CloudRecoMan: Cloud adoption made easy

A platform for assisting small and medium enterprises to adopt cloud solutions

Christos Mettouris*, Evangelia Vanezi, George Zampas, Stylianos Argyrou

Department of Computer Science, University of Cyprus, 2109 Nicosia, Cyprus, {mettour, evanez01, gzampa01, argyrou.stylianos}@cs.ucy.ac.cy

Achilleas Achilleos

Department of Electrical and Computer Engineering and Informatics, Frederick University, 7 Yianni Frederickou Street, 1036, Nicosia, Cyprus, com.aa@frederick.ac.cy

Avgoustinos Constantinides

IBSCY Ltd, Office 11, 1st Floor, Harmony House, 4 Theklas Lisioti street, CY-3030, Limassol, Cyprus, avgoustinos@ibs.com.cy

George A. Papadopoulos

Department of Computer Science, University of Cyprus, 2109 Nicosia, Cyprus, george@cs.ucy.ac.cy

Although the adoption and use of cloud computing in large, medium and small enterprises is increasing every year, statistics show that medium and especially small enterprises have a much slower adoption than large enterprises. This suggests that various types of issues may affect SMEs in making the choice to transit towards the public cloud. Based on experts and literature, small to medium enterprises' adoption of cloud is impacted by lack of resources, concerns on security and data control and uncontrolled increasing cloud costs due to the lack of cloud experience. A platform is not yet available that would assist non-technical SME managers in specifying non-technical business requirements in order to receive assistance in the form of recommendations of cloud resources, cloud applications and deployment infrastructures. The CloudRecoMan platform aims to enable this through the development of a web-based Cloud recommender and deployment platform.

CCS CONCEPTS • Information systems • Information systems applications • Computing platforms

Additional Keywords and Phrases: Cloud Recommender System, Cloud deployment platform, Cloud computing

1 INTRODUCTION

Cloud computing exploits the technological advancements in network, virtualisation and web services in order to provide the capability to deliver hardware and software services using a faster, simplified and more efficient business model. However, the process of starting using and/or migrating to the cloud is difficult, especially for SMEs where generally little

* Place the footnote text for the author (if applicable) here.

or no expertise exists in using the technical tools and platforms offered by cloud providers. Most importantly, business administrators and managers struggle to visualise the implications in terms of measurable threats and benefits from application transition and/or migration to the Cloud. Thus, while Clouds have been used for shared email environments, shared storage systems and similar purposes, there has been only a slow take-up of Cloud technology for real business applications.

A recent analysis presented in [1] shows that there are significant differences across countries regarding cloud adoption: in Sweden (75%), Finland (75%), the Netherlands (65%) and Denmark (65%) at least 65% of enterprises use cloud computing, while in Greece (22%), Romania (14%) and Bulgaria (13%) less than 25% of enterprises do so [1]. Cyprus lies in the middle (50%), while only a few years back, in 2018, the percentage was 27. Among the enterprises using cloud computing, about 79% relied on a cloud solution for their e-mail: instead of setting up a server infrastructure for their e-mail system, these enterprises opted for a cloud solution [1].

According to [1], the use of cloud computing was particularly high in large enterprises: 72% used cloud computing in 2021, which is a 7% increase compared with 2020. In terms of medium-sized enterprises, 53% used cloud computing in 2021 (compared to 46% in 2020), while in small enterprises, the use of cloud computing increased by 5% to 38% [1]. Although the adoption and use of cloud computing is increasing every year in large, medium and small enterprises, it is evident that medium and especially small enterprises have a much slower adoption than large enterprises. This strongly suggests that there are some issues that affect SMEs in making the choice to transit towards the public cloud. Therefore, it is particularly interesting to research and identify the reasons why the percentage drops as the size of the enterprise decreases.

Based on the experience and expertise of IBSCY Ltd[1] and the supporting literature [2], small to medium enterprises' adoption of cloud is impacted by lack of resources (due to shortage of expert or trained staff and time) [3], concerns on security and data control [4] and uncontrolled increasing cloud costs due to the lack of cloud experience [5]. The results of an online survey we have conducted in Cyprus from December 2020 to April 2021 are similar: from the 9 representatives of small enterprises participating, only 3 stated that they are experts or have adequate knowledge in cloud computing, with the rest noting very little (3) or basic (3) knowledge. From the latter group, 2 did not use cloud computing at all stating the "lack of resources" as the reason. The same 2 participants stated that they would be willing to switch their operations to the cloud if they found that to be advantageous and easy to install/use. To the question "what issues did you have while adopting the cloud?", participants mentioned security concerns, lack of knowledge, data control concerns, (difficult) use of the cloud and (problems with) internet access.

