



LUND UNIVERSITY

Spot Pricing in the Cloud Ecosystem: A Comparative Investigation

Li, Zheng; Zhang, He; O'Brien, Liam; Jiang, Shu; Zhou, You; Kihl, Maria; Ranjan, Rajiv

Published in:

Journal of Systems and Software

DOI:

[10.1016/j.jss.2015.10.042](https://doi.org/10.1016/j.jss.2015.10.042)

2016

Document Version:

Peer reviewed version (aka post-print)

[Link to publication](#)

Citation for published version (APA):

Li, Z., Zhang, H., O'Brien, L., Jiang, S., Zhou, Y., Kihl, M., & Ranjan, R. (2016). Spot Pricing in the Cloud Ecosystem: A Comparative Investigation. *Journal of Systems and Software*, 114(1), 1-19. <https://doi.org/10.1016/j.jss.2015.10.042>

Total number of authors:

7

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Spot Pricing in the Cloud Ecosystem: A Comparative Investigation

Zheng Li^a, He Zhang^b, Liam O'Brien^c, Shu Jiang^b, You Zhou^b, Maria Kihl^a, Rajiv Ranjan^d

^aDepartment of Electrical and Information Technology, Lund University, Lund, Sweden

^bState Key Laboratory of Novel Software Technology, Software Institute, Nanjing University, Jiangsu, China

^cGeoscience Australia, Canberra, ACT, Australia

^dSchool of Computing Science, Newcastle University, Newcastle, United Kingdom

Abstract

Background: Spot pricing is considered as a significant supplement for building a full-fledged market economy for the Cloud ecosystem. However, it seems that both providers and consumers are still hesitating to enter the Cloud spot market. The relevant academic community also has conflicting opinions about Cloud spot pricing in terms of revenue generation.

Aim: This work aims to systematically identify, assess, synthesize and report the published evidence in favor of or against spot-price scheme compared with fixed-price scheme of Cloud computing, so as to help relieve the aforementioned conflict.

Method: We employed the systematic literature review (SLR) method to collect and investigate the empirical studies of Cloud spot pricing indexed by major electronic libraries.

Results: This SLR identified 61 primary studies that either delivered discussions or conducted experiments to perform comparison between spot pricing and fixed pricing in the Cloud domain. The reported benefits and limitations were summarized to facilitate cost-benefit analysis of being a Cloud spot pricing player, while four types of theories were distinguished to help both researchers and practitioners better understand the Cloud spot market.

Conclusions: This SLR shows that the academic community strongly advocates the emerging Cloud spot market. Although there is still a lack of practical and easily deployed market-driven mechanisms, the overall findings of our work indicate that spot pricing plays a promising role in the sustainability of Cloud resource exploitation.

Keywords: Cloud Computing, Cloud Ecosystem, Cloud Spot Pricing, Comparative Evidence, Systematic Literature Review

1. Introduction

Cloud computing has been increasingly accepted in industry not only for benefiting Cloud providers by creating more business opportunities, but also for relieving Cloud consumers of purchasing, installing, and maintaining local compute resources. To guarantee a successful and sustainable Cloud ecosystem, suitable pricing techniques must be developed and implemented [19]. When it comes to trading Cloud resources, fixed pricing is the dominant strategy in the Cloud market nowadays [XL13], [1].¹ In particular, the most common pricing scheme (namely *pay as you go*) is for on-demand Cloud services, where employing a unit of service is charged a fixed price per unit of time [AKK12]. Given the normally unpredictable and stochastic demand, however, there would always be unused resources in the virtually infinite compute capacity of the Cloud. To help further and better utilize the idle compute resources, a promising approach is to provide spot resources at a reduced price so as to attract more demands with toleration of service delay and interruptions [WQH⁺13]. In fact,

a commercial spot market has been established when a spot instance service was launched by Amazon in December 2009 [SYG13]. Given the de facto spot price traces that are generally far below the on-demand prices, spot pricing is claimed to be the most cost-effective scheme among the existing options for Cloud consumers. More importantly, Amazon's offering of spot service has been regarded as the first step toward a full-fledged market economy for Cloud computing [AKK12].

Unfortunately, there seems to be a lack of confidence in becoming Cloud spot pricing players in industry. Both providers and consumers are still hesitating to enter the Cloud spot market. For instance, the overwhelming majority of the existing Cloud providers have not employed the spot pricing scheme yet [ZG11], and the only currently available provider Amazon is still using contests to encourage more spot applications [3]. The possible reasons for not joining the Cloud spot market could exist behind the limitations of spot pricing. Unlike the static and straightforward pricing schemes of on-demand and reserved Cloud services, the market-driven mechanism for pricing spot service would be complicated for both implementation and understanding. Moreover, since the overall supply and demand of spot resources are both uncertain during runtime, spot service consumers would have to suffer from the irregular fluctuations in service price and availability.

Meanwhile, there are also conflicting opinions in the relevant

Email address: hezhang@nju.edu.cn (He Zhang)

¹We use two types of bibliography formats: the alphabetic format denotes the Cloud service evaluation studies (primary studies) of the SLR, while the numeric format (present in the "References" section) refers to the other references for this article.

academic community. As mentioned previously, offering spot resources has been viewed as an effective approach to attracting more consumers, fully utilizing the Cloud resources, and generating more revenue [WQH⁺13]. Nevertheless, some theoretical analysis and simulation argued that directly using fixed prices would bring higher expected revenues for providers than employing a hybrid (fixed + spot) pricing scheme [AKK12]. It is difficult to give a quick judgment on this even if referring to Amazon's spot service as a concrete example, because the public can obtain little information except the short-term history of spot prices.

To help alleviate the conflict in such a background and help understand whether or not it is reasonable to employ spot pricing for Cloud computing, we conducted a systematic literature review (SLR) [13] in order to rigorously identify, assess, and synthesize empirical evidence in favor of or against Cloud spot pricing. In addition to analyzing the benefits and limitations of spot pricing, we also distinguished between different theories [12] proposed to describe/predict prices or explain/prescribe pricing mechanisms in the Cloud spot market. Furthermore, we particularly investigated the fault-tolerance techniques developed to address the limitations of Cloud spot pricing.

Accordingly, the contributions of this work are mainly threefold. Firstly, the systematically summarized discussions and empirical evidence can help both Cloud providers and consumers gain a quick impression of the pros and cons of the spot pricing scheme. Moreover, this report is further able to act as a checklist to facilitate cost-benefit analysis of offering/employing spot services. Secondly, the four types of relevant theories can help both researchers and practitioners better understand the Cloud spot market. In practice, the price prediction techniques involved in the predictive theories would be particularly valuable for Cloud consumers to make proper bidings, while the various prescriptive theories would be able to inspire Cloud providers to develop/improve their spot pricing mechanisms. In academia, researchers may refer to the existing theories to cross check and review new studies on Cloud spot pricing. Thirdly, the collected fault-tolerance techniques aim to give Cloud consumers an overview about how to achieve trade-offs between economic benefits and service availability.

The remainder of this paper is organized as follows. Section 2 briefly introduces the related work on Cloud spot pricing. Section 3 specifies the SLR method and logistics in our study. Section 4 reports an overview of the reviewed studies, while Section 5 presents our main findings from this SLR by answering the predefined research questions. A set of possible threats to the validity of this study are highlighted in Section 6. Conclusions and some future work are discussed in Section 7.

2. Related work

It has been identified that adequate pricing techniques would play a key role in the success of Cloud computing in practice [19]. In the de facto Cloud market, different providers have employed different strategies to attract consumers and sell their Cloud services. In general, there are three typical pricing

schemes mainly based on Amazon's specification [2], as listed below.

- **On-Demand Service Pricing scheme:** Cloud consumers pay a fixed cost per service unit on an hourly basis without upfront fee and long-term commitment. An analogy of this pricing scheme can be paying per view from a video on demand (VOD) service through the Internet.
- **Reserved Service Pricing scheme:** Cloud consumers pay an upfront fixed fee to ensure discounted hourly pricing for a long-term commitment of service availability (e.g., 1 year, 3 years). An analogy of this pricing scheme can be signing a two-year subscription of mobile service to receive cheaper data plans with a free phone.
- **Spot Service Pricing scheme:** Cloud consumers bid on spare resources and employ them whenever the bid exceeds the current spot price, while the employed service will be interrupted when the spot price exceeds the current bid. An analogy of this pricing scheme can be the dynamic pricing in the electricity distribution industry.

Although the fixed pricing schemes are dominant approaches to trading Cloud resources nowadays [19], spot pricing has been broadly agreed as a significant supplement for building a full-fledged market economy for the Cloud ecosystem [AKK12]. In fact, a wide consensus on efficient management of resources in our society is to avoid fixed pricing [17]. As such, spot pricing (also dynamic pricing) methods have been increasingly developed and adopted in various industries ranging from airlines to electric utilities [7, 10]. Nevertheless, considering the unique characteristics of Cloud computing like location independence, resource virtualization and rapid elasticity, it could be improper to directly confirm the benefits of Cloud spot pricing by analogy with that in other industries.

As a matter of fact, the recent intensive investigations into Cloud spot pricing (e.g., [JTB13, LLLZ13, ZG13]) have delivered diverse and even contradictory statements. In particular, the simulation in the study [AKK12] indicates that spot pricing scheme would generate lower expected revenue for Cloud providers, which conflicts with the aforementioned consensus in most cases. Several survey studies tried to give an overview of, and comparison between different Cloud pricing schemes [1, 18]. Unfortunately, none of them emphasized empirical evidence for the comparison between spot pricing and fixed pricing. Furthermore, these survey studies delivered confusing terminology and classifications to readers. For example, various theoretical and mathematical models were directly treated as different Cloud pricing schemes, whereas those models are essentially for explaining or implementing the spot pricing scheme.

Our work focuses on the empirical evidence of spot pricing in the Cloud industry. In addition to outlining an overview of the benefits and limitations of Cloud spot pricing, we also report a theory-based classification of the relevant studies for researchers' and practitioners' reference.




Fig. 1: Review procedure of this SLR.

3. Review method

The comprehensive guidelines for performing SLR have been specified by Kitchenham and Charters [13]. Here we adapted the guidelines to this work and followed a three-stage review procedure, as illustrated in Fig. 1.

In particular, given the existing experiences of SLR [4, 16], we also emphasized the pilot review during Planning Review stage. In fact, initially reading some relevant studies would be crucial for understanding the domain knowledge and justifying the SLR work, which essentially brought the research questions. Furthermore, during the development of the review protocol, the pilot review can help gradually improve search strategy, refine inclusion/exclusion criteria, and verify data extraction schema by collecting pilot data. The remainder of this section specifies more details we prepared for conducting this SLR.

3.1. Research questions

There are inconsistent and even conflicting claims with regard to the benefits of applying spot pricing in the Cloud market. In order to objectively judge those claims, we decided to summarize and analyze the existing evidence, and the work was driven by the primary research question RQ1.

RQ1: *What empirical evidence has been presented in the literature regarding the benefits and limitations of spot pricing in comparison to fixed pricing in the Cloud market?*

Moreover, as mentioned previously, the de facto mechanism for pricing spot resources in the Cloud market seems still to be a mystery. To help understand the Cloud spot market, it is worth further investigating the backend theories proposed in the relevant studies. In information systems, there always exists multiple types of theories exposing assumptions under different viewpoints, and different types of theories all can be valuable [12]. Therefore, according to the classification of theory in [12], we also distinguish between four types of theories related to the Cloud spot market, as specified in Table 1.

Given the pre-clarified theory types, we further defined the following research question to drive the corresponding investigation:

RQ2: *What theories have been proposed to describe/predict spot prices or explain/prescribe spot pricing in the Cloud market?*

Considering that the limitations of spot pricing also imply research opportunities, there could already exist studies focusing on the potential cures. Thus, we also define a secondary research question of RQ1:

RQ3: *What techniques have been developed to address the limitations of Cloud spot pricing?*

Note that the research questions for this SLR were framed obeying the PICO (Population, Intervention, Comparison, Outcome) criteria [13]. In detail, the population here refers to the participants in the Cloud market; the intervention is the Cloud spot pricing including its backend theories; the comparison intervention is the static strategies for pricing Cloud resources; and the outcomes that are of interest to this systematic review indicate the usefulness of spot pricing and fixed pricing within the Cloud computing domain.

