

Cloud-based Platform for Personalization in a Wellness Management Ecosystem: Why, What, and How

Pei-Yun S. Hsueh, Raymund J.R. Lin, Mark J.H. Hsiao, Liangzhao Zeng, Sreeram Ramakrishnan, and Henry Chang, *Member, IEEE*

Abstract—

Offering personalized services through dynamically formed ecosystems is essential to personal wellness management. In this paper, we present the design of a cloud-enabled platform to facilitate the collection and delivery of evidence for personalization in a multi-provider ecosystem environment. In addition, the platform also provides essential building blocks of personalization services: smarter analytics for active personalization and dynamic provisioning. While the former common service takes charge of inferring user wellness risks from multiple data sources on the fly and making risk-driven recommendations, the latter common service determines optimal platform pricing and resource allocation given the constraint of acceptable quality of service.

***Index Terms—*computer-tailored lifestyle intervention, cloud computing, personalization, dynamic pricing, active learning, feature selection, sample selection**

I. INTRODUCTION

WELLNESS management is essential for keeping aging populations for years to come. Large-scale national studies have presented evidence of the importance of lifestyle intervention (e.g., dietary intake planning and exercise) on chronic disease management and health promotion programs. For example, the Finland National Type II Diabetes Prevention Programme (FIN-D2D) (Saaristo et al., 2007) is a 5-yr study across 1.5 million subjects, which tests whether lifestyle intervention can prevent Type II diabetes from onset or at least delay its progress in high-risk subjects. Shortly before the FIN-D2D study, the Center of Disease Control (CDC) has also initiated a Diabetes Prevention Program (DPP, 2002) to trace the effect of lifestyle intervention on high-risk subjects who are obese or have pre-diabetes. Both studies yielded positive evidences of lifestyle intervention in high-risk subjects.

However, there exists a danger in believing that lifestyle intervention can be structured in a way that is universally

applicable. The planning of lifestyle interventions in each step of the wellness management process can vary among individuals. While existing providers are expanding service choices to cover various user needs, there is still a long tail of demand unsatisfied. The success of the previous intervention studies often relies on some sort of user-tailoring. For one, the delivery of lifestyle intervention in DPP relies on intensive individual counseling and motivational support on diet, exercise, and behavior modification. Moreover, there also exist dangers in applying universal standards on self-motivation, screening high-risk subjects, and early diagnosis and treatment. Take the Medical Research Council study (Alderman et al., 1993) as an example. Researchers compared the effectiveness of using a universal threshold (e.g., systolic pressures at 135 mm Hg) and a personalized threshold on determining the need of antihypertensive therapies. The results suggested that the personalized threshold would prevent 53 events per 1,000 persons treated, yielding a treated-to-benefited ratio 44 times better than the universal threshold.

Despite all the positive evidence on the need of attending to individual differences, personalizing services has remained as a concept. The large number of variations in risk thresholds and lifestyle interventions pose challenges to existing providers. Recently computer-tailored programs have been introduced. Initial evidence has been on the positive side, especially for weight reduction (Lee & Lustria, 2009; Kroeze et al., 2008). However, previous computer tailoring methods are often limited to computer-generated reminders for caregivers or tailored education materials.

Meanwhile, there are more and more information about a person's wellness conditions, e.g., monitoring vital signs, are available electronically. Consumer reports show a growing trend in the market of mobile patient monitoring, which is predicted to reach \$1.9 billion by 2014 (Juniper Research, 2010). Previously, these records are left unanalyzed and therefore led to ill-informed decision making. In some cases, the analytical decision could have foreseen problems and introduced appropriate interventions. To put personalized intervention services to work, this paper sets out to understand the design and enablement requirements of personalized services in a multi-provider environment like wellness management.

Manuscript received August 30, 2010. This work was supported in part by the Ministry of Economic Affairs, Taiwan.

P. S. Hsueh, L. Zeng, S. Ramakrishnan and H. Chang are with the IBM T.J. Watson Research Center, Hawthorne, NY 10583 USA.

R. J. Lin and M.J. Hsiao are with the IBM Taiwan Research Collaboratory, Taipei, Taiwan.

II. WELLNESS ECOSYSTEM IN THE CLOUD ENVIRONMENT

A. *Ecosystem design for personalization services*

Unlike healthcare services which center around large organizations, most of the wellness service providers are small- or mid-size businesses (SMBs) (such as fitness centers, nutrition clinics, and physical exam centers). To provide a total solution of health promotion – from physical examination to follow-up interventions and feedbacks – it is necessary for them to work together. Therefore, being able to share service information in a standard way is important.

The creation of personalization services has taken the need of sharing and standardization to a new level. Take the following scenario as example. A diabetic patient came to the physical exam center for her check-up through her family physician's referral. She has also been advised to participate in a health promotion program (including fitness and dietary suggestions). Now, her case manager has to help her create an intervention plan that suits her wellness risk profile. First, her wellness risk profile is created by scanning through population wellness records to identify risk groups that match her wellness status and risk factors. Second, her feedback reward is determined by comparing her progress with cohorts who started with similar wellness status. Finally, follow-up suggestions for monitoring and interventions are created according to the outcomes of the identified cohorts.

