

Gerard Briscoe and Alexandros Marinos Community cloud computing

Conference paper

Original citation:

Originally presented at First International Conference on Cloud Computing, 1-4 December 2009, Beijing, China.

This version available at: <http://eprints.lse.ac.uk/26516/>

Available in LSE Research Online: December 2009

Funding provided by European Union, Open Philosophies for Associative Autopoietic Digital Ecosystems (OPAALS)

© 2009 Gerard Briscoe & Alexandros Marinos

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Community Cloud Computing

Alexandros Marinos
Department of Computing
University of Surrey
United Kingdom
e-mail: a.marinos@surrey.ac.uk

Gerard Briscoe
Department of Media and Communications
London School of Economics and Political Science
United Kingdom
e-mail: g.briscoe@lse.ac.uk

Abstract— Cloud Computing is rising fast, with its data centres growing at an unprecedented rate. However, this has come with concerns over privacy, efficiency at the expense of resilience, and environmental sustainability, because of the dependence on Cloud vendors such as Google, Amazon and Microsoft. Our response is an alternative model for the Cloud conceptualisation, providing a paradigm for Clouds in the community, utilising networked personal computers for liberation from the centralised vendor model. Community Cloud Computing (C3) offers an alternative architecture, created by combing the Cloud with paradigms from Grid Computing, principles from Digital Ecosystems, and sustainability from Green Computing, while remaining true to the original vision of the Internet. It is more technically challenging than Cloud Computing, having to deal with distributed computing issues, including heterogeneous nodes, varying quality of service, and additional security constraints. However, these are not insurmountable challenges, and with the need to retain control over our digital lives and the potential environmental consequences, it is a challenge we must pursue.

Index Terms— Cloud Computing, Community Cloud, Community Cloud Computing, Green Computing, Sustainability.

I. INTRODUCTION

The recent development of Cloud Computing provides a compelling value proposition for organisations to outsource their Information and Communications Technology (ICT) infrastructure [1]. However, there are growing concerns over the control ceded to large Cloud vendors [2], especially the lack of information privacy [3]. Also, the data centres required for Cloud Computing are growing exponentially [4], creating an ever-increasing *carbon footprint* and therefore raising environmental concerns [5], [6].

The distributed resource provision from Grid Computing, distributed control from Digital Ecosystems, and sustainability from Green Computing, can remedy these concerns. So, Cloud Computing combined with these approaches would provide a compelling socio-technical conceptualisation for sustainable distributed computing, utilising the spare resources of networked personal computers collectively to provide the facilities of a virtual *data centre* and form a Community Cloud. Therefore, essentially reformulating the Internet to reflect its current uses and scale, while maintaining the original intentions [7] for sustainability in the face of adversity. Including extra capabilities embedded into the infrastructure which would become as fundamental and invisible as moving packets is today.

II. CLOUD COMPUTING

Cloud Computing is the use of Internet-based technologies for the provision of services [1], originating from the *cloud* as a metaphor for the Internet, based on depictions in computer network diagrams to abstract the complex infrastructure it conceals [8]. It can also be seen as a commercial evolution of the academic-oriented Grid Computing [9], succeeding where Utility Computing struggled [10], [11], while making greater use of the self-management advances of Autonomic Computing [12]. It offers the illusion of infinite computing resources available on demand, with the elimination of upfront commitment from users, and payment for the use of computing resources on a short-term basis as needed [3]. Furthermore, it does not require the node providing a service to be present once its service is deployed [3]. It is being promoted as the cutting-edge of scalable web application development [3], in which dynamically scalable and often virtualised resources are provided as a service over the Internet [13], [1], [14], [15], with users having no knowledge of, expertise in, or control over the technology infrastructure of the Cloud supporting them [16]. It currently has significant momentum in two extremes of the web development industry [3], [1]: the consumer web technology incumbents who have resource surpluses in their vast *data centres*¹, and various consumers and start-ups that do not have access to such computational resources. Cloud Computing conceptually incorporates Software-as-a-Service (SaaS) [18], Web 2.0 [19] and other technologies with reliance on the Internet, providing common business applications online through web browsers to satisfy the computing needs of users, while the software and data are stored on the servers.

Figure 1 shows the typical configuration of Cloud Computing at run-time when consumers visit an application served by the central Cloud, which is housed in one or more data centres [20]. *Green* symbolises resource *consumption*, and *yellow* resource *provision*. The role of *coordinator* for resource provision is designated by *red*, and is centrally controlled. Even if the central node is implemented as a distributed grid, which is the usual incarnation of a data centre, control is still centralised. Providers, who are the controllers, are usually companies with other web activities that require large computing

¹ A *data centre* is a facility, with the necessary security devices and environmental systems (e.g. air conditioning and fire suppression), for housing a *server farm*, a collection of computer servers that can accomplish server needs far beyond the capability of one machine [17].




Figure 1. Cloud Computing: Typical configuration when consumers visit an application served by the central Cloud, which is housed in one or more data centres [20]. Green symbolises resource consumption, and yellow resource provision. The role of coordinator for resource provision is designated by red, and is centrally controlled.

resources, and in their efforts to scale their primary businesses have gained considerable expertise and hardware. For them, Cloud Computing is a way to resell these as a new product while expanding into a new market. Consumers include everyday users, Small and Medium sized Enterprises (SMEs), and ambitious start-ups whose innovation potentially threatens the incumbent providers.

A. Layers of Abstraction

While there is a significant *buzz* around Cloud Computing, there is little clarity over which offerings qualify or their interrelation. The key to resolving this confusion is the realisation that the various offerings fall into different levels of abstraction, as shown in Figure 2, aimed at different market segments.

1) *Infrastructure-as-a-Service (IaaS)* [21]: At the most basic level of Cloud Computing offerings, there are providers such as Amazon [22] and Mosso [23], who provide *machine instances* to developers. These instances essentially behave like dedicated servers that are controlled by the developers, who therefore have full responsibility for their operation. So, once a machine reaches its performance limits, the developers have to manually instantiate another machine and scale their application out to it. This service is intended for developers who can write arbitrary software on top of the infrastructure with only small compromises in their development methodology.

2) *Platform-as-a-Service (PaaS)* [24]: One level of abstraction above, services like Google App Engine [25] provide a programming environment that abstracts machine instances and other technical details from developers. The programs are executed over data centres, not concerning the developers with matters of allocation. In exchange for this, the developers have to handle some constraints that the environment imposes on their application design, for example the use of *key-value stores*² instead of *relational databases*.