A platform is not yet available that would allow non-technical users (e.g. managers) to specify non-technical business requirements in the form of a company profile in order to receive recommendations of cloud resources, cloud applications and deployment infrastructures. The CloudRecoMan project aims to enable this through the development of a web-based Cloud recommender and deployment platform. It aims to deliver a cloud provider agnostic approach that allows end-users to specify their business requirements in an abstract form and, with the help of a cloud applications knowledge database, transform them into concrete requirements that can be processed and transformed into a cloud deployment solution.

CloudRecoMan aims to deliver a platform that allows identifying and recommending the optimal cloud deployment solution, based on the company profile and similar company profiles defined in the database. This enables overcoming the vendor "lock-in" characteristic of public cloud providers. The platform aims at satisfying a plethora of diverse client

[1] A leading provider of total IT solutions and IT services, specializing in the areas of cloud services and applications, systems integration, IT infrastructure, collaboration, management and security solutions. Currently, IBSCY Ltd is one of the largest Microsoft cloud providers in Cyprus and in the greater region. www.ibs.com.cy/en/cloud-services/microsoft-azure, www.ibs.com.cy/en/blog.

requirements such as deployment infrastructure (private, public, or hybrid), cost, location, etc., while at the same time supporting, optimising and automating the deployment process.

The CloudRecoMan platform is consisted of 4 software components:

- The *Recommender System (RS)* enables the computation of recommendations of cloud resources, cloud applications and deployment plans to the user. The recommendations are provided as concrete cloud requirements, e.g. hardware, business applications and networking. The recommendation algorithm relies on the company profile, as well as on explicit user feedback data.
- The *Company Profile Management System (CPMS)* enables the definition of the business requirements by the client and the storage and management of a company profile and related history data.
- The *Cloudifier* supports querying cloud provider APIs, based on the concrete requirements generated by the recommender component. It receives as input the cloud solution that was recommended and ultimately selected by the user and executes the deployment plan.
- The *Visualization component* allows for two visualizations: 1. a map graph depicting the geographical location of cloud servers of public cloud providers and 2. a tree/dendrogram graph showing all available services of well-known public cloud providers, categorized by cloud provider name, thematic areas applications (e.g. Machine Learning, AI, Analytics etc.) and name of service.

The paper is structured as follows: Section 2 discusses related work. In Section 3, the CloudRecoMan platform architecture is presented and discussed. Sections 4 to 7 present the 4 CloudRecoMan software components: the RS in Section 4, the CPMS in Section 5, the Cloudifier in Section 6 and the Visualization component in Section 7. Section 8 discusses the integration of the 4 components. The paper completes with conclusions and future work in Section 9.

2 RELATED WORK

In [6] a cloud service selection framework is presented, that uses a RS that assists a user to select the optimal services from different cloud providers that matches requirements of the user. The RS creates ranks of different services with providers and presents them to users so that they can select the appropriate or optimal services [6].

The CloudCmp framework [7] assists potential cloud customers to estimate the performance and costs of running a legacy application on a cloud without actually deploying the application. The framework first characterizes the services offered by various cloud providers into a set of common service interfaces, and benchmarks the performance and costs of these services [7]. It then expresses an application’s workload using the interfaces, and estimates the application’s performance and costs based on the benchmarking results. Benchmarking results indicate that cloud providers differ significantly in performance and costs of the services they provide, and one provider is unlikely to ace all services [7].

Research projects in data management domain have tended to focus on the improved presentation and categorisation of data in Clouds to aid integration with Cloud services. A good example of such work can be seen in the cloudTM project [8]. Here the project is focused on creating a data centric middleware in order to aid better identification of data and its requirements to aid better efficiency and fault tolerance in the Cloud.

Modelling projects are focusing on the use of models to support specific challenges such as the migration of legacy systems to the Cloud. In the Artist project, models are used to describe and wrap legacy systems to aid migration [9]. Other projects are looking to existing standards to aid the model-based management of Clouds, such as the Mosaic Cloud project that has embraced ontologies as central to their modelling solution [10]. PaaSage [11] takes a step back from providing a purely technical interface to the PaaSage application developer user and encourages the user to model his/her requirements before technical integration takes place. User defined application models expressed in CAMEL allow for a richer expression of application and business application end-user requirements translating down to how the Cloud is managed in terms of resource usage. Advancing research in MODAClouds and projects such as MOSAIC [10] inform work in PaaSage which groups domain specific standards in models at all stages of application use in the Cloud. The models are used from modelling through to execution in the Cloud; they will ensure that the user is presented with a consistent model of application activity based on his/her original requirements during design and deployment.