To support such a scenario, we need to engage many distinctive types of participants other than service providers, e.g., end users, health promotion enablers and sponsors, data providers, and wellness risk researchers.

The concept of ecosystem, therefore, comes into the center stage. An ecosystem, in this context, is an economic community, which produces goods and services of value to customers, who are themselves members of the ecosystem (Moore, 1996; WEF, 2008). Previous examples in government services and global integrated enterprise services have demonstrated the effectiveness and efficiency of IT-facilitated ecosystems on resource allocation. In this study, we then further extend the concept of ecosystem into the domain of wellness management.

B. *Ecosystem enablement of personalization services*

The formation and management of ecosystems require support from an on-demand service integration platform. To achieve this, the common services provided on the platform need to satisfy the following criteria: (a) accessible from anywhere, at anytime (such that participants in the ecosystem can leverage each other's services); (b) standardized interface (such that pre-structured service template of one participant can be shared and reused by others in their own context); (c) scalable and elastic (such that ecosystem participants can guarantee in-time services to a large number of users when they are in need); (d) rapid prototyping of ecosystem services; and (e) ecosystem service provisioning and management.

In this paper, we design a wellness cloud, which is a dynamic infrastructure pattern that follows service-oriented approaches to deliver services of multiple ISVs, service providers, and other stakeholders in the domain of wellness management. We follow the dynamic business application SaaS (Software-as-a-Service) model of software deployment (Forrester Research, 2008), where an application is hosted as a service provided to customers across the Internet and shared in a multi-tenant environment that comes with dynamic orchestration and provisioning capabilities. (See Section III for detailed descriptions of its building blocks.)

In addition, personalized service providers need to dynamically select and compose user-specific service plans from the pre-structured service templates offered by multiple providers. The development of such personalized services typically involves two major tasks: (1) tailoring user-specific lifestyle interventions from pre-structured applications or services according to the user's need and (2) ensuring smooth service delivery in dynamically formed ecosystems.

This has incurred requirements beyond current cloud platform support. Specifically, the platform needs to provide common services that are essential to personalization. For Task (1), the platform needs an analytics service that is capable of learning user needs and adapting services to users and a data service to handle pre-population and data exchange between services. For Task (2), this platform needs a dynamic provisioning common service to determine the optimal resource allocation and platform pricing.

C. *Roadmap to the provisioning of personalization services*

Having identified the requirements of ecosystem design and enablement, we need to put these into the context of wellness management for consideration. The remaining of the paper is then organized as following. In Section III, we describe our cloud-based platform support for dynamically composing wellness services in a multi-provider wellness management ecosystem. Then, in Section IV and V, we introduce the two common services needed for the development and the provisioning of personalized services. Finally, in Section VI, we summarize the needed support from wellness cloud and explore future directions.

The end goal is to construct a personalization engine that can allow ISVs, service providers or stakeholders, with permission from the target users, to plug in the analytical model of personal wellness risks and to learn what to recommend to users based on the user's wellness risks. Because a personalized wellness management ecosystem encompasses many distinctive types of participants, it is also important to identify pricing strategies that can be used to streamline the process and redistribute revenues.

III. WELLNESS CLOUD: OPEN PLATFORM SUPPORT ON GREENOLIVE

As discussed in Section II, wellness management entails some unique requirements on the platform, such as standardization, knowledge sharing, dynamicity support, and scalability. Existing publicly available cloud platforms, such as Amazon Web Service and Force.com, are not designed to meet the needs in this domain. In order to tackle these challenges, we propose GreenOlive cloud platform (as depicted in Fig. 1), which consists of three layers: (i) infrastructure layer, (ii) platform layer, and (iii) application layer.

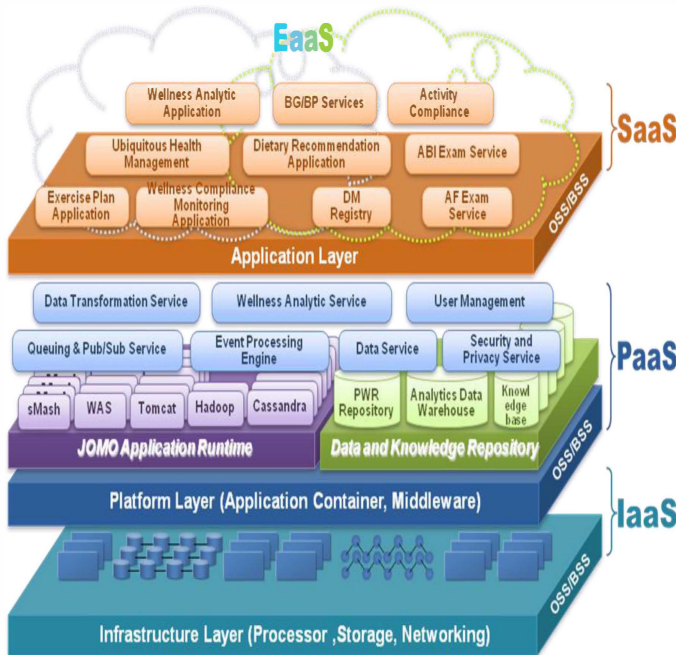


Fig 1. GreenOlive Cloud Platform for Wellness Management

(i) The infrastructure layer is designed to distribute computational resources, including processing time, storage space, and networking bandwidth, over web services. The services are named as infrastructure-as-a-service (IaaS).