3) *Software-as-a-Service (SaaS)* [18]: At the consumer-facing level are the most popular examples of Cloud Computing, with well-defined applications offering




Figure 2. Abstractions of Cloud Computing: While there is a significant buzz around Cloud Computing, there is little clarity over which offerings qualify or their interrelation. The key to resolving this confusion is the realisation that the various offerings fall into different levels of abstraction, aimed at different market segments.

users online resources and storage. This differentiates SaaS from traditional websites or web applications which do not interface with user information (e.g. documents) or do so in a limited manner. Popular examples include Microsoft's (Windows Live) Hotmail, office suites such as Google Docs and Zoho, and online business software such as Salesforce.com.

To better understand Cloud Computing we can categorise the roles of the various actors. The *vendor* as resource provider has already been discussed. The application *developers* utilise the resources provided, building services for the *end users*. This separation of roles helps define the stakeholders and their differing interests. However, actors can take on multiple roles, with *vendors* also developing services for the *end users*, or *developers* utilising the services of others to build their own services. Yet, within each Cloud the role of provider, and therefore controller, can only be occupied by the *vendor* providing the Cloud.

B. Concerns

The Cloud Computing model is not without concerns, as others have noted [29], [3], and we consider the following as primary:

1) *Failure of Monocultures*: The uptime³ of Cloud Computing based solutions is an advantage, when compared to businesses running their own infrastructure, but often overlooked is the co-occurrence of downtime in vendor-driven *monocultures*. The use of globally decentralised *data centres* for vendor Clouds minimises failure, aiding its adoption. However, when a cloud fails, there is a cascade effect crippling all organisations dependent on that Cloud, and all those dependent upon them. This was illustrated by the Amazon (S3) Cloud outage [31], which disabled several other dependent businesses. So, failures are now system-wide, instead of being partial or

² A distributed storage system for structured data that focuses on scalability, at the expense of the other benefits of relational databases [26], e.g. Google's BigTable [27] and Amazon's SimpleDB [28].

³ Uptime is a measure of the time a computer system has been running, i.e. up. It came into use to describe the opposite of downtime, times when a system was not operational [30].

localised. Therefore, the efficiencies gained from centralising infrastructure for Cloud Computing are increasingly at the expense of the Internet's resilience.

2) *Convenience vs Control*: The growing popularity of Cloud Computing comes from its convenience, but also brings vendor control, an issue of ever-increasing concern. For example, Google Apps for in-house e-mail typically provides higher uptime [32], but its failure [33] highlights the issue of lock-in that comes from depending on vendor Clouds. The even greater concern is the loss of information privacy, with vendors having full access to the resources stored on their Clouds. So much so the British government is considering a 'G Cloud' for government business applications [34]. In particularly sensitive cases of SMEs and start-ups, the provider-consumer relationship that Cloud Computing fosters between the owners of resources and their users could potentially be detrimental, as there is a potential conflict of interest for the providers. They profit by providing resources to up-and-coming players, but also wish to maintain dominant positions in their consumer-facing industries.

3) *Environmental Impact*: The other major concern is the ever-increasing *carbon footprint* from the *exponential growth* [4] of the *data centres* required for Cloud Computing. With the industry expected to exceed the airline industry by 2020 [6], raising sustainability concerns [5]. The industry is being motivated to address the problem by legislation [6], [35], the operational limit of power grids (being unable to power anymore servers in their data centres) [36], and the potential financial benefits of increased efficiency [37], [6]. Their primary solution is the use of *virtualisation*⁴ to maximise resource utilisation [39], but the problem remains [40], [41].

While these issues are endemic to Cloud Computing, they are not flaws in the Cloud conceptualisation, but the vendor provision and implementation of Clouds [25], [22], [42]. There are attempts to address some of these concerns, such as a portability layer between vendor Clouds to avoid lock-in [43]. However, this will not alleviate issues such as inter-Cloud latency [44]. An open source implementation of the Amazon (EC2) Cloud [22], called Eucalyptus [45], allows a data centre to execute code compatible with Amazon's Cloud. Allowing for the creation of *private internal* Clouds, avoiding vendor lock-in and providing information privacy, but only for those with their own data centre and so is not really Cloud Computing (which by definition is to avoid owning data centres [1]). Therefore, vendor Clouds remain synonymous with Cloud Computing [13], [1], [14], [15]. Our response is an alternative model for the Cloud conceptualisation, created by combining the Cloud with paradigms from Grid Computing, principles from Digital Ecosystems, and sustainability from Green Computing, while remaining true to the original vision of the Internet [46].

III. GRID COMPUTING: DISTRIBUTING PROVISION

Grid Computing is a form of distributed computing in which a *virtual super computer* is composed from a cluster




Figure 3. Grid Computing: Typical configuration in which resource provision is managed by a group of distributed nodes [47]. Green symbolises resource consumption, and yellow resource provision. The role of coordinator for resource provision is designated by red, and is centrally controlled.

of networked, loosely coupled computers, acting in concert to perform very large tasks [47]. It has been applied to computationally intensive scientific, mathematical, and academic problems through volunteer computing, and used in commercial enterprise for such diverse applications as drug discovery, economic forecasting, seismic analysis, and back-office processing to support e-commerce and web services [47].

What distinguishes Grid Computing from *cluster computing* is being more loosely coupled, heterogeneous, and geographically dispersed [47]. Also, grids are often constructed with general-purpose grid software libraries and middleware, dividing and apportioning pieces of a program to potentially thousands of computers [47]. However, what distinguishes Cloud Computing from Grid Computing is being web-centric, despite some of its definitions being conceptually similar (such as computing resources being consumed as electricity is from power grids) [9].