3.2. Research scope

It has been identified that there are unique characteristics of and insights into Cloud computing compared to other computing paradigms [11]. Therefore, although extensive research efforts on pricing in communication networks and Internet can be found in the literature, we only concentrate on the pricing studies in the domain of Cloud computing.

When it comes to the participants in the Cloud market, we are concerned with Cloud providers and consumers without considering their external-party connections like with power suppliers. In particular, Cloud consumers can be Cloud service brokers, secondary service providers, or end users. As for Cloud providers, recall that Amazon is the de facto vendor in the spot market at the time of writing. To make our work closer to the real situation in industry, we only consider the pricing scenario where a single Cloud provider sells spot resources to its consumers. Although a multi-provider spot market will very possibly appear in the future, we are not interested in the evidence for a concern that is not yet a problem. Moreover, different empirical evidence for different pricing scenarios could not be compatible or comparable to each other.

Table 1: Four theory types related to the Cloud spot market.

Theory Type	Specification
Descriptive Theory	This theory provides a description of the spot prices based on observations and/or statistical analyses, which essentially considers the backend spot pricing mechanism as a black box.
Explanatory Theory	This theory provides an explanation/clarification of the spot pricing problem, while the explanation/clarification is usually intended to promote solutions to the spot pricing problem.
Predictive Theory	This theory can simulate/generate spot prices without being aware of the causal demand/supply details, which essentially considers the backend spot pricing mechanism as a black box.
Prescriptive Theory	This theory provides an explicit prescription (e.g., methods, techniques, principles, functions, or a combination of them) for realizing the backend spot pricing mechanism.

3.3. Search strategy

It has been identified that the rigor of the search strategy is crucial for its corresponding systematic reviews [20]. A rigorous strategy can minimize the bias and increase the confidence of repeatable search process and consistent search results. To achieve this, we emphasize the following steps: setting a precise publication time span, employing popular electronic libraries, designing a search string carefully, and using a manual search to supplement the automated search.

3.3.1. Publication time span

As the commercial Cloud spot market was established in December of 2009 [SYG13], we particularly focused on the literature published from the beginning of 2010. Given such a concrete case in industry, the collected evidence would be more valuable and convincing for their practical motivations and validations in the relevant studies. Meanwhile, considering the possible delay of publishing, we only explored studies published before mid 2015. In other words, we restricted the publication time span to being between **January 1st, 2010** and **June 30th, 2015**.

3.3.2. Electronic data sources

According to the statistics of the literature search engines [20], and following the referential experiences reported in the existing SLR protocols and reports, we also employed five popular electronic libraries to achieve a broad enough coverage of relevant primary studies. Note that, as described in Section 3.3.4, the selected studies found by the automated search essentially acted as secondary data sources for the subsequent manual search (i.e. reference snowballing).

- ACM Digital Library (<http://dl.acm.org/>)
- Google Scholar (<http://scholar.google.com>)
- IEEE Xplore (<http://ieeexplore.ieee.org>)
- ScienceDirect (<http://www.sciencedirect.com>)
- SpringerLink (<http://www.springer.com>)

3.3.3. Search string

We mainly followed the approach proposed in [15] to compose the search string. Firstly, we derived major terms from the aforementioned PICO-based research questions, as listed below.

Population: cloud

Intervention: spot pricing




Fig. 2: Manual search steps (adapted from [6]).

Comparison: fixed pricing

Outcomes: benefit, limitation, revenue, cost

Secondly, we identified the alternative spellings and synonyms for these major terms, and also supplemented the terms by checking the keywords in the pilot relevant studies.

Thirdly, we used the Boolean OR and AND to properly integrate all the determined terms into a string for the automated search:

cloud AND (“dynamic pricing” OR “spot pricing” OR “market pricing” OR “dynamic price” OR “spot price” OR “market price” OR “static pricing” OR “fixed pricing” OR “on-demand pricing” OR “static price” OR “fixed price” OR “on-demand price” OR “reserve price”) AND (profit OR profitable OR revenue OR cost OR limit OR limitation OR benefit OR benefiting)

3.3.4. Search process

The search process is roughly an automated search followed by a manual search. The automated search is to run the search string through the search engines of those five popular electronic data sources. After collecting the results from automated search, we unfolded the manual search by implementing three steps, as illustrated in Fig. 2. Note that the manual search here essentially acts as data immersion [6] that facilitates data extraction of this SLR, which also confirms the importance of pilot review (cf. Fig. 1).

- (1) *Quick Scanning:* An initial selection from the automated search results was performed by scanning their titles, keywords and abstracts against the inclusion/exclusion criteria.
- (2) *Thorough Reading:* The full texts of initially selected publications were read entirely to further decide their rele-

vance to this SLR. In particular, the unsure papers were discussed in our team meetings.

- (3) *Reference Snowballing*: To explore the potential relevant studies to the most degree, we reread the Related Work section of each selected study and then supplemented a reference snowballing [14]. New papers identified during reference snowballing were finally determined by repeating the previous steps.

3.4. Inclusion and exclusion criteria

The inclusion and exclusion criteria were mainly shaped from the research scope (cf. Section 3.2) of this SLR. In detail, the criteria are specified as below:

Inclusion Criteria:

- (1) Publications that investigate the de facto spot market of Cloud computing.
- (2) Publications that propose new mechanisms to support pricing/allocating/provisioning Cloud spot resources.
- (3) Publications that perform comparisons between spot pricing and fixed pricing in the Cloud market.
- (4) Publications that develop techniques to address the limitations of Cloud spot pricing.
- (5) Publications that study requesting/bidding/applying spot resources together with discussions about the Cloud spot market/mechanism/prices.
- (6) An overall condition for above inclusion criteria is that the published studies consider a single-provider pricing scenario only.
- (7) Moreover, the above inclusion criteria apply only to regular academic publications (Full journal/conference/workshop papers, technical reports, and book chapters).

Exclusion Criteria:

- (1) Publications that emphasize external connections other than Cloud provider and consumer when investigating Cloud resource pricing. To the best of our knowledge, the external connections are normally with power suppliers.
- (2) Publications that consider the competition among multiple Cloud providers, or allow one consumer to employ spot resources from multiple Cloud providers.
- (3) Publications that do not include any discussion about benefits or limitations of Cloud spot pricing.
- (4) Publications that discuss pricing for non-Cloud computing paradigms.
- (5) Publications that describe Cloud-related studies with respect to the fixed pricing only.
- (6) Publications that investigate Cloud pricing in a generic sense without distinguishing between spot and fixed mechanisms.
- (7) Publications that are previous versions of the later published work.

- (8) Survey papers (secondary studies) that do not contribute the first-hand evidence.
- (9) In addition, short/position papers, demo, posters, extended abstracts or industry publications are all excluded.

3.5. Study quality assessment

Recall that it is usually difficult to distinguish between the methodological rigor and reporting quality of a research study [9]. For the purpose of conciseness, and also by referring to the existing SLR reports, we proposed to use five assessment questions as Part I to examine a relevant study's overall quality.

Part I:

- **QA1:** Is there a rationale for why this study was undertaken?
- **QA2:** Is there a description and justification for distinguishing this study from the related work?
- **QA3:** Is there a description and justification for how this study was conducted (the research design)?
- **QA4:** Is there a clear statement of the findings in this study?
- **QA5:** Does this study present sufficient data/analysis to support/justify the findings?

Considering that this SLR work is about comparative evidence of Cloud spot pricing, we further defined three assessment questions as Part II to measure the strength of the evidence contributed by a primary study. In other words, we can use the score of Part II to reflect how relevant the primary study is to our SLR work.

Part II:

- **QA6:** Is there any description or justification of benefits and limitations of Cloud spot pricing?
- **QA7:** Is there any comparison between spot pricing and fixed pricing in the Cloud market?
- **QA8:** Does this study present sufficient data/analysis to support/justify the comparison?

Answers to these quality assessment questions were assigned numerical scores 0 ('no'), 0.5 ('to some extent'), or 1 ('yes'). In particular, we scored answers to the evidence strength questions 0.5 if the corresponding study only focused on Amazon, and 1 if the study discussed Cloud spot pricing in a broad sense. In addition, we noted the information locations within a paper according to which the assessments were made. As such, we easily double checked the quality scores in our meetings and resolved disagreements.

3.6. Data extraction

To collect useful data to answer the previously defined research questions, we employed a data extraction schema when reviewing the primary studies, as specified in Table 2. We also defined a set of data extraction questions to help clarify the data element to be collected. Note that, the data extraction questions

Table 2: Data extraction schema.

ID	Data Element	Data Extraction Question	Corresponding Research Question
(1)	Reference key	N/A (To help readers quickly find a particular study.)	N/A (Metadata)
(2)	Publication title	What is the title of the publication?	
(3)	Author	What is/are the authors' name(s)?	
(4)	Affiliation	What is/are the authors' affiliation(s)?	
(5)	Publication year	In which year was the evaluation work published?	
(6)	Venue type	What type of the venue does the publication have? (Journal, Conference, Workshop, Book Chapter, or Technical Report)	
(7)	Venue name	Where is the publication's venue? (name of the journal, conference, workshop, or institute)	
(8)	Resource type	What type of Cloud spot resource is considered in this study?	N/A (Context data)
(9)	Demand distribution	What distribution/scenario of consumer demand is considered in this study?	
(10)	Application	What application is considered to consume the Cloud spot resource?	
(11)	Provider benefits	What are the benefits of using spot pricing for Cloud provider?	RQ1
(12)	Provider limitations	What are the limitations of using spot pricing for Cloud provider?	
(13)	Consumer benefits	What are the benefits of using spot pricing for Cloud consumer?	
(14)	Consumer limitations	What are the limitations of using spot pricing for Cloud consumer?	
(15)	Descriptive theory	What descriptive theory about Cloud spot pricing is proposed in this study?	RQ2
(16)	Explanatory theory	What explanatory theory about Cloud spot pricing is proposed in this study?	
(17)	Predictive theory	What predictive theory about Cloud spot pricing is proposed in this study?	
(18)	Prescriptive theory	What prescriptive theory about Cloud spot pricing is proposed in this study?	
(19)	Techniques	What techniques are proposed to address the limitations of Cloud spot pricing?	RQ3
(20)	Addressed limitations	What limitations of Cloud spot pricing is supposed to be addressed in this study?	

are essentially driven by the research questions of this SLR, while the context data extraction aims to help judge the strength of the identified evidence.

The data extraction work was conducted mainly by three research students under the supervision of three researchers. The extracted raw data were kept in a spreadsheet for subsequent analysis.² Moreover, the data elements' locations in their corresponding publications were also marked, so as to help us quickly backtrack and validate these data collected by different reviewers.

3.7. Data synthesis and aggregation

Since the raw data we collected were mostly qualitative descriptions, we employed the approach of thematic analysis [6] to realize data synthesis and aggregation for our SLR. As suggested by Cruzes and Dybå [6], the process of thematic synthesis and aggregation drives different forms of the data with an increasing level of abstraction, as shown in Fig. 3. The four types of data forms are briefly explained below.

- **Text** refers to the raw data with qualitative descriptions.
- **Codes** are descriptive labels that represent different segments of the raw data.
- **Themes** categorize the initial codes into a smaller set of concentrated-meaning units.




Fig. 3: Data evolution during thematic synthesis and aggregation.

- **Model** here can also denote taxonomy or theory that portrays a big picture consisting of higher-order themes and their relationships.

4. Overview of the reviewed studies

Following the search process (cf. Section 3.3.4), we identified 61 relevant primary studies in total. In detail, the automated search initially retrieved 213, 99, 162, and 335 papers from ACM, IEEE Xplore, ScienceDirect, and SpringerLink digital libraries respectively. As for the Google Scholar, we only explored the first 20 pages of the returned results (i.e. 200 records). The manual search through Quick Skimming and Thorough Reading eventually identified 59 publications, while two more valid studies were found in the Reference Snowballing stage, as also shown in Fig. 2.

Through grouping the primary studies according to their publishing time, we show the study distribution over the past four

²The extracted raw data have been shared online: <https://docs.google.com/spreadsheets/d/1vOcuiIHH-eNgIz0r-0b4h3qwkZgTQqV0JhqLaf1e0pw>




Fig. 4: Study distribution over the publication years. (*Note that we only explored primary studies published before mid 2015.)