(ii) The platform layer offers Platform-as-a-Service (PaaS). This layer contains an application runtime environment dubbed Jomo, a collection of knowledge and data repositories, and a collection of common services that run on top of Jomo and use the knowledge/data stored in the repositories. Each of the common services provides a set of APIs that allow independent software vendors (ISVs) to develop new services on their own.

The common services cover a wide area of functionalities, including those that are essential to personalization, e.g., wellness analytics services which generate new guidelines or new insights of existing guidelines for wellness management; ecosystem provisioning services which optimize resource allocation and platform pricing.

Other common services include security and privacy services which protects the integrity of user data and privacy, data transformation services which transform data/events from

different sources to a standard format (like Clinical Document Architecture, abbreviated CDA), data store services which provide APIs for accessing knowledge and data repositories (for instance, Personal Wellness Record, abbreviated PWR), event processing services which process events in a real-time fashion (like Monitor as a Service, abbreviated MaaS), queuing & pub/sub services which provides communication channels between services, and user services which manage user identity.

(iii) The application layer adopts the single-app SaaS model (Forrester Research, 2008). Applications are developed using the open APIs of the common services on the platform layer.

Finally, with support from all the three layers, the platform provides ecosystem-as-a-service, which is an application composition/mash-up mechanism that can facilitate service integration in an ecosystem.

IV. SMARTER ANALYTICS FOR ACTIVE PERSONALIZATION

One of the personalization-essential common services on the Platform layer is the wellness analytics service. Over the past years, we have seen tremendous interest and urgent need in data analytics in various domains. Despite the plethora of commercial software in the market, there are still challenges on the accessibility, reusability, scalability and efficiency of these well developed statistical algorithms.

Many believe that Cloud computing is a compelling new paradigm that can lead to IT revolution and reshape the entire ICT industry. In Section III, we presented the three layers of support on the cloud platform. In this section, we introduce a cloud-based implementation of wellness analytic services, i.e., the Open Analytics Services (OAS), using the IBM Research Compute Cloud (RC2). To fulfill the requirement of accessibility and standard interface for knowledge sharing, we adopt the RESTful (REpresentational State Transfer) HTTP protocol to deliver analytical functionalities for personalization. In addition, we follow a standard that is commonly accepted by the vendor-independent data mining consortium, i.e., the Predictive Model Markup Language (PMML 4.0) standard (Guazzelli et al., 2009). The standardization effort has enforced compatibility of the learned models, making it possible to exchange these models between services or applications.

A. Grab-and-Build Open Analytics Services

OAS offers an agile, flexible, and robust framework for online data analytics. The goal of this framework is threefold: (1) to support the dynamic requirements of online wellness risk analytics for both wellness domain experts and developers; (2) to allow the interoperability of models between different applications and platforms; and (3) to enable the publishing and sharing of valuable models.

To support online analytics, the OAS functions serve as a common service on the PaaS layer, providing various statistical functionalities and data mining algorithms. As we do not expect personalized service developers to be domain experts who are

necessarily highly tech-savvy in programming, we need to provide easy access to the analytics tool over web services in a cloud environment. The analytics engine in OAS is thus wrapped with a RESTful API for easy access. RESTful HTTP protocol aim to express every service delivered over the web as a set of resources using universal resource identifiers (URI). By following the REST architectural principle, HTTP methods can be used to create, read, update, and delete the OAS resources. Together with JSON (JavaScript Object Notation), which is a lightweight data-interchange format, the client application can easily interact with OAS to present analytics results, or acquire information for further processing (e.g., visualization).

Once the analysis is done and the analytics model is built, we face the issue of how to allow models to be interoperable between different analytics applications and platforms. Here we adopt the PMML 4.0, which is a XML-based markup language and has been broadly supported by data mining community and industry, to alleviate the interoperability problems. Once the trained predictive models are expressed in PMML and deployed on cloud, users can again invoke a RESTful-based web service to acquire the learned model and apply it to the applications of their interest, with a different set of testing data. This is to say, if a company wants to switch from one analytical model to another, vendor lock-in or system migration costs would not be an issue.

The combination of RESTful analytics API and PMML open standard provides the opportunity for domain experts to deliver wellness risk models as a cost-effective service, i.e., Analytics-as-a-Service (AaaS). AaaS allows domain experts or data owners to publish and share their research results freely, benefiting from an environment in which flexibility and interoperability are truly attainable. General users who are not domain experts can also join any dynamically formed ecosystem and leverage the latest wellness evidence made available by ecosystem partners on the cloud.

B. Smarter analytics for active personalization

In the last version of Green Olive platform design, we have presented several wellness analytics services needed for chronic disease management (GreenOlive, 2010). In this study, we further equip the platform with common services of smarter personalization analytics that are essential to personalization, i.e., suggesting lifestyle interventions with respect to users' need. Specifically, we offer personalization services over the OAS framework to complete the two tasks in Section II.B: (1) Profiling the target user's wellness risks from personal wellness records, and (2) match-making recommendable services and adapting services according to the user's wellness profile.

