsourcing their IT functionality to the grid, perhaps as much as 30% [10]. As a result,

¹The most general market equilibrium model, not surprisingly, is called the general equilibrium model. In this model all exogenous prices (on goods) are variable and equilibrium requires that all markets clear under independent decisions made by agents.

grid computing is now widely deployed by organizations and provides seamless temporary processing capacity expansion to handle peak period demand on e-commerce servers, distributed gaming, and content storage and distribution [11]. This trend is evident through reports such as “Grid Computing: A Vertical Market Perspective 2005-2010” which claimed that the worldwide Grid spending would increase from \$714.9 million in 2005 to approximately \$19.2 billion in 2010 [12]. In fact, [13] has argued that information technology’s shift from a fragmented, locally owned, capital asset model to a unified and centralized grid computing based service model will be momentous.

Grid computing enables commoditization of computing resources [14], resulting in a market place for trading these resources [15]². In this market place owners of computing resources can sell the use of their idle resources to owners of computing tasks, whose computing requirements surpass their own computing resources. A natural way for efficient management of these resources in a grid computing environment is to use an economic framework [15], which assumes that agents participating in grid computing are not altruistic and that their incentives matter—an assumption that we stick to for the rest of the paper. A variety of economic models that have been used in the context of selling goods and services are also relevant to grid computing: each of these entails a different model for setting the price of the resources or goods based on supply-and-demand and their utility to the agent. These works include monopoly markets [16-18], competitive markets [2,15,19] as well as auctions [11,20].

From the perspective of an agent (for example, a manager of computing resources in a firm) it is critical that trade takes place once it has decided to participate in a grid market. That is, the agent would like to know upfront whether trades would indeed take place and that the grid markets exist. Interestingly, researchers so far have implicitly assumed that such markets exist. This in turn assumes that grid markets are identical to other commodity markets—an assumption that may not be true, which may cast doubt on the validity of these results. It is not necessary that such grid markets have to exist since grid resources and the environment within which they get traded differ fundamentally from traditional commodities such as grain, iron ore and other raw materials for which there are well established commodity markets. Previous research on grid market structures does not provide any guidance or approach to ascertain the existence of grid markets which deal with commodities that are unique and

have no commonality with traditional resources [21,22].

Although grid computing markets can be thought of as a way for exchanging commodities, these markets have additional constraints on what bundles can be bought and sold at a certain time. These bundles may arise due to requirements on multiple resources, such as disk space, memory bandwidth, etc. in addition to the usual requirement of computing power. For example, a grid computing bundle may require a certain amount of disk space in addition to computing power. This suggests that certain complementarity and substitution effects may exist among various resources in the grid computing market. These can be referred to as grid computing specific constraints. These constraints may significantly alter the assumptions typically used for demonstrating existence of market equilibria [23]. That is, there is a possibility that markets may not even exist under grid computing specific constraints and, if they do exist, they may require additional conditions on agents’ utilities as well as on commodity spaces.

Thus, there is a significant gap in the existing literature due to this assumption of existence of grid computing markets. Our work fills this gap by studying the existence of market equilibria under grid computing specific constraints. Such a study is helpful to IT managers as it would unequivocally assure that grid markets are feasible and importantly would provide guidance to Grid market operators such as France Telecom for Mobile Grids, Sun Microsystems for Computational and Data Grids, etc. in developing practical mechanisms that are critical to the development of robust grid markets in a real world competitive environment within which these firms operate [24].

1.2. Our Work

We earlier alluded to how grid computing markets are different from regular commodity markets. We now expand on this further by highlighting key characteristics of grid computing markets.

First, in traditional commodity markets, typically one agent always has the same resource in surplus while another agent always demands the same resource. In contrast, in grid computing an owner who has surplus computational resource at a particular time may run short and hence demand the same resource at another time. That is, in such a market place, depending on the resources owned and computing needs, it is possible that agents change their roles over time, and an agent can be both a buyer and a seller.

Second, the key aspect of distributed ownership poses a challenge to realizing the full potential of grid computing. That is, not only the agents are in competition with each other, an independent agent who has excess com-

²Amazon’s S3 is an example of such commoditization of computing resources [9]. We note that, similar to regular markets, service (*i.e.*, product) differentiation is possible in the context of grid computing. Therefore, it would require a corresponding extension to our basic model. This is beyond the scope of this paper.

puting resources may have no incentive to offer it for use by another independent agent who faces a shortage of such resources. For instance, if agents care about money they would like to be compensated for offering the use of their excess resources—in particular, if they need to cover the costs of operating these resources. On the other hand, if the agents care about exchanging resources they would like to know that the resource exchanges are fair—in this case, one way to ensure fairness is the existence of implicit prices that underlie trades. In addition, not only the owners of each type of these resources may have different usage or access policies, cost structures and varying availabilities but also the agent who lacks enough resources also has requirements, which may not allow for a perfect exchange to occur.

Finally, these resources could be scattered across multiple geographically distributed organizations with different usage rates, varying availability and separate cost functions; and managing these resources is a complex activity. Thus, there is a need for robust economic models that realistically model the interaction among self-interested agents in the grid computing marketplace. This is an integral part of a robust environment that, engenders complex resource management and scheduling issues that are central to grid computing.

With these distinctions in mind, we are interested in exploring two fundamental theoretical questions regarding the existence of grid computing markets. These, in our opinion, should be resolved before the properties of the market mechanism as well as its implementation aspects are thoroughly investigated. First, would regular markets even exist under grid computing? Second, if so, under what conditions? One may ask as to why we even care—that is, what is different about grid computing that may change the structure of the economic interaction among agents trading for resources in a market place.

Indeed, there are three basic factors that differentiate trading for grid resources from trading in regular markets.

First, computing requires many resources together, such as microprocessors, local memories (e.g., dynamic access memories), mass storage (e.g., hard disks), etc. For example, a task requiring m units of compute power (through microprocessors) may also require at least n units of hard disk space. In economic terminology, these resources are *complements*.

Second, computing needs may be spread over time. For example, an agent may need a total of m units of compute power over two distinct time slots (e.g., day and night) but may not care how the compute power is distributed over the two time slots. So, given m_1 and m_2 as the assigned compute powers in the two time slots, the agents would require that $m = m_1 + m_2$, with minimal restrictions of the actual values of m_1 and m_2 . That is, in

economic terminology, m_1 and m_2 are substitutes.

Third, agents may have *job must be finished* requirements. For example, an agent may have a strict requirement of at least m units of compute power. Without meeting this requirement, it may not participate in the market. Note that most of the requirements of a grid computing environment can be expressed in terms of these factors. For example, a quality of service requirement of a certain (maximum) response time may be translated into a combination of job-must-be-finished requirements and complementarity conditions.

The study of the existence of a regular market would require examining partial equilibrium concept, where all prices other than the good being studied are assumed to remain fixed, under the three additional sets of constraints outlined above. However, we study the general equilibrium model instead not only for its generality but also to cover the special case of grid computing as an exchange economy due to the observation that regular markets for grid computing have not become as popular as originally hoped for [15]. This is evident from the “Popcorn Market” experience [25], where agents were not interested in using real money due to various reasons, such as complexity of the markets, the lack of value creation from users’ perspective, and technical reasons such as security concerns, etc. This brings forth the need for studying alternative market mechanisms, such as an exchange economy [23,26], to study the interaction among agents in a grid market place³. Thus, for most of the paper we stick to the general equilibrium concept and examine its application to regular markets in Section 4 as an example.

To study the general equilibrium concept for grid computing, we present a discrete time model [11] that allows for substitution across time slots. We assume that both supply and demand are deterministic, and look at two standard concepts of market equilibria. The first concept is the most general model used in the study of market equilibria—the less known price quasiequilibrium with transfers [23]. The second concept is the well-known (but specific) Walrasian equilibrium [23]. We study both of these concepts to ensure a full breadth of coverage of market equilibrium models. These equilibria depend on underlying assumptions on agents’ utilities, and may be implemented in fundamentally different ways. These equilibrium concepts are further defined in detail in Section 3, where we describe why studying these different equilibria is instructive. Given that the managers of computing resources are eventually interested in the mechanism through which they would exchange resources, they

³That is, a market place where agents with surplus computing resources would come to the market place and exchange their resources with those owners who need them, as long as it is mutually beneficial to do so over time—note that this would work as long as agents change roles over time.

would care about not only whether markets would exist but also which form they would eventually take.

Under some standard assumptions (e.g., convexity, etc.) on agents' utilities, we show that these equilibria do exist in the presence of the first two grid computing specific requirements of complementarity and substitution, given that some basic necessary conditions that capture the third requirement of job-must-be-finished at an aggregate level. Specifically, similar to the typical approaches to demonstrating existence of market equilibria [23], we show that convexity of agents' utilities is essential to showing existence of market equilibrium. Further, we show that the sufficient conditions get stricter as we move from the price quasiequilibrium with transfers to the Walrasian equilibrium. The former requires the agents utilities to be convex and locally-satiated, whereas the latter requires them to be continuous, strictly convex, and strongly monotone⁴. However, we also show that the somewhat strict sufficient conditions for the existence of Walrasian equilibrium are not necessary, and provide some illustrative examples that demonstrate the same even when all the conditions ensuring its existence are not met.