Fig. 5: Study distribution over the evidence strength.

years in Fig. 4. The rise in the number of publications indicates an increasing research interest in investigating the Cloud spot pricing. Furthermore, in spite of the various conferences related to Cloud computing, more than 20% of the relevant studies were published in highly ranked journals, and the majority of venues (six out of eight) are IEEE/ACM transactions. Thus, it is clear that the academic community has recognized spot pricing as a significantly noteworthy research topic for satisfying the needs of Cloud industry.

According to the quality assessment (cf. detailed scores in Appendix A), the selected studies were overall well conducted and reported. This confirms that our selection criteria have excluded possibly low-quality studies (e.g., short papers without sufficient evaluation). As for the strength of evidence, we further portray the study distribution over different scores in Fig. 5. Given the *Exclusion Criterion* (3), all the relevant studies have discussed the benefits or limitations of Cloud spot pricing. However, not all of the studies performed comparison between spot and fixed pricing in the Cloud market. In particular, the scores also reveal that Amazon has naturally acted as the focus of and a concrete sample in many research works. More than half of the discussions and comparisons directly aimed at the de facto spot market. Only 11 studies further used simulation or statistical analysis to justify their comparisons, and most of them directly employed Amazon’s spot price traces.

5. Review results and discussions

The results and discussions of this SLR are organized following the sequence of answers to the three research questions.




Fig. 6: The benefits of employing spot pricing for Cloud providers. Note that one primary study may have specified more than one benefits.

Corresponding to the data evolution process (cf. Fig. 3), we categorize the benefits and limitations of spot pricing for Cloud providers and consumers respectively in Section 5.1, summarize four different theories for understanding the Cloud spot market in Section 5.2, and list currently developed techniques for addressing the limitations of Cloud spot pricing in Section 5.3.

5.1. Benefits and limitations of spot pricing (RQ1)

5.1.1. Benefits for Cloud providers

There are 34 out of the 61 primary studies recognizing the benefits of spot pricing for Cloud providers. We summarize the identified benefits into seven categories, as shown in Fig. 6.

Driven by the economics rules, dynamic pricing has been considered to be inherently more efficient than static pricing for allocating resources [KS13]. The efficiency of resource usage would be automatically maximized under auction-based provisioning mechanisms [TSK11], because the Cloud resources can be matched to the consumers who have the highest valuation [ZG13]. In other words, the spot pricing scheme can help effectively discover the market value of Cloud resources [SZW⁺14], especially when allocating relatively limited resources to a potentially large number of demands.

In practice, the efficiency of spot pricing for Cloud providers is mainly embodied in resource utilization and profit maximization. When it comes to the resource utilization, a clear consensus is that the spot pricing scheme can improve data centers’ utilization by attracting more consumption, because it allows Cloud providers to sell spare compute capacity at lower prices which would otherwise be idle [AKK12, AKY10, AVCP14, AVCP15, BYBYST13, CVK14, DCL13, GCWZ15, KS13, Liu11, LTH14, MD11, SD15, SMG14, SZW⁺14, TSK11, WQH⁺13, YAK12, ZLH⁺14]. More importantly, since fixed

pricing cannot effectively reflect the underlying trends in the under- and over-demand scenarios [WRM12, WTS⁺13, ZG13], spot pricing has been recognized as being a promising alternative to better cope with unpredictable demands and unbalance problem in the Cloud market [XL13, ZGBX11]. Through demand shifting, spot pricing can smooth out some of the computation requests with monetary incentives and lead to a more efficient use of Cloud infrastructure [KDF⁺11, Liu11, Wee11].

As for the profit maximization, the efficiency may be further distinguished between cost saving and revenue earning. By increasing the usage of spare resources, on the one hand, spot pricing can help reduce the costs associated with idle infrastructure [BYBYST13, KS14, MHL⁺14]; on the other hand, spot pricing can help accommodate various consumers to generate higher revenues [SZW⁺14, XL13, ZG11]. Moreover, it was believed that spot resources essentially provide elasticity to the fixed-price resources without real harm to Cloud providers' main offering [BYBYST13]. Although Abhishek et al. [AKK12] argued that fixed pricing is good enough in terms of revenue generation, more empirical studies have shown that under particular mechanisms the spot pricing outperforms fixed pricing by achieving more revenue for a Cloud provider [KDF⁺11, KLCD14, ZG13, ZLH⁺14], especially when allocating competitive resources to a large number of consumers [WRM12].

Furthermore, spot pricing also has marginal benefits for Cloud providers. For example, the spot-price scheme gives Cloud providers a price-setting power to reclaim their compute capacity when necessary [XL13], which can also be regarded as a freedom to dispatch Cloud resources to dynamic demands [TSK11]. In particular, this control power prevents Cloud users from monopolizing resources through the spot market [MHL⁺14]. In some circumstances, as a marketing strategy, Cloud providers may employ imitative spot pricing to create an impression of constant change in a booming spot market, while intentionally masking the truth of low demand and price inactivity [BYBYST13].

5.1.2. Limitations for Cloud providers

Unlike the aforementioned benefits, the limitations of spot pricing for Cloud providers are considered by 11 studies only, and the concerns can be classified into five categories, as shown in Fig. 7. Compared to the straightforward fixed pricing, the spot pricing's major drawback for Cloud providers is the difficulties and challenges in developing a suitable market-driven mechanism [ZLH⁺14]. Suppose the spot market is formulated as auctions, no matter whether targeting revenue maximization or resource allocation, the corresponding optimization problem would be NP-hard or NP-complete [SYG13, WRM12, ZG13, ZGBX11]. In addition, implementing such a market-driven mechanisms could require more effort and managerial overheads [XL13].

Furthermore, untruthful bidding and mutual cooperation can cause a cyclical effect and eventually make a market-driven strategy turn to be unprofitable for Cloud providers [KS13, AVCP14]. In fact, the de facto spot pricing in the current Cloud market has been considered to be too primitive to guarantee




Fig. 7: The limitations of employing spot pricing for Cloud providers, discussed by 34 out of the 61 studies. Note that one primary study may have considered more than one limitations.

truthful bidding and fair resource allocation [WRM12]. In particular, Abhishek et al. [AKK12] believed that operating a spot market could create price discrimination and generate less expected revenues than using fixed prices.

There is also a special concern about Amazon's policy of not charging customers for the interrupted partial hours, if the spot service usage is terminated by Amazon [3]. In the worst case, the uncharged time rises up to 30% of a spot instance's total run time [DTM12], which would become a burden and reduce the provider's profit [SMG14]. However, since this is Amazon's own policy, we do not treat this concern as a limitation of spot pricing in a generic sense.

5.1.3. Benefits for Cloud consumers

As shown in Fig. 8, the benefits for Cloud consumers were discussed by 53 out of the 61 studies, and most of the discussions concentrated on the cost effectiveness of using spot services. Although the de facto spot service of Amazon may still unexpectedly charge more than the on-demand option [CKLN11], it has been widely accepted that spot resources are on average cheaper (e.g., [JCC⁺11, KS15a, MD11, PRB14, SYG13]). Such a consensus is mainly based on the investigation into the current Cloud spot market. For example, the quantitative analyses of Amazon's price history show that consumers can expect to save more than half the expense if replacing on-demand instances with the spot ones [BYBYST13, JTB13, SYG13]. The empirical studies deliver even more encouraging results: with proper bids, the total cost of employing spot resources can be maintained between 13% and 36% of using the equivalent on-demand resources [LLLZ13, MVB10].

Meanwhile, the spot pricing scheme has been claimed to be able to improve customer satisfaction on overall Cloud performance [XL13, ZLH⁺14]. Firstly, if there is a shortage of on-demand resource, Cloud users can shorten their waiting time by finding alternative resources in the spot market [KLCD14].




Fig. 8: The benefits of employing spot pricing for Cloud consumers, discussed by 53 of the 61 studies. Note that one primary study may have specified more than one benefits.

Secondly, spot instances have been considered as accelerators for small-scale jobs. For example, by employing suitable fault-tolerant technique, the study [JT13] shows that the application turnaround time can be shortened by up to 58% and lowering the monetary cost by up to 59%; by employing additional spot nodes in the MapReduce process, the study [CCS⁺10] shows that the speedup for the overall MapReduce time of some workloads can exceed 200% with an extra monetary cost of 42% [CCS⁺10]. Thirdly, if a spot market implements the mechanism of combinatorial auction, trading different types of Cloud resources would be more efficient and more convenient for the consumers with complex requirements [CVK14].

Moreover, the free market structure behind spot pricing can ensure fair interactions between Cloud providers and consumers [Tai12]. The market-driven dynamic resource provisioning also enables higher social welfare for the entire Cloud ecosystem [SZW⁺14, ZLH⁺14].

5.1.4. Limitations for Cloud consumers

Similar to the benefits for Cloud consumers, the limitations of spot pricing were highlighted also by most of the selected studies (50 out of 61), as distributed in Fig. 9. The most significant limitation consumers have to face is the compromised availability/reliability of Cloud spot services. Given the nature of spot pricing [3], a Cloud provider has the control over terminating spot services through price adjustment. In other words, a spot service can become unavailable at any time without any notice due to the supply and demand fluctuations [AVCP14, AVCP15, CKLN11, DCL13, KM13, VB12, VGB11, WRM12, YAK12, ZGBX11]. In fact, Amazon's spot price history indicates that consumers may still experience out-




Fig. 9: The limitations of employing spot pricing for Cloud consumers, discussed by 50 out of the 61 studies. Note that one primary study may have considered more than one limitations.

of-bid events even if their bid is as high as the regular prices [CKLN11, JTB13, LTH14].

Consequently, it will become hard to analyze the availability level of services or applications built on top of spot instances [GCWZ15], and the completion time of job processing could be seriously delayed [JCC⁺11, CCS⁺10, JLYS14]. From the end user's point of view, the quality of those spot-instance-based services/applications would have been deteriorated due to their uncertain performance, which inevitably leads to negative impacts on the economic advantages of using spot service for two reasons: Firstly, the quality deterioration might incur penalties for failing to meet performance and availability objectives [HSSI15, MD11]. Secondly, if not used carefully, the cheaper spot resources with frequent interruptions would eventually be more expensive than fix-priced, on-demand instances with respect to the overall cost of a Cloud-based application [CKLN11, DCK14, SMG14].

As such, spot pricing has been considered not to be a suitable scheme for workloads with little flexibility [XL13], and spot resources have been particularly reported as inappropriate for supporting long running jobs [CKLN11, KM13].

Additionally, spot pricing seems complicated for Cloud consumers to understand psychologically [XL13], and therefore the inherent complexity could make it difficult to carry out a wise bidding [AKY10, CLT14, JTB13, KS13, KS15b, LTH14, MSJ14, ZG11]. Indeed, Cloud consumers might always face a dilemma when making bidding decisions. On one hand, bidding low prices can result in extremely long execution time without reducing much monetary cost [AKY10]. On the other hand, bidding high prices can lead to a large cost increase without decreasing much computation time [KS15a]. In particular, bidding above on-demand price is claimed to be useless, al-




Fig. 10: Spot service consumption chain in the Cloud ecosystem.

though it might be intuitively attractive [GCWZ15, KS15b].

5.1.5. Discussion

For the convenience of reading, we concisely highlight the collected benefits and limitations of Cloud spot pricing in Table 3.

Through identifying these opinions about Cloud spot pricing, we find a clear consumption chain in the Cloud ecosystem, as shown in Fig. 10. In addition to the investigations from the provider’s and consumer’s perspectives, the existing studies were also concerned with concrete applications, and a service broker (or secondary provider) that acted as a middleware between spot resources and their end consumption.

Although the relevant studies were unfolded from various aspects of the spot service consumption chain, a large degree of consensus has been reached on the benefits and limitations of Cloud spot pricing.

When it comes to the benefits of spot pricing, although different beforehand assumptions may make different studies incomparable, the majority of empirical evidence advocates that spot pricing would be more profitable. A conflicting opinion is that the fixed-price scheme can generate higher expected revenue for Cloud providers [AKK12]. Since suitable mechanisms should be a prerequisite of implementing spot pricing, it is possible that the simulation results could have been flawed by the improper assumptions implementation in [AKK12]. In fact, by analogy, the spot service pricing has been successfully applied to different industrial fields [7].