IV. DIGITAL ECOSYSTEMS: DISTRIBUTING CONTROL

Digital Ecosystems are distributed adaptive open socio-technical systems, with properties of self-organisation, scalability and sustainability, inspired by natural ecosystems [48], [49]. Emerging as a novel approach to the catalysis of sustainable regional development driven by SMEs [50]. Aiming to help local economic actors become active players in globalisation [51], valorising their local culture and vocations, and enabling them to interact and create value networks at the global level [52]. Increasingly this approach, dubbed *glocalisation*, is being considered a successful strategy of globalisation that preserves regional growth and identity [53], [54], [55], and has been embraced by the mayors and decision-makers of thousands of municipalities [56]. The community focused on the deployment of Digital Ecosystems, REgions for Digital Ecosystems Network (REDEN) [50], is supported by projects such as the Digital Ecosystems Network of regions for (4) DissEmination and Knowledge Deployment (DEN4DEK) [57]. This thematic network that aims to share experiences and disseminate knowledge to let regions effectively deploy of Digital Ecosystems at all levels (economic, social, technical and political) to produce real impacts in the economic activities of European regions through the improvement of SME business environments.

⁴ Virtualisation is the creation of a virtual version of a resource, such as a server, which can then be stored, migrated, duplicated, and instantiated as needed, improving scalability and work load management [38].

In a traditional market-based economy, made up of sellers and buyers, the parties exchange property, while in a new network-based economy, made up of servers and clients, the parties share access to services and experiences [58]. Digital Ecosystems aim to support network-based economies reliant on next-generation ICT that will extend the Service-Oriented Architecture (SOA) concept [59] with the automatic combining of available and applicable services in a scalable architecture, to meet business user requests for applications that facilitate business processes. Digital Ecosystems research is yet to consider scalable resource provision, and therefore risks being subsumed into vendor Clouds at the infrastructure level, while striving for decentralisation at the service level. So, the realisation of their vision requires a form of Cloud Computing, but with their principle of community-based infrastructure where individual users share ownership [48].

V. GREEN COMPUTING: GROWING SUSTAINABLY

Green Computing is the efficient use of computing resources, with the primary objective being to account for the *triple bottom line*⁵, an expanded spectrum of values and criteria for measuring organisational (and societal) success [61]. Given computing systems existed before concern over their environmental impact, it has generally been implemented retroactively, but is now being considered at the development phase [61]. It is systemic in nature, because ever-increasingly sophisticated modern computer systems rely upon people, networks and hardware. So, the elements of a *green* solution may comprise items such as end user satisfaction, management restructuring, regulatory compliance, disposal of electronic waste, telecommuting, virtualisation of server resources, energy use, thin client solutions and return on investment [61].

One of the greatest environmental concerns of the industry is their data centres [41], which have increased in number over time as business demands have increased, with facilities housing a rising amount of evermore powerful equipment [17]. As data centres run into limits related to power, cooling and space, their ever-increasing operation has created a noticeable impact on power grids [36]. To the extent that data centre efficiency has become an important global issue, leading to the creation of the Green Grid [62], an international non-profit organisation mandating an increase in the energy efficiency of data centres. Their approach, virtualisation, has improved efficiency [40], [41], but is optimising a flawed model that does not consider the whole system, where resource provision is disconnected from resource consumption. For example, competing vendors must host significant redundancy in their data centres to manage usage spikes and maintain the illusion of infinite resources. So, we would argue that an alternative more systemic approach is required, where resource consumption and provision are connected, to minimise the environmental impact and allow sustainable growth.

⁵ The triple bottom line (*people, planet, profit*) [60].

VI. COMMUNITY CLOUD

C3 arises from concerns over Cloud Computing, specifically control by vendors and lack of environmental sustainability. The Community Cloud aspires to combine distributed resource provision from Grid Computing, distributed control from Digital Ecosystems and sustainability from Green Computing, with the use cases of Cloud Computing, while making greater use of self-management advances from Autonomic Computing. Replacing vendor Clouds by shaping the underutilised resources of user machines to form a Community Cloud, with nodes potentially fulfilling all roles, *consumer*, *producer*, and most importantly *coordinator*, as shown in Figure 4.




Figure 4. Community Cloud: Created from shaping the underutilised resources of user machines, with nodes potentially fulfilling all roles, consumer, producer, and most importantly coordinator. Green symbolises resource consumption, yellow resource provision, and red resource coordination.

A. Conceptualisation

The conceptualisation of the Community Cloud draws upon Cloud Computing [20], Grid Computing [9], Digital Ecosystems [48], Green Computing [63] and Autonomic Computing [12]. A paradigm for Cloud Computing in the *community*, without dependence on Cloud vendors, such as Google, Amazon, or Microsoft.

1) *Openness*: Removing dependence on vendors makes the Community Cloud the open equivalent to vendor Clouds, and therefore identifies a new dimension in the open versus proprietary struggle [64] that has emerged in code, standards and data, but has yet to be expressed in the realm of hosted services.

2) *Community*: The Community Cloud is as much a social structure as a technology paradigm [65], because of the community ownership of the infrastructure. Carrying with it a degree of economic scalability, without which there would be diminished competition and potential stifling of innovation as risked in vendor Clouds.

3) *Individual Autonomy*: In the Community Cloud, nodes have their own utility functions in contrast with data centres, in which dedicated machines execute software as instructed. So, with nodes expected to act in their own self-interest, centralised control would be impractical, as with consumer electronics like game consoles [66]. Attempts to control user machines counter to their self-interest results in cracked systems, from black market hardware

modifications and *arms races* over hacking and securing the software (routinely lost by the vendors) [66]. In the Community Cloud, where no concrete vendors exist, it is even more important to avoid antagonising the users, instead embracing their self interest and harnessing it for the benefit of the community with measures such as a *community currency*.

4) *Identity*: In the Community Cloud each user would inherently possess a unique identity, which combined with the structure of the Community Cloud should lead to an inversion of the currently predominant membership model. So, instead of users registering for each website (or service) anew, they could simply add the website to their identity and grant access. Allowing users to have multiple services connected to their identity, instead of creating new identities for each service. This relationship is reminiscent of recent application platforms, such as Facebook's f8 and Apple's App Store, but decentralised in nature and so free from vendor control. Also, allowing for the reuse of the connections between users, akin to Google's Friend Connect, instead of reestablishing them for each new application.

5) *Graceful Failures*: The Community Cloud is not owned or controlled by any one organisation, and therefore not dependent on the lifespan or failure of any one organisation. It therefore ought to be robust and resilient to failure, and immune to the system-wide cascade failures of vendor Clouds, because of the diversity of its supporting nodes. When occasionally failing doing so gracefully, non-destructively, and with minimal downtime, as the unaffected nodes mobilise to compensate for the failure.