Finally, we provide some extensions to our basic general equilibrium model that demonstrate the robustness of our results and make our work more complete. First, we extend our results to the specific case of regular competitive markets, where money is also a commodity. Recall that this was our objective to begin with. Second, we relax our assumption of deterministic supply and demand and show that our result would hold even when these are stochastic—a more realistic scenario. Third, we establish a finite non-cooperative basis for the competitive outcome which shows that our results hold even in the case of finite markets where agents may have market power.

We note that the ideas of grid computing as well as market equilibrium are not new. The novelty of this paper lies in the modeling of grid computing as a commodity market (and an exchange economy) for reasons outlined above. Further, our theoretical results, which not only apply to grid computing as a regular market but also to grid computing as an exchange economy, fill a gap in the literature by ensuring existence of these markets under grid computing specific constraints. We believe this gap to be a major one since previous literature examined characteristics and operations of grid markets assuming these markets exist, which may not have been true after all.

1.3. Related Work

The early work in grid computing has primarily focused

⁴Many of the technical terms, such as convexity, local non-satiation, monotonicity, etc., are explained in Section 3.3.

on scheduling problems using a conventional static approach where a scheduling mechanism decides which users resource request would be executed using which owners resource. Typically, a cost function was optimized in making the scheduling decision [27,28]. However, the cost optimization was normally based on system-wide considerations rather than individual users requirements, such as delivery time frame. Further, this treatment implied that all user requests are the same thereby preventing individual users to negotiate based on their demand, priority and budget availability. In addition, despite recent attempts to model the scheduling problem in a non-cooperative way [29], this line of research has not considered an economic framework using which agents could trade resources.

Grid computing economics has an extensive literature. [15] was one of the first to look at economic mechanisms for grid computing, and introduced a simple economic framework. The framework allowed for the users (consumers) and owners (producers) of these resources to have different goals, objectives, and strategies; and discussed how they can be mapped to grid environment and existing grid tools and architectures. It also enabled the use of a variety of mechanisms and market models for setting prices of these resources based on their supply and demand and their value to the user. However, though [15] identified the application of competitive markets and auctions to grid computing it did not provide any specific mechanisms. Also, it did not make an attempt to analytically model the interactions between the agents in these diverse market structures. Finally, it did not examine the economic properties of the above mentioned market structures, namely whether they are allocatively efficient or incentive compatible.

Many papers [11,19,20,25] have looked at application of auctions to grid computing. [25] was one of the first papers to look at the application of auctions to grid computing. It described a market where people could trade with virtual money and sellers would simply auction their goods to the highest bidder. However, it looked at computing power in isolation and ignored the effects due to complementarity and substitution. [11] addressed these deficiencies and presented a unified framework that addressed complementarities among various grid resources as well as substitution effects across the time dimension. It formulated the problem as a combinatorial call auction and presented a portfolio of three solution approaches that trade off economic properties, such as allocative efficiency, incentive compatibility, and fairness in allocation, with computational efficiency. However, although the mechanism accounted for quality and time attributes and enabled the simultaneous trading of multiple buyers and sellers, it did not account for competition on the sellers' side as all orders were aggregated to one

virtual order. Moreover, the mechanism did not take co-allocation constraints into account. However, a major shortcoming of this work is that it does not raise the issues of efficient mechanism design and incentive alignment of the participating agents. All these deficiencies were addressed in [20]. Our work borrows the latest modeling techniques for auctions from [11] and [20], and applies them to commodity markets, in particular to exchange economies.

[19] looked at the efficiency of resource allocation under two different market conditions: commodities markets and auctions. It compared both market strategies in terms of price stability, market equilibrium, consumer efficiency, and producer efficiency. It then showed that commodities markets are a better choice for controlling grid resources than previously defined auction strategies, such as [25]. This motivated us to look at analytical properties of commodity markets. In this line of research, some recent work has also developed stochastic market equilibrium models for QoS assurance in grid computing environment [2].

Computational grids differ from other high performance computing systems in that the nodes where computations are carried out exhibit varying degree of heterogeneity. Therefore computing jobs in such grids generate a stochastic demand for computing resources. Another stream of recent research has explicitly considered the stochastic nature of the demand for computing resources. [2] developed a competitive model of the interaction between computing resources of the grid service providers and the ubiquitous applications embedded in the subscribers mobile devices. It presented the equilibrium behavior of the model when the interaction results in stochastic demand which is satisfied according to guaranteed quality of service standards. Further, given that efficiently balancing the load on the grid implied by these jobs is important, [29] developed a game theoretic solution to the load balancing problem.

Finally, many papers have focused on grid computing from a monopolist's point of view [16-18] and view grid computing as a utility that a monopolist may be interested in selling. These have looked at determining cost structures in providing utility computing services and focus on developing pricing mechanisms in a stochastic environment. Given our focus on competitive markets, our work does not have much in common with this body of work, except that some of these ideas could be borrowed for the stochastic extension of our work.

None of the above papers has examined the existence of market equilibrium, the main focus of our paper. Our work is influenced by [26] that modeled a peer-to-peer file sharing system as an exchanged economy, and established not only the existence of market equilibrium but also a non-cooperative (finite) framework that would

converge to a competitive equilibrium in the limit. Given that grid computing can also be viewed as a peer-to-peer computing system, our work can be viewed as the extension of [26] to grid computing.

1.4. Organization

The paper is organized as follows. In Section 2 we present our basic model. In Section 3 we present our theoretical results as well as numerical examples. In Section 4 we present various extensions to our results. Finally, Section 5 concludes.

2. The Model

The main idea is to introduce the time dimension with time as a discrete parameter and use it to define goods [11]. The duration T (for example, a day) is divided into J non-overlapping chunks (for example, a day can be divided into 24 one-hour slots). Then a good may be defined as one computation-unit-slot. Another example could be a storage-unit-slot.

Then, the problem we are trying to solve is: *given J slots, G type of goods (i.e., a total of $L = G * J$ goods), and I agents, what is the set of conditions that ensure the existence of an exchange economy?* We denote by w_{igj} as agent i 's initial endowment of good g in slot j . Similarly, by x_{igj} , we denote the final allocation of good g for agent i in slot j . Note that one can think of $l \in \{1, \dots, L\}$ as the index traversing a vector of goods that results by vectorizing the matrix of goods indexed by rows $g \in \{1, \dots, G\}$ and columns $j \in \{1, \dots, J\}$.

Having defined goods this way, now we can have a regular exchange economy with agents holding initial endowments of goods and possible zero preference for certain goods. The following, somewhat trivial, example shows the existence of one exchange economy.

Example 1. *Say there are J agents that own one unit of computation unit each and want J computation units over one slot each, with no overlaps among computing requirements. That is, agent i has a computing requirement of J during time slot i only. Then this is a feasible problem, with the prices of all goods being equal. Agent i would be the only consumer of good i and all agents would supply good i .*

Before we start a formal treatment of the grid computing markets, some initial observations are as follows. First, an agent needs to have some goods to sell in this exchange economy. In order to have enough buying power in this economy, it should have some combination of initial endowments that provides enough "cash" or money that allows the agent to acquire the required allocations. As an example, the agents can either have a lot of goods with low demands or a few goods with high demands. Note that this cash acts as a medium of transfer,

and may be real or virtual depending on the kind (regular vs. exchange) of economy under study. Second, typically an agent would be interested in participating only if its initial endowments don't satisfy its required allocation. However, it may want to participate even if its initial endowments satisfy its requirements, especially if participating in the economy improves its utility in some way. For example, this may happen if the agent already has its required computing allocations in terms of its initial endowments in the first slot, and the agent is indifferent between getting the same allocations in the first and second slots. Then, if the agent can improve its cash balance (e.g., in a regular market) by selling its endowments in the first slot and buying the same amount of allocations in the second slot, its overall utility can be improved, and it would participate in the market.

3. Equilibrium with Grid Constraints

As discussed earlier, there may be some exceptions to a general exchange economy due to how grid markets may work. We are interested in how these exceptions affect the existence of an equilibrium. We begin by looking at these exceptions in Section 3.1. We then derive some necessary conditions for the existence of various market equilibria in Section 3.2. Next, we provide sufficient conditions for the existence of various market equilibria in Section 3.3. Finally, Section 3.4 wraps up this section with some illustrative examples of existence of Walrasian equilibrium.