Similarly, the discussions about limitations of offering spot pricing were also intensive. A typical limitation is that the complexity in backend mechanisms could prevent both Cloud providers and consumers from joining the spot market. Furthermore, Cloud consumers are particularly concerned with unpredictably frequent interruptions when using spot services, although in some optimistic scenarios spot instances could be as reliable as standard instances.

5.2. Theories of Cloud spot pricing (RQ2)

5.2.1. Descriptive Theories

Descriptive theories only focus on the spot prices or the variation in spot prices (cf. Table 1). Given the observations on Amazon’s price history, a major standpoint is that there is no

law behind the spot price variation [MD11, SZL12], and correspondingly the spot prices are totally uncertain [CKLN11] or randomly generated within a particular band [BYBYST13, MVB10]. The spot service reliability varies strongly depending on bid price, instance type, and availability zone [AKY10, BYBYST13, LTH14, RFB⁺15, KS15a].

To simplify their research work, some studies assumed that spot prices followed a mathematical distribution like normal distribution [MD11] or standard Gaussian distribution [MSJ14], and the spot price variations followed a semi-Markovian process [SZL12, GCWZ15]. However, several statistical analyses have revealed that the normal distribution does not fit well the price history [ZPL⁺12], while the Mixture of Gaussians (MoG) distribution delivers better price approximation compared to the other distribution models [JTB13]. In addition to the complex MoG distribution, the lognormal distribution can also adequately model the spot-price distribution [KS15b].

5.2.2. Explanatory theories

When explaining the spot pricing problem, this type of theories proposed by different studies vary significantly.

Following the initial description about Cloud spot market [3], the most straightforward idea is to consider spot prices as results from auctions [AVCP15, CVK14, KS14, KS15b, MHL⁺14, PRB14, SD15]. For example, since different consumers’ bids are unknown to each other and the identical spot resources are sold at an identical price, the Cloud spot market was viewed as a continuous sealed-bid uniform price action model [ZGBX11]; considering that only losers are allowed to repetitively submit new bids while winners have to remain with their positions, Song et al. [SYG13] treated the spot market as a modified version of the repeated single-price auction; suppose the top- N bidders win and the spot price equals to the highest unsuccessful bid ($(N + 1)$ st bid), the spot market can be modeled as a $(N + 1)$ st price auction with multiple goods [BYBYST13, KS15b]; while if taking into account the nature of diverse demands of different numbers and types of Cloud resource, the combinational auction was claimed to be best suited for representing the allocation of Cloud spot resources [WRM12, ZG13, SZW⁺14].

Consequently, game theories have been further used to model the competitions that may happen within auctions. For instance, the Prisoner Dilemma game and the Generalized Nash Equilibrium (GNE) game was employed to formulate the conflicts between a provider and its consumers [DCL13, KS13, SS12]; and particularly, inspired by the attacker-defender scenarios, the situation of a single Cloud provider with multiple consumers was modeled as a Stackelberg game [DCL13].

The existing studies also tried to specify the Cloud spot market from engineering viewpoints. For example, Abhishek et al. modeled the spot market as k parallel $M/M/1$ queues [AKK12], and another work regarded the market as a faulty machine driven by a semi-Markovian process with up and down states [SZL12].

From the perspective of Cloud provider, the spot pricing problem has been often translated into revenue maximization

Table 3: Benefits and limitations of Cloud spot pricing.

Cloud Spot Pricing	Benefits	Limitations
For Providers	<ul style="list-style-type: none"> • Efficient resource utilization <ul style="list-style-type: none"> – Inherently efficient market-driven mechanism – Improving utilization of idle infrastructure – Reflecting trends in changing demand and supply – Balancing between changing demand and supply • Profit Maximization <ul style="list-style-type: none"> – Saving cost associated with idle infrastructure – Earning revenue by attracting more consumers • Control power over unpredictable demand 	<ul style="list-style-type: none"> • Difficulty in developing a suitable mechanism • Leading to untruthful and cooperative bidding • Requiring managerial overhead • Creating price discriminations • Generating lower expected revenue*
For Consumers	<ul style="list-style-type: none"> • High cost effectiveness • Improving customer satisfaction <ul style="list-style-type: none"> – Shortening waiting time for alternative resources – Accelerating small-scale jobs – Efficient & Convenient for complex requirements – improving degree of overall user satisfaction • Higher social welfare for the entire Cloud system • Fair interactions between supplier & consumer 	<ul style="list-style-type: none"> • Compromised service availability/reliability <ul style="list-style-type: none"> – Increasing job completion time – Deteriorating quality of application – Incurring high penalty and/or cost • Difficulty in making bidding wisely • Psychologically complicated pricing logic

*Only one study claims that fixed-price scheme can generate higher expected revenue [AKK12].

problems [WQH⁺13]. In particular, the revenue maximization was further formulated as a finite-horizon stochastic dynamic program [XL13]; the virtual resource allocation was formed into a multi-dimensional knapsack problem (MKP) with the additional constraints [CVK14]; while Kantere et al. modeled optimal pricing as an optimal control problem with a finite horizon, and the demand curve was modeled by employing second order differential equations with constant parameters [KDF⁺11].

Nevertheless, a reverse engineering study doubted about the market-driven mechanism behind the de facto Cloud spot market [BYBYST13]. Given the nearly identical trend of price changing history of different spot instance types at various regions, it was believed that the spot prices were driven by a dynamic reserve price algorithms rather than depending on the market activities claimed by Amazon. The pricing algorithm upgrade might be rolled out in one data center first, followed by other data centers [KS14].

5.2.3. Predictive theories

By using the empirical data from Amazon’s price history, the regressive process seems to be a common technique to build predictive theories. Through reverse engineering, an auto-regressive algorithm AR(1) was developed to simulate spot price generation, and the generation results showed positive match with Amazon’s price traces [BYBYST13]. Unlike AR(1) that tries to keep a linear relation between service availability and prices, the autocorrelation function (ACF)-based algorithm focused on the relation between price variation and the corresponding time difference, which also delivered encouraging

prediction performance [MD11]. On the contrary, the analysis of spot price predictability claimed that using an auto-regressive moving average model could not yield satisfactory accuracy of price approximation [ZPL⁺12]. To achieve better accuracy, a Multi-Layer-Perceptron (MLP) model with a moving simulation mode was developed for short-term price prediction only [WTS⁺13].

In addition to direct spot price predictions, the predictive theory can also be spot instance failure estimations. For example, by modeling the failure probability of a spot instance based on a semi-Markovian process of spot prices, the spot instance’s failure probability for the next bidding interval can be estimated under a specific bid, which essentially implies the price prediction [GCWZ15].

5.2.4. Prescriptive theories

The prescriptive theories in this case are normally proposed solutions to the previously formulated spot pricing problems.

Corresponding to the auction models, the provisioning and allocation of spot resources were mostly emphasized for proposing pricing mechanisms, such as CA-PROVISION [ZG13], greedy allocation algorithm, collusion-resistant algorithm [CVK14, WRM12], and a multi-round combinatorial auction framework [SZW⁺14]; while Song et al. [SYG13] developed an order-statistic-based online pricing (OSOP) algorithm to determine spot prices without considering consumer behaviors.

With respect to the game theories, equilibrium-related mathematical methods were directly adapted from the economics field to the Cloud case (e.g., solving variational inequality

Table 4: Detailed theories of Cloud spot pricing.

Cloud Spot Pricing	About the Past	About the Future
Black-box Theories	Descriptive theory: <ul style="list-style-type: none"> • Price variations are uncertain • High price variability between zones and between instance types • Price variations follow a semi-Markovian process • Prices follow the normal distribution • Prices follow the standard Gaussian distribution • Prices follow a mixture of Gaussians distribution • Prices follow the lognormal distribution 	Predictive Theory: <ul style="list-style-type: none"> • Linear function of service availability and price • Autocorrelation function of price variation and time difference • Multi-Layer-Perceptron model with a moving simulation mode • Semi-Markovian process-based failure estimation
White-box Theories	Explanatory theory: <ul style="list-style-type: none"> • Auction models <ul style="list-style-type: none"> – Continuous sealed-bid uniform price auction – Repeated single-price auction – $(N + 1)$st price auction with multiple goods – Combinational auction • Game theories <ul style="list-style-type: none"> – Prisoner Dilemma game – Generalized Nash Equilibrium game – Stackelberg game • Mechanical models <ul style="list-style-type: none"> – K parallel $M/M/1$ queues – A faulty machine • Revenue maximization problems <ul style="list-style-type: none"> – Finite horizon stochastic dynamic problem – Optimal control problem – Multi-dimensional knapsack problem (MKP) • Unknown dynamic reserve price algorithm 	Prescriptive theory: <ul style="list-style-type: none"> • Resource provisioning & allocation algorithm (for auction models) <ul style="list-style-type: none"> – CA-PROVISION – Greedy allocation + collusion-resistant algorithm – Order-statistic-based online pricing algorithm – Multi-round combinatorial auction framework • Equilibrium-related methods from Economics (for game theories) • Optimization solutions (for revenue maximization problems) <ul style="list-style-type: none"> – Lyapunov optimization algorithm – Hamilton-Jacobi condition based approach – Mixed-integer nonlinear programming – N-armed bandit ϵ-greedy approach

[DCL13]).

When it comes to solving the revenue maximization problem, Wang et al. [WQH⁺13] adopted a Lyapunov optimization algorithm with the basic idea of minimizing a bound on the drift-plus-penalty term; Xu and Li [XL13] proposed an optimal solution based on a demand model, the Hamilton-Jacobi conditions, and a standard numerical approach; and Kantere et al. [KDF⁺11] employed mixed-integer nonlinear programming to optimize the service pricing; and in one of the experimental scenarios, the provider was assumed to set spot prices using an N -armed bandit ϵ -greedy approach that repeatedly select the greedy pricing action among a set of N actions [AVCP14].

Overall, the sophisticated prescriptive theories here confirm that developing a practical spot pricing mechanism would still be challenging, because both Cloud providers and consumers often prefer simplicity in practice [KS14].

5.2.5. Discussion

Similarly, we further highlight the identified theories of Cloud spot pricing in Table 4. In essence, the four types of theories come from four different angles of the spot market. On the one hand, the descriptive and predictive theories are more consumer-oriented, while the explanatory and prescriptive the-

ories are more provider-oriented. On the other hand, the descriptive and explanatory theories are developed mainly to reveal the knowledge of the existing Cloud spot prices, while the predictive and prescriptive theories are proposed mainly to aim at future spot prices.

When it comes to the research method, two opposite approaches can be distinguished in the relevant studies: the black-box investigations are usually based on the statistical analyses of Amazon’s spot price traces, while the white-box investigations are mostly based on theoretical models of market participant activities. In particular, behind the explanatory theories, we find two popular classes of domain-specific techniques employed from the economics and engineering disciplines respectively. From the perspective of economics, Cloud spot pricing has been treated as various auctions and games to reach some equilibrium (e.g., the Prisoner Dilemma game [DCL13]), as shown in Fig. 11. From the perspective of engineering, researchers have modeled Cloud spot pricing as electronical dynamics with input/output factors (e.g., k parallel $M/M/1$ queues [AKK12]) or mechanical dynamics within a state space (e.g., a faulty machine [SZL12]), as illustrated in Fig. 12.

After all, it is clear that there is no one-size-fits-all theory for revealing the secrets of Cloud spot pricing. As a result, we




Fig. 11: Cloud spot pricing from the perspective of economics.

