Cloud Computing and Software as a Service in BMI

Timoteus Ziminski

tbz@engr.uconn.edu

May 7, 2011

1 Introduction

The term "cloud computing" is one of the currently most frequently used buzz words and most controversially discussed topics in the world of information technologies. While we agree that a good portion of critique regarding the manner in which cloud computing is marketed is certainly not out of place (not convinced readers are invited to take a brief excursion through the sales brochures of the countless cloud solution vendors) and while we equally realize that to a certain extent the cloud is a re-branding of distributed computing concepts that have been present for years, we also argue that it is a field that spans over a large variety of useful applications and that has significantly matured in the past years.

The scope of this paper is twofold; the first part explores the concept of cloud computing by providing a definition, discussing the main characteristics of cloud computing and giving an overview of the cloud computing solution offers of major providers on the market. Furthermore, we discuss technological advances that have increased the usability of the concept in the past years. The second part is centered around Biomedical Informatics (BMI), a highly complex research area which aims to facilitating the management and usage of biomedical data, information and knowledge in areas such as clinical research, public health, biomedicine or clinical care. BMI includes computational application fields which are inherently rich in facets and present various challenges of technical and organizational nature. Our goal is to identify a selection of those challenges and to explore how strengths of cloud computing can be leveraged for implementing related solutions.

2 Cloud Computing

In this section we introduce a definition and core characteristics of the cloud computing paradigm, as well as discuss some popular platforms offered by well known IT solutions vendors. Furthermore, we explore the technical foundation of the concept, namely advances in hardware, software and network technologies.

2.1 What is the Cloud?

Many definitions of cloud computing can be found [1] at different degrees of formalism and detail. For the purpose of this paper we will refer to the comparatively sophisticated definition that is provided by the National Institute of Standards and Definitions (NIST): "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."[2] What we want to summarize from this definition is that at the core of cloud computing is the idea of abstract resources that hide computational details and complexity from their users.

Although the cloud computing paradigm is still considered to be subject to transformations by some sources [3], several very mature and usable platforms became apparent in recent time. Once cloud services are established they become a product that is easy to offer to customers, which supports the general interest in of the paradigm.

2.2 Providers

In this section we explore the the cloud computing portfolios of some major information technology providers. We only discuss a small selection of available products without claiming completeness. As we go through the applications available on the market, the goal is to develop an understanding of the different types of cloud computing applications, which we will then subsequently categorize.

2.2.1 Google

Google has been a strong advocate of cloud computing applications for over ten years, often impressively demonstrating what can be implemented in terms of comfortable user interfaces in a purely browser based setting. Examples of popular applications by Google are the office tools Google Docs, Google Calendar and Google Reader. Noteworthy communication and collaboration tools are Gmail, Postini, Google Groups and the currently retired project Google Wave.

Finally, the company also offers Google Apps, a framework that allows developers and system administrators to create a common work environment based on the above applications. The set of available applications can be extended either in-house or through the Google Apps Marketplace. Google Apps is available in several payed and free versions.

Of particular interest to the BMI scope of this paper is the Personal Health Record (PHR) Google Health, which since 2008 allows users to consolidate health care and wellness information from disparate source into one record. The product is exposing an API with support for clients in multiple programming languages and uses a subset of the Continuity of Care Record (CCR) as data format. Those properties make Google Health an accessible and worthwhile study object for the integration of health care applications. An overview of most of Google's cloud services can be found at [4], Google Health is accessible at [5].

2.2.2 Amazon

As an extension to its core business, the online vendor Amazon has been very active in the field of cloud computing since 2002. Under the name Amazon Web Services (AWS), the company offers a collection of cloud applications of which Amazon Elastic Cloud (EC2) and Amazon Simple Storage System (S3) are the most popular ones.

EC2 is a virtualization solution offering customers to rent a virtual computer (an "instance") which can be equipped with a variety of operating systems and essentially serve for any required computational task. The Amazon S3 service is a cloud file storage and distribution system supporting multiple communication protocols. An overview of Amazon's cloud services can be found at [6].

2.2.3 IBM

Another relevant actor on the cloud computing market is IBM. Next to a whole series of business focused applications such as SmartCloud Services, LotusLive and Blueworks, the company also particularly focuses on some cloud based BMI applications. One example is the Collaborative Care Electronic Medical Record (EMR) and decision supports system, which has been already deployed into real life clinical care environments. Another project of interest is the attempt to leverage the Watson artificial intelligence technology for clinical decision support, in which the challenging storage of data will be approached with cloud computing solutions [7].