3.1. The Grid Constraints

First, note that the goods representing the same type of good across time slots (*i.e.*, x_{ig^*}) are substitutes, and all the agent may care about is that it gets enough of a particular type of good over all time slots. For example, x_{ig1} may substitute for x_{ig2} and vice versa—the computing cycles in slot 1 are a substitute for the computing cycles in slot 2.

This substitution may have underlying coefficients of substitution embedded in the utility functions. That is, a unit of x_{ig1} may not weigh the same as a unit of x_{ig2} . However, for simplicity we assume that the coefficients of substitution are all one. This is particularly true for CPU time due to the sequential nature of programs. Further, the issue of substitution of other goods becomes somewhat irrelevant due to the complementarity conditions outlined below. For example, for disk capacity, 1 GB in period 5 and 1 GB in period 6 may not be the same as 2 GB in period 5 on its own—however, the capacity breakdown may suffice in conjunction with the complementarity requirements associated with CPU time.

Second, there may be intra-slot complementarities be-

tween (type of) goods. That is, there may be dependencies among different types of goods. This is reflected as,

Requirement 1. $x_{igj} \geq k_{igj} x_{ilj}$, $\forall g \in \{1, \dots, G\}$, $\forall j \in \{1, \dots, J\}$, where the k_{igj} 's are constants.

For example, given x_{i1j} , x_{i2j} as the acquired (type of) good-type 1 and 2 in slot j , agent i may want $x_{i1j} = k_{i2j} x_{i2j}$ ⁵. Specifically, an agent may want one computational unit with three storage units in the same slot.

Third, agents may have overall job-must-be-finished requirements. That is, an agent may have a minimum requirement on a type of good over all slots. This is reflected as,

Requirement 2. $\sum_{j=1}^J x_{igj} \geq r_{ig}$, $\forall i \in \{1, \dots, I\}$, $\forall g \in \{1, \dots, G\}$.

That is, for agent i , the sum of type-of-good g (e.g., computation units) over the time slots should be greater than its requirement for good g . For example, a requirement for agent i may be $\sum_{j=1}^J x_{i1j} \geq r_{i1}$. An agent may not participate in the exchange economy if its job is not guaranteed to finish. By deciding not to participate it may break the equilibrium.

Given Requirements 1 and 2, and using $k_{igj} = k_{ig}$, consumer i 's problem may be written as

$$\max_{\substack{x_i \in X_i; p \cdot x_i \leq p \cdot w_i \\ \sum_{j=1}^J x_{igj} \geq r_{ig}, \forall g \in \{1, \dots, G\} \\ x_{igj} \geq k_{ig} x_{ilj}, \forall g \in \{1, \dots, G\}, \forall j \in \{1, \dots, J\}}} u_i(x_{i1}), \tag{1}$$

where $k_{ig} \in \mathbb{R}_+$. Note that, given the agent may be able to improve its utility via inter-timeslot substitution, this property is embedded in the agent utility function $u_i(x_i)$. Note that, if the inequality $x_{igj} \geq k_{ig} x_{ilj}$ were an equality, consumer i 's problem (1) can be modified to

$$\max_{\substack{x_i \in X_i; p \cdot x_i \leq p \cdot w_i \\ \sum_{j=1}^J x_{i1j} \geq r_{i1}}} u_i(x_{i1}), \tag{2}$$

assuming that the Requirements 2 are related through k_{ig} and are automatically satisfied. That is,

$$\sum_{j=1}^J x_{i1j} \geq r_{i1} \Rightarrow \sum_{j=1}^J x_{igj} \geq r_{ig}, \quad g \in \{2, \dots, G\}.$$

At this point we have captured many of the important idiosyncrasies of the grid computing environment, such as inter-slot substitution among same type of good, intra-slot complementarities among different types of goods, and overall job-must-be-finished requirements. Many more characteristics can be added to this mix. These include intra-slot substitution among different types of goods, inter-slot complementarities, per-slot job-must-be-finished requirements, and technical requirements such as interoperability, security, etc. Given the intricacies of grid computing, we note that at this point our aim is not to be comprehensive in our representation but to be

⁵This rules out results ensuring uniqueness of equilibrium [23,F] since these results rely on the goods being gross substitutes.

realistic. This allows us to derive useful results without overly complicated models. We come back to these additional characteristics in Section 4.1, and show that they do not change the basic nature of our results.

In particular, we define three equilibrium concepts that work under different assumptions on how the markets function. These show that we move from centralized trade to fully decentralized trade by using increasingly strict requirements on users' preferences. In this process, we start with the most general form of price quasiequilibrium with transfers and end up with the well known form of Walrasian equilibrium. In what follows we not only discuss what these equilibria are but also point out as to why managers of a firm's computing resources should be interested in them.

We start with the price quasiequilibrium with transfers that requires the agents' utilities to be convex and locally-nonsatiated. Note that convexity in preferences implies that an agent prefers averages to extremes—it is a generalization of the assumption of “diminishing marginal rates of substitution, whereas local non-satiation implies that one can always do a little bit better, even if one is restricted to only small changes in the consumption bundle.”⁶ These requirements are the least constraining and allow for the broadest set of user utilities to be considered. However, this equilibrium assumes the existence of a central planner that would re-distribute wealth through lumpsum transfers that add up to implicit budget balances. There are two potential issues with this concept. First, it assumes the presence of a central planner (e.g., an auctioneer), which may not be present under all situations—in particular, in a completely distributed setting. Second, it does not specify a functional form of how these lumpsum transfers are calculated, which may pose distributional issues in a market.

We then look at the concept of Walrasian quasiequilibrium that extends the concept of price quasiequilibrium with transfers to specify a functional form of how these lumpsum transfers are calculated. These lumpsum transfers are defined as implicit budget balances that correspond directly to agents' expectations. However, this concept still requires the presence of a central planner. That is, this concept may be applied to some market

⁶Mathematically, convexity is defined as follows: given x, y , and z in X such that $x \succeq z$ and $y \succeq z$, then $tx + (1-t)y \succeq z$ for all $0 \leq t \leq 1$. For strong convexity, the definition changes to: given $x \neq y$, and z in X such that $x \succeq z$ and $y \succeq z$, then $tx + (1-t)y \succeq z$ for all $0 \leq t \leq 1$.

Mathematically, non-satiation is defined as follows: given any x in X and any $\varepsilon > 0$, there is some y in X with $|x - y| < \varepsilon$ such that $y \succ x$.

Note that when we write x is preferred (using the relation \succeq) over y , we mean “the consumer thinks that the bundle x is at least as good as the bundle y ”.

settings, such as auctions but not to others, such as commodity markets. This extension is achieved by constraining agents' utilities and requires them to be continuous, in addition to the existing requirements of convexity and local-nonsatiation.

Finally, we look at the concept of Walrasian equilibrium that further relaxes the remaining requirement of the presence of a central planner, and allows for fully distributed trade. That is, it applies to all different kinds of market forms. This concept maintains the continuity requirement of Walrasian quasiequilibrium. However, it requires the agents' utilities to be strictly convex and strictly monotone instead of convexity and local-nonsatiation. Note that monotonicity implies that “at least as much of everything is at least as good”, whereas strong monotonicity implies that at least as much of every good, and strictly more of some good, is strictly better⁷.

We now capture the necessary and sufficient conditions for the exchange economy to exist.

3.2. Necessary Conditions

We start with the necessary conditions. These are grid computing specific, and do not exist in regular exchange economies. These conditions, if not met, will render the question of the existence of exchange economies moot. From the perspective of a manager of computing resources this is the first test of whether trading of such resources would be possible in a grid computing market place.

The following necessary condition deals with the job-must-be-finished requirement (Requirement 2), as follows.

Lemma 1. *An equilibrium exists only if there is enough supply to satisfy demand over the whole duration T . That is,*

$$\sum_{i=1}^I \sum_{j=1}^J w_{ij} \geq \sum_{i=1}^I r_{ig}, \forall g \in \{1, \dots, G\}. \quad (3)$$

Proof: If the above condition is not satisfied then there exists a $\tilde{g} \in \{1, \dots, G\}$ such that

$$\sum_{i=1}^I \sum_{j=1}^J w_{i\tilde{g}j} < \sum_{i=1}^I r_{i\tilde{g}}. \quad (4)$$

However, if all agents have Requirement 2, then

$$\sum_{j=1}^J x_{i\tilde{g}j} \geq r_{i\tilde{g}}, \forall i \in \{1, \dots, I\},$$

which upon summing over i gives

$$\sum_{i=1}^I \sum_{j=1}^J w_{i\tilde{g}j} \geq \sum_{i=1}^I r_{i\tilde{g}},$$

⁷Mathematically, monotonicity is defined as follows: if $x \succeq y$ then $x \succeq y$. For strong monotonicity, the definition changes to: if $x \succeq y$ and $x \neq y$, then $x \succ y$.

a contradiction to (9). \square

This intuitive condition essentially says that there is enough aggregate supply over all slots to satisfy the aggregate demand over all slots. An illustrative example is as follows.