This involves addressing the following challenges. First, because a person's wellness risk is often multi-fold, we need to infer the multi-faceted characteristics of wellness risks for further mitigation. One possibility is to position the target user into various risk groups by scanning through the target wellness records and identifying relevant risk factors. However, this entails another question. As the wellness records are coming from various sources and the diagnosis of some of the wellness

risks can be uncertain, it is dangerous to make wellness-related decisions based on the records alone. Finally, because new data records can be added at anytime, we need to handle the constant online update of wellness risk profiles given the new records.

Multi-faceted wellness profiling

The first challenge is to filter statistical noises from the vast number of irrelevant features and diagnosis results that can preempt the signals from true risk factors. With the expansion of personal wellness record (PWR) and corresponding feature space, it is now imperative to develop an automatic mechanism that can identify risk factors, i.e., a portion of features that are correlated strongly with patients' wellness status, for future monitoring and intervention.

However, prior data mining research focuses on finding common risk factors for common diseases and yields limited success due to the statistical noise problem. It is thus essential to develop PWR-wide risk stratification services that can scan across the selected PWR databases to further narrow in on a subset (or subsets) of high-risk subjects and the risk factors associated with each subset. Each grouping of a subset of cases (including the information of relevant risk factors) then composes one risk group.

In fact, many feature/sample evaluation, subset section, and ensemble modeling algorithms have been developed in the field of machine learning and pattern recognition. For example, feature filter (Das, 2001) and wrapper (Kohavi and John, 1996) for feature selection; bagging (Sutton et al., 2005) and boosting (O'Sullivan et al., 2000) for ensemble modeling. Our design does not assume any specific algorithm to be used. Instead, we present a summarization framework (as shown on the left side of Fig.2) to identify the most representative risk groups, which minimize the overall distance between the predicted and diagnosed risk level of the target wellness risk.

At run-time, when a target user's wellness record is available for wellness risk profiling, the PWR-wide risk stratification service will be activated to compare the user's records against the various risk groups', so that the user's wellness risks can be profiled, and follow-up interventions can be recommended by making changes on the most important risk factors.




Fig 2. Active Characterization of Personal Wellness Status

Health risk assessment quality control

Having identified user-relevant risk groups from population PWR, the second challenge is to provide data quality assurance by quantifying the sources of deviation of risk prediction. To achieve this, we have designed a data quality monitor to determine what risk groups are safe to be included for the prediction task. Similar to Hsueh et al. (2009), we characterize data- and case-based prediction errors in three ways: risk group noise, case ambiguity, and noise-adjusted case ambiguity.

- 1) **Risk group noise** quantifies data-based errors, e.g., the average deviation of the prediction results obtained with one target risk group. Ref. Formula (1), wherein g_i refers to the i th risk group, $\hat{Risk}(g_i)$ and $Risk(g_i)$ are the predicted and actual diagnosed risk level (between -1 and 1), and N is the number of cases whose risk levels this risk group has predicted on,

$$Noise(g_i) = \sum_{n=1}^N \frac{|\hat{Risk}(g_i) - Risk(g_i)|}{N} \quad (1)$$

- 2) **Case ambiguity** quantifies case-based errors, e.g., the average deviation of all predictions yielded on one single case, based on all relevant risk groups. Ref. Formula (2)-(5), wherein $\theta_+(Case_j)$ is the proportion of positive predictions and $\theta_-(Case_j)$ is that of negative predictions.

$$Amb(case_j) = \frac{Risk(case_j)}{Risk_{max}} H(\theta(Case_j)) \quad (2)$$

$$H(\theta(Case_j)) = -(\theta_+(Case_j) \log(\theta_+(Case_j)) + \theta_-(Case_j) \log(\theta_-(Case_j))) \quad (3)$$

$$\theta_+(Case_j) = Risk_+(Case_j) / Risk(Case_j) \quad (4)$$

$$\theta_-(Case_j) = Risk_-(Case_j) / Risk(Case_j) \quad (5)$$

- 3) **Noise-adjusted case ambiguity** is the case ambiguity score reweighted with respect to the noise level of each risk group involved in the determination of case ambiguity. Ref. Formula (6)-(7) wherein $w(case_j)$ is the aggregated noise level of all the risk group predictions made on the target case, and $\theta(ij)$ is an indicator function of whether the prediction of a case agrees with its diagnosis.

$$Amb_{adj}(Case_j) = \frac{w(Case_j)}{\sum_j w(Case_j)} Amb(Case_j) \quad (6)$$

$$w(Case_j) = \sum_j Noise(g_i) \left(\frac{1}{e}\right)^{\theta(ij)} \quad (7)$$

With the aid of the data quality monitor, it is now possible to design a pro-active learner to request for more cases or more features to be included, based on the data monitoring results. Fig. 2 has demonstrated the combination of the risk stratification and the data monitoring mechanism.

Online monitoring of risk profiles

The recommender monitors the quality of the risk predictions using the three different quality metrics, i.e., risk group noise, case ambiguity, noise-adjusted adjustment. Previous experiments in crowd sourcing data quality have shown that selecting annotations based on a simple linear combination of data quality metrics will yield better learner performance than considering each of the metrics alone.