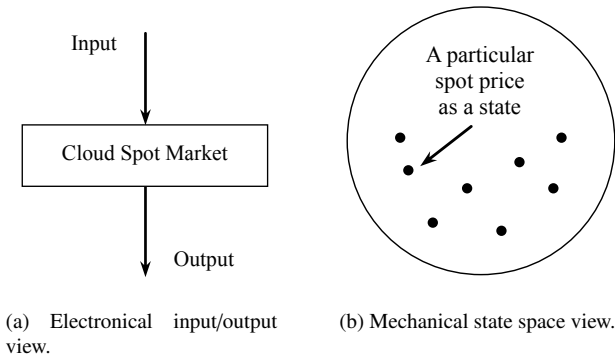


Fig. 12: Cloud spot pricing from the perspective of engineering.

find that different types of theories could supplement each other for understanding the de facto pricing mechanism. For example, as a descriptive theory, the likelihood of employing reserve prices was claimed to be evidence against Amazon’s market-driven mechanism [BYBYST13]; nevertheless, considering the investment cost mentioned by some explanatory theories, setting a reserve price would be reasonable for auctioneers to prevent losses when demand was very low [ZG13]. However, we admit that none of the economic models can explain the similar price fluctuations of Amazon’s different spot resources.

Within the same theory type, we find that it is hard to aggregate the collected data due to the diverse and sometimes contradictory statements. Yet a positive side effect of the contradictory theories is that people can use them to cross assess different primary studies. For example, since the systematic predictability analysis proved that normal distribution was inappropriate for spot price approximation [ZPL⁺12], the assumption of normally distributed prices could have flawed the corresponding work on spot service bidding [MD11].

In particular, given the rational and intelligent participants in the Cloud market, it has been identified that economic models would be more suitable than conventional models in the context of Cloud spot pricing [PR14]. On the other hand, the engineering models could suffer from simplified while improper assumptions. For example, the study [AKK12] assumed that there were always jobs coming up, and its simulation suggested that a well-chosen fixed price would beat spot prices in terms of revenue maximization. In practice, however, there could be jobs attracted by spot prices only, and for these jobs a “well-chosen”

fixed price would not be easily chosen well in advance. Thus, the conclusions in the study [AKK12] might have been flawed by its unsuitable assumptions, which confirms our previous discussions in Section 5.1.5.

5.3. Techniques to address limitations of Cloud spot pricing (RQ3)

Given the consensus on the major drawbacks of spot pricing for Cloud consumers, most of the relevant studies attempted to address the compromised availability of spot services. We have identified six typical fault-tolerance techniques developed from the consumers’ perspective, and two improvement approaches proposed from the provider’s point of view. Although some papers report composite or integrated fault-tolerance solutions (e.g., a combination of four OS-level mechanisms is proposed in [HSSI15]), we focus on the individual techniques only, as briefly described in the following subsections.

5.3.1. Checkpointing

Checkpointing seems to be the most popular fault-tolerance technique to boost overall computing performance and productivity [JLYS14], and sometimes to save the cost [PRB14], when employing spot resources. In brief, the checkpointing technique allows consumers to save their intermediate work when spot service is terminated. A set of practical checkpointing strategies developed by the relevant studies are:

- *Adaptive Scheme*: Checkpoints are taken or skipped at a 10-minute frequency based on the estimation of expected recovery time in case of future service interruptions [YAK12].
- *Application-centric Scheme*: This scheme is based on a sophisticated event generation system that delivers decision to either take checkpoints or terminate spot instances [KM13].
- *Enhanced Adaptive Incremental Scheme*: This is a revised version of the adaptive scheme based on an adjusted Markov model that takes into account both service revocations and hardware failures [JT13].
- *Hourly Scheme*: Checkpoints are taken periodically prior to the beginning of each spot service hour [VB12, YAK12].
- *Rising edge-driven Scheme*: Checkpoints are taken after every increase (rising edge) of spot prices even if the out-of-bid event does not happen [LLLZ13, YAK12].

5.3.2. Duplication / Redundancy

It has been identified that Cloud spot service may only be suitable for short-term jobs due to the possible frequent interruptions [CKLN11]. To increase the chance of satisfying longer jobs’ deadline constraints, the duplication technique creates one replica of each job that could run for more than one hour, and the replicas are supposed to be deployed with different instance type/datacenter combinations [VB12]. Considering that each

zone has a strong dependency on its own price history, the study [MHL⁺14] particularly suggested the across-zone redundancy as a complementary fault-tolerance mechanism.

5.3.3. Lineage-based Recovery

Lineage-based recovery works by re-executing the computation of the failed nodes [DCK14]. Compared to the other fault-tolerance techniques, lineage-based recovery generally incurs little overhead during normal execution. However, it does not work if many spot instances fail at the same time. Therefore, this technique suggests employing different instance types to reduce the risk that all spot instances fail simultaneously.

5.3.4. Migration

The migration technique can be viewed as an improvement of checkpointing. By using the checkpointing technique, the suspended jobs have to wait until the spot service is resumed. To avoid waiting for re-acquiring the same spot resources, the migration technique suggests rebidding at a comparable per-core price for different types of instances even from a different datacenter [JCCY13, YAK12]. If successful, the whole image of a spot instance will be migrated from its original physical host to another host that runs virtual machines with other prices and even with other pricing schemes [CLT14].

5.3.5. Nested Virtualization

Nested virtualization indicates the architecture where a virtual machine runs inside another virtual machine and a job runs inside the nested virtual machine [HSSI15]. In essence, the purpose of nested virtualization is to facilitate the aforementioned migration, by having the complete control over the nested virtual machine without requiring any privileged access to the native one.

5.3.6. Stop-and-Redo Model

Unlike the lineage-based recovery technique, the stop-and-redo model will restart a job from scratch even if only one spot instance fails, which is usually for a traditional database system deployed in spot markets [DCK14]. Thus, the best application strategy driven by the stop-and-redo model is to use as many spot instances as necessary to finish the job in less than one hour.

5.3.7. Service Scaling Down

Recall that the termination of spot instances would not only interrupt the consumers' applications but also reduce the provider's profit (cf. Section 5.1.2). From the provider's perspective, the study [DTM12] suggested scaling down a spot service proportionally to the increase in its price, so as to improve the spot service's availability. In other words, when necessary to free some compute resources, the provider may sacrifice some capacity of its service for not terminating the relevant spot instances.

5.3.8. Hybrid Spot Instance

To relieve the overhead burden of not charging the last partial hour, the study [SMG14] proposed a Hybrid Spot Instance approach that could share the fault tolerance cost between the service provider and consumers, and meanwhile could make the spot service more reliable. The Hybrid Spot Instance approach is essentially based on the aforementioned checkpointing technique, and it allows stretching the user bid till checkpointing is done when an out-of-bid situation occurs.

5.3.9. Discussion

As mentioned previously, unexpected service interruption is the most significant concern for consumers in the spot market. Then, fault-tolerance techniques would be crucial to help increase practitioners' confidence in using Cloud spot services. In fact, although the widely accepted use cases of spot service are small-scale, flexible and delay-tolerant jobs, suitable fault-tolerance techniques have been demonstrated to be able to keep the spot service's availability level as high as using on-demand instances [GCWZ15, OP12], and to make spot instances possible to support highly-reliable always-on online applications [GCWZ15, HSSI15].

However, the existing fault-tolerance techniques are not free [SMG14]. For example, checkpointing would inevitably occupy extra fixed-price resource (e.g., storage), while duplication and migration would require more compute resources. Moreover, a rough trend we have found is: the more sophisticated the approach is, the more cost it may incur. Therefore, the target of service availability should be well balanced, otherwise in the worst case the fault-tolerance cost could overwhelm the economic benefits from employing Cloud spot service.

Ideally, it is also possible to ask the provider like Amazon to modify the spot market and improve its spot service's reliability [MHL⁺14], for example, using the hybrid spot mechanism to allow finalizing a checkpoint before any termination. From the provider's perspective, unfortunately, such market modifications would not be desirable if they lead to fewer consumers entering the usually more profitable on-demand market.

6. Threats to validity

Although we strove to perform the review activities as rigorously and objectively as possible, the findings of this SLR study might still have been affected by certain limitations, as listed below. Readers may need to consider these validity threats when applying the reported findings to their own work.

6.1. Completeness

It is possible that our paper selection does not cover all the relevant studies. First of all, as revealed in Section 5.1.5, different researchers could investigate Cloud spot pricing from different angles with various terms and concepts. As a result, the studies that have irregular descriptions of their empirical investigations might have been missed out. Second, to balance the possible workload with coverage, we searched five most popular electronic libraries instead of looking up every possible literature resource. To alleviate these two issues, we employed a

manual search by snowballing the references to further identify the possibly missed publications.

Last but not the least, we performed automated search through the metadata (namely titles, keywords and abstracts) rather than full text of the papers. However, it has been identified that the search precision would be reduced dramatically by scanning full text [8], and the automated search could inevitably miss relevant studies due to the limitations of search engines [5]. Therefore, we consider this point as a known observation instead of a limitation of this SLR study.

6.2. Reviewers reliability

Recall that our work concentrates on the empirical evidence of spot pricing in the Cloud computing domain only (cf. Section 3.2). When collecting and synthesizing evidence, for this comparative investigation, we tried to distinguish the statements about spot pricing in the Cloud-specific sense from those in the generic sense. Thus, different reviewers involved in this study could have different opinions in some cases.

To reduce the possible bias, we cross-reviewed the collected data, and further discussed unsure issues in our group meetings. Since this study is an international collaboration with team members separated in separated geographical locations, we largely resorted to Skype and telephone for our group meetings. Although the modern communication facilitates are convenient, the different time zones still made it difficult for us to meet on a frequent basis, which might also incur biased review process and results. In fact, compared to our previous practices, we put more effort on organizing meetings and making our discussions more efficient in this study. To the best of our knowledge, few SLR studies have reported their experiences of remote collaboration on conducting SLR. Such a challenge could bring new research opportunities in our EBSE community.

6.3. Evidence strength

When assessing the quality of the relevant studies (cf. Section 3.5), we supposed that their empirical work like simulations and data analyses were valid, and the corresponding statements and findings were trustable. Nevertheless, it is possible that the primary studies could have been flawed by improper assumptions or unsuitable modeling of the Cloud spot market. This could also be the reason why there are contradictory statements about Cloud spot pricing. For example, the conclusion drawn by study [AKK12] about the spot pricing scheme is contrary against the other studies.

The ideal approach to judging the evidence strength would be to replicate and compare those empirical investigations. However, the replication verification is out of the scope of this paper. Thus, we highlight this possible limitation only as a reminder for readers to take into account evidence strength when interpreting the results of this SLR study.

6.4. Industrial Gap

As previously mentioned, there are few Cloud providers using the spot pricing scheme to offer service in industry,

and the only one provider Amazon might have employed artificial algorithm instead of real market-driven mechanism to fluctuate spot prices [BYBYST13]. Such a phenomenon could bring bias in the existing studies and enlarge the gap between academic outcomes and industrial needs. For example, when proposing descriptive and predictive theories (cf. Section 5.2.1 and 5.2.3), nearly all the relevant studies took or had to take Amazon's price trace to perform empirical investigations. As for the explanatory and prescriptive theories (cf. Section 5.2.2 and 5.2.4), there is little evidence to show that those sophisticated pricing system design and assumptions are practically aligned with the real market [KS14].

We have carefully balanced our research scope to reduce the impact of this threat. On one hand, to avoid discussions sticking to Amazon, we tried to distinguish the research focus between on general pricing mechanism and on Amazon's own policies (e.g., Section 5.1.4). On the other hand, to make our work closer to the current situation in industry, we have excluded studies on procurement auction (or reverse auction) that considers competition among Cloud providers (cf. Section 3.2).

In addition, we believe that such an industrial gap also reminds our community of potential research opportunities in this domain. For example, the industry still desires a widely-accepted and easily-implementable pricing mechanism for the Cloud spot market.

7. Conclusions and future work

Appropriate pricing schemes and techniques are crucial for developing and maintaining a successful and sustainable Cloud ecosystem. Through collecting, assessing and analyzing the relevant evidence from 61 primary studies, our systematic literature review (SLR) shows that the academic community strongly advocates the emerging Cloud spot market. Most of the relevant studies reported encouraging discussions about and/or analyses of applying spot pricing to Cloud computing. Although the complexity in backend mechanisms could be a long-term obstacle for both offering and employing Cloud spot services, a set of fault-tolerance techniques have been developed to help reduce some limitations from the consumers' side. In fact, only one study favored fixed pricing from the perspective of Cloud providers according to its queuing theory-based analysis and simulation [AKK12]. However, we have doubts about such work not only because of the opposite majority opinion, but also because its improper assumptions could have weakened the validity of the simulation result.