6) *Convenience and Control*: The Community Cloud, unlike vendor Clouds, has no inherent conflict between convenience and control, resulting from its community ownership providing distributed control, which would be more democratic. However, whether the Community Cloud can provide technically quality equivalent or superior to its centralised counterparts is an issue that will require further research.

7) *Community Currency*: The Community Cloud would require its own currency to support the sharing of resources, a *community currency*, which in economics is a medium (currency), not backed by a central authority (e.g. national government), for exchanging goods and services within a community [67]. It does not need to be restricted geographically, despite sometimes being called a local currency [68]. An example is the Fureai kippu system in Japan, which issues credits in exchange for assistance to senior citizens [69]. Family members living far from their parents can earn credits by offering assistance to the elderly in their local community, which can then be transferred to their parents and redeemed by them for local assistance [69].

8) *Quality of Service*: Ensuring acceptable quality of service (QoS) in a heterogeneous system will be a challenge. Not least because achieving and maintaining the different aspects of QoS will require reaching *critical mass* in the participating nodes and available services. Thankfully, the *community currency* could support long-term promises by resource providers and allow the higher quality providers, through market forces, to command a

higher price for their service provision. Interestingly, the Community Cloud could provide a better QoS than vendor Clouds, utilising time-based and geographical variations advantageously in the dynamic scaling of resource provision.

9) *Environmental Sustainability*: We expect the Community Cloud to have a smaller *carbon footprint* than vendor Clouds, on the assumption that making use of underutilised user machines requires less energy than the dedicated data centres required for vendor Clouds. The server farms within data centres are an intensive form of computing resource provision, while the Community Cloud is more organic, growing and shrinking in a symbiotic relationship to support the demands of the community, which in turn supports it.

10) *Service Composition*: The great promise of service-oriented computing is that the *marginal cost* of creating the n-th application will be virtually zero, as all the software required already exists to satisfy the requirements of other applications. Only their composition and orchestration are required to produce a new application [70], [71]. Within vendor Clouds it is possible to make services that expose themselves for composition and compose these services, allowing the hosting of a complete service-oriented architecture [20]. However, current service composition technologies have not gained widespread adoption [72]. Digital Ecosystems advocate service composability to avoid centralised control by large service providers, because easy service composition allows coalitions of SMEs to compete simply by composing simpler services into more complex services that only large enterprises would otherwise be able to deliver [52]. So, we should extend decentralisation beyond resource provision and up to the service layer, to enable service composition within the Community Cloud.

B. Architecture

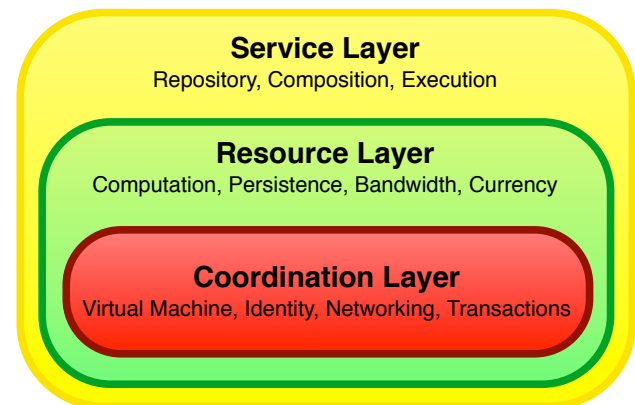


Figure 5. Community Cloud Computing: An architecture in which the most fundamental layer deals with distributing coordination. One layer above, resource provision and consumption are arranged on top of the coordination framework. Finally, the service layer is where resources are combined into end-user accessible services, to then themselves be composed into higher-level services.

The method of materialising the Community Cloud is the distribution of its server functionality amongst a population of nodes provided by user machines, shaping their underutilised resources into a *virtual data centre*. While straightforward in principle, it poses challenges on

many different levels. So, an architecture for C3 can be divided into three layers, dealing with these challenges iteratively. The most fundamental layer deals with distributing *coordination*, which is taken for granted in homogeneous data centres where good connectivity, constant presence and centralised infrastructure can be assumed. One layer above, *resource* provision and consumption are arranged on top of the coordination framework. Easy in the homogeneous grid of a data centre where all nodes have the same interests, but more challenging in a distributed heterogeneous environment. Finally, the *service* layer is where resources are combined into end-user accessible services, to then themselves be composed into higher-level services.

1) *Coordination Layer*: To achieve coordination, the nodes need to be deployed as isolated *virtual machines*, forming a fully distributed P2P⁶ network that can provide support for distributed identity, trust, and transactions.

a) *Virtual Machines (VMs)*: Executing arbitrary code in the machine of a resource-providing user would require a *sandbox*⁷ for the guest code, a VM⁸ to protect the host. The role of the VM is to make system resources *safely* available to the Community Cloud, upon which Cloud processes could be run safely (without danger to the host machine). Fortunately, feasibility has been proven with heavyweight VMs such as the Java Virtual Machine, lightweight JavaScript VMs present in most modern web browsers, and new approaches such as Google's Native Client. Furthermore, the age [76] of *multi-core processors*⁹ has resulted in unused and underutilised cores being commonplace in modern personal computers [78], which lend themselves well to the deployment and background execution of Community Cloud facing VMs. Regarding deployment, users would be required to maintain an active browser window or tab, or install a dedicated application. While the first would not require installation privileges, the later would with the benefit of greater functionality. However, more likely a hybrid of both would occur, facilitating the availability and advantages of each in different scenarios.

b) *Distributed Identity*: In distributed systems with variable node reliability, historical context is logically required to have certainty of node interactions. Fundamental to this context is the ability to identify nodes and therefore reference previous interactions. However, current identification schemes have *identity providers* controlling

provision. Such as in the DNS¹⁰, which while nominally distributed, remains under centralised control both technologically and organisationally, permitting numerous distortions in the network. Including *domain squatting*¹¹, abuses by domain registrars [81], subjection to political control [82], [83] and risks to the infrastructure being compromised [84]. Identity in the Community Cloud has to arise naturally from the structure of the network, based on the relation of nodes to each other, so that it can scale and expand without centralised control. We can utilise the property that a large enough identifier-space is unlikely to suffer collisions. For example, the Git distributed version control system [85] assigns a universal identifier to each new commission, without coordination with other repositories. Analogously, assuming each node independently produces a private-public key pair, the probability of public key collision is negligible. Also, from the human identification of nodes we can utilise the property that each node, despite formal identity, possesses a unique position in the network, i.e. set of connections to other nodes. So, combining these two properties provides reasonable certainty for a *distributed identity model* where universal identification can be accomplished without centralised mediation, but this is still an active area of research.