Example 2. Say $I = 2, J = 2, G = 1; w_{11*} = \{2, 6\}, w_{21*} = \{1, 6\};$ and $r_{11} = 9, r_{21} = 8.$ Then, given the total supply of $15(= w_{111} + w_{112} + w_{211} + w_{212}),$ there is no way to satisfy the total demand of $17(= r_{11} + r_{21}),$ and no exchange economy can exist. However, it is possible to satisfy the total demand if $w_{21*} = \{3, 6\}.$ Then exchange economies can exist—for example, a possible solution is $x_{11*} = \{5, 4\}$ as well as $x_{21*} = \{0, 8\}.$

3.3. Sufficient Conditions

We now present sufficient conditions for existence of various market equilibria. These results essential show that grid constraints don't matter as long as agents' utilities are convex. The only requirement is that the intuitive necessary condition (3), which is based on Requirement 2, is satisfied. We note that the nature of these results, and hence the proofs, do not change if the overall job-must-be-finished condition (Requirement 2) is replaced by the corresponding per-slot condition (6). The only difference is that the necessary condition, Lemma 1, is replaced by Lemma 2. Formally stated, in the form introduced in [23, 17.B.1], we have a private ownership economy specified by⁸

$$\left(\left\{ (X_i, \pm_i) \right\}_{i=1}^I, Y_1, \{W_i\}_{i=1}^I \right),$$

where $x_i \in X_i = \mathbb{R}_+^L$ and $w_i \in W_i = \mathbb{R}_+^L$ are the final and initial allocations, respectively. $y_1 \in Y_1 = -\mathbb{R}_+^L$ is the variable representing free disposal (*i.e.*, a quantity that can be disposed off for free, assuming that is possible). Finally $u_i(x_i)$ is the utility function associated with the underlying preference relation $\succeq_i.$

We first define a *price quasiequilibrium with transfers* (see Definition [23, 16.D.1]): this is the most general form, with the least strict requirements of convexity and local nonsatiation on the agents' utilities. That is, it makes no assumption about continuity and monotonicity of preferences.

Definition 1. In the economy defined above, an allocation (x^*, y_1^*) and a price vector $p = (p_1, \dots, p_L) \neq 0$ constitute a *price quasiequilibrium with transfers* if there is an assignment of wealth levels (tw_1, \dots, tw_I) with $\sum_i tw_i = p \cdot (\sum_i w_i + y_1^*)$ such that

- $p \cdot y_1 \leq p \cdot y_1^*.$

⁸Under a private ownership economy, individuals, rather than government, own and operate their self owned businesses or property. They make the majority of, or all decisions for their business with little or no government regulation or interference.

- For every $i,$ if $x_i \succ_i x_i^*$ then $p \cdot x_i \geq tw_i.$
- $\sum_{i=1}^I x_i^* = \sum_{i=1}^I w_i + y_1^*,$

where \succ_i is the strict preference relation. Then we state the following result that follows Proposition [23, 16.D.1].

Proposition 1. Consider an economy specified by

$\left(\left\{ (X_i, \succeq_i) \right\}_{i=1}^I, Y_1, \{w_i\}_{i=1}^I \right),$ and suppose that every preference relation \succeq_i is convex and locally nonsatiated. Then, given Requirements 1 and 2, if (3) is satisfied, a price quasiequilibrium with transfers exists.

Proof: This proof is along the same lines as the one for Proposition [23, 16.D.1] (a version of the second fundamental welfare theorem: any Pareto optimal allocation (x^*, y_1^*) is a price quasiequilibrium with transfers), which essentially involves showing the existence of a price hyperplane for convex and nonsatiated utilities. We don't reproduce it here, except to highlight what is different.

In this case, Requirements 1 and 2 can be included in agents' optimization problems as linear constraints, as in (1). If u_i are convex functions, then consumer i 's problem is still convex since the feasible set is the intersection of convex sets [30] (recall that linear constraints—either inequalities or equalities—are allowed). This means that the analysis in Proposition [23, 16.D.1] will still apply: we just have to guarantee that Pareto solutions are feasible. This, in turn, relies upon showing that the collective feasible set (*i.e.*, intersection of individual feasible sets, which are derived from respective constraints) of outcomes is not empty—note that this is also a necessary condition for the existence of an equilibrium.

But, this is straightforward, since the non-emptiness of individual feasible sets ensures the non-emptiness of the collective feasible set. Recall that, given any price vector $p,$ agent i 's choice is independent of agent j 's choice [23]. The independence of agents' choices is a characteristic of decentralized economies where agents do not have market power. Given a market clearing price " p " the agents optimize for their utilities independently of each other and hence their choices are independent of each other. That is, given X_i as agent i 's feasible set in own variables, i 's overall feasible set $(\subset \mathbb{R}_+^L)$ is $\mathbb{R}^L \times \dots \times X_i \times \dots \times \mathbb{R}^L.$ This implies that, if we focus on agent i 's variables, agent j 's feasible set is $\mathbb{R}_+^L,$ which when intersected with agent i 's feasible set gives back agent i 's feasible set. Thus the intersection of overall feasible sets gives the union of the feasible sets in respective variables, each of which is non-empty. Note that it is possible that the feasible region contains only the initial endowments. This is acceptable since a single point is a (closed) convex set.

Finally, any market solution not only satisfies all the linear constraints for all the agents but also is Pareto su-

terior to initial endowments. \square

A price quasiequilibrium with transfers assumes the existence of a central planner that would assign wealth transfers among agents after trade has taken place. Further, it does not specify how the wealth transfers are distributed among the agents. The absence of this mechanism may not be acceptable in practical settings, where rational agents (or firm managers) would like to understand how these wealth transfers are related to their initial endowments as well as prices (implicit or explicit) of goods. Thus, we would like a distributed solution that does not require the presence of the central planner—a solution that endogenizes the wealth transfer. We would also like the solution to specify the wealth transfer in a form that is acceptable to rational agents.

A partial solution, Walrasian (or competitive) quasiequilibrium, which solves the second issue, is a special case of the price quasiequilibrium with transfers, with $tw_i = p \cdot (w_i + \theta_{i1}y_1^*)$, where θ_{i1} is the share of agent i of free disposal—note that the transfers are now basically the implicit budget balances of the agents. Compared to the price quasiequilibrium with transfers, this concept places the additional restriction of continuity (in addition to convexity and local nonsatiation) on agents' preferences. That is, it requires the agents' preferences to satisfy the properties of continuity, convexity, and nonsatiation.

We now state the following result that guarantees the existence of a Walrasian quasiequilibrium (see Proposition [23, 17.BB.2])⁹:

Proposition 2. *Suppose that for an economy with $I > 0$ consumers, we have*

- For every i
 - $X_i \subset \mathbb{R}^L$ is closed and convex;
 - \succsim_i is a rational, continuous, locally nonsatiated, and convex preference relation defined on X_i ;
 - $w_i \geq \hat{x}_i$ for some $\hat{x}_i \in X_i$.
- The set of feasible allocations

$$A = \left\{ (x, y) \in \mathbb{R}^{LI} \times \mathbb{R}^{L1} : x_i \in X_i \right. \\ \left. \text{for all } i, y_1 \in Y_1, \text{ and } \sum_i x_i \leq \sum_i w_i + y_1 \right\}$$

is compact.

Then, Requirements 1 and 2, if (3) is satisfied, a Walrasian quasiequilibrium exists.

Proof: The proof is along the same lines as the one for Proposition [23, 17.BB.2], which essentially involves showing the existence of an equilibrium price through a fixed-point argument. We do not reproduce it here, except to highlight the differences, and to show that the differences do not alter the end result.

⁹Note that Proposition [23, BB.2] extends Proposition [23, D.1] to a Walrasian quasiequilibrium.

First, since intersection of closed and convex sets remains closed and convex [33], the inclusion of the linear constraints (closed and convex sets) implied by Requirements 1 and 2 creates closed and convex choice sets $\hat{X}_i \subset X_i \subset \mathbb{R}^L$, and the analysis in Proposition [23, 17.BB.2] still applies.

Second, not only Pareto solutions are feasible but also any market solution satisfies all the linear constraints for all the agents and is Pareto superior to initial endowments, as argued in the proof of Proposition 1. \square

Since a Walrasian quasiequilibrium restricts transfers, the conditions for Proposition 2 are stronger than the conditions for Proposition 1. That is, the existence of a Walrasian quasiequilibrium suffices to ensure the existence of a price quasiequilibrium with transfers. However, the reverse is not true—the existence of a price quasiequilibrium with transfers does not automatically imply the existence of a Walrasian quasiequilibrium.

A Walrasian quasiequilibrium still assumes the existence of a central planner that would assign wealth transfers among agents after trade has taken place. However, such a central planner may not be present in many settings. Then we would like a distributed solution that does not require the presence of the central planner—a solution that endogenizes the wealth transfer. That is, we would like the existence of equilibria among independent agents that are restricted by their implicit budget balance based on corresponding initial endowments.