Information about IBM's cloud computing portfolio can be found at [8], resources about the Collaborative Care EMR can be retrieved from [9].

2.2.4 Microsoft

Microsoft is also very engaged in the cloud computing market with a wide range of products; some examples are Windows Azure, which offers a cloud application development and hosting environment, WindowsServer Hyper-V, which is a server virtualization service, and Office 365, a cloud version of the popular Microsoft Office Suite. The company also runs the PHR Microsoft HealthVault, which can be particularly interesting for a setting where the integration of health and fitness related devices, such as glucose meters or exercise treadmills, is of importance.

An overview of Microsoft's cloud services can be found at [10], the HealthVault service is accessible at [11].

2.2.5 Importance to IT Market

The discussion of cloud computing is so vivid not least because the estimates of its potential monetary impact on businesses are quite significant depending on which sources one consults. Although the numbers tend to vary and the situation is complicated by not all companies reporting their outcomes in relation to cloud computing, there is a general agreement on predicting at least a basic significance in the years to come. Some available examples are:

- Citigroup estimates Amazon will generate $650 million revenues through AWS in 2011[12].
- IBM intends to generate revenues of $7.0 billion from cloud computing by 2015 (over 5% of total revenues)[13].
- Microsoft will be spending 90% of its $9.6 billion research and development budget on cloud strategies in 2011[14].
- Gartner predicts the global cloud market to grow to$102.1 billion net this year from $68.3 billion in 2010.

2.3 Characteristics

The list of services presented in the previous section is far from complete, since we only focused on the most popular services provided the major companies. Even on the contrary, there exists an enormous range of products covering a variety of tasks such as collaborative software development and testing, professional training, file archiving and data analysis, collaborative document processing, and a large number of multi media applications.

However, for the purpose of our next task the above list is sufficient, as we want to explore the different cloud service models. We can easily identify three style in the set of presented services regarding the level of infrastructural detail they expose: while several of them to some extent behave the same way a standalone desktop software installation would behave (e.g., Google Calendar, Office 356), others are providing a more complete framework that allows embedding and configuration of user specific applications (e.g., Google Apps, Azure), and again other are exposing a virtual image of hardware to the user (e.g., AWS, Hyper-V). Based on this difference in detail exposure, the NIST cloud computing definition [2] distinguishes between the following cloud computing service models with the noted key functionalities:

- Software as a Service (SaaS)
  - Provides applications
  - Enables cross-device and cross-platform capabilities
- Platform as a Service (PaaS)
  - Access to the cloud hosting environment
  - Deployment of applications
- Infrastructure as a Service (IaaS)
  - Provides access to basic resources
  - CPU cycles
  - Network connections
  - File storage space

A fourth model, Data as a Service (DaaS), is mentioned in some sources [1], although it can conceptually be incorporated in the three core models defined above.

Another important difference between cloud computing services is the deployment model [2]. Cloud computing does not necessarily mean that resources are moved outside of organizational boundaries. More concretely, the paradigm does not only support the idea of a public cloud, which will be typically reachable through the Internet and pursue a business model of some nature, but also copes well with the notion of private clouds that are hosted in local data centers and are accessible only through organizational intranets. Furthermore, one can also imagine community clouds that span over multiple organizations or hybrid clouds that interconnect multiple clouds and have and enforce mixed forms of access control on their resources.

Finally we can recognize key a set of characteristics of cloud computing services as such [2]: cloud computing typically provides on-demand self service, which means that adjustments to the service (e.g., registration, configuration, extension) are mostly automatic and can therefore be executed programmatically. Service will usually depend on broad network access through standard mechanisms (e.g., Web Services). This requires on one hand a constantly available network connection with sufficient quality of service, on the other hand it also opens the door to using very thin hardware clients.

Resource pooling plays an important role in the background of a typical cloud computing service. However, the user has in most cases no control of any physical resource instances and is protected from the complexity of that the process. The complexity can be in fact very significant if one takes into account that cloud services operate under rapid elasticity terms, meaning that they present their resources as unlimited and allow the user to grow and shrink capacities according to requirements that potentially change in real time. This is a key feature of cloud computing, since the paradigm was born from the need to cater towards application fields with very high punctual peaks in infrastructural requirements (e.g., web stores on Christmas or Black Friday). To match up with the elasticity feature, cloud services are often implementing a form of service measurement (e.g., "gigabytes transferred", "IPs allocated"), that allows to bill customers only for what they exactly use. This directly means that taking advantage of cloud services requires little up-front capital.