3 CLOUDRECOMAN ARCHITECTURE

Figure 1 depicts the architecture of the CloudRecoMan platform. The platform backend is consisted of the Company Profile Management System (CPMS), the Recommender System (RS) the Cloudifier and the database (DB). The CPMS manages the company profiles and interacts with the DB and the CloudRecoMan platform frontend which uses to receive users' input and display data to them.

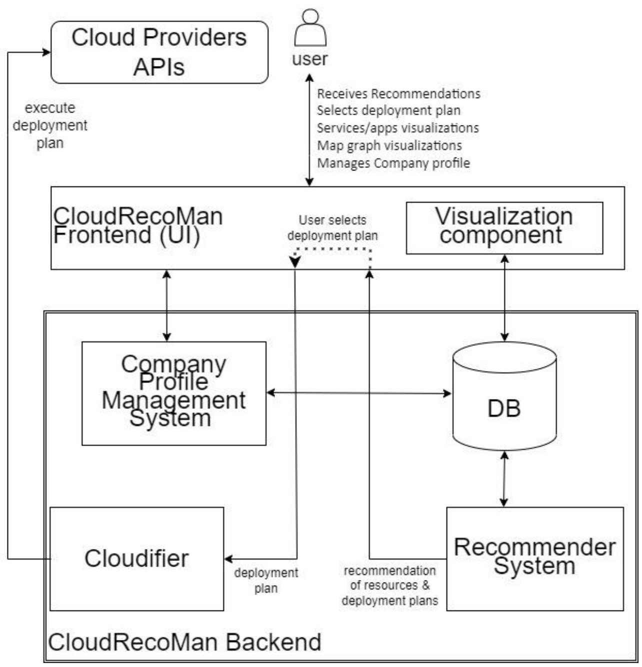

Figure 1: CloudRecoMan platform architecture.

The RS retrieves company data, user interaction data and user preferences data (see Section 4) from the DB and uses them to compute recommendations of cloud resources and deployment packages and plans for the respective company of the user. The recommendations are presented to users through the platform's frontend User Interface (UI). At this point, the user is able to select a recommended deployment plan and send it to the Cloudifier for execution on the respective cloud provider, using the corresponding cloud provider's API.

In terms of the CloudRecoMan platform frontend, it includes the respective UIs for the CPMS and the RS as discussed above, as well as the Visualization component. The latter displays cloud providers' services and applications visualizations, as well as map graph visualizations (see Section 7).

4 RECOMMENDER SYSTEM

Recommender systems (RSs) use a variety of filtering techniques and recommendation methods to provide personalised recommendations to their users. RSs are important due to the information overload that modern life experiences in all fields, and their importance will be further increased as this overload is expanding exponentially.

RSs use information retrieved from the user profile, from user's usage history, as well as information related to the items to be recommended in order to compute recommendations. The user profile (in our work we use the company profile) includes valuable information regarding the user, his/her personality, preferences, habits and more. Such information could be used by RSs to filter out any recommendations not suitable for the user. User's behavioural data are data produced based on user actions (behaviour). User actions reflect user preferences and needs, thus analysing such actions during the recommendation computation process is valuable. For example, e-commerce RSs may use any browser history information and logging data to opine about what the user likes. Finally, information related to the items to be recommended is used. Such information depends on the particular items that are recommended. For example, in movie recommendations where items under study are movies, such information may include the title of the movie, its genre, its duration, the actors etc. The most well-known recommendation approaches are Collaborative Filtering (CF), Content-based filtering and Hybrid recommendation techniques. While traditional recommender systems use limited or none contextual information to produce recommendations, recently the term Context-Aware Recommender Systems (CARS) appeared in the literature that denoted systems that aim in using many contextual parameters to improve recommendations [12].

In the context of the CloudRecoMan platform users are non-technical persons (e.g. managers) and the user profile used in the recommendation process is the company profile, as the recommendations need to best suit the company characteristics, needs and line of work. Users through the CPMS specify in the platform non-technical business requirements and company information, thus forming the company profile and defining the company data. In addition, users are enabled to rate recommended (or not) cloud resources and deployment plans to show their interest in them. This is a form of *explicit user feedback* data. RSs use explicit user feedback data on items by users, such as product ratings in e-commerce scenarios, to elicit and model user preferences and offer personalised recommendations of products users would enjoy [13, 14]. In case explicit feedback is not available, RSs use *implicit user feedback* on items by tracking user behaviour, such as user transaction data (purchase history), clickstream data, click-through rate (CTR) and browser history information [15, 16, 17, 18]. Implicit techniques have been used in RSs for products in online retail stores (e-commerce websites), as well as movies, music, journal articles, web documents and content, online news articles, books, television programs, and other RSs applications.