The solution, a Walrasian equilibrium, is the most restrictive form from the perspective of agents' preferences. This is defined as follows (see Definition [23, 17.B.1]).

Definition 2. *In this economy, an allocation (x^*, y_1^*) and a price vector $p = (p_1, \dots, p_L)$ constitute a Walrasian equilibrium if*

- $y_1^* \leq 0, p \cdot y_1^* = 0, \text{ and } p \geq 0$ ¹⁰.
- $x_i^* = x_i(p, p \cdot w_i), \forall i$, where $x_i(\cdot)$ is consumer i 's Walrasian demand function. That is, x_i^* solves

$$\max_{x_i \in X_i : p \cdot x_i \leq p \cdot w_i} u_i(x_i). \tag{5}$$

- $\sum_{i=1}^I x_i^* = \sum_{i=1}^I w_i + y_1^*$.

Note that a Walrasian equilibrium is the same as the Walrasian quasiequilibrium except for the stronger profit maximizing condition (5)¹¹. Also note that Proposition 1 does not guarantee the existence of a Walrasian equilibrium, whereas Proposition 2 would guarantee the existence of a Walrasian equilibrium (Proposition [23, 17.BB.1]) if the Walrasian quasiequilibrium satisfies the

¹⁰From Exercise [23, B.1], this is equivalent to $p \cdot y_1 \leq p \cdot y_1^*$, the first condition in the definition of a price quasiequilibrium with transfers.

¹¹ x_i^* solves $\max_{x_i \in X_i : p \cdot x_i \leq p \cdot w_i} u_i(x_i)$ is the same as: for every i , if $x_i \succ_i x_i^*$ then $p \cdot x_i > p \cdot w_i$. This is stronger than: for every i , if $x_i \succ_i x_i^*$ then $p \cdot x_i > tw_i$.

cheaper consumption condition (Definition [23, 17.BB.2]).

To show the existence of a Walrasian equilibrium, an approach, based on [23, 17.B.1], is as follows. We first define the excess demand functions for the consumers as

$$z_i(p) = x_i(p, p \cdot w_i) - w_i, i \in \{1, \dots, I\}$$

and the excess demand function of the economy as

$$z(p) = \sum_{i=1}^I z_i(p).$$

Then the following conditions guarantee the existence of a Walrasian equilibria by guaranteeing that the system of equations $z(p) = 0$ has a solution (see Propositions [23, 17.B.2] and [23, 17.C.1]). Note that, compared to the Walrasian quasiequilibrium, while it keeps the requirement of continuity, it replaces the requirements of convexity and non-satiation with the stricter requirements of strict convexity and strong monotonicity. That is, it requires the agents' preferences to satisfy the properties of continuity, strict convexity and strong monotonicity.

Proposition 3. *Suppose that, for every consumer i , $X_i = \mathbb{R}_+^L$ and \succeq_i is continuous, strictly convex, and strongly monotone. Also suppose that $\sum_{i=1}^I w_i \gg 0$. Then $z(p)$, defined for all $p \gg 0$ satisfies*

- $z(\cdot)$ is continuous.
- $z(\cdot)$ is homogeneous of degree zero.
- $p \cdot z(p) = 0$ for all p .
- There is an $s > 0$ such that $z_l(p) > -s$ for every commodity l and all p .
- If $p^n \rightarrow p$, where $p \neq 0$ and $p_l = 0$ for some l , then $\max z_1(p^n), \dots, z_L(p^n) \rightarrow \infty$.

Further suppose, given Requirements 1 and 2, (3) is satisfied. Then the system of equations $z(p) = 0$ has a solution—that is, a Walrasian equilibrium exists.

Proof: The proof is along the same lines as the one for Proposition [23, 17.B.2], which essentially involves showing the existence of an equilibrium price through a fixed-point argument. We do not reproduce it here, except to highlight the differences, and to show that the differences do not alter the end result.

First, the only potential issue in the application of Proposition [23, 17.B.2] followed by Proposition [23, 17.C.1] is the continuity of the excess demand function $z(\cdot)$ (it can be trivially verified that all the other properties of $z(\cdot)$ in Proposition [23, 17.B.2] remain unchanged). This continuity is established as follows: the agents' problems remain strictly convex despite the inclusion of the linear constraints implied by Requirements 1 and 2, and the Walrasian demand correspondences $x_i(p, p \cdot w)$ remain single valued, implying continuity of the excess demand function $z(\cdot)$ (Proposition [23, 3.AA.1]).

Second, not only Pareto solutions are feasible but also any market solution satisfies all the linear constraints for all the agents and is Pareto superior to initial endow-

ments, as argued in the proof of Proposition 1. \square

Recall that, since a Walrasian equilibrium uses the strictest conditions, the existence of a Walrasian equilibrium suffices to ensure the existence of a Walrasian quasiequilibrium, and hence a price quasiequilibrium with transfers. One may wonder as to why we have provided three different (though related) results on the existence of a market equilibrium under an exchange economy. It is obvious that these results deal with increasingly strict conditions on users' preferences. That is, the set of utilities satisfying these conditions gets smaller as the conditions get stricter.

3.4. Illustrative Examples

Let us now examine as to how reasonable each of these assumptions are in the grid computing setting. It is reasonable to assume that utilities are continuous. However, under the restrictions posed by the grid economy, utilities may not be strongly monotone and/or strongly convex. That is, it may be more appropriate to work with monotone (i.e., increasing) and convex utilities (e.g., linear functions, as in Examples 3 and 4), and allow the possibility of $p_l = 0$ and $z_l(p) < 0$ for some l . Then, in terms of theoretical results, only Propositions 1 and 2 would be useful, and existence of Walrasian equilibria needs to be demonstrated through examples, as in Examples 3 and 4 that follow.

We next look at some representative examples of an exchange economy in the grid computing setting. These examples demonstrate that the strong assumptions (i.e., continuity, strong convexity, and strong monotonicity) of Proposition 3 are not necessary to demonstrate the existence of a Walrasian equilibrium, and Walrasian equilibrium can exist in presence of continuous, convex, and monotone (i.e., non-decreasing) utilities as well.

The following example shows the existence of an exchange economy in the absence of complementarity requirement (Requirement 1).

Example 3. *Say $I = 2$, $J = 2$, $G = 1$; $w_{1*} = \{2, 6\}$, $w_{2*} = \{3, 6\}$; $r_{11} = 9$, $r_{21} = 8$; and $u_1 = x_{111} + 2x_{112}$, $u_2 = 2x_{211} + x_{212}$. Then the agents' respective problems are given by*

$$\begin{aligned} \max & u_1 \\ \text{s.t.} & \sum_{j=1}^2 p_j x_{1j} \leq \sum_{j=1}^2 p_j w_{1j} \\ & \sum_{j=1}^2 x_{1j} \geq 9 \\ & x_{1j} \geq 0, j \in \{1, 2\} \end{aligned}$$

and

$$\begin{aligned} \max & u_2 \\ \text{s.t.} & \sum_{j=1}^2 p_j x_{2j} \leq \sum_{j=1}^2 p_j w_{2j} \\ & \sum_{j=1}^2 x_{2j} \geq 8 \\ & x_{2j} \geq 0, j \in \{1, 2\} \end{aligned}$$

A possible solution is $x_{1*} = \{0, 9\}$ as well as $x_{2*} = \{5, 3\}$, with supporting prices $p_* = \{1.5, 1\}$, where p_{gj} is the price of good g in slot j .

Any scalar multiple of p_1^* would also give the same solution, as in any exchange economy [23]. This solution is unique in the sense that all equilibrium allocation have $x_{11^*} = \{0, 9\}$, $x_{21^*} = \{5, 3\}$, with $p_{11}/p_{12} = 1.5$. Note that each customer's demand is homogenous of degree zero in the price vector p_1^* ; that is, if prices double, then the consumer's wealth also doubles and his budget set remains unchanged. Thus, from Definition [23, B.1], we see that if p_1^* is a Walrasian equilibrium price vector, then so is αp_1^* . In short, only the relative prices are determined in an equilibrium.

Example 3 can be extended to include complementarities, as demonstrated in the following, more comprehensive example.