When moving to an online setting, there are new requirements being introduced into the scene. The system needs to collect annotations to establish a pool of cases with ground truth labels. Then, with each risk group found, the quality metrics in the risk model pool will be updated once. Based on the new information, the system needs to determine which case to be selected "next" for analysis and which risk model to annotate the chosen example on the fly.

The above requirements call for an online learning algorithm, which can dynamically update the weightings of the different quality metrics while maintaining the best learner performance on the target risk prediction task. We take a sliding window approach. At time t and $t + 1$, we will extract summary statistics and save them into information vectors $X(t)$ and $X(t+1)$. Given a state pair $\langle y(t-1), y(t) \rangle$, the scoring function is then defined as follows, i.e., Formula (8).

$$g(\langle y_{t-1}, y_t \rangle) = \lambda_1 Q_{noise}(X_t, X_{t+1}) + \lambda_2 Q_{amb}(X_t, X_{t+1}) + \lambda_3 Q_{ambadj}(X_t, X_{t+1}), \quad (8)$$

where the weights of the matrix, $\lambda = \langle \lambda_1, \lambda_2, \lambda_3 \rangle \in \mathfrak{R}_n$, needs to be learned by the system. The task can then be cast as finding the state $Y(t+1)$ that optimizes the scoring function g .

V. PLATFORM SERVICE PRICING

Another personalization-essential common service on the Platform layer is the dynamic provisioning service. "Personalization" allows users to dynamically customize their wellness services (monthly or weekly, depending on granularity). On the one hand, personalization has created a dynamic resource allocation issue, since SaaS providers' demands for platform resources are constantly changing. On the other hand, as the wellness service users come and go and the industry is still in the inception phase, the SaaS providers are expecting increasing but dynamic usage of the services. Therefore, among the pricing schemes, *Pay-as-you-go* (Sandholm, 2002) has become the main one for the open platform services, for that it resonates with the core concept of Cloud computing such as *utility computing* (Sandholm, 2002).

Unfortunately, while multiple SaaS providers are competing for the platform resources, it is often difficult to guarantee Quality of Service (QoS). Imagine the following scenario. Multiple SaaS providers issue virtual machine (VM) resource requests at the same time, and the platform service provider is unable to fulfill the requests due to resource shortage. This failure may cause serious QoS problem if the wellness service is time critical, for instance, an alerting service of hypertensive

conditions.

It is sometimes assumed that Cloud has unlimited resources, but in reality, the resources are bounded. It is important for SaaS providers to allocate enough resources to ensure QoS at SaaS level. Although it is also possible to delegate the responsibility of determining the right quantity of bounded resources to the platform service provider, the complexity of QoS at platform level is often unmanageable due to the wide varieties of parties and services on the platform.

In this study, we assume the responsibilities of ensuring QoS at SaaS level are delegated to SaaS providers and allow them to "lock" platform resources. Therefore, some pricing mechanisms must be included in the platform to support resource locking.

Before performing detailed analysis of the pricing mechanisms, we should revisit the definition of the levels of platform services in Fig. 2 and Fig. 3. The IaaS provider provides VMs to PaaS and SaaS, while PaaS provider provides common wellness services, such as PWR, CDA accelerator, AaaS and monitoring service, to SaaS providers. Each end user basically faces one or more SaaS providers. However, some SaaS providers, like provider *B* in Fig. 3, are capable of delivering services via integrating many other SaaS providers.

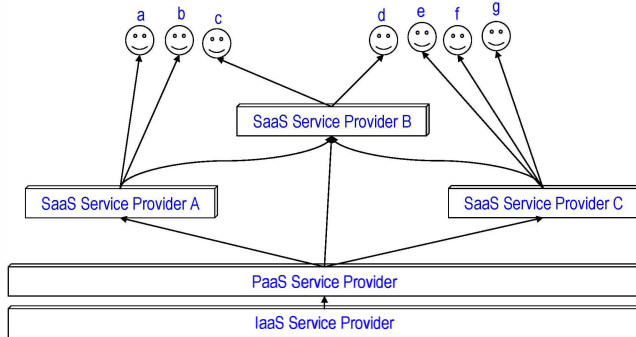


Fig 3. Layers of Service Provision

In sum, the problems to be solved in the context of platform service pricing include:

1. What is the appropriate pricing structure and mechanism for the proposed open platform?
2. In some critical cases, PaaS/IaaS computing resources need to be locked by some SaaS providers to ensure QoS at SaaS level. What is the pricing mechanism for the resources locked?

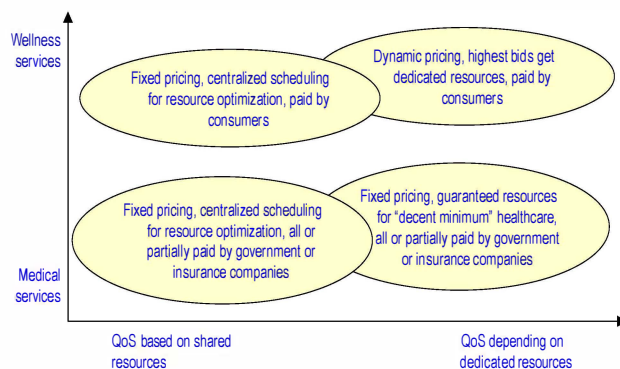


Fig 4. Pricing Structure and Resource Allocation

To justify the pricing mechanism used in the open platform, we must first analyze the context in which pricing is needed. As shown in Fig. 4, payers of the services could be an important factor to distinguish wellness services from medical services. Wellness services are mainly paid by consumers, while medical services are often partially paid by government or insurance companies to guarantee tax payers a "decent yet minimum" healthcare measure. Given that wellness services are mostly self-paid, people willing to pay high prices to wellness services thus are expected to get high QoS. Under this assumption, pricing mechanisms like auctions or automated negotiations could then provide incentive-compatible optimization of pricing and resource allocation (locking). Its consequence is *dynamic pricing*, meaning that the price changes according to the level of demand.