2.4 Technologies

After exploring the conceptual details of cloud computing in the previous sections, we now analyze a set of technological progress aspects that enabled cloud computing to show significant improvements in the past years. This is of particular interest to our overview of the field, since it is our belief that the "re-branding" of distributed computing concepts to the cloud computing paradigm and the increased effort of marketing it was an effect of reaching a threshold in usability caused by those aspects. We identify progress in three areas: hardware, networking, and software.

2.4.1 Progress in Hardware Technologies

As a result of the topics discussed in Section 2.2.5, cloud computing calls for reliable, cost-efficient, and powerful data centers. Certainly the general progress in hardware technologies (i.e., availability of larger and faster disc space and memory as well as improvements on energy efficiency) is weighting in positively on this aspect. However, as more important factor we identify the developments in hardware supported virtualization, hardware multiplexing architectures, and clustering techniques that cater to the scalable nature of cloud computing.

On the side of client hardware we observe a rapidly growing popularity of very thin yet sufficiently smart devices such as mobile phones, tablet computers, and sub-notebooks. Scenarios where multiple of those devices are being used by one person become very common and create additional synchronization issues that can be addressed by data storage in the cloud. While not in the scope of our discussion, we also want to want to point out the increasing importance of those devices in clinical care and refer the reader to [15] for additional information on the topic.

2.4.2 Progress in Networking Technologies

Applications that make use of cloud services are dependent on a fast and reliable connections to the providing data centers. When the link to the cloud is disconnected, their functionality is limited in the best case and completely disturbed in the worst. However, networking technologies have over the last years reached the point of allowing even mobile broadband connections (e.g., 4G networks) which allows the assumption that we are well capable of building sufficiently performance networks to support cloud computing.

2.4.3 Progress in Software Technologies

In regards of software technologies we identify progress on multiple fields as relevant to the increasing usability and importance of cloud computing. First, as noted above, virtualization plays an important role on the with in the cloud data centers. Therefore, the development and increasing performance of virtualization software products such as VMware, VirualBox, Virtual PC, and Xen has created the prerequisites of building those centers.

Second, from an organizational perspective, Service-oriented architectures (SOA) have promoted the decomposition of IT processes and their encapsulation into loosely coupled services with defined interface interfaces as presented in Figure 1. The implementation of an SOA (with technologies such as REST, Web Services or CORBA), creates the potential for replacement of locally bound services with options from the cloud. Third,

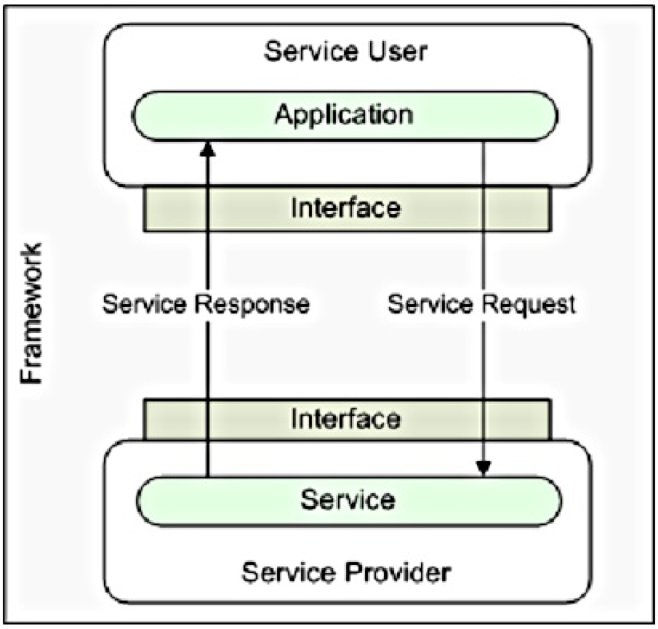

Figure 1: Service Interaction in a Service-oriented Architecture.

from the perspective of front-ends, cloud application developers can leverage increasingly sophisticated web browsers (e.g., IE9, Firefox 4, Chrome, Safari 5) that fully support features such as HTML5, JavaScript/AJAX, Flash, and Silverlight. Furthermore, the operating systems that are being shipped with mobile devices such as Android, iOS, and Chromium OS are natively aware and supportive of cloud applications.