c) *Networking*: At this level, nodes should be interconnected to form a P2P network. Engineered to provide high resilience while avoiding single points of control and failure, which would make decentralised super-peer based control mechanisms [86] insufficient. Newer P2P designs [87] offer sufficient guarantees of distribution, immunity to super-peer failure, and resistance to enforced control. For example, in the Distributed Virtual Super-Peer (DVSP) model a collection of peers logically combine to form a virtual super-peer [87], which dynamically changes over time to facilitate fluctuating demands.

d) *Distributed Transactions*: A key element of distributed coordination is the ability of nodes to jointly participate in transactions that influence their individual state. Appropriately annotated business processes can be executed over a distributed network with a transactional model maintaining the ACID¹² properties on behalf of the initiator [89]. Newer transaction models maintain these properties while increasing efficiency and concurrency. Other directions of research include relaxing these properties to maximise concurrency [90]. Others still, focus on distributing the coordination of transactions [87]. A feature vital for C3, as distributed transaction capabilities are

⁶ Peer-to-peer (P2P) computing or networking is a distributed application architecture that partitions tasks or work loads between service peers. Peers are equally privileged participants in the application, and are said to form a peer-to-peer network of nodes [73].

⁷ A sandbox is a security mechanism for safely running programs, often used to execute untested code, or untrusted programs from unverified third-parties, suppliers and untrusted users [74].

⁸ A virtual machine is a software implementation of a machine (computer) that executes programs like a real machine [75].

⁹ A multi-core processor is an integrated circuit to which two or more processors have been attached for enhanced performance, reduced power consumption, and more efficient simultaneous processing of multiple tasks [77].

¹⁰ The Domain Name System (DNS) is a hierarchical naming-space for computers, services, and other resources participating in the Internet. It translates *domain names* meaningful to humans into their counterpart numerical identifiers associated with networking equipment to locate and address these devices world-wide [79]. So, translating human-friendly computer hostnames into Internet Protocol (IP) addresses, e.g. `www.example.com` translates to `208.77.188.166`.

¹¹ Domain squatting (also known as cybersquatting) is registering, trafficking in, or using a domain name in bad faith, with the intent to profit from the goodwill of a trademark belonging to someone else. The cybersquatter then offers to sell the domain to the person or company who owns a trademark contained within the name at an inflated price [80].

¹² ACID (Atomicity, Consistency, Isolation, Durability) is a set of properties that guarantee transactions are processed reliably [88].

fundamental to permitting multi-party service composition without centralised mediation.

2) *Resource layer*: With the networking infrastructure now in place, we can consider the first consumer-facing uses for the *virtual data centre* of the Community Cloud. Offering the usage experience of Cloud Computing on the PaaS layer and above, because Cloud Computing is about using resources from the Cloud. So, Utility Computing scenarios [91], such as access to raw storage and computation, should be available at the PaaS layer. Access to these abstract resources for service deployment would then provide the SaaS layer.

a) *Distributed Computation*: The field has a successful history of centrally controlled incarnations [92]. However, C3 should also take inspiration from Grid Computing and Digital Ecosystems to provide distributed coordination of the computational capabilities that nodes offer to the Community Cloud.

b) *Distributed Persistence*: The Community Cloud would naturally require storage on its participating nodes, taking advantage of the ever-increasing surplus on most¹³ personal computers [94]. However, the method of information storage in the Community Cloud is an issue with multiple aspects. First, information can be file-based or structured. Second, while constant and instant availability can be crucial, there are scenarios in which recall times can be relaxed. Such varying requirements call for a combination of approaches, including distributed storage [95], distributed databases [96] and key-value stores [26]. Information privacy in the Community Cloud should be provided by the encryption of user information when on remote nodes, only being unencrypted when accessed by the user, allowing for the secure and distributed storage of information.

c) *Bandwidth Management*: The Community Cloud would probably require more bandwidth at the user nodes than vendor Clouds, but can take advantage of the ever-increasing bandwidth and deployment of broadband [97]. Also, P2P protocols such as BitTorrent [98] make the distribution of information over networks much less bandwidth-intensive for content providers, accomplished by using the downloading peers as repeaters of the information they receive. C3 should adopt such approaches to ensure the efficient use of available network bandwidth, avoiding fluctuations and sudden rises in demand (e.g. the Slashdot effect¹⁴) burdening parts of the network.

d) *Community Currency*: An important theme in the Community Cloud is that of nodes being contributors as well as consumers, which would require a *community currency* (redeemable against resources in the community) to reward users for offering resources [100]. This would also allow traditional Cloud vendors to participate by offering their resources to the Community Cloud to gather consid-

erable *community currency*, which they can then monetise against participants running a *community currency* deficit (i.e. contributing less than they consume). The relative cost of resources (storage, computation, bandwidth) should fluctuate based on market demand, not least because of the impracticality of predicting or hard-coding such ratios. So, a node of the network would gather *community currency* by performing tasks for the community, which its user could then use to access resources of the Community Cloud.

e) *Resource Repository*: Given that each node providing resources has a different location in the network and quality characteristics, a distributed resource repository would be required that could respond to queries for resources according to desired performance profiles. Such a query would have to consider historical performance, current availability, projected cost and geographical distribution of the nodes to be returned. A constraint optimisation problem, the results returned would be a set of nodes that fit the required profile, proportionally to the availability of suitable nodes.