Table 1 SLA Contents

SLA contents (example):	
1.	All events are processed by MaaS within N seconds, compliance rate $R\%$.
2.	All analytic services will respond to requests within N seconds, compliance rate $R\%$.
3.	PWR data warehouse services will respond to requests within N seconds, compliance rate $R\%$.
4.	Each CDA conversion transaction must be completed in N seconds
5.	Each CDA transfer transaction must be completed in N seconds
6.	All allocated VMs are equipped with X gigabytes memory and Y gigabytes disk space.
7.	Platform must guarantee successful VM allocation requests for a total number of Z VMs.
8.	Inbound bandwidth p mega bytes
9.	Outbound bandwidth q mega bytes

Given the number of user subscriptions to a certain SaaS, the SaaS provider could bid for dedicated PaaS/IaaS resources monthly, allowing the PaaS/IaaS provider to commit to an SLA (Service Level Agreement) that ensures preferred QoS at SaaS level. However, SLA negotiations are *multi-issue negotiations* (Klein et al., 2003; Lin and Chou, 2004), and when multiple SaaS providers are negotiating for SLAs monthly, it becomes a complex multi-party multi-issue negotiation. Since auctions often refer to single-issue (price) negotiations, a multi-issue negotiation mechanism is considered much more difficult than auctions. The following SLA contents (Table 1) has shown that multiple variables (issues) need to be negotiated and settled before an agreement can be reached. Some of the issues shown in Table 1 are from PaaS level, and some of the others from IaaS level. In the following paragraphs, we will first analyze the pricing mechanisms at both IaaS level and PaaS level and then discuss the required technology components supporting multi-party multi-issue negotiations.

Table 2 shows available pricing mechanisms at IaaS level. Since the fundamental "products" provided by IaaS are virtual machines, the related QoS issues in Table 1 are issue 6 and 7. If dedicated VM resources are required to ensure QoS at the SaaS level, a dynamic pricing mechanism will be invoked to determine QoS at the IaaS level. The negotiation-based resource optimization approach allows SaaS providers to "lock in" a number of Z base VMs. In other words, once the QoS is agreed, the platform must grant VM allocation requests up to a number of Z. The platform does not guarantee successful VM requests beyond this level. The dynamic allocation of VMs is supported by the Jomo technology described in Cramton et al. (2006).

Table 2 IaaS Pricing Mechanisms

Descriptions	Policies	Pricing policy	Resource allocation policy	Quality assurance
Pricing mechanisms				
Fixed pricing		Fixed price per VM	First come first served	May fail to acquire new VMs
Fixed pricing, centralized resource optimization		Fixed price per VM	Resource requests are queued, high priority VM requests (loading high priority services) served first	May fail to acquire new VMs
Dynamic pricing, negotiation-based resource optimization		Dynamic price for base VMs, fixed price for extra VMs	Resource requests (for base VMs) are queued, resource allocations are optimized by auctions or negotiations	Guaranteed allocation for base VMs, may fail to acquire extra VMs

The platform will support two service levels in the future.

Standard service level: At this service level, all computing resources are shared. The pricing mechanism is pay-as-you-go. However, the platform does not guarantee successful resource allocation. In other words, SaaS providers will experience delays or denial of services during peak periods.

VIP service level: Dedicated computing resources will be allocated to SaaS providers at this service level. Therefore, the system would enforce stricter time and space constraints to ensure QoS at SaaS level. However, the SaaS providers must compete and negotiate for the dedicated computing resources. Prices of the resources are dynamically determined: During off-peak periods, services are priced at the standard service level; the prices can be raised during peak periods.

At the PaaS level, several common services are provided to SaaS providers. Each service can be priced differently with respect to its characteristics. As shown in

Table 3, Personal Wellness Record (PWR) service is used to store personal wellness records, and the price will be determined by the usage of disk space. In addition, data services like the conversion and transmission of CDA and AaaS are charged per use. Finally, MaaS is charged per message. However, there can be extra payments for the same service if SaaS providers request for dedicated computing resources. Dynamic pricing mechanisms are used to ensure optimized pricing and resource allocation.