2.5 Overview

We finish our background-focused discussion of cloud computing by providing the high level conceptual schema in Figure 2. In summary, the cloud consists of services of different types (SaaS, Iaas, PaaS), that are provided in an abstracted and complexity hiding way to a large number and variety of clients. We conclude that although it is not living up to the commercial hype around it, the cloud computing paradigm nonetheless represents a powerful tool that can be leveraged for various architectural scenarios.

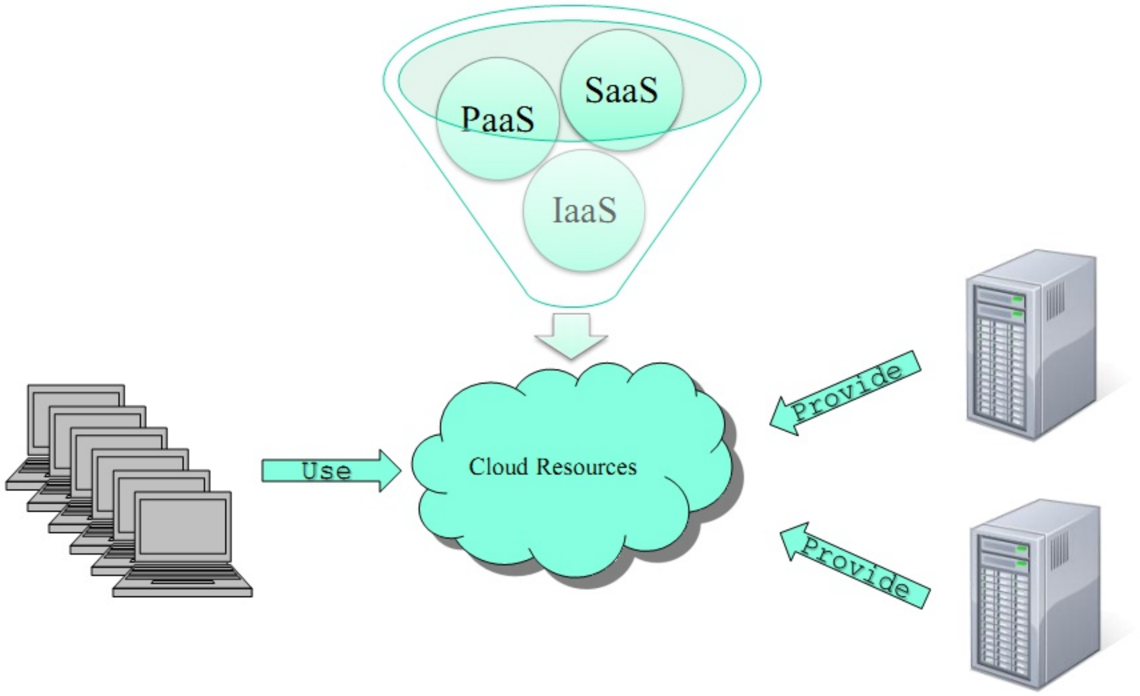

Figure 2: Schematic Overview of the Cloud Computing Paradigm.

3 BMI in the Cloud

In this section we aim to relate cloud computing to the Biomedical Informatics domain. In order to achieve this, we chose the following perspective for the remainder of this section: given the properties identified in Section 1 and an understanding of cloud computing as a tool suitable for solving certain types of computational problems, we traverse through a set of BMI application fields, explore their typical challenges and discuss how they could be addressed with cloud computing.

3.1 Application Fields

We explore the four BMI disciplines defined and extensively discussed by [16] in our analysis: Imaging Informatics, Clinical Informatics, Bioinformatics, and Public Health Informatics. Furthermore, we add Clinical Research Informatics and Health Care Training to the list, as we found those areas to be of relevance while studying the topic. Figure

3 gives an overview of the discipline set that has served as the outline for our work. For additional details on each of the areas we refer the reader to [16].