Meanwhile, this SLR also reveals that there is still a lack of practical market-driven mechanisms to support Cloud spot pricing. A possible dilemma is: economic models would be able to better reflect the Cloud spot market, while the corresponding optimization problems could be NP-hard or NP-complete. As for the de facto spot market with respect to Amazon, although reserve price is reasonable from the economic angle, there is little evidence clearing the doubt about Amazon's spot-price scheme. By synthesizing four types of relevant theories, we confirm the existence of a large gap between the sophisticated pricing models imposed by the existing studies and the real

needs in reality, because both providers and consumers would prefer simple mechanism in practice.

Overall, the findings of this SLR have raised our confidence in the research topic of Cloud spot pricing. We would also expect this report to be able to increase practitioners' confidence in joining the Cloud spot market. Our future work will then be unfolded in two directions. The first, we will focus on the reusability of the extracted data, e.g., developing a factor checklist for cost-benefit analysis of being a spot pricing player. The second, we will put more effort on investigating practical and easily deployed spot mechanisms.

Appendix A. Quality assessment scores of the primary studies

The quality assessment has generated a set of detailed scores for each primary study, as listed in Table A.5.

Appendix B. Selected primary studies

- [AKK12] Vineet Abhishek, Ian A. Kash, and Peter Key. Fixed and market pricing for Cloud services. In *Proceedings of the 7th Workshop on the Economics of Networks, Systems, and Computation (NetEcon 2012)*, pages 157–162, Orlando, FL, USA, 20 March 2012. IEEE Computer Society.
- [AKY10] Artur Andrzejak, Derrick Kondo, and Sangho Yi. Decision model for Cloud computing under SLA constraints. In *Proceedings of the 18th Annual IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2010)*, pages 257–266, Miami, FL, USA, 17–19 August 2010. IEEE Computer Society.
- [AVCP14] Marco Abundo, Valerio Di Valerio, Valeria Cardellini, and Francesco Lo Presti. Bidding strategies in QoS-aware Cloud systems based on N-Armed Bandit problems. In *Proceedings of the 3rd IEEE Symposium on Network Cloud Computing and Applications (NCCA 2014)*, pages 38–45, Rome, Italy, 5–7 February 2014. IEEE Computer Society.
- [AVCP15] Marco Abundo, Valerio Di Valerio, Valeria Cardellini, and Francesco Lo Presti. QoS-aware bidding strategies for VM spot instances: A reinforcement learning approach applied to periodic long running jobs. In *Proceedings of the 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM 2015)*, pages 53–61, Ottawa, Ontario, Canada, 11–15 May 2015. International Federation for Information Processing.
- [BSZ⁺15] Carsten Binnig, Abdallah Salama, Erfan Zamanian, Muhammad El-Hindi, Sebastian Feil, and Tobias Ziegler. Spotgres - parallel data analytics on spot instances. In *Proceedings of the 31st IEEE International Conference on Data Engineering Workshops (ICDEW 2015)*, pages 14–21, Seoul, Korea, 13–17 April 2015. IEEE Press.
- [BYBYST13] Orna Agmon Ben-Yehuda, Muli Ben-Yehuda, Assaf Schuster, and Dan Tsafir. Deconstructing Amazon EC2 spot instance pricing. *ACM Transactions on Economics and Computation*, 1(3):1–20, September 2013.
- [CCS⁺10] Navraj Chohan, Claris Castillo, Mike Spreitzer, Malgorzata Steinder, Asser Tantawi, and Chandra Krintz. See spot run: Using spot instances for MapReduce workflows. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing (HotCloud 2010)*, pages 1–7, Boston, MA, USA, 22 June 2010. USENIX Association.
- [CKLN11] Sivadon Chaisiri, Rakpong Kaewpuang, Bu-Sung Lee, and Dusit Niyato. Cost minimization for provisioning virtual servers in Amazon elastic compute Cloud. In *Proceedings of the 19th Annual IEEE International Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2011)*, pages 85–95, Singapore, 25–27 July 2011. IEEE Computer Society.
- [CLT14] Changbing Chen, Bu Sung Lee, and Xueyan Tang. Improving Hadoop monetary efficiency in the Cloud using spot instances. In *Proceedings of the 6th IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2014)*, pages 312–319, Singapore, 15–18 December 2014. IEEE Computer Society.
- [CVK14] Sergei Chichin, Quoc Bao Vo, and Ryszard Kowalczyk. Truthful market-based trading of Cloud services with reservation price. In *Proceedings of the 2014 IEEE International Conference on Services Computing (SCC 2014)*, pages 27–34, Anchorage, AK, 27 June - 2 July 2014. IEEE Computer Society.
- [DCK14] Elkhan Dadashov, Ugur Cetintemel, and Tim Kraska. Putting analytics on the spot: Or how to lower the cost for analytics. *IEEE Internet Computing*, 18(5):70–73, September–October 2014.
- [DCL13] Valerio Di Valerio, Valeria Cardellini, and Francesco Lo Presti. Optimal pricing and service provisioning strategies in Cloud systems: A Stackelberg game approach. In *Proceedings of the 6th IEEE International Conference on Cloud Computing (CLOUD 2013)*, pages 115–122, Santa Clara, CA, USA, 28 June - 3 July 2013. IEEE Computer Society.
- [DTM12] Wesam Dawoud, Ibrahim Takouna, and Christoph Meinel. Reliable approach to sell the spare capacity in the cloud. In *Proceedings of the 3rd International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING 2012)*, pages 229–236, Nice, France, 22–27 July 2012. IARIA.
- [GCWZ15] Weichao Guo, Kang Chen, Yongwei Wu, and Weimin Zheng. Bidding for highly available services with low price in spot instance market. In *Proceedings of the 24th International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC 2015)*, pages 191–202, Portland, Oregon, USA, 15–19 June 2015. ACM Press.
- [HSSI15] Xin He, Prashant Shenoy, Ramesh Sitaraman, and David Irwin. Cutting the cost of hosting online services using Cloud spot markets. In *Proceedings of the 24th International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC 2015)*, pages 207–218, Portland, Oregon, USA, 15–19 June 2015. ACM Press.
- [JCC⁺11] Daeyong Jung, SungHo Chin, KwangSik Chung, HeonChang Yu, and JoonMin Gil. An efficient checkpointing scheme using price history of spot instances in Cloud computing environment. In *Proceedings of the 8th IFIP International Conference on Network and Parallel Computing (NPC 2011)*, pages 185–200, Changsha, China, 21–23 October 2011. Springer-Verlag.
- [JCCY13] Daeyong Jung, SungHo Chin, Kwang Sik Chung, and Heon-Chang Yu. VM migration for fault tolerance in spot instance based Cloud computing. In *Proceedings of the 8th International Conference on Grid and Pervasive Computing (GPC 2013)*, pages 142–151, Seoul, Korea, 9–11 May 2013. Springer-Verlag.
- [JLYS14] Daeyong Jung, JongBeom Lim, Heonchang Yu, and Taeweon Suh. Estimated interval-based checkpointing (EIC) on spot instances in Cloud computing. *Journal of Applied Mathematics*, 2014:217547, May 2014.
- [JT13] Itthichok Jangjaimon and Nian-Feng Tzeng. Effective cost reduction for elastic Clouds under spot instance pricing through adaptive checkpointing. *IEEE Transactions on Computers*, in press, December 2013.
- [JTB13] Bahman Javadi, Ruppa K. Thulasiram, and Rajkumar Buyya. Characterizing spot price dynamics in public Cloud environments. *Future Generation Computer Systems*, 29(4):988–999, June 2013.
- [KDF⁺11] Verena Kantere, Debabrata Dash, Grégory François, Sofia Kyriakopoulou, and Anastasia Ailamaki. Optimal service pricing for a Cloud cache. *IEEE Transactions on Knowledge and Data Engineering*, 23(9):1345–1358, September 2011.
- [KLCD14] Murat Karakus, Zengxiang Li, Wentong Cai, and TA Nguyen Binh Duong. OMTIR: Open market for trading idle Cloud resources. In *Proceedings of the 6th IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2014)*, pages 719–722, Singapore, 15–18 December 2014. IEEE Computer Society.

Table A.5: Detailed score card for the quality assessment of the 61 primary studies.

Paper	Quality of the Study					Sum (1-5)	Strength of the Evidence			Sum (6-8)
	QA1	QA2	QA3	QA4	QA5		QA6	QA7	QA8	
[AKK12]	1	1	1	1	0.5	4.5	1	1	1	3
[AKY10]	1	0.5	1	1	1	4.5	1	0	0	1
[AVCP14]	0.5	1	1	1	1	4.5	1	0.5	0	1.5
[AVCP15]	1	1	0.5	1	1	4.5	1	0	0	1
[BSZ ⁺ 15]	1	0.5	0.5	1	1	4	1	0.5	0	1.5
[BYBYST13]	1	1	1	1	1	5	1	1	0.5	2.5
[CCS ⁺ 10]	1	0	1	1	1	4	0.5	0.5	0	1
[CKLN11]	0.5	1	0.5	1	1	4	1	0.5	0.5	2
[CLT14]	1	1	1	1	1	5	1	1	1	3
[CVK14]	1	1	0.5	1	1	4.5	1	0	0	1
[DCK14]	1	0	0	0	0	1	1	0.5	0	1.5
[DCL13]	1	0.5	1	1	1	4.5	1	1	0	2
[DTM12]	1	1	1	1	1	5	0.5	0.5	0	1
[GCWZ15]	1	1	0.5	1	1	4.5	1	1	1	3
[HSSI15]	1	0.5	0.5	1	1	4	1	0.5	0	1.5
[JCC ⁺ 11]	1	1	1	1	1	5	1	0.5	0	1.5
[JCCY13]	1	1	1	1	1	5	0.5	0.5	0	1
[JLYS14]	1	1	1	1	0.5	4.5	1	0.5	0	1.5
[JT13]	1	1	1	1	1	5	1	0.5	0	1.5
[JTB13]	1	1	1	1	1	5	0.5	0	0	0.5
[KDF ⁺ 11]	0.5	1	0.5	1	1	4	1	1	1	3
[KLCD14]	0.5	0.5	0	1	1	3.5	1	1	1	3
[KM13]	1	1	1	1	1	5	1	0.5	0	1.5
[KS13]	1	0	1	1	1	4	0.5	0	0	0.5
[KS14]	0.5	1	0.5	1	1	4	1	1	1	3
[KS15a]	1	0.5	0.5	1	1	4	1	0	0	1
[KS15b]	1	0	1	1	1	4	1	0.5	0	1.5
[Liu11]	1	0.5	1	1	1	4.5	1	0	0	1
[LLZ13]	1	1	1	1	1	5	0.5	0.5	0.5	1.5
[LTH14]	1	1	0.5	1	1	4.5	1	0.5	0	1.5
[MD11]	1	1	1	1	1	5	1	0.5	0.5	2
[MHL ⁺ 14]	1	1	1	1	1	5	1	0.5	0	1.5
[MSJ14]	1	1	1	1	1	5	1	1	0	2
[MVB10]	1	1	1	1	1	5	0.5	0.5	0.5	1.5
[OP12]	1	1	0.5	1	1	4.5	0.5	0.5	0.5	1.5
[PRB14]	0.5	0	1	1	1	3.5	1	1	1	3
[QWHD12]	1	0.5	0.5	1	1	4	1	1	1	3
[RFB ⁺ 15]	0.5	0.5	1	1	1	4	1	1	1	3
[SD15]	1	0	0.5	1	1	3.5	1	1	0	2
[SLG ⁺ 15]	1	1	1	1	1	5	1	0	0	1
[SMG14]	1	0	0	1	1	3	1	0	0	1
[SS12]	1	0.5	1	1	1	4.5	1	1	0.5	2.5
[SYG13]	1	1	1	1	1	5	0.5	0.5	0	1
[SZL12]	1	1	1	1	1	5	1	1	0	2
[SZW ⁺ 14]	0.5	1	1	0.5	1	4	1	1	0	2
[Tai12]	1	1	1	1	1	5	1	0.5	0	1.5
[TSK11]	1	1	1	1	1	5	1	0	0	1
[TYL12]	1	1	1	1	1	5	0.5	0	0	0.5
[VB12]	1	1	1	1	1	5	0.5	0.5	0	1
[VGB11]	1	1	1	1	1	5	0.5	0.5	0	1

continued on next page

Table A.5 continued from previous page

Paper	Quality of the Study					Sum (1-5)	Strength of the Evidence			Sum (6-8)
	QA1	QA2	QA3	QA4	QA5		QA6	QA7	QA8	
[Wee11]	1	0.5	1	1	1	4.5	0.5	0.5	0.5	1.5
[WQH ⁺ 13]	1	1	0.5	1	1	4.5	0.5	0	0	0.5
[WRM12]	1	1	1	1	1	5	0.5	0.5	0	1
[WTS ⁺ 13]	1	0	1	1	1	4	0.5	0	0	0.5
[XL13]	1	1	0.5	1	1	4.5	1	1	0	2
[YAK12]	1	1	1	1	1	5	1	0.5	0	1.5
[ZG11]	1	1	0.5	1	1	4.5	1	1	0	2
[ZG13]	1	1	1	1	1	5	1	1	1	3
[ZGBX11]	1	1	1	0.5	1	4.5	1	1	0	2
[ZLH ⁺ 14]	1	1	0.5	1	1	4.5	1	1	1	3
[ZPL ⁺ 12]	1	1	1	1	1	5	0.5	0.5	0	1
Total	57	47.5	49.5	59	59	273	52.5	34.5	15	102
Average	0.93	0.78	0.81	0.97	0.97	4.47	0.86	0.57	0.25	1.67