Table 3 Pricing at PaaS level

	Standard SLA with share resources	VIP SLA with dedicated resources
PWR	Pay per storage unit, fixed price	Pay per storage unit, fixed price Extra payment for dedicated computing power, dynamic pricing
CDA	Pay per use, fixed price	Pay per use, fixed price Extra payment for dedicated computing power, dynamic pricing
AaaS	Pay per use, fixed price	Pay per use, fixed price Extra payment for dedicated computing power, dynamic pricing
MaaS	Pay per message, fixed price	Pay per message, fixed price Extra payment for dedicated computing power, dynamic pricing

One advanced topic in resource-allocation negotiation is *combinatorial negotiation and auction* (Sandholm, 2002; Cramton et al., 2006). For the SaaS providers who are in need of computing resources, they may be willing to pay the premium rate. Take the scenario in Fig. 5 for example, the marketing campaign will only be executed if four week-day computing "shifts" in August are secured. Losing any week-day allocation will risk its QoS and therefore intolerant. If an SaaS provider is willing to pay more than the rest of SaaS providers to secure the resources, the platform should be able to allow the negotiations.

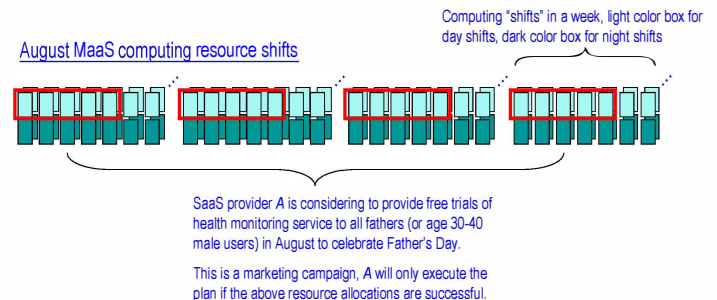


Fig 5. Combinatorial Negotiations/Auctions

Table 4 summarizes five levels of resource optimization, SLA and pricing. It also shows the corresponding technologies to support the pricing mechanism. The ultimate goal of the platform is to optimize and dynamic pricing. Developing technology components for multi-party multi-issue negotiation is especially challenging since it requires the knowledge of *game theory* (Harsanyi and Selten, 1988). The field of research is known as *mechanism design* (Naor et al., 1999; Feigenbaum and Shenker, 2002; Parkes et al., 2008). Since all SaaS providers will devise their best strategies to compete for resources, how the platform elicits private information and aggregates various preferences by designing a "strategy-proof" negotiation mechanism is then the next main research question.

One final issue is how to guarantee end-to-end QoS, especially when multiple parties are involved in providing an integrated service to end users. Since it is assumed that SaaS providers must ensure QoS at SaaS level, the responsibilities of ensuring end-to-end QoS by confirming propagated QoS requirements of collaborating partners are delegated to SaaS integrators (for instance, provider B in Fig. 3).

Table 4 Pricing and Resource Optimization Levels

	Level description (pricing, SLA, and resource optimization)	Technology required
Level 5	Pay-as-you-go, dynamic price for base VMs, supporting VIP SLA, fine-grained combinatorial resource optimization	Jomo, multi-party multi-issue negotiation given fine-grained issues and issue-interdependency
Level 4	Pay-as-you-go, dynamic price for base VMs, supporting VIP SLA, combinatorial resource optimization	Jomo, multi-party multi-issue negotiation given issue-interdependency
Level 3	Pay-as-you-go, dynamic price for base VMs, supporting VIP SLA, negotiation-based resource optimization	Jomo, multi-party multi-issue negotiation
Level 2	Pay-as-you-go, fixed price, standard SLA, centralized resource optimization	Jomo, queuing and scheduling
Level 1	Pay-as-you-go, fixed price, standard SLA, first come first served	Jomo

VI. CONCLUSION & FUTURE WORK

Wellness management ecosystems encompass distinctive types of participants, e.g., end users, health promotion program operators and sponsors, wellness service providers, data providers and researchers. To support the formation and management of personalized services, in this study, we design wellness cloud to operate across three distinctive layers, including infrastructure, platform, and application. It provides virtualization and elastic computing capabilities on the IaaS layer, common services on the PaaS layer, and an ecosystem application mash-up mechanism on the SaaS layer.

One important decision to be made is to what extent the platform should be opened – i.e., structured to encourage participation (Eisenmann et al., 2008). Currently, the GreenOlive platform is designed to open to developers/domain experts. It is expected to spur adoption by harnessing network effects, reducing concerns of lock-in, and encouraging differentiated services that meet the long tail of demand.

In particular, we focus on two open common services: smarter analytics for active personalization and dynamic provisioning. While the former service is taking charge of inferring user wellness risks and making risk-driven recommendations, the latter service determines optimal platform pricing and resource allocation given the constraint of acceptable quality of service. What is the best way to integrate these common services remains as the main challenge to this work. The ultimate goal is to construct a personalization engine that can allow any provider to plug in their analytics model and data, which are likely to be provided by their preferred ecosystem partners, and develop their personalized services that are easily accessible, scalable, and elastic in no time.