3) *Service Layer*: Cloud Computing represents a new era for service-oriented architectures, making services explicitly dependent on other resource providers instead of building on self-sufficient resource locations. C3 makes this more explicit, breaking down the stand-alone service paradigm, with any service by default being composed of resources contributed by multiple participants. So, the following sections define the core infrastructural services that the Community Cloud would need to provide.

a) *Distributed Service Repository (DSR)*: The service repository of the Community Cloud must provide persistence, as with traditional service repositories [101], for the pointers to services and their semantic descriptions. To support the absence of service-producing nodes during service execution, there must also be persistence of the executable code of services. Naturally, the implementation of a distributed service repository is made easier by the availability of the distributed storage infrastructure of the Community Cloud.

b) *Service Deployment and Execution*: When a service is required, but is not currently instantiated on a suitable node, a copy should be retrieved from the DSR and instantiated as necessary, allowing for flexible responsiveness and resilience to unpredictable traffic spikes. As nodes are opportunistically interested in executing services to gather *community currency* for their users, so developers should note the resource cost of their services in their descriptions, allowing for pre-execution resource budgeting, and post-execution *community currency* payments. Being in a developer's own interest to mark resource costs correctly, because over-budgeting would burden their users and under-budgeting would cause premature service termination. Additionally, developers could add a subsidy to promote their services. Remote service execution would need to be secured against potentially compromised nodes, perhaps through encrypted processing schemes [102]. Otherwise, such nodes while unable to access a complete traffic log of the services they execute, could potentially access the business logic; and we would be replacing the vendor introspection problem, with an anyone intro-

¹³ The only exception is the recent arrival of Solid-State Drives (SSDs), popular for mobile devices because of their lack of moving parts, growing in use as their size and price reach that of traditional Hard Disk Drives (HDDs) [93].

¹⁴ The Slashdot effect, also known as slashdotting, is the phenomenon of a popular website linking to a smaller site, causing the smaller site to slow down or even temporarily close due to the increased traffic [99].

spection problem. Since delivering a service over large distances in the network comes at a potentially high cost, the lack of a central well-connected server calls for a fundamental paradigm shift, from *pull*-oriented approaches to hybrid *push/pull*-oriented approaches. So, instead of the *pull*-oriented approach of supplying services only upon request [103], service provision should also follow a *push*-oriented approach of preemptive deployment to strategically suitable nodes, including modifying their deployment profile based on the traffic patterns they face at run time.

c) *Programming Paradigm*: A key innovation of Cloud Computing in its PaaS incarnation, is the offering of a well-specified context (programming paradigm) within which the services should be executed [20]. The programming paradigm that produces these services is also important to C3, because it forms a contract between the service developers and resource providers. The current state-of-the-art requires manipulation of source code in which each line is context dependent, and so a single intended change may necessitates significant alterations at different locations in the codebase. A paradigm shift to *declarative generative programming* [104] would be greatly beneficial, avoiding the need to manually manage cascading changes to the codebase. As the requirements behind a service would be made explicit and executable, and being human readable could therefore be manipulated directly as stand-alone artifacts. Additionally, barriers to service composition would be significantly decreased [105], beneficial to C3 and beyond.

C. Distributed Innovation

When considering the Community Cloud over time, current software distribution models would cause problems. Should the infrastructure be dependent on a single provider for updates, they would become a single point of control, and possibly failure. Entrusting a single provider with the power to control the evolution of the architecture, even if they are considered benevolent, risks the development goals becoming misaligned with the community. Therefore, the Community Cloud should follow an *evolutionary software distribution model*. Extending an already-growing trend of using distributed code repositories such as Git [85] and Mercurial [106], over centralised code repositories such as Subversion [107] and CVS [108]. So, modifications to services, including infrastructural ones, should be distributed locally to migrate over the Community Cloud from where they are deployed, making use of the existing relationships between users. Users or their nodes (by default) could even choose to follow the updates that other trusted peers adopt. Therefore, new versions of a service would compete with older versions, and where superior (fitter) would distribute more widely, spreading further across the Community Cloud. So, updates to services would permeate through the network, in a distributed but regulated manner. We could even consider the updates to services, as the release of patches (modifications), allowing for frequent, smaller and iterative releases more akin to an *evolutionary software distribution model*. Potential speciation (branching) would encourage developers to coordinate their releases and ensure their patches are viable across different branches. Obviously, the ability to undo

patches and step back through versions of infrastructural services would be necessary to maintain the Community Cloud. Still, without a more granular approach to conflict resolution from different patching sources, poor developer relations could risk fragmentation of the codebase and network. So, an alternative non-centralised *software innovation model* would be required, such as the *declarative generative programming paradigm* [104] mentioned.

VII. IN THE COMMUNITY CLOUD

While we have covered the fundamental motivations and architecture of the Community Cloud, its practical application may still be unclear. So, this section discusses the cases of Wikipedia and YouTube, where the application of C3 would yield significant benefits, because they have unstable funding models, require increasing scalability, and are community oriented.

A. Wikipedia

Wikipedia suffers from an ever-increasing demand for resources and bandwidth, without a stable supporting revenue source [109]. Their current funding model requires continuous monetary donations for the maintenance and expansion of their infrastructure [110]. The alternative being contentious advertising revenues [109], which caused a long-standing conflict within their community [111]. While it would provide a more scalable funding model, some fear it would compromise the content and/or the public trust in the content [112]. Alternatively, the Community Cloud could provide a self-sustaining scalable resource provision model, without risk of compromising the content or public trust in the content, because it would be compatible with their communal nature (unlike their current *data centre* model), with their user base accomplishing the resource provision they require.

Were Wikipedia to adopt C3, it would be distributed throughout the Community Cloud alongside other services. With the core operations of Wikipedia, providing web-pages and executing server-side scripts, being handled as service requests. Participants would use their *community currency* to interact with Wikipedia, performing a search or retrieving a page, while gaining *community currency* for helping to host Wikipedia across the Community Cloud. More complicated tasks, such as editing a Wikipedia web-page, would require an update to the distributed storage of the Community Cloud, achieved by transmitting the new data through its network of nodes, most likely using an eventual consistency model [90].

B. YouTube

YouTube requires a significant bandwidth for content distribution, significant computational resources for video transcoding, and is yet to settle on a profitable business model [113], [114]. In the Community Cloud, websites like YouTube would also have a self-sustaining scalable resource provision model, which would significantly reduce the income required for them to turn a profit.

Were YouTube to adopt C3, it would also be distributed throughout the Community Cloud alongside other services. Updates such as commenting on a YouTube video, would similarly need to propagate through the

distributed persistence layer. So, the community would provide the bandwidth for content distribution, and the computational resources for video transcoding, required for YouTube's service. The QoS requirements for YouTube are significantly different to those of Wikipedia, because while constant throughput is desirable for video streaming, occasional packet loss is tolerable. Also, YouTube's streaming of live events has necessitated the services of bespoke content distribution networks [115], a type of service for which the Community Cloud would naturally excel.