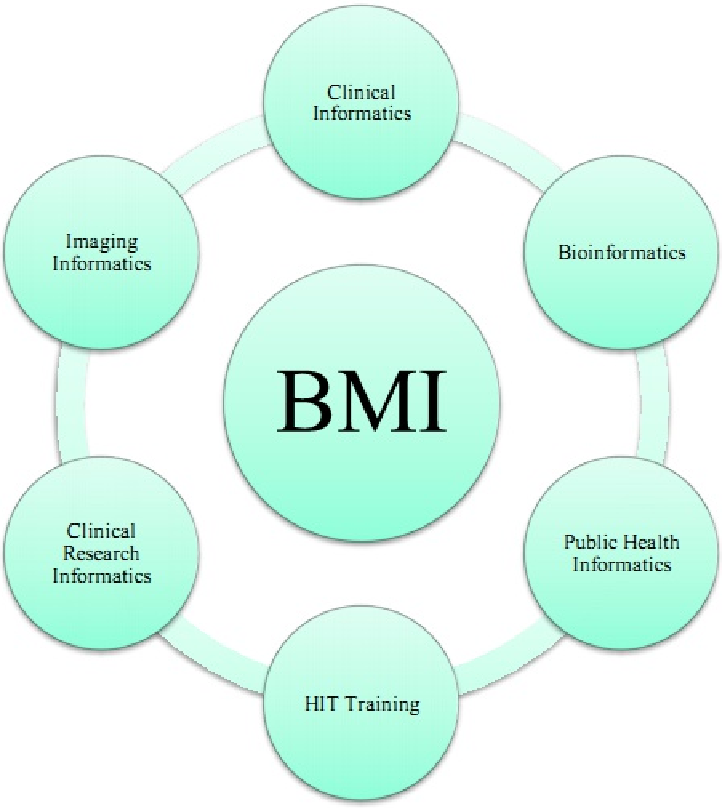

Figure 3: Biomedical Informatics Disciplines

3.1.1 Clinical Informatics

The field of Clinical Informatics is facing significant challenges in efficiently documenting and organizing provision of care to patients in all possible settings (e.g., inpatients, outpatients, supported living patients, patients with chronic conditions). Patient records in electronic form, such as Electronic Medical Records (EMR) or Personal Health Records (PHR) have been identified as valuable tools for this task. However, as we think of implementing those solutions on a national scale, those approaches come with an inherit need for efficient and reliable large-scale data storage and exchange.

Another big topic in Clinical Informatics is enabling collaboration between health care professionals (e.g., for tasks such as team-based collaboration within one point of care or transition of care between multiple points of care), health care organizations (e.g., hospitals, clinics, practices, labs, pharmacies, nursing homes), and, on a more abstract level, between whole health care domains (e.g., providers, payers, researchers).

When looking for cloud computing solutions for those tasks, we find that there are already some full-fledged applications on the market such as Google Health or Microsoft HealthVault (see Section 2.1). Regarding EMRs, we also begin to see some serious initiatives such as the Collaborative Care project (see Section 2.2.2) We expect further development in the cloud EMR sector, since the model comes with some significant benefits.

First, as we have seen in Section 1 cloud computing platforms typically foster modular approaches, which brings configurable plug-in architectures with of support various extensions (e.g., adapters for health information exchange, decision support, note processing and voice recognition, patient education) into scope. Second, we can imagine deployment models that streamline the implementation of health information exchange by installing regional clouds and enabling communication through hybrid clouds. While the paradigm itself does not solve the interoperability issues involved with this topic, we would expect that it can be used a vehicle to promote related solutions. Third, cloud EMRs could be used to efficiently implement the application service provider model [17] which is important in the context of implementing EMRs in small practices that do not have the financial means to hire technical technical know-how for software and hardware maintenance.

Yet another example from the Clinical Informatics discipline is data management and data analysis of laboratory data (e.g., from a hospital pathology lab). Currently there is a lack of affordable off-the shelve solutions for this task, so that in-house solutions are created and potentially cause new data silos. Since cloud computing is a strong paradigm for in this context, it could be used to provide affordable generic solutions. We conclude