REFERENCES

- [1] T. Saaristo, M. Peltonen, S. Keinänen-Kiukaanniemi, M. Vanhala, J. Saltevo, L. Niskanen, H. Oksa, E. Korpi-Hyovalti, J. Tuomilehto and F.-D. S. Group, "National type 2 diabetes prevention programme in Finland: FIN-D2D," *Int. J Circumpolar Health*, pp.101-112.
- [2] Diabetes Prevention Program (DPP), "National Institute of Diabetes and Digestive and Kidney Diseases," US Department of Health and Human Services, 2008.
- [3] M. H. Alderman, "Blood Pressure Management: Individualized Treatment Based on Absolute Risk and the Potential for Benefit," *Annals of Internal Medicine*, vol. 119(4), pp.329-335, 1993.
- [4] J.A. Lee and M.L.A. Lustria, "Measuring effects of web-based, tailored health interventions programs for nutrition and physical activity on clinical outcomes: A systemic review," *ASIS&T Annual Meeting*, 2009.
- [5] W. Kroeze, A. Werkman and J. Brug, "A systematic review of randomized trials on the effectiveness of computer-tailored education on physical activity and dietary behaviors," *Annals of Behavioral Medicine*, pp. 205-223, 2008.
- [6] Juniper Research, "Mobile Healthcare Opportunities- Monitoring, Applications & mHealth Strategies 2010 - 2014," 2010.
- [7] J. F. Moore, "The Death of Competition: Leadership & Strategy in the Age of Business Ecosystems," New York: HarperBusiness, 1996.
- [8] World Economic Forum, "ICT for Economic Growth: A Dynamic Ecosystem Driving The Global Recovery," 2008.
- [9] S. Ried, J. R. Rymer, and R. Iqbal, "Forrester's SaaS Maturity Model: Transforming Vendor Strategy While Managing Customer Expectations", Forrester Research, 2008.
- [10] A. Guazzelli, M. Zeller, W. Chen, and G. Williams. "PMML: An Open Standard for Sharing Models," *The R Journal*, Vol. 1(1), May 2009.
- [11] L. Zeng, P. Hsueh, and H. Chang, "Greenolive: an Open Platform for Wellness Management Ecosystem," *Proceedings of IEEE/INFORMS International Conference on Service Operations and Logistics, and Informatics (SOLI 2010)*, 2010.
- [12] S. Das, "Filters, Wrappers and a Boosting-Based Hybrid for Feature Selection," *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, pp. 74-81, 2001.
- [13] R. Kohavi and G. H. John, "Wrappers for Feature Subset Selection," *Artificial Intelligence*, pp. 273-324, 1996.
- [14] C. Sutton, M. Sindelar, and A. McCallum, "Feature Bagging: Preventing Weight Undertraining in Structured Discriminative Learning," *Center for Intelligent Information Retrieval, U. of Massachusetts*, 2005
- [15] J. O'Sullivan, J. Langford, R. Caruna and A. Blum, "FeatureBoost: A MetaLearning Algorithm that Improves Model Robustness," *Proc. of the 7th International Conference on Machine Learning*, pp. 703-710, 2000.
- [16] P., Hsueh, P. Melville, and V. Sindhwani, "Data Quality from Crowdsourcing: A Study of Annotation Selection Criteria," *NAACL'09: Active Learning Workshop*, 2009.
- [17] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin and M. Zaharia, "Above the Clouds: A Berkeley View of Cloud Computing," 2009, Available: <http://www.csee.usf.edu/~anda/PDS/papers/EECS-2009>
- [18] J. Yang, J. Qiu, and Y. Li, "A Profile-based Approach to Just-in-time Scalability for Cloud Applications," *IEEE International Conference on Cloud Computing*, 2009
- [19] M. Naor, B. Pinkas, and R. Sumner, "Privacy preserving auctions and mechanism design," *Proceedings of the 1st ACM conference on Electronic commerce*, pp.129 – 139, 1999
- [20] J. Feigenbaum and S. Shenker, "Distributed algorithmic mechanism design: recent results and future directions," *Proceedings of the 6th international workshop on Discrete algorithms and methods for mobile computing and communications*, 2002.
- [21] D. Parkes, M. O. Rabin, S. M. Shieber, and C. Thorpe, "Practical secrecy-preserving, verifiably correct and trustworthy auctions," *Electronic Commerce Research and Applications*, vol.7, issue 3, pp. 294-312, 2008.
- [22] R. J. Lin and S. T. Chou, "Mediating a Bilateral Multi-Issue Negotiation," *Electronic Commerce Research and Applications*, vol.3, issue 2, pp. 126-138, 2004
- [23] M. Klein, P. Faratin, H. Sayama and Y. Bar-Yam, "Protocols for Negotiating Complex Contracts," *IEEE Intelligent Systems*, vol. 18, no. 6, pp. 32-38, Nov./Dec. 2003
- [24] T. Sandholm, "Algorithm for optimal winner determination in combinatorial auctions," *Artificial Intelligence*, vol. 135, issue 1-2, pp. 1-54, Feb. 2002.
- [25] P. Cramton, Y. Shoham, and R. Steinberg. "Combinatorial Auctions," Cambridge, MA: MIT Press, 2006
- [26] J. C. Harsanyi and R. Selten. "A General Theory of Equilibrium Selection in Games," the MIT Press book, 1988, Available: <http://ideas.repec.org/b/mtp/titles/0262582384.html>.
- [27] T. R. Eisenmann, G. Parker, G., and M.V. Alstyne. "Opening Platforms: How, When and Why?" Gawer, A. (Ed.) *Platforms, Markets and Innovation*, Edward Elgar: Cheltenham, UK, 2008.