We have discussed Wikipedia and YouTube in the Community Cloud, but other sites such as arXiv and Facebook would equally benefit. As C3's organisational model for resource provision moves the cost of service provision to the user base, effectively creating a micro-payment scheme, which would dramatically lower the barrier of entry for innovative start-ups.

VIII. CONCLUSIONS

We have presented the Community Cloud as an alternative to Cloud Computing, created from blending its usage scenarios with paradigms from Grid Computing, principles from Digital Ecosystems, self-management from Autonomic Computing, and sustainability from Green Computing. So, C3 utilises the spare resources of networked personal computers to provide the facilities of data centres, such that the community provides the computing power for the Cloud they wish to use. A socio-technical conceptualisation for sustainable distributed computing.

While the Open Cloud Manifesto [116] is well intentioned, its promotion of open standards for vendor Cloud interoperability has proved difficult [117]. We believe it will continue to prove difficult until a viable alternative, such as C3, is developed. Furthermore, we hope that the Community Cloud will encourage innovation in vendor Clouds, forming a relationship analogous to the creative tension between open source and proprietary software.

In the future we will continue to refine the various elements of C3, such as suitable mechanisms for a *community currency*, distributed alternatives to DNS, DVSPs, RESTful Clouds, *declarative generative programming paradigms*, distributed innovation, and the environmental impact of the Community Cloud relative to vendor Clouds.

ACKNOWLEDGEMENTS

We would like to thank for comments and helpful discussions Paulo Siqueira of the Instituto de Pesquisas em Tecnologia e Inovacao, Eva Tallaksen and Alexander Deriziotis.

REFERENCES

- [1] M. Haynie, "Enterprise cloud services: Deriving business value from Cloud Computing," Micro Focus, Tech. Rep., 2009.
- [2] R. L. Peter Lucas, Joseph Ballay, "The wrong cloud," MAYA Design, Inc, Tech. Rep., 2009. [Online]. Available: http://www.maya.com/file_download/126/The%20Wrong%20Cloud.pdf
- [3] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "Above the Clouds: A Berkeley view of Cloud Computing," *University of California, Berkeley*, 2009. [Online]. Available: <http://d1smfj0g31qzek.cloudfront.net/abovetheclouds.pdf>

- [4] J. Hayes, "Cred - or croak?" IET Knowledge Network, Tech. Rep., 2008. [Online]. Available: <http://kn.theiet.org/magazine/issues/0820/cred-croak-0820.cfm?SaveToPDF>
- [5] P. Mckenna, "Can we stop the internet destroying our planet?" *New Scientist*, vol. 197, no. 2637, pp. 20–21, 2008.
- [6] J. Kaplan, W. Forrest, and N. Kindler, "Revolutionizing data center energy efficiency," McKinsey & Company, Tech. Rep., 2008. [Online]. Available: http://www.mckinsey.com/clientservice/bto/pointofview/pdf/Revolutionizing_Data_Center_Efficiency.pdf
- [7] B. Leiner, V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts, and S. Wolff, "A brief history of the internet," Institute for Information Systems and Computer Media, Tech. Rep., 2001. [Online]. Available: http://www.iicm.tugraz.at/thesis/cguet_diss/literatur/Kapitel02/References/Leiner_et_al_2000/brief.html?timestamp=1197467969844
- [8] J. Scanlon and B. Wiens, "The internet cloud," The Industry Standard, Tech. Rep., 1999. [Online]. Available: <http://www.thestandard.com/article/0,1902,5466,00.html>
- [9] I. Foster, Y. Zhao, I. Raicu, and S. Lu, "Cloud Computing and Grid Computing 360-degree compared," in *Grid Computing Environments Workshop*, 2008, pp. 1–10.
- [10] T. Foremski, "Sun services CTO says utility computing acceptance is slow going," ZDNet, CBS Interactive, Tech. Rep., 2006. [Online]. Available: <http://blogs.zdnet.com/Foremski/?p=33>
- [11] A. Orłowski, "The Cell chip - how will MS and Intel face the music?" The Register, Tech. Rep., 2005. [Online]. Available: http://www.theregister.co.uk/2005/02/03/cell_analysis_part_two/
- [12] J. Kephart, D. Chess, I. Center, and N. Hawthorne, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [13] G. Gruman and E. Knorr, "What Cloud Computing really means," InfoWorld Inc., Tech. Rep., 2008. [Online]. Available: http://www.infoworld.com/article/08/04/07/15FE-cloud-computing-reality_1.html
- [14] Gartner, "Cloud Computing will be as influential as e-business," Gartner, Tech. Rep., 2008. [Online]. Available: <http://www.gartner.com/it/page.jsp?id=707508>
- [15] P. Gaw, "What's the difference between Cloud Computing and SaaS?" Proofpoint, Tech. Rep., 2008. [Online]. Available: <http://blog.fortiva.com/fortivablog/2008/05/what-is-the-dif.html>
- [16] K. Danielson, "Distinguishing Cloud Computing from Utility Computing," ebizQ, Tech. Rep., 2008. [Online]. Available: http://www.ebizq.net/blogs/saasweek/2008/03/distinguishing_cloud_computing/
- [17] M. Arregoces and M. Portolani, *Data center fundamentals*. Cisco Press, 2003.
- [18] M. Turner, D. Budgen, and P. Brereton, "Turning software into a service," *Computer*, vol. 36, no. 10, pp. 38–44, 2003.
- [19] T. O'Reilly, "What is Web 2.0: Design patterns and business models for the next generation of software," O'Reilly Media, Tech. Rep., 2008. [Online]. Available: <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>
- [20] R. Buyya, C. Yeo, and S. Venugopal, "Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities," in *Conference on High Performance Computing and Communications*. IEEE, 2008.
- [21] A. Newman, A. Steinberg, and J. Thomas, *Enterprise 2.0 Implementation*. McGraw-Hill Osborne Media, 2008.
- [22] Amazon, "Amazon Elastic Compute Cloud (EC2)," Amazon Web Services LLC, Tech. Rep., 2009. [Online]. Available: <http://aws.amazon.com/ec2/>
- [23] Mosso, "Deploy and scale websites, servers and storage in minutes," Rackspace, Tech. Rep., 2009. [Online]. Available: <http://www.mosso.com/>
- [24] R. Buyya, C. Yeo, and S. Venugopal, "Market-oriented Cloud Computing: Vision, hype, and reality for delivering it services as computing utilities," in *High Performance Computing and Communications*. IEEE Press, 2008.
- [25] Google, "Google App Engine: Run your web apps on Google's infrastructure," Google, Tech. Rep., 2009. [Online]. Available: <http://code.google.com/appengine/>
- [26] T. Bain, "Is the relational database doomed?" *ReadWriteWeb.com*, 2008. [Online]. Available: http://www.readwriteweb.com/archives/is_the_relational_database_doomed.php
- [27] F. Chang, J. Dean, S. Ghemawat, W. Hsieh, D. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. Gruber, "Bigtable: A distributed storage system for structured data," in *USENIX Symposium on Operating Systems Design and Implementation*, 2006.
- [28] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels,