Example 4. Say $I = 2$, $J = 2$, $G = 2$; $w_{11^*} = \{2, 6\}$, $w_{12^*} = \{2, 2\}$, $w_{21^*} = \{3, 6\}$, $w_{22^*} = \{2, 3\}$; $r_{11} = 9$, $r_{21} = 8$; $x_{111} \geq 2x_{121}$, $x_{112} \geq 3x_{122}$, $x_{211} \geq 4x_{221}$, $x_{212} \geq 5x_{222}$; and $u_1 = x_{111} + 2x_{112} + x_{121} + x_{122}$, $u_2 = 2x_{211} + x_{212} + x_{221} + x_{222}$. Then the agents' respective problems are given by

$$\max_{\substack{\sum_{g=1}^2 \sum_{j=1}^2 p_{gj} x_{1gj} \leq \sum_{g=1}^2 \sum_{j=1}^2 p_{gj} w_{1gj} \\ \sum_{j=1}^2 x_{11j} \geq 9 \\ x_{1gj} \geq 0, g \in \{1, 2\}, j \in \{1, 2\}}} u_1$$

and

$$\max_{\substack{\sum_{g=1}^2 \sum_{j=1}^2 p_{gj} x_{2gj} \leq \sum_{g=1}^2 \sum_{j=1}^2 p_{gj} w_{2gj} \\ \sum_{j=1}^2 x_{21j} \geq 8 \\ x_{2gj} \geq 0, g \in \{1, 2\}, j \in \{1, 2\}}} u_2$$

A possible solution is $x_{11^*} = \{0, 9\}$, $x_{12^*} = \{0, 3\}$, $x_{21^*} = \{5, 3\}$, $x_{22^*} = \{1.25, 0.6\}$, with supporting prices $p_{1^*} = \{1.5, 1\}$ and $p_{2^*} = \{0, 0\}$.

That is, the prices of the second (complementary) good don't matter and, in addition, the complementarity conditions bind (due to the utility functions' dependence on the second good). These prices may be derived through numerical computations that model price dynamics. Interestingly enough, it does not matter how the ratio $p_{11}/p_{12} = 1.5$ may be reached either from above or below. Note that the upper bounds on the first type of good are provided only through the respective prices, whereas the upper bounds on the second type of good may also be provided by the complementary conditions. The intra time slot complementarity conditions imply that it may suffice for every agent to effectively optimize over only one type of good, and establish conditions for the existence of Walrasian equilibria. All other goods can then be exchanged for free, provided that the necessary conditions for resource availability are met.

In general, given that the utility depends on all goods in an strictly increasing (*i.e.*, strictly monotone) manner, we can easily show that non-zero prices are necessary only on the *primary* variables, which are defined the variables that have no other upper bounds than the en-

dowment driven wealth effects. Strictly positive prices are necessary on the primary variables since these prices result in the creation of upper bounds through the endowment driven wealth effects. If any of these prices were zero, these variables will increase without bound due to the strict monotonicity of the utility. Further, strictly positive prices are not necessary on the *secondary* variables, which are defined as variables that have other upper bounds (either exogenous or through complementarity conditions with the primary variables), as demonstrated by Example 4.

4. Extensions

So far we have looked at a competitive exchange economy with deterministic supply and demand. This representation, though representative, may not be comprehensive. First, there may be many more characteristics of grid computing environments that are not captured in Section 2. Second, agents' demands may not be deterministic but stochastic. Third, we would like to know whether the results apply to regular markets where real money is also a commodity. Fourth, we would like to know if these results apply to finite markets where many of the appealing properties of competitive markets do not hold. In this section, we look at corresponding extensions to our basic model. The first extension examines the inclusion of additional, grid computing related, constraints to our model. The next two extend our results to the case of uncertain supply and demand as well as to regular markets with money. The final extension looks at a non-cooperative basis for our results. In totality these extensions enable a firm's manager to understand the intricacies, nuances and abstractions of the grid market place as they emerge in real world contexts.

4.1. Other Grid Computing Characteristics

As we mentioned in Section 2, our treatment of the characteristics of grid computing environment was realistic enough to develop relevant results. We now examine other relevant characteristics and argue that our basic results hold.

First, the job-must-be-finished requirements in Section 2 may not be comprehensive. That is, in addition to the across-slot job-must-be-finished requirements, an agent may have per-slot requirements

$$x_{igj} \geq r_{igj}, \forall i \in \{1, \dots, I\}, \forall g \in \{1, \dots, G\}, \forall j \in \{1, \dots, J\}. \quad (6)$$

This results in the following necessary condition deals with the stricter, per-slot requirement (6), and essentially says that there is enough supply in every slot to meet the total per slot demand.

Lemma 2. An equilibrium exists only if there is enough supply to satisfy demand in every slot. That is,

$$\sum_{i=1}^I w_{igj} \geq \sum_{i=1}^I r_{igj}, \forall g \in \{1, \dots, G\}, \forall j \in \{1, \dots, J\}.$$

Proof: If the above condition is not satisfied then there exist $\tilde{g} \in \{1, \dots, G\}, \tilde{j} \in \{1, \dots, J\}$ such that

$$\sum_{i=1}^I w_{i\tilde{g}\tilde{j}} < \sum_{i=1}^I r_{i\tilde{g}\tilde{j}}. \tag{7}$$

However, if all agents have Requirement 2, then

$$\sum_{i=1}^I w_{i\tilde{g}\tilde{j}} \geq \sum_{i=1}^I r_{i\tilde{g}\tilde{j}},$$

which contradicts (10). □

An illustrative example is as follows.

Example 5. Say $I = 2, J = 2, G = 1$; $w_{11*} = \{2, 6\}, w_{21*} = \{1, 6\}$; and $r_{11*} = \{4, 4\}, r_{21*} = \{0, 7\}$. Then, there is no way to satisfy $3 = w_{111} + w_{211} \geq x_{111} \geq r_{111} = 4$, and no exchange economy can exist. However, a solution is feasible if $w_{21*} = \{2, 6\}$. In this case, one solution is $x_{11*} = r_{11*} = \{4, 4\}$ as well as $x_{21*} = r_{21*} = \{0, 7\}$.

These per-slot requirements do not change the nature of the problem, and the nature of the proofs for main results (Propositions 1 - 3), except that they create additional necessary conditions (Lemma 2 in Section 3.2). That is, the only difference is that the necessary condition, Lemma 1, is replaced by Lemma 2. Therefore, without loss of generality, we stick to the overall job-must-be-finished conditions for the remainder of this paper.

Further, the job-must-be-finished requirements imply an “all or nothing” approach. However, for certain type of jobs it is possible that there is value from partial fulfillment. Addressing this requires relaxation (or removal) of the job-must-be-finished constraints, which makes the agents’ optimization problems less constrained, and our task of ensuring existence easier. In summary, our formulation allows exploration of this avenue as well.

Second, the substitution constraints in Section 2 may not be comprehensive. For example, in addition to the inter-time slot substitution between same type of good there may be substitution between different types of goods. We have not included this due to the following reasons. To begin with, substitution between different type of goods is possible but not very common. For example, a requirement for CPU capacity may not be easily replaced with a requirement for disk capacity. Next, it is indeed possible to include additional constraints that capture substitution between different types of goods. However, it does not change the nature of the proofs for our main results (*i.e.*, Propositions 1 - 3) —the only difference would be that an agent will just have more linear constraints in its optimization problem (1). Therefore, without loss of generality, we stick to the inter-slot substitution conditions between the same type of goods for

the remainder of this paper.

Third, the complementarity constraints in Requirement 1 may not be comprehensive. For example, these conditions may be across time slots—*i.e.*, $x_{i\tilde{g}\tilde{j}} \geq k_{i\tilde{g}\tilde{j}} x_{i1\tilde{j}}, g \in \{1, \dots, G\}; \tilde{j}, \hat{j} \in \{1, \dots, J\}$. Further, there may be additional complementarity requirements between different (e.g., second and third) type of goods. There are many ways to address this. To start with, complementarity conditions need not be comprehensive. For example, in grid computing, the complementarity conditions are typically tied to CPU cycles, the first type of good. Next, it is possible to introduce additional complementarity conditions not only between all types of goods but also across time slots. This does not change the nature of results established under the original, intra-slot complementarity conditions—the only difference would be that an agent will just have more linear constraints in its optimization problem (1). Finally, we could replace the inequalities by equalities, which implicitly makes the conditions comprehensive. Again, in this case, though the linear constraints are tighter, the nature of the proofs doesn’t change. Therefore, without loss of generality, we stick to the intra-slot complementarity conditions for the remainder of this paper.

Finally, note that grid computing may have many abstractions, including, economic, operational, and technical. The technical and operational requirements can be further broken down as follows. There are infrastructure issues (e.g., identifying and integrating with suitable computer clusters), interoperability issues (e.g., lack of standards, portability of applications, virtualization, etc.), and quality-of-service issues (e.g., service level agreements). Then, there are security (e.g., private and public clouds) as well as communications issues (for example, some programs require more communication than others). We have tried to model the interactions at the economic level while representing some essential requirements from other levels. Our aim is not to be comprehensive but to be realistic is our representation. Further, many of these technical characteristics (e.g., infrastructure issues, interoperability, security, etc.) are necessary from the perspective of final implementation, and can be modeled as combinations of existing types of constraints, without limiting our models in any way.