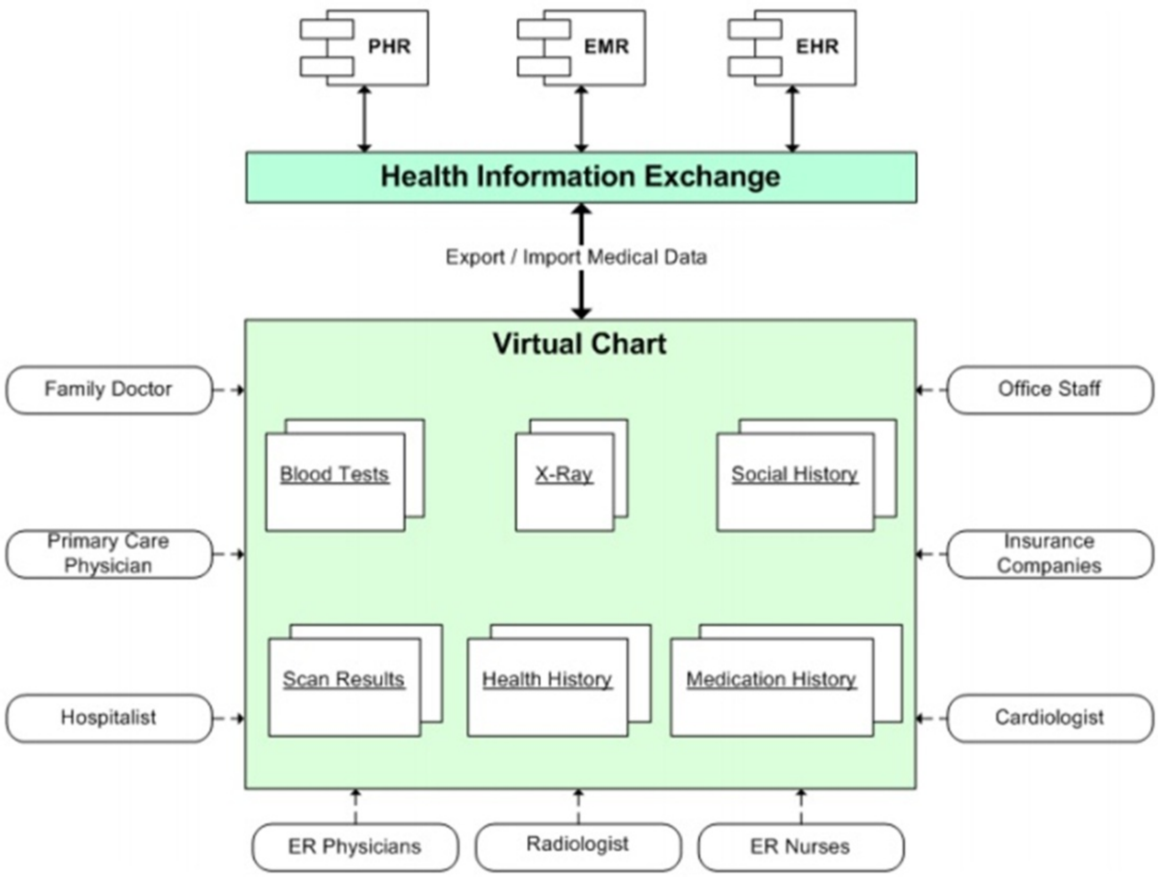

Figure 4: Virtual Patient Chart.

our discussion of Clinical Informatics with a look at one of its most challenging concepts, the virtual patient chart (see Figure4), which organizes the collaboration between the disparate actors of clinical care in a virtual distributed document. Although it is a frequently discussed topic, we so far do not have a convincing idea of how build it. Following thoughts come into mind when looking at the challenges of the virtual chart in the context of cloud computing:

- Cloud computing services can provide very abstract storage resources
- Space and performance are presented as limitless and are automatically adjustable
- Data is stored in a highly accessible fashion
- We are starting to have experience real-life proof platforms for real-time collaboration on documents

Although those points are certainly worth exploring, we will discuss in Section 3.1.5 why cloud computing is not the silver bullet solution to the problem.

3.1.2 Bioinformatics and Imaging Informatics

The field of Bioinformatics supports areas such as DNA sequencing, genomics, protein analysis, and mutation analysis. In addition to very specific hardware for those tasks, it mainly requires enormous storage space (up to the scale of petabytes) and CPU cycles.

Institutions wanting to conduct Bioinformatics research have to face the problem of hiring staff that is capable of managing data of such dimensions, since even implementing backup solutions becomes a challenging operation in such an environment. We again see potential to approach the problem with cloud computing, namely accumulating the data processing know-how at service providers and offering the necessary functionalities with a “pay as you go” model to interested institutions. The goal is to remove the burden of handling the data sets from the research teams, so that they can focus on their research rather than worrying about maintenance tasks. The issue of providing CPU cycles as resource is more related to the filed of Grid computing and out of scope of our analysis, however we note that cloud services could be used to increase the discoverability of suitable Grids.

Due to very similar requirements, the same thoughts apply to the field of Imaging Informatics, so that we do not discuss the area for itself at this point.