- [KM13] Sunirmal Khatua and Nandini Mukherjee. Application-centric resource provisioning for Amazon EC2 spot instances. In *Proceedings of the 19th International Conference on Parallel Processing (Euro-Par 2013)*, pages 267–278, Aachen, Germany, 26-30 August 2013. Springer-Verlag.
- [KS13] Sowmya Karunakaran and Rangaraja P. Sundarraj. On using prisoner dilemma model to explain bidding decision for computing resources on the Cloud. In *Proceedings of the 13th International Conference on Group Decision and Negotiation (GDN 2013)*, pages 206–215, Stockholm, Sweden, 17-20 June 2013.
- [KS14] Vedsar Kushwaha and Yogesh Simmhan. Cloudy with a spot of opportunity: Analysis of spot-priced VMs for practical job scheduling. In *Proceedings of the 2014 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM 2014)*, pages 33–40, Bangalore, India, 15-17 October 2014. IEEE Press.
- [KS15a] Bogumił Kamiński and Przemysław Szufel. On optimization of simulation execution on Amazon EC2 spot market. *Simulation Modelling Practice and Theory*, in press, 2015.
- [KS15b] Sowmya Karunakaran and Rangaraja P. Sundarraj. Bidding strategies for spot instances in Cloud computing markets. *IEEE Internet Computing*, 19(3):32–40, May-June 2015.
- [Liu11] Huan Liu. Cutting MapReduce cost with spot market. In *Proceedings of the 3rd USENIX Conference on Hot Topics in Cloud Computing (HotCloud 2011)*, pages 1–6, Portland, OR, USA, 14-15 June 2011. USENIX Association.
- [LLLZ13] Luke M. Leslie, Young Choon Lee, Peng Lu, and Albert Y. Zomaya. Exploiting performance and cost diversity in the Cloud. In *Proceedings of the 6th IEEE International Conference Cloud Computing (CLOUD 2013)*, pages 107–114, Santa Clara, CA, USA, 28 June - 3 July 2013. IEEE Computer Society.
- [LTH14] Seung-Hwan Lim, Gautam S. Thakur, and James L. Horey. Analyzing reliability of virtual machine instances with dynamic pricing in the public Cloud. In *Proceedings of the 28th IEEE International Parallel & Distributed Processing Symposium Workshops (IPDPSW 2014)*, pages 885–893, Phoenix, Arizona, USA, 19-23 May 2014. IEEE Computer Society.
- [MD11] Michele Mazzucco and Marlon Dumas. Achieving performance and availability guarantees with spot instances. In *Proceedings of the 13th IEEE International Conference on High Performance Computing and Communications (HPCC 2011)*, pages 296–303, Banff, Canada, 2-4 September 2011. IEEE Computer Society.
- [MHL⁺14] Aniruddha Marathe, Rachel Harris, David Lowenthal, Bronis R. de Supinski, Barry Rountree, and Martin Schulz. Exploiting redundancy for cost-effective, time-constrained execution of HPC applications on Amazon EC2. In *Proceedings of the 23rd International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC 2014)*, pages 279–290, Vancouver, Canada, 23-27 June 2014. ACM Press.
- [MSJ14] Ishai Menache, Ohad Shamir, and Navendu Jain. On-demand, spot, or both: Dynamic resource allocation for executing batch jobs in the Cloud. In *Proceedings of the 11th International Conference on Autonomic Computing (ICAC 2014)*, pages 177–187, Philadelphia, PA, 18-20 June 2014. USENIX Association.
- [MVB10] Michael Mattess, Christian Vecchiola, and Rajkumar Buyya. Managing peak loads by leasing Cloud infrastructure services from a spot market. In *Proceedings of the 12th IEEE International Conference on High Performance Computing and Communications (HPCC 2010)*, pages 180–188, Melbourne, VIC, Australia, 1-3 September 2010. IEEE Computer Society.
- [OP12] Simon Ostermann and Radu Prodan. Impact of variable priced Cloud resources on scientific workflow scheduling. In *Proceedings of the 18th International Conference on Parallel Processing (Euro-Par 2012)*, pages 350–362, Rhodes Island, Greece, 27-31 August 2012. Springer-Verlag.
- [PRB14] Deepak Poola, Kotagiri Ramamohanarao, and Rajkumar Buyya. Fault-tolerant workflow scheduling using spot instances on Clouds. In *Proceedings of the 14th International Conference on Computational Science (ICCS 2014)*, pages 523–533, Cairns, Australia, 10-12 June 2014. Elsevier.
- [QWHD12] Huazheng Qin, Xing Wu, Ji Hou, and Wanchun Dou. Self-adaptive cloud pricing strategies with markov prediction and data mining method. In *Proceedings of the 2012 IEEE International Conference on Cloud and Service Computing (CSC 2012)*, pages 219–226, Shanghai, China, 22-24 November 2012. IEEE Computer Society.
- [RFB⁺15] Maristella Ribas, C. G. Furtado, Giovanni Barroso, Alberto S. Lima, Neuman Souza, and Antônio Moura. Modeling the use of spot instances for cost reduction in cloud computing adoption using a Petri net framework. In *Proceedings of the 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM 2015)*, pages 1428–1433, Ottawa, Ontario, Canada, 11-15 May 2015. International Federation for Information Processing.
- [SD15] Vivek Kumar Singh and Kaushik Dutta. Dynamic price prediction for Amazon spot instances. In *Proceedings of the 48th Hawaii International Conference on System Sciences (HICSS 2015)*, pages 1513–1520, HI, USA, 5-8 January 2015. IEEE Computer Society.
- [SLG⁺15] Prateek Sharma, Stephen Lee, Tian Guo, David Irwin, and Prashant Shenoy. SpotCheck: Designing a derivative IaaS Cloud on the spot market. In *Proceedings of the 10th European Conference on Computer Systems (EuroSys 2015)*, pages 1–15, Bordeaux, France, 21-24 April 2015. ACM Press.

- [SMG14] Naidila Sadashiv, Dilip Kumar S. M., and R. S. Goudar. Hybrid spot instance based resource provisioning strategy in dynamic Cloud environment. In *Proceedings of the 2014 International Conference on High Performance Computing and Applications (ICHPCA 2014)*, pages 1–6, Bhubaneswar, Odisha, India, 22–24 December 2014. IEEE Press.
- [SS12] K. Sowmya and R. P. Sundarraj. Strategic bidding for Cloud resources under dynamic pricing schemes. In *Proceedings of the 2012 International Symposium on Cloud and Services Computing (ISCOS 2012)*, pages 25–30, Mangalore, India, 17–18 December 2012. IEEE Computer Society.
- [SYG13] Kai Song, Yuan Yao, and Leana Golubchik. Exploring the profit-reliability trade-off in Amazon’s spot instance market: A better pricing mechanism. In *Proceedings of the 21st IEEE/ACM International Symposium on Quality of Service (IWQoS 2013)*, pages 1–10, Montreal, Canada, 3–4 June 2013. IEEE Press.
- [SZL12] Yang Song, Murtaza Zafer, and Kang-Won Lee. Optimal bidding in spot instance market. In *Proceedings of the 31st Annual IEEE International Conference on Computer Communications (INFOCOM 2012)*, pages 190–198, Orlando, Florida, USA, 25–30 March 2012. IEEE Press.
- [SZW⁺14] Weijie Shi, Linquan Zhang, Chuan Wu, Zongpeng Li, and Francis C.M. Lau. An online auction framework for dynamic resource provisioning in Cloud computing. In *Proceedings of the 2014 ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2014)*, pages 71–83, Austin, Texas, USA, 16–20 June 2014. ACM Press.
- [Tai12] Moussa Taifi. Banking on decoupling: Budget-driven sustainability for HPC applications on EC2 spot instances. In *Proceedings of the 31st IEEE International Symposium on Reliable Distributed Systems (SRDS 2012)*, pages 442–447, Irvine, CA, USA, 8–11 October 2012. IEEE Computer Society.
- [TSK11] Moussa Taifi, Justin Y. Shi, and Abdallah Khreishah. SpotMPI: A framework for auction-based HPC computing using Amazon spot instances. In *Proceedings of the 11th International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP 2011)*, pages 109–120, Melbourne Australia, 24–26 October 2011. Springer-Verlag.
- [TYL12] ShaoJie Tang, Jing Yuan, and Xiang-Yang Li. Towards optimal bidding strategy for Amazon EC2 Cloud spot instances. In *Proceedings of the 5th IEEE International Conference on Cloud Computing (CLOUD 2012)*, pages 91–98, Honolulu, Hawaii, USA, 24–29 June 2012. IEEE Computer Society.
- [VB12] William Voorsluys and Rajkumar Buyya. Reliable provisioning of spot instances for compute-intensive applications. In *Proceedings of the 26th IEEE International Conference on Advanced Information Networking and Applications (AINA 2012)*, pages 542–549, Fukuoka, Japan, 26–29 March 2012. IEEE Computer Society.
- [VGB11] William Voorsluys, Saurabh Kumar Garg, and Rajkumar Buyya. Provisioning spot market Cloud resources to create cost-effective virtual clusters. In *Proceedings of the 11th International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP 2011)*, pages 395–408, Melbourne Australia, 24–26 October 2011. Springer-Verlag.
- [Wee11] Sewook Wee. Debunking real-time pricing in Cloud computing. In *Proceedings of the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2011)*, pages 585–590, Newport Beach, CA, USA, 23–26 May 2011. IEEE Computer Society.
- [WQH⁺13] Peijian Wang, Yong Qi, Dou Hui, Lei Rao, and Xue Lin. Present or future: Optimal pricing for spot instances. In *Proceedings of the 33rd International Conference on Distributed Computing Systems (ICDCS 2013)*, pages 410–419, Philadelphia, USA, 8–11 July 2013. IEEE Computer Society.
- [WRM12] Qian Wang, Kui Ren, and Xiaoqiao Meng. When Cloud meets eBay: Towards effective pricing for Cloud computing. In *Proceedings of the 31st Annual IEEE International Conference on Computer Communications (INFOCOM 2012)*, pages 936–944, Orlando, FL, USA, 25–30 March 2012. IEEE Press.
- [WTS⁺13] Richard M. Wallace, Volodymyr Turchenko, Mehdi Sheikhalshahi, Iryna Turchenko, Vladyslav Shults, José Luis Vazquez-Poletti, and Lucio Grandinetti. Applications of neural-based spot market prediction for Cloud computing. In *Proceedings of the 7th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS 2013)*, pages 710–716, Berlin, Germany, 12–14 September 2013. IEEE Press.
- [XL13] Hong Xu and Baochun Li. Dynamic Cloud pricing for revenue maximization. *IEEE Transactions on Cloud Computing*, 1(2):158–171, July–December 2013.
- [YAK12] Sangho Yi, Artur Andrzejak, and Derrick Kondo. Monetary cost-aware checkpointing and migration on Amazon Cloud spot instances. *IEEE Transactions on Services Computing*, 5(4):512–524, October–December 2012.
- [ZG11] Sharrukh Zaman and Daniel Grosu. Efficient bidding for virtual machine instances in Clouds. In *Proceedings of the 4th IEEE International Conference on Cloud Computing (CLOUD 2011)*, pages 41–48, Washington, DC, USA, 4–9 July 2011. IEEE Computer Society.
- [ZG13] Sharrukh Zaman and Daniel Grosu. A combinatorial auction-based mechanism for dynamic VM provisioning and allocation in Clouds. *IEEE Transactions on Cloud Computing*, 1(2):129–141, July–December 2013.
- [ZGBX11] Qi Zhang, Eren Gürses, Raouf Boutaba, and Jin Xiao. Dynamic resource allocation for spot markets in Clouds. In *Proceedings of the 11th USENIX Workshop on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services (Hot-ICE 2011)*, pages 1–6, Boston, MA, USA, 29 March 2011. USENIX Association.
- [ZLH⁺14] Yonglong Zhang, Bin Li, Zhiqiu Huang, Jin Wang, Junwu Zhu, and Huanfeng Peng. Strategy-proof auction mechanism with group price for virtual machine allocation in Clouds. In *Proceedings of the 2nd International Conference on Advanced Cloud and Big Data (CBD 2014)*, pages 60–68, Huangshan, China, 20–22 November 2014. IEEE Press.
- [ZPL⁺12] Han Zhao, Miao Pan, Xinxin Liu, Xiaolin Li, and Yuguang Fang. Optimal resource rental planning for elastic applications in Cloud market. In *Proceedings of the 26th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2012)*, pages 808–819, Shanghai, China, 21–25 May 2012. IEEE Computer Society.