4.2. Stochastic Supply and Demand

So far we have focused on markets where the supply and demand is known in advance. In realistic situations, there is always some uncertainty about these quantities. For example, an agent may not know for sure how much supply it will actually bring to the market. Similarly, it may not know the actual demand it may actual face. Then, from the perspective of a manager of computing

resources, a natural question is whether our results on existence of market equilibria can be extended to the setting where supply and demand are uncertain.

To answer this, we follow the treatment in Section [23, 19.C]. We first define the “state” of the world as one of the (finite) possible realizations of supply and demand. The state of an agent may be defined as $s_i \in \{1, \dots, S_i\}$, where S_i is the maximum possible number of states for agent i . Then, the state of the world may be defined as $s \in \{1, \dots, S\}$, where $S = S_1 * \dots * S_I$. Now, we can define a (state-)contingent commodity as a combination of the actual commodity $l \in \{1, \dots, L\}$ and state $s \in \{1, \dots, S\}$. The allocation for agent i is now given by $x_i \in \mathbb{R}_+^{LS}$. Similarly, the initial allocation and free disposal are denoted by $w_i \in \mathbb{R}_+^{LS}$ and $y_i \in -\mathbb{R}_+^{LS}$.

The only remaining issue is of defining agents’ preferences. We say that

$$x_i \succ_i x'_i \text{ iff } \sum_s \pi_{si} u_{si}(x_i) \geq \sum_s \pi_{si} u_{si}(x'_i), \quad (8)$$

where π_{si} and u_{si} are state dependent (could be objective or subjective) probability distributions and utility functions, respectively, for agent i . The summation here essentially represents expectations. Then, the following definition describes the *Arrow-Debreu* equilibrium, an extension of the Walrasian equilibrium.

Definition 3. *An allocation*

$$(x_1^*, \dots, x_I^*, -y_1^*) \in X_1 \times \dots \times X_I \times (-Y_1) \subset \mathbb{R}_+^{LS(I+1)}$$

and a system of prices for the contingent commodities $p = (p_{11}, \dots, p_{LS}) \in \mathbb{R}^{LS}$ constitute an *Arrow-Debreu equilibrium* if:

- $p \cdot y_1^* \geq p \cdot y_1$ for all $y_1 \in Y_1$.
- For every i , x_i^* is maximal for \succeq_i in the budget set

$$\{x_i \in X_i : p \cdot x_i \leq p \cdot w_i + \theta_i p \cdot y_1\}.$$

- $\sum_i x_i^* = y_1^* + \sum_i w_i$.

Then, from Section [23, 19.C], Proposition 3 can be extended in a straightforward way to show the existence of an Arrow-Debreu equilibrium. That is, the answer to the question above is yes—our results hold even in the presence of uncertainty. In the present context, the usual convexity assumption takes on an interpretation in terms of risk aversion. For example, in the expected utility setting, the preference relation \succeq_i is convex if the Bernoulli utilities $u_{si}(x_i)$ are concave. Further, the Pareto optimality implication of Arrow-Debreu equilibrium says, effectively, that the possibility of trading in contingent commodities leads, in equilibrium, to an efficient allocation of risk.

4.3. Regular Markets

So far we have focused on exchange markets where

agents trade based on implicit prices making bartering of goods rational. This may apply to certain markets such as P2P computing. However, in many realistic situations, such as cloud computing, agents may want to trade with real money. Thus, from the perspective of a manager of computing resources, it is important to know whether such regular markets would exist. In this section, we look at the case where money is also a commodity—recall that this was the main focus of our paper to begin with. Given our focus on a more general model of exchange economies, this result simply becomes an application of our more general results.

That is, we look at the case where an agent’s allocation is given by $x_i(m_i, x_i)$, where $m_i \in \mathbb{R}_+$ is the money component, and $x_i \in X_i = \mathbb{R}_+^L$ is as defined earlier (see Section 3). The utility of the agent is now given by the well-known quasilinear form

$u_i(x_i) = u_i(m_i, x_i) = m_i + u_i(x_i)$, where $u_i(x_i)$ is as defined earlier. Further, we can define the initial allocation as $w_i = (wm_i, w_i)$, where wm_i is agent i ’s initial allocation of money and $w_i \in W_i = \mathbb{R}_+^L$ is as defined earlier (see Section 3); and free disposal as $y_i = (ym_i, y_i)$, where ym_i is the free disposal of money and y_i is as defined earlier (see Section 3).

Note that, by defining x_i, w_i, y_i in this manner, we can use all the existing machinery developed so far—all we need to do is to replace all the existing terms for the final allocations, initial allocations, and free disposal with these new definitions, and all results in Section 3 (in particular, Proposition 3) carry forward in a straightforward way to the case of regular markets where money is also a commodity.

This extension works since $u_i(\cdot)$ is a linear combination of m_i and $u_i(\cdot)$: if $u_i(\cdot)$ is concave in x_i then $u_i(\cdot)$ is concave in x_i . That is, the crucial convexity assumption on preferences is preserved in regular markets. In fact, it is now easier to show the existence of market equilibria since these equilibria are now partial equilibria (see Chapter [10,23]) which are a special case of general equilibria we have looked at so far.

This result is key from the perspective of previous research on grid computing as regular markets [15,19], given that previous research examined properties of these markets with the implicit assumption that these markets exist under the grid specific constraints outlined in this paper. Our work verifies that the assumption holds true—a key contribution to grid computing markets literature.

4.4. Non Cooperative Basis

So far we have treated agents as price takers. This assumption holds as long as there are infinite number of agents in the market place and an individual agent cannot influence the price of goods through its action. From the perspective of a manager of computing resources, a na-

tural question is: will trade happen in presence of a finite number of rational agents? That is, is there a model that captures the interaction in the presence of finite number of agents—a model that converges to the competitive market in the limit? Given the finite number of agents, it is no longer possible to assume that they are price takers. In fact, in any finite market, rational agents would influence the eventual outcome through their actions.

We now look at one such (possible) non-cooperative basis as presented in [31]. We note that this does not have to be the *only* possible basis. Other representations may still be possible, such as dynamic bargaining and matching games [32]. However, this representation suffices for our purposes in demonstrating the existence of a non-cooperative basis that leads to our results in the limit.

In this model, money (as a medium of exchange) plays an essential role. Therefore, we focus on regular markets instead of pure exchange economies. We use the definitions of Section 4.3, and define the following non-cooperative game [31]. Let us imagine L trading posts, one for each of the L goods. Each agent i makes bids on all the goods by allocating amounts $b_i = \{b_i^j\}_{j=1}^J$ —this is also agent i 's strategy. Then the price of the j^{th} commodity is determined by

$$p_j = \frac{\sum_{i=1}^I b_i^j}{\sum_{i=1}^I w_i^j}.$$

Next, the amount of good j received by agent i is given by

$$x_i^j = \begin{cases} \frac{b_i^j}{p_j} & \text{if } p_j > 0 \\ 0 & \text{if } p_j = 0. \end{cases}$$

Finally, his final amount of money is given by

$$m_i = wm_i - \sum_{j=1}^L b_i^j + \sum_{j=1}^L w_i^j p_j.$$

Under this setup, agents i chooses its bid b_i in response to other agents bids $\{b_j\}_{j \neq i}$ such that its utility $u_i(x_i)$ is maximized under the results allocations $\{x_i\}_{i=1}^I$. Then, we get the following result that is a straightforward extension of Theorem 1 in [31]—again, the key to this result is the crucial convexity assumption on preferences.

Proposition 4. *For each agent $i \in \{1, \dots, I\}$, let u_i be continuous, concave, and non-decreasing. For each good $j \in \{1, \dots, L\}$, let there be at least two agents with*

$wm_i > 0$ whose utility in good j is strictly increasing. Then a non-cooperative equilibrium exists.

This result essentially establishes the existence of a Nash equilibrium for the non-cooperative game modeling a finite market with finite number of agents. This Nash equilibrium can further be shown to converge to a competitive limit (*i.e.*, Walrasian equilibrium) indicated in Proposition 3 as the number of agents goes to infinity (e.g., see Theorem 2 in [31]). That is, the competitive limit of this finite market is the same as the solution of the competitive market. Thus, we have shown the existence of a non-cooperative basis for regular markets for grid computing.

5. Conclusions

In this paper, we have examined the existence of grid computing markets. Due to similarities with other commodity markets, this existence has been implicitly assumed by prior literature. However, grid computing markets, though similar to commodity markets, have various distinguishing characteristics, including substitutions effects, complementarities, and job-must-be-finished requirements. These characteristics, if not properly examined, could result in non-existence of markets, as established by various examples. Thus, our study fills a much needed gap by establishing intuitive conditions under which these markets would exist.