3.1.3 Public Health Informatics

The discipline of Public Health Informatics aims to facilitate the tracking of health trends in the population on different levels (i.e., regional level, state-wide level, nation-wide level). The goal is to gain knowledge that can be used for disease control and prevention as well as to organize emergency response procedures (e.g., in the case of an pandemic or a natural disaster).

The main challenges in this this area are the collection, storage and analysis of public health data (i.e., large scale data warehousing). While we repeat our argument of cloud computing being very suitable for data management regarding the two later tasks, we see a form of chain reaction being beneficial for the first one: the implementation cloud based virtual charts could be further leveraged for Public Health Informatics by implementing data collection through inter-cloud communication.

3.1.4 Clinical Research Informatics

For achieving their goals, researches in clinical research require extensive access to repositories and databases holding patient data and study outcomes in suitable formats. They will then typically either perform meta-studies (i.e., studies of studies) or attempt to build patient cohorts suitable for clinical trials. Data acquisition, data de-identification, data sharing and data querying are challenges for building and managing those repositories in the first place. Once the trial stage is reached, data collection, data management and finally data analysis become relevant tasks.

Cloud computing could be used here to outsource the responsibility for regulation conform de-identification and data processing. Again we see the opportunity to have technical and regulatory knowledge accumulate on the provider side, where the cloud service not only support the main target, but also become a motor for the promotion of best practices.

3.1.5 Health Care Training

Training is an important and often underestimated component within BMI. Challenges include to adequately train health care professionals and students of related fields. This includes components such as providing medical education (e.g., transfer of new knowledge, adherence to best practices) or preparing for the usage of software systems (e.g., EMRs, Practice Management Systems, e-Prescribing Systems).

The virtualization level that can be achieved by cloud computing turns has proven to be a very powerful tool in designing scalable, flexible and cost efficient training environments. The main advantage here is the "disposable" nature of the virtual platforms, which can be leveraged to simulate a large amount of situations with minimal effort. Furthermore, training can be provided through thin clients (i.e., independently of the exact nature of available physical workstations), which helps service providers to minimize their cost and simplifies maintaining a large course portfolio.

3.2 Limitations

As we discussed in Section 2.5, cloud computing can be used for approaching BMI challenges, but it is naturally not an universal solution to all encountered problems. Quite contrary, cloud computing and SaaS are even associated to some risks and limitations that we need to recognize in the context of BMI applications. Namely, we identify the following problematic areas: trust and security concerns, reliability and availability risks, performance issues, lack of standardization, and communication complexity. The remainder of this section provides details on those areas.

3.2.1 Trust and Security

BMI applications typically make use of very sensitive data that is subject to many restrictive policies. When using third party services, we have to be aware that data is being loaded outside of organizational boundaries, which is by definition a security risk. Even worse, it is likely to be simply forbidden by privacy regulations. Not only would loading data into a third party cloud mean that it becomes available to this party, but also that it can potentially be stored anywhere in the world, since cloud computing does not assume any limitations regarding physical locations [18]. Algorithms that aim to overcome those issues by suitable encryption, are currently still experiencing performance issues. Furthermore, we have to recognize that access control concepts are not a natively supported concept of cloud computing. Although nothing prevents service providers to implement models such as role based access control (RBAC), it is not something that comes without effort.

3.2.2 Reliability and Availability

BMI systems are frequently of critical importance to patients, particularly clinical care systems can directly influence decisions that have potentially crucial impact on the success of treatment and the well-being of patients. Giving the responsibility for the systems to third parties also means to a accept a risk of system outages and performance problems. Although this issue can equally occur with in-house data centers, one has to be aware that cloud computing is not fail proof, as the recent outage of Amazon Web Services has demonstrated vividly [19]. For many BMI settings this is not an acceptable scenarios, and while we can certainly think of options how to overcome those limitations (e.g., local caching and fail-overs), we will in this case immediately sacrifice the low complicity we hoped to gain from cloud computing.

3.2.3 Performance

Although cloud computing has demonstrated some very mature applications with smart front-ends and although low latency Internet can usually be assumed for the application environment, cloud applications remain vulnerable to potential fluctuations in quality of service. In a domain such as clinical care with hard time limitations within the daily workflows, this can have a harsh impact on the acceptance of the applications.

3.2.4 Standardization

Cloud Computing does not give any immediate advantages regarding interoperability since data models, standards, vocabularies, and ontologies are a completely separate issue. Even worse, there are no binding standards for cloud interfaces, which means the risk of vendor-lock. There are currently multiple national and international organizations [20] working on establishing standards, but an universal agreement is not in sight.