Appendix C. Explanation of the typical excluded studies

See Table C.6. To save space, we only show several typical excluded publications. Most excluded papers were discussed in our group meetings. Note that the studies listed here may also be used as a clue for readers to further identify useful information.

Appendix D. Typical excluded primary studies

- [CC14] Se-Hak Chun and Byong-Sam Choi. Service models and pricing schemes for cloud computing. *Cluster Computing*, 17(2):529–535, June 2014.
- [dSdMSNG⁺12] Francisco Airton Pereira da Silva, Paulo Anselmo da Mota Silveira Neto, Vinicius Cardoso Garcia, Rodrigo Elia Assad, and Fernando Antonio Mota Trinta. Accounting models for Cloud computing: A systematic mapping study. In *Proceedings of the 2012 International Conference on Grid Computing and Applications (GCA 2012)*, pages 3–9, Las Vegas, USA, 16–19 July 2012.
- [KKP15] Sowmya Karunakaran, Venkataraghavan Krishnaswamy, and Sundarraj Rangaraja P. Business view of Cloud: Decisions, models and opportunities – a classification and review of research. *Management Research Review*, 38(6):582–604, June 2015.

Table C.6: Explanation of several typical excluded studies

Paper	Brief Explanation	Corresponding Exclusion Criteria
[CC14]	This paper analyzes two types of fixed pricing schemes.	(5)
[dSdMSNG ⁺ 12]	This is a secondary study on accounting models for Cloud computing.	(8)
[KKP15]	This paper performs a survey type of work on Cloud computing from the business view.	(8)
[KM13]	This is a poster presentation without clear specification about validation.	(9)
[MAT10]	This paper emphasizes the imperfect competition among multiple Cloud providers.	(2)
[PR14]	This paper investigates reverse auction that involves multiple providers.	(2)
[PRC13]	This paper focuses on the energy consuming when discussing pricing for Cloud computing.	(1)
[PS12]	This paper studies load balancing of resource requirements in a federated Cloud environment.	(2)
[STT ⁺ 12]	This paper discusses Cloud resource pricing in a generic sense.	(6)
[TWQY12]	This is a short paper without clear specification about validation.	(9)
[VDB11]	This paper investigates the pricing mechanism for Grid computing systems.	(4)
[ZG10]	This is an old-version duplication of the selected study [ZG13].	(7)
[ZG12]	This paper focuses on the mechanism implementation without clearly discussing spot pricing.	(3)
[ZLW ⁺ 14]	This paper considers job pricing instead of Cloud service pricing.	(4)
[ZSS ⁺ 15]	This paper does not specify any benefit or limitation of Cloud spot pricing.	(3)

- [KM13] Sunirmal Khatua and Nandini Mukherjee. A novel check-pointing scheme for Amazon EC2 spot instances. In *Proceedings of the 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2013)*, pages 180–181, Delft, the Netherlands, 13-16 May 2013. IEEE Computer Society.
- [MAT10] Thomas Meinl, Arun Anandasivam, and Michiaki Tatsubori. Enabling Cloud service reservation with derivatives and yield management. In *Proceedings of the 12th IEEE Conference on Commerce and Enterprise Computing (CEC 2010)*, pages 150–155, Shanghai, China, 10-12 November 2010. IEEE Computer Society.
- [PR14] Abhinandan S. Prasad and Shrisha Rao. A mechanism design approach to resource procurement in Cloud computing. *IEEE Transactions on Computers*, 63(1):17–30, January 2014.
- [PRC13] Marco Polverini, Shaolei Ren, and Antonio Cianfrani. Capacity provisioning and pricing for Cloud computing with energy capping. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton 2013)*, pages 413–420, Monticello, IL, USA, 2-4 October 2013. IEEE.
- [PS12] Komal Singh Patel and A. K. Sarje. VM provisioning method to improve the profit and SLA violation of Cloud service providers. In *Proceedings of the 1st IEEE International Conference on Cloud Computing in Emerging Markets (CCEM 2012)*, pages 1–5, Bangalore, India, 11-12 October 2012. IEEE Computer Society.
- [STT⁺12] Bhanu Sharma, Ruppa K. Thulasiram, Parimala Thulasiraman, Saurabh K. Garg, and Rajkumar Buyya. Pricing Cloud compute commodities: A novel financial economic model. In *Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2012)*, pages 451–457, Ottawa, ON, Canada, 13-16 May 2012. IEEE Computer Society.
- [TWQY12] Cheng Tian, Ying Wang, Feng Qi, and Bo Yin. Decision model for provisioning virtual resources in Amazon EC2. In *Proceedings of the 8th International Conference on Network and Service Management (CNSM 2012)*, pages 159–163, Las Vegas, USA, 22-26 October 2012. International Federation for Information Processing.
- [VDB11] Kurt Vanmechelen, Wim Depoorter, and Jan Broeckhove. Combining futures and spot markets: A hybrid market approach to economic Grid resource management. *Journal of Grid Computing*, 9(1):81–94, March 2011.
- [ZG10] Sharrukh Zaman and Daniel Grosu. Combinatorial auction-based allocation of virtual machine instances in Clouds. In *Proceedings of the 2nd IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2010)*, pages 127–134, Indianapolis, USA, 30 November - 3 December 2010. IEEE Computer Society.
- [ZG12] Sharrukh Zaman and Daniel Grosu. An online mechanism for dynamic VM provisioning and allocation in Clouds. In *Proceedings of the 5th IEEE International Conference on Cloud Computing (CLOUD 2012)*, pages 253–260, Honolulu, HI, USA, 24-29 June 2012. IEEE Computer Society.
- [ZLW⁺14] Jian Zhao, Hongxing Li, Chuan Wu, Zongpeng Li, Zhizhong Zhang, and Francis C.M. Lau. Dynamic pricing and profit maximization for the Cloud with geo-distributed data centers. In *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM 2014)*, pages 118–126, Toronto, Canada, 27 April - 2 May 2014. IEEE Press.
- [ZSS⁺15] Ao Zhou, Qibo Sun, Lei Sun, Jinglin Li, and Fangchun Yang. Maximizing the profits of Cloud service providers via dynamic virtual resource renting approach. *EURASIP Journal on Wireless Communications and Networking*, pages 1–12, March 2015.

References

- [1] M. Al-Roomi, S. Al-Ebrahim, S. Buqrais, I. Ahmad, Cloud computing pricing models: A survey, *International Journal of Grid and Distributed Computing* 6 (2013) 93–106.
- [2] Amazon, Amazon EC2 instances purchasing options, <https://aws.amazon.com/ec2/purchasing-options/>, 2015.
- [3] Amazon, Amazon EC2 spot instances, <https://aws.amazon.com/ec2/purchasing-options/spot-instances/>, 2015.
- [4] M.A. Babar, H. Zhang, Systematic literature reviews in software engineering: Preliminary results from interviews with researchers, in: *Proceedings of the 3rd International Symposium on Empirical Software Engineering and Measurement (ESEM 2009)*, IEEE Computer Society, Lake Buena Vista, Florida, USA, 2009, pp. 346–355.
- [5] P. Brereton, B.A. Kitchenham, D. Budgen, M. Turner, M. Khalil, Lessons from applying the systematic literature review process within the software engineering domain, *Journal of Systems and Software* 80 (2007) 571–583.

- [6] D.S. Cruzes, T. Dybå, Recommended steps for thematic synthesis in software engineering, in: Proceedings of the 5th International Symposium on Empirical Software Engineering and Measurement (ESEM 2011), IEEE Computer Society, Banff, AB, Canada, 2011, pp. 275–284.
- [7] R. Desiraju, S.M. Shugan, Strategic service pricing and yield management, *Journal of Marketing* 63 (2006) 44–56.
- [8] O. Dieste, A. Grimán, N. Juristo, Developing search strategies for detecting relevant experiments, *Empirical Software Engineering* 14 (2009) 513–539.
- [9] T. Dybå, T. Dingsøy, Strength of evidence in systematic reviews in software engineering, in: Proceedings of the 2nd International Symposium on Empirical Software Engineering and Measurement (ESEM 2008), ACM Press, Kaiserslautern, Germany, 2008, pp. 178–187.
- [10] W. Elmaghraby, P. Keskinocak, Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions, *Management Science* 49 (2003) 1287–1309.
- [11] I. Foster, Y. Zhao, I. Raicu, S. Lu, Cloud computing and Grid computing 360-degree compared, in: Proceedings of the 2008 Grid Computing Environments Workshop (GCE 2008), IEEE Computer Society, Austin, TX, USA, 2008, pp. 1–10.
- [12] S. Gregor, The nature of theory in information systems, *MIS Quarterly* 30 (2006) 611–642.
- [13] B.A. Kitchenham, S. Charters, Guidelines for Performing Systematic Literature Reviews in Software Engineering, Technical Report EBSE 2007-001, School of Computer Science and Mathematics, Keele University and Durham University, 2007.
- [14] B.A. Kitchenham, Z. Li, A. Burn, Validating search processes in systematic literature reviews, in: Proceedings of the 1st International Workshop on Evidential Assessment of Software Technologies (EAST 2011) in conjunction with ENASE 2011, SciTePress, Beijing, China, 2011, pp. 3–9.
- [15] B.A. Kitchenham, E. Mendes, G. Travassos, Protocol for systematic review of within- and cross-company estimation models, <https://community.dur.ac.uk/ebse/study.php?type=protocol&id=5>, 2006.
- [16] Z. Li, H. Zhang, L. O’Brien, R. Cai, S. Flint, On evaluating commercial Cloud services: A systematic review, *Journal of Systems and Software* 86 (2013) 2371–2393.
- [17] R. Meir, J.S. Rosenschein, Avoid fixed pricing: Consume less, earn more, make clients happy, in: Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), ACM Press, Saint Paul, Minnesota, USA, 2013, pp. 239–246.
- [18] P. Samimi, A. Patel, Review of pricing models for Grid & Cloud computing, in: Proceedings of the 2011 IEEE Symposium on Computers & Informatics (ISCI 2011), IEEE Press, Kuala Lumpur, 2011, pp. 634–639.
- [19] C. Weinhardt, A. Anandasivam, B. Blau, N. Borissov, T. Meinel, W. Michalk, J. Stößer, Cloud computing a classification, business models, and research directions, *Business and Information Systems Engineering* 1 (2009) 391–399.
- [20] H. Zhang, M.A. Babar, P. Tell, Identifying relevant studies in software engineering, *Information and Software Technology* 53 (2011) 625–637.