In this process, we examine three different market equilibria that capture different market conditions and hence require different conditions on the agents' utilities. The tradeoff is essentially between the complexity of the markets and the strictness of requirements on utilities. On one end of the spectrum is the price quasiequilibrium with transfers. This is the most restrictive from a market perspective since it requires the presence of a central planner. Further, though it assumes eventual wealth transfers, it does not fully specify how this wealth transfer would be determined. By placing these restrictions it reduces the strictness requirements on agents' utilities in that they are simply required to be convex and non-satiated. On the other end of the spectrum is the Walrasian (or market) equilibrium. It is the least restrictive from a market perspective since it not only does not need a central planner but also it links the eventual wealth transfers to market outcomes. It buys the market simplicity by restricting the agents' utilities to continuous, strictly convex and strongly monotone.

We believe that we have taken the first step in establishing conditions under which agents will participate in regular markets (and exchange economies) for grid resources. These results may also be applicable to other settings, such as cloud computing [6], where distributed computing resources may be traded. We have also in-

cluded various extensions, such as uncertainty in supply and demand, regular competitive markets where money is also a commodity, non-cooperative basis for the competitive markets, etc.

However, we recognize that our work is by no means complete. We envisage the following extensions as part of future work. First, we have looked at a finite horizon, revolving model. This model allows us to study a spot market. A natural extension would be to study an infinite horizon model that describes actual markets in a more realistic manner. It would also be desirable to combine this with stochasticity, and analyze an infinite horizon, continuous model in which supply and demand are stochastic. Second, beyond the marketplace for grid computing resources, a second integral component of the grid computing eco-system is the scheduling and resource management of the grid resources. We would like to perform an integrated analysis of the economics of the grid marketplace and the resource management of the same, and study how the optimal choices in these impact each other.

REFERENCES

- [1] V. Hingne, A. Joshi, T. Finin, H. Kargupta and E. Houstis, "Towards a Pervasive Grid," *Proceedings of 17th International Parallel and Distributed Processing Symposium*, Nice, 22-26 April 2003.
- [2] N. Roy, S. K. Das, K. Basu and M. Kumar, "Enhancing Availability of Grid Computational Services to Ubiquitous Computing Applications," *IEEE Transactions on Parallel & Distributed Systems*, Vol. 20, No. 7, 2009, pp. 953-967. [doi:10.1109/TPDS.2009.15](https://doi.org/10.1109/TPDS.2009.15)
- [3] F. Berman, G. Fox and A. J. G. Hey, "Grid Computing: Making the Global Infrastructure a Reality," John Wiley and Sons, Hoboken, 2003.
- [4] L. F. G. Sarmenta, "Bayanihan: Web-Based Volunteer Computing Using Java," *Lecture Notes in Computer Science* 1368, Springer-Verlag, Berlin, 1998, pp. 444-461.
- [5] R. Chollmeier, "A Definition of Peer-to-Peer Networking for the Classification of Peer-to-Peer Architectures and Applications," *Proceedings of the First International Conference on Peer-to-Peer Computing*, Washington DC, 27-29 August 2001.
- [6] B. Hayes, "Cloud Computing," *Communications of the ACM*, Vol. 51, No. 7, 2008, pp. 9-11. [doi:10.1145/1364782.1364786](https://doi.org/10.1145/1364782.1364786)
- [7] I. Foster, "What Is the Grid? A Three Point Checklist," Unpublished.
- [8] CIO Update, "Finding the Business Case for Grids," CIO Update, Columbus, 2005. <http://www.cioupdate.com/trends/article.php/3553301/Finding-the-Business-Case-for-Grids.htm>
- [9] M. Palankar, A. Lamnitchi, M. Ripeanu and S. Garfinkel, "Amazon S3 for Science Grids: A Viable Solution?" DADC 08, Boston, 2008.
- [10] D. Minoli, "A Networking Approach to Grid Computing," John Wiley, Hoboken, 2005.
- [11] R. Bapna, S. Das, R. Gardfinkel and J. Stallaert, "A Market Design for Grid Computing," *Inform's Journal on Computing*, Vol. 20, No. 1, 2008, pp. 100-111. [doi:10.1287/ijoc.1070.0221](https://doi.org/10.1287/ijoc.1070.0221)
- [12] Insight Research Corporation, "Grid Computing: A Vertical Market Perspective 2005-2010," Insight Research Corporation, Dallas, 2005
- [13] N. Carr, "The End of Corporate Computing," *MIT Sloan Management Review*, Vol. 46, No. 3, 2005, pp. 67-73.
- [14] A. Davies, "Computational Intermediation and the Evolution of Computation as a Commodity," *Applied Economics*, Vol. 36, No. 1, 2004, pp. 1131-1142. [doi:10.1080/0003684042000247334](https://doi.org/10.1080/0003684042000247334)
- [15] R. Buyya, H. Stockinger, J. Giddy and D. Abramson, "Economic Models for Management of Resources in Grid Computing," Unpublished, 2002.
- [16] E. Afgan and P. Bangalore, "Computation Cost in Grid Computing Environments," *29th International Conference on Software Engineering Workshops*, Minneapolis, 19-27 May 2007.
- [17] G. A. Paleologo, "Price-at-Risk: A Methodology for Pricing Utility Computing Services," *IBM Systems Journal*, Vol. 43, No. 1, 2004, pp. 20-31. [doi:10.1147/sj.431.0020](https://doi.org/10.1147/sj.431.0020)
- [18] M. A. Rappa, "The Utility Business Model and the Future of Computing Services," *IBM System Journal*, Vol. 43, No. 1, 2004, pp. 32-42. [doi:10.1147/sj.431.0032](https://doi.org/10.1147/sj.431.0032)
- [19] R. Wolski, J. S. Plank, J. Brevik and T. Bryan, "Analyzing Market-Based Resource Allocation Strategies for the Computational Grid," *The International Journal of High Performance Computing Applications*, Vol. 15, No. 3, 2001, pp. 258-281. [doi:10.1177/109434200101500305](https://doi.org/10.1177/109434200101500305)
- [20] B. Schnizler, D. Neumann, D. Veit and C. Weinhardt, "A Multiattribute Combinatorial Exchange for Trading Grid Resources," *Proceedings of the Research Symposium on Emerging Electronic Markets*, Amsterdam, 2-3 September, 2005.
- [21] D. Neumann, J. Ster, C. Weinhardt and J. Nimis, "A Framework for Commercial Grids Economic and Technical Challenges," *Journal of Grid Computing*, Vol. 6, No. 3, 2008, pp. 325-347. [doi:10.1007/s10723-008-9105-0](https://doi.org/10.1007/s10723-008-9105-0)
- [22] D. J. Veit and W. Gentzsch, "Grid Economics and Business Models," *Journal of Grid Computing*, Vol. 6, 2008, pp. 215-217. [doi:10.1007/s10723-008-9106-z](https://doi.org/10.1007/s10723-008-9106-z)
- [23] A. Mas-Colell, M. D. Whinston and J. R. Green, "Microeconomic Theory," Oxford University Press, Oxford, 1995.
- [24] M. Mossmann, J. Ster, D. Neumann, D. Savourey and R. Krishnashwamy, "A Combinatorial Exchange for Complex Grid Services," Unpublished, 2007.
- [25] O. Regev and N. Nisan, "The Popcorn Market: An Online Market for Computational Resources," ICE, Charleston, 1998.
- [26] C. Aperjis and R. Johari, "A Peer-to-Peer System as an Exchange Economy," *Proceedings of Gamenets*, Pisa, October 2006.

- [27] S. Chapin, J. Karpovich and A. Grimshaw, "The Legion Resource Management System," *Proceedings of the 5th Workshop on Job Scheduling Strategies for Parallel Processing*, San Juan, 16 April 1999.
[doi:10.1007/3-540-47954-6_9](https://doi.org/10.1007/3-540-47954-6_9)
- [28] K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith and S. Tuecke, "A Resource Management Architecture for Metacomputing Systems," *IPPS/SPDP'98 Workshop on Job Scheduling Strategies for Parallel Processing*, Orlando, 30 March 1998, pp. 62-82.
- [29] R. Subrata, A. Y. Zomaya and B. Landfeldt, "Game-Theoretic Approach for Load Balancing in Computational Grids," *IEEE Transactions on Parallel & Distributed Systems*, Vol. 19, No. 1, 2008, pp. 66-76.
[doi:10.1109/TPDS.2007.70710](https://doi.org/10.1109/TPDS.2007.70710)
- [30] S. Boyd and L. Vanderberghe, "Convex Optimization," Cambridge University Press, Cambridge, 2004.
- [31] L. Shapley and M. Shubik, "Trade Using One Commodity as a Means of Payment," *The Journal of Political Economy*, Vol. 85, No. 5, 1997, pp. 937-968.
[doi:10.1086/260616](https://doi.org/10.1086/260616)
- [32] D. Gale, "Strategic Foundations of General Equilibrium," Cambridge University Press, Cambridge, 2000.
[doi:10.1017/CBO9780511492310](https://doi.org/10.1017/CBO9780511492310)
- [33] J. E. Marsden and M. J. Hoffman, "Elementary Classical Analysis," W.H. Freeman and Co., New York, 1993.