We propose a Hybrid RS that consists of: a knowledge-based RS deployed initially, followed by a CF RS that offers refined recommendations. Figure 2 shows the recommendation process workflow. Initially the user completes the company profile. At this point, due to the cold start problem that CF RSs face, the knowledge-based RS is deployed. Recommendations at this point may not be of the desired accuracy. The cold start problem is an inherent problem of RSs which dictates that, for RSs to be able to produce meaningful and accurate recommendations, user interaction with items first needs to take place. Knowledge-based RSs have functional knowledge on how a particular item (e.g. cloud resource) meets a particular user need (company need), and can therefore reason about whether an item could meet a user need and should therefore be recommended [19]. The user profile can be any knowledge structure that supports this process: from a simple query to a more detailed representation of a user's needs [19]. In our case, the user query is the company profile, which defines aspects of the company such as sector, company size (number of employees), number of IT personnel, years of operation, whether they are currently using the cloud, current knowledge of cloud computing, etc. Based on the company data, recommendations are computed. The reasoning approach of the Knowledge-based RS relies on using empirical rules provided by one of the most respected cloud experts in Cyprus, IBSCY Ltd, that is derived from several years of experience in providing cloud services to their customers. Any data provided for the implementation of the knowledge-based RS were fully anonymized.

After the user has received knowledge-based recommendations based on the company profile, he/she is able to rate them from 1 to 5 (1 being the least desirable and 5 the most desirable). The user is also able to rate other cloud resources and deployment plans not currently recommended to them. Using this explicit user feedback data, the CF RS is able to better refine the recommendations.

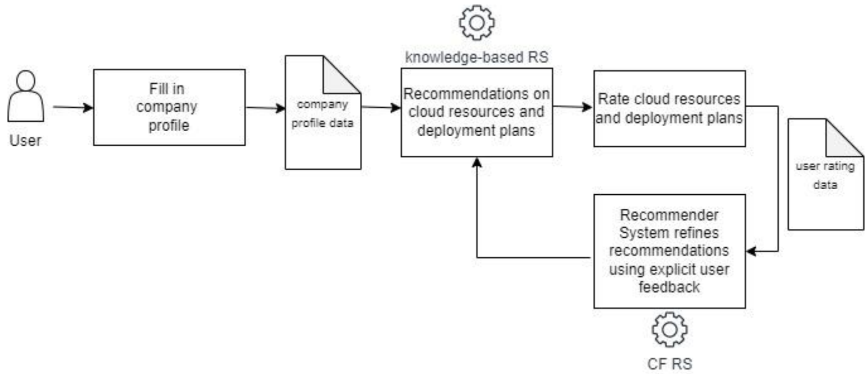

Figure 2: CloudRecoMan recommendation process workflow.

Collaborative Filtering recommends items that similar users to the active user have highly rated (hence like). There are two types of methods followed in CF, the memory-based or neighbourhood methods that use similarity functions (Pearson Correlation or Cosine Distance) to compute the user's neighbourhood, and the model-based methods that use user feedback on items (i.e. ratings) to learn a model for the user that is then used for computing recommendations [20]. The assumption made in CF is that those who agreed in the past tend to agree also in the future. Thus, since the active user's neighbours tend to agree with him/her (they are similar users), the items these neighbours like the most might be included highly in the list of preferences of the active user as well, and hence can be recommended. The biggest advantage of CF is that it does not depend on any system representations of the items to be recommended and thus can function well with complex items such as music and movies [21].

In the context of CloudRecoMan, CF is used as a complementary method to the knowledge-based method described above. The knowledge-based method recommends cloud resources and deployment plans based exclusively on the company profile. However, if there exist resources or deployment plans that would potentially work for a company whose profile though does not suggest them, then using the knowledge-based approach solely, these resources and plans would never be recommended. On the contrary, CF may provide such recommendations in the case where another company that happens to be a neighbouring company to the active company has assigned high ratings to these resources and plans.