3.2.5 Overhead

Technologies such as Web Service and SOAP RPC which are used in cloud computing might be still too heavy-weight for building a collaborative exchange network that is aiming to span over the nationwide health care system. From this ultra large scale system [21] perspective we might need to rely on even more abstract concepts such as HTTP/S REST.

4 Conclusion

In this paper with have explored the cloud computing paradigm, its key characterises, and its main service models. We presented a selection of currently available applications and discussed the technologies related to the field. Once we established an understanding of the paradigm as an architectural tool and described its capabilities, we identified challenges from BMI disciplines that could be approached by cloud computing. In this process we discovered that cloud computing is a model that conceptually aligns with many problems encountered in BMI, so that it can be regarded as a worthwhile study object. However, we also recognized that it is prone to multiple serious limitations that need to be taken into account when approaching practical problems. Additional research topics of interest split roughly into two fields. On the one hand there is work attempting to improve cloud concepts (e.g., implementation of trust and security, improvement of scheduling and virtualization, enabling of mobile could computing, and inter-cloud communication). On the other hand, there are many project working on leveraging cloud computing for tasks such as data analysis and management, data modeling, text recognition, and voice recognition. A good entry point to explore the topics can be found at [22].

References

[1] L. Wang, J. Tao, M. Kunze, AC Castellanos, D. Kramer, and W. Karl. Scientific cloud computing: Early definition and experience. In *Proceedings of the 10th IEEE International Conference on High Performance Computing and Communications, Austin, TX*, pages 825–830, 2008.

[2] Peter Mell and Timothy Grance. The NIST Definition of Cloud Computing (Draft) Recommendations of the National Institute of Standards and Technology. *Nist Special Publication*, 145:1–2, 2011.

[3] R. Buyya, C.S. Yeo, S. Venugopal, J. Broberg, and I. Brandic. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6):599–616, 2009.

[4] Google Products. `http://www.google.de/options/`.

[5] Google Health. `http://www.google.com/health`.

[6] Amazon AWS. `http://aws.amazon.com/`.

[7] IBM. IBM to collaborate with Nuance to apply IBM's "Watson" analytics technology to healthcare. `http://www-03.ibm.com/press/us/en/pressrelease/33726.wss`.

[8] IBM Cloud Computing. `http://www.ibm.com/cloud-computing/`.

[9] IBM. Sharp Community Medical Group taps IBM and Activehealth Collaborative Care solution to help doctors deliver better patient care. `http://www-03.ibm.com/press/us/en/pressrelease/32278.wss`.

[10] Microsoft Cloud Power. `http://www.micrsoft.com/cloud`.

[11] Microsoft HealthVault. `http://www.healthvault.com`.

[12] Andrew R Hickey. Amazon Cloud Revenue Could Exceed $500 Million In 2010: Report. *CRN*, 2010.

[13] Seeking Alpha. Smart Cloud to Boost IBM's Cloud. `http://seekingalpha.com/article/262389-smart-cloud-to-boost-ibm-s-cloud`.

[14] Aydan Eksi. Microsoft says to spend 90% of R&D on Cloud Strategy. *Bloomberg*, 2011.

[15] John D. Halamka MD. Cool technology of the week. `http://geekdoctor.blogspot.com/2011/04/cool-technology-of-week_15.html`.

[16] E.H. Shortliffe and J.J. Cimino. *Biomedical Informatics: Computer Applications in Health Care and Biomedicine*. Springer Verlag, 2006.

[17] K. Wager and P Lee, F. amd Glaser. *Health Care Information Systems: A Practical Approach for Health Care Management*. Jossey-Bass, 2009.

[18] T. Velte, A. Velte, T.J. Velte, and R.C. Elsenpeter. *Cloud computing: a practical approach*. McGraw-Hill Osborne Media, 2009.

[19] Dan Mitchell. Amazon's cloud nightmare. *CNNMoney.com*, 2011.

[20] Cloud Standards Wiki. `http://cloud-standards.org/wiki/`.

[21] P. Feiler, R.P. Gabriel, J. Goodenough, R. Linger, T. Longstaff, R. Kazman, M. Klein, L. Northrop, D. Schmidt, K. Sullivan, et al. Ultra-large-scale systems: The software challenge of the future. *Software Engineering*, 2006.

[22] National Science Foundation. National science foundation awards millions to fourteen universities for cloud computing research. `http://www.nsf.gov/news/news_summ.jsp?cntn_id=114686`.