- "Dynamo: Amazon's highly available key-value store," in *Symposium on Operating Systems Principles*. ACM, 2007, pp. 205–220.
- [29] B. Johnson, "Cloud Computing is a trap, warns GNU founder Richard Stallman," *The Guardian*, Tech. Rep., 2008. [Online]. Available: <http://www.guardian.co.uk/technology/2008/sep/29/cloud.computing.richard.stallman>
- [30] J. McCabe, *Network analysis, architecture, and design*. Morgan Kaufmann, 2007. [Online]. Available: http://books.google.co.uk/books?id=iddGPgR48_MC
- [31] A. Modine, "Web startups crumble under Amazon S3 outage," *The Register*, Tech. Rep., 2008. [Online]. Available: http://www.theregister.co.uk/2008/02/15/amazon_s3_outage_feb_2008/
- [32] J. Montgomery, "Google Apps sees 99.9reliability," *Tech.Blorge*, Tech. Rep., 2008. [Online]. Available: <http://tech.blorge.com/Structure:%20/2008/11/02/google-apps-sees-999-uptime-proves-cloud-reliability/>
- [33] J. Perez, "Google Apps customers miffed over downtime," *IDG News Service*, Tech. Rep., 2007. [Online]. Available: http://www.pcworld.com/businesscenter/article/130234/google_apps_customers_miffed_over_downtime.html
- [34] Kable, "Carter recommends 'g cloud' for gov it," *The Register*, 2009. [Online]. Available: http://www.channelregister.co.uk/2009/06/17/government_cloud_computing/
- [35] Environmental Protection Agency, "EPA report to congress on server and data center energy efficiency," US Congress, Tech. Rep., 2007.
- [36] R. Miller, "NSA maxes out Baltimore power grid," *Data Center Knowledge*, Tech. Rep., 2006. [Online]. Available: <http://www.datacenterknowledge.com/archives/2006/08/06/nsa-maxes-out-baltimore-power-grid/>
- [37] K. McIsaac, "The data centre goes green, the CFO saves money," *Intelligent Business Research Services*, Tech. Rep., 2007.
- [38] C. Wolf and E. Halter, *Virtualization: from the desktop to the enterprise*. Apress, 2005.
- [39] R. Talaber, T. Brey, and L. Lamers, "Using virtualization to improve data center efficiency," *The Green Grid*, Tech. Rep., 2009.
- [40] K. Brill, "The invisible crisis in the data center: The economic meltdown of Moore's law," *Uptime Institute*, Tech. Rep., 2007.
- [41] J. Brodtkin, "Gartner in 'green' data centre warning," *Techworld*, 2008. [Online]. Available: <http://www.techworld.com/green-it/news/index.cfm?newsid=106292>
- [42] Microsoft, "Azure services platform," *Micrsoft*, Tech. Rep., 2009. [Online]. Available: <http://www.microsoft.com/azure/>
- [43] C. Metz, "The Meta Cloud - flying data centers enter fourth dimension," *The Register*, Tech. Rep., 2009. [Online]. Available: http://www.theregister.co.uk/2009/02/24/the_meta_cloud/
- [44] T. Kulmala, "The cloud's hidden lock-in: Latency," *Archivd*, 2009. [Online]. Available: <http://blog.archivd.com/1/post/2009/04/the-clouds-hidden-lock-in-latency.html>
- [45] D. Nurmi, R. Wolski, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov, "The Eucalyptus open-source cloud-computing system," in *Cloud Computing and Its Applications*, 2008.
- [46] J. Abbate, *Inventing the internet*. MIT press, 1999.
- [47] I. Foster and C. Kesselman, *The grid: blueprint for a new computing infrastructure*. Morgan Kaufmann, 2004.
- [48] G. Briscoe and P. De Wilde, "Digital Ecosystems: Evolving service-oriented architectures," in *Conference on Bio Inspired Models of Network, Information and Computing Systems*. IEEE Press, 2006. [Online]. Available: <http://arxiv.org/abs/0712.4102>
- [49] G. Briscoe, "Digital ecosystems," Ph.D. dissertation, Imperial College London, 2009.
- [50] L. Rivera León. Regions for Digital Ecosystems Network (REDEN). [Online]. Available: <http://reden.opaals.org/>
- [51] P. Dini, G. Lombardo, R. Mansell, A. Razavi, S. Moschoyiannis, P. Krause, A. Nicolai, and L. Rivera León, "Beyond interoperability to digital ecosystems: regional innovation and socio-economic development led by SMEs," *International Journal of Technological Learning, Innovation and Development*, vol. 1, pp. 410–426, 2008.
- [52] F. Nachira, A. Nicolai, P. Dini, M. Le Louarn, and L. Rivera León, Eds., *Digital Business Ecosystems*. European Commission, 2007.
- [53] R. Robertson, "Globalisation or glocalisation," *Journal of International Communication*, vol. 1, pp. 33–52, 1994.
- [54] E. Swyngedouw, "The mammon quest. 'Glocalisation', interspatial competition and the monetary order: the construction of new scales," in *Cities and regions in the new Europe: The Global-local Interplay and Spatial Development Strategies*, M. Dunford and G. Kafkalas, Eds. Wiley, 1992, pp. 39–67.
- [55] H. Khondker, "Glocalization as globalization: Evolution of a sociological concept," *Bangladesh e-Journal of Sociology*, vol. 1, pp. 1–9, 2004.
- [56] Glocal Forum and CERFE, "The glocalization manifesto," *The Glocal Forum*, Tech. Rep., 2004. [Online]. Available: http://www.glocalforum.org/mediagallery/mediaDownload.php?mm=warehouse/documents/the_glocalization_manifesto.pdf
- [57] L. Rivera León. Digital Ecosystems Network of