Model-based CF approaches - or latent factor models - use machine learning techniques such as Matrix Factorisation to decompose a multidimensional user-item-context matrix to learn latent factors for each user and item in the data to compute recommendations [22, 23]. While model-based methods have been proven more accurate and superior to neighbourhood methods as they allow the incorporation of additional information, such as implicit feedback, to be utilised in the recommendation process [24], in our work we are using memory-based (neighbourhood) methods due to their simplicity and also the nature of our recommendation problem. Resource and deployment plan recommendations in cloud settings constitute a reduced recommendation problem, in comparison to a full product recommender system where users and products to recommend are millions, and the latent factors for the users and items are too many. In our case, we estimate to have a few hundreds of different company profiles, while the number of cloud resources and deployment plans to recommend are fewer too. As a first implementation of the CloudRecoMan RS, the algorithm we have selected to implement is the correlation-based approach [20].

The top 6 recommendations of cloud resources and deployment plans are displayed to the user by our Hybrid RS as follows: 3 from the knowledge-based RS and 3 from the CF RS. If one of the two recommenders has less recommendations to display, then the overall recommendations shown to the user are also less than 6.

5 COMPANY PROFILE MANAGEMENT SYSTEM

The CPMS's aim is to manage the company profiles on the backend, by storing them and retrieving them from the DB, while at the same time to display all the information to users and to provide an easy way for users to set up company profiles, in the respective frontend UIs. This information includes the company profile data, and the user rating data.

Users should first set-up a company profile in order for the recommendation process to be initiated. The platform allows for multiple company profiles to be set-up by the same user. Each company will get their own customised recommendations based on their own profile data, independently of whether they are handled by the same user or not. The same applies for the CF RS, the explicit user feedback data acquired by the platform is linked not only to the user generating the interaction data, but also to the respective company profile. Users can view and manage all company profiles created by them in the "Manage Companies" UI (Figure 3). For each company profile, the following options are provided: select the specific company to view its profile data (that includes past ratings of resources and plans), view the recommendations for cloud resources and deployment plans for that company, view the already deployed plans, edit the company profile data, and delete the company profile. The latter two functionalities, together with the functionality of entirely deleting a user account and all its data, contribute to the conformation of the CloudRecoMan platform with the GDPR. Figure 4 displays part of the company profile set-up form, with both basic information and cloud related information respectively.

6 CLOUDIFIER

The Cloudifier has been developed by using Pulumi [25], an open source "infrastructure as code"[2] tool for creating, deploying, and managing cloud infrastructure. Pulumi works with traditional infrastructures like VMs, networks, and databases, in addition to modern architectures, including containers, Kubernetes clusters, and serverless functions. The Cloudifier receives as input the selection of the user from the recommended cloud resources and deployment plans and creates the cloud infrastructure. At this stage, the Cloudifier supports only Azure for creating and managing cloud infrastructure. It is in our future work plans that other cloud providers will be added as an option, such as AWS, Google Cloud, etc. Regarding the Azure integration, other tools could have been used such as the Azure API Management and Terraform, but we found Pulumi, along with Azure command-line interface (Azure CLI) [26] to be the easiest and cleanest to use. The Azure CLI is a command-line interface for creating and managing Azure resources of users.

[2] Infrastructure as code (IaC) is the process of provisioning and managing infrastructure defined through code, as opposed to a manual process.

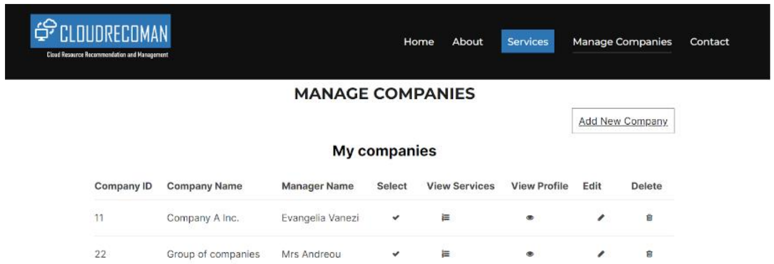

Figure 3: User interface to manage companies.

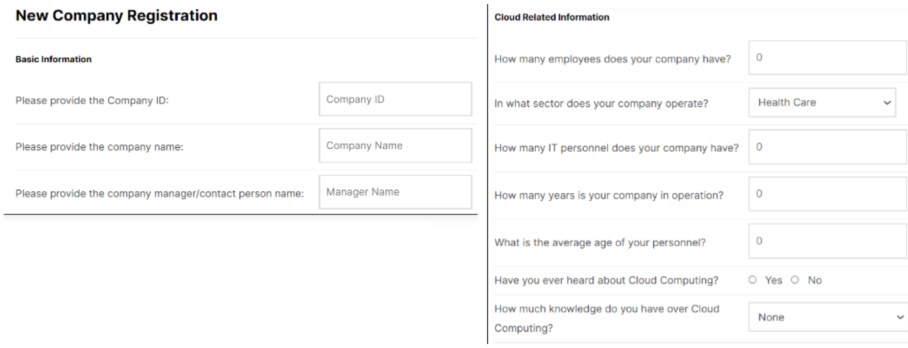

Figure 4: Screens of the company profile.

When the user selects a recommendation, the Cloudifier initiates the deployment process of the cloud infrastructure. The first step of the execution is to connect to the selected cloud provider's API through the user's account using the appropriate CLI (for Azure it is the Azure CLI) as shown in Figure 5 (user's account info has been deleted).

Figure 5: Connecting to user's account on Azure using Azure CLI.

After a successful connection with the CLI, the Cloudifier updates the configuration file of the Pulumi so that the tool is able to build the cloud infrastructure that the user has selected. The final step is the execution of Pulumi's main typescript program using a script that contains the command "pulumi up": the algorithm parses the data from the configuration file and updates the needed variables, creates a new resource group with the name of the user's company and builds the cloud infrastructure. Figure 6 shows a visualisation of the resources built using Cloudifier in an Azure user account: a new Resource Group and a Virtual Machine was created with the same name.

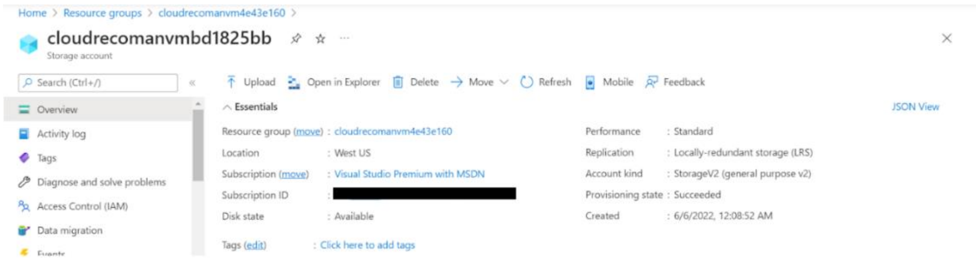

Figure 6: Visualization of the resources built using Cloudifier in an Azure test account.

7 VISUALIZATION COMPONENT

7.1 Map Graph

The Physical location of each server plays a significant role when choosing a cloud provider. The closer the server to the user's machine the faster the data transfer speed is. This is due to the physical distance between the server and the user's location. Therefore, for business owners the best approach is to select a cloud provider that can offer servers in the geographic area of most customers so that they can have quick access to the company's data. Servers' physical location also plays a role in the regulations that needs to be followed. For instance, different locations provide different privacy measures since laws in some countries allow server access to authorities for monitoring purposes.

In the context of the visualization component, a map graph has been implemented showing all the available server locations based on each cloud provider. The locations for each server have been obtained from the official website of five well-known cloud providers: AWS, Alibaba, Azure, Firebase, Oracle. The graph shows with different colours the locations of the servers for each specific cloud provider, as shown in Figure 7, left part. The graph is interactive; the user can hover on top of the legend and the map will update to show only the locations for the server selected. Moreover, by hovering over a specific location on the map, the graph depicts the city represented by that location and the cloud provider who owns that specific server. The graph has been implemented using Datawrapper, an online tool for creating visualisations without coding required. Datawrapper allows for free publication of the final chart and provides a link to embed the visualisation on a specific platform. Once the visualisation has been tested, a WordPress plugin has been created for the inclusion of the graph on the website.

7.2 Tree/Dendrogram Graph

The tree/dendrogram graph (Figure 7, right part) shows all the available services of each cloud provider. This graph is split into the following categories: cloud provider name, thematic areas applications (e.g. Machine Learning, AI, Analytics etc.) and name of service. This graph is used as a tool for the user to observe whether his/her selected choice of cloud provider has the services needed, or to show the complexity of such systems and promote the use of the application provided. The graph has been implemented using the D3JS library in Javascript. D3JS is a library dedicated to manipulating documents based on different kinds of data. Such data can be geographical data, time series data or even comma separated values. This specific visualisation has been obtained by adapting a pre-existing template to suit the needs of the application. The data have been obtained through the official websites of the cloud service providers.

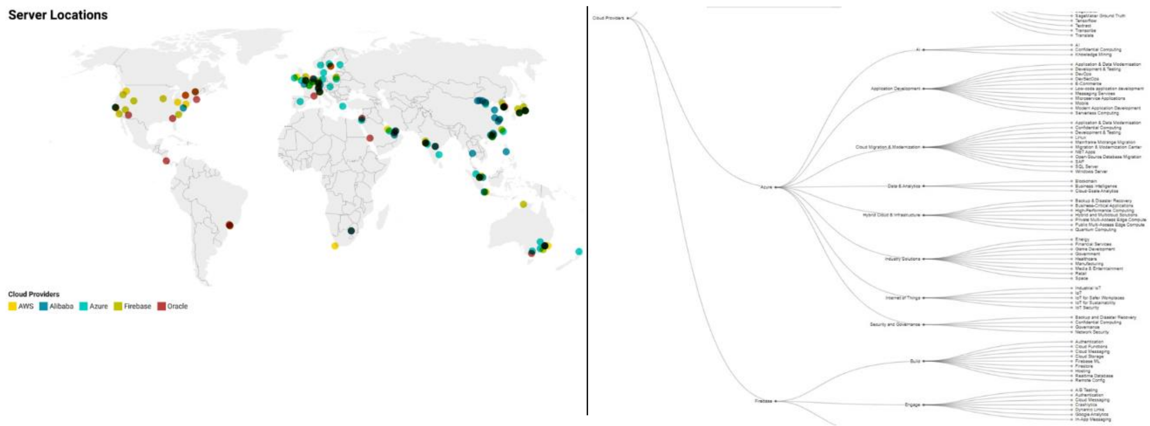

Figure 7: left: Graph map, right: Tree/Dendrogram.

8 INTEGRATION

The CloudRecoMan platform was developed on top of the well-known, widely used, open source WordPress[3] platform. WordPress was used as the basis for development, where all platform components (the RS, the CPMS, the Cloudifier and the visualizer) were developed as plugins. WordPress is nowadays used for the development of millions of website and web platforms and has an active supporting community. In terms of integration, the components being developed as plugins have the advantage of easy installation on the platform, portability and easier maintenance.

9 CONCLUSIONS AND FUTURE WORK

This paper presents the CloudRecoMan platform and its components. The platform will next be evaluated in real settings. We aim to invite company owners, managers and personnel from IBSCY Ltd’s clientele to an evaluation process, where evaluation participants will be asked to freely use the CloudRecoMan platform for a period of 1 month and report their feedback. We aim to evaluate the platform’s effectiveness in providing meaningful recommendations, its efficiency in terms of enabling the users complete their tasks quickly, usability and ease of use. The results of the evaluation will drive the platform’s refinement, as well as functionality updates and improvements. Special attention will be given to the improvement of the RS based on the evaluation results. The algorithms used, especially the correlation-based approach in the CF RS is subject to improvements as it was selected, not because of its efficiency in recommendations’ accuracy, but in terms of its simplicity, problem complexity, as well as due to the limited data available at the current phase (model-based approaches require more data to provide accurate results). We expect that the CF RS will offer adequate results.

As future work, we aim to investigate whether more advanced model-based CF recommendation techniques can be used in the case where the recommendation problem at hand increases in complexity. This depends on the amount of companies registered in the platform and the amount of cloud resources and deployment plans defined in the platform. In case these parameters greatly increase, then the recommendation problem will be increased in complexity, allowing us to utilize more complex recommendation techniques such as Matrix Factorization (see Section 4).

The Cloudifier can be extended to monitor and report the deployment status based on sensors (e.g., CPU, RAM) provided by the platform, as well as perform scaling actions when the deployment rules are violated, e.g., high CPU load.

[3] https://wordpress.com/

ACKNOWLEDGMENTS

The CloudRecoMan project is funded by the University of Cyprus, under the "ONISILOS" funding scheme for postdoctoral researchers, 1st call for proposals.

REFERENCES

[1] Eurostat, "Cloud computing - statistics on the use by enterprises", Online: https://ec.europa.eu/eurostat/statistics-explained/index.php/Cloud_computing_-_statistics_on_the_use_by_enterprises#Enterprises_using_cloud_computing [Accessed June 2022]

[2] RightScale, "State of the cloud report, 2017" Online: https://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2017-statecloud-survey#cloud-workloads [Accessed June 2022]

[3] Tammi S. Feltham, Glenn Feltham, and James J. Barnett. 2005. The dependence of family businesses on a single decision-maker. Journal of small business management, 43(1), 1-15.

[4] Nelson Gonzalez, Charles Miers, Fernando Redígolo, Marcos Simplício, Tereza Carvalho, Mats Näslund, and Makan Pourzandi. 2012. A quantitative analysis of current security concerns and solutions for cloud computing. Escola Politécnica at the University of São Paulo (EPUSP), São Paulo, Brazil.

[5] Tim Mazzarol. 2003. A model of small business HR growth management. International Journal of Entrepreneurial Behavior & Research, 9(1), 27-49.

[6] Seung-Min Han, Mohammad Mehedi Hassan, Chang-Woo Yoon, Hyun-Woo Lee, and Eui-Nam Huh. 2009. Efficient Service Recommendation System for Cloud Computing Market. In: Ślęzak, D., Kim, Th., Yau, S.S., Gervasi, O., Kang, BH. (eds) Grid and Distributed Computing. GDC 2009. Communications in Computer and Information Science, vol 63. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-10549-4_14

[7] Ang Li, Xiaowei Yang, Srikanth Kandula, and Ming Zhang. 2010. CloudCmp: shopping for a cloud made easy. In Proceedings of the 2nd USENIX conference on Hot topics in cloud computing (HotCloud'10). USENIX Association, USA, 5.

[8] cloudTM EU project homepage, Online: http://www.cloudtm.eu/. [Accessed June 2022]

[9] Artist EU project Homepage, Online: https://cordis.europa.eu/project/rcn/105117_en.html. [Accessed April 2022]

[10] Mosaic EU project homepage, Online: http://www.mosaic-cloud.eu/. [Accessed March 2022]

[11] PaaSage EU project homepage, Online: https://paasage.ercim.eu/. [Accessed March 2022]

[12] Gediminas Adomavicius and Alexander Tuzhilin. 2011. A.: Context-Aware Recommender Systems. In: Ricci F., Rokach L., Shapira B., Kantor P. (eds) Recommender Systems Handbook. Springer, Boston, MA, pp. 217-253.

[13] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In Proc. of the 8th IEEE International Conference on Data Mining, pp. 263–272.

[14] Wai Tik So and Katsutoshi Yada. 2017. A Framework of Recommendation System Based on In-store Behavior. in Proceedings of the 4th Multidisciplinary International Social Networks Conference, New York, NY, USA, p. 33:1–33:4.

[15] Edward Rolando Núñez-Valdéz, Juan Manuel Cueva Lovelle, Oscar Sanjuán Martínez, Vicente García-Díaz, Patricia Ordoñez de Pablos, and Carlos Enrique Montenegro Marín. 2012. Implicit feedback techniques on recommender systems applied to electronic books. Computers in Human Behavior, Volume 28, Issue 4.

[16] Ladislav Peska. 2016. Using the Context of User Feedback in Recommender Systems. In Proceedings of the 11th Doctoral Workshop on Mathematical and Engineering Methods in Computer Science, MEMICS 2016, pp. 1-12.

[17] Byoungju Yang, Sangkeun Lee, Sungchan Park, Sang-goo Lee. 2012. Exploiting Various Implicit Feedback for Collaborative Filtering. In Proceedings of the 21st International Conference Companion on World Wide Web, WWW '12 Companion, ACM, New York, NY, USA, pp. 639–640.

[18] Xing Yi, Liangjie Hong, Erheng Zhong, Nathan Nan Liu, Suju Rajan. 2014. Beyond Clicks: Dwell Time for Personalization. In: Proceedings of the 8th ACM Conference on Recommender Systems, RecSys '14, ACM, New York, NY, USA, pp. 113–120.

[19] Pearl Pu, Li Chen, and Rong Hu. 2012. Evaluating recommender systems from the user's perspective: survey of the state of the art. User Model User-Adap Inter 22, 317–355. https://doi.org/10.1007/s11257-011-9115-7

[20] Gediminas Adomavicius, Ramesh Sankaranarayanan, Shahana Sen, and Alexander Tuzhilin. 2005. Incorporating contextual information in recommender systems using a multidimensional approach. ACM Transactions on Information Systems (TOIS), vol. 23 103–145.

[21] Robin Burke. 2002. Hybrid Recommender Systems: Survey and Experiments, User Modeling and User-Adapted Interaction, vol. 12, no. 4, pp. 331–370.

[22] Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. 2010. Multiverse Recommendation: N-dimensional Tensor Factorization for Context-aware Collaborative Filtering. In Proceedings of the Fourth ACM Conference on Recommender Systems, New York, NY, USA, pp. 79–86.

[23] Trung V. Nguyen, Alexandros Karatzoglou, and Linas Baltrunas. 2014. Gaussian Process Factorization Machines for Context-aware Recommendations. In Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval, New York, NY, USA, pp. 63–72.

[24] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. In Computer, vol. 42, pp. 30-37. DOI 10.1109/MC.2009.263.

[25] Pulumi, Online: Available: https://www.pulumi.com/docs. [Accessed June 2022]

[26] Azure CLI, Online: Available: https://docs.microsoft.com/en-us/cli/azure. [Accessed June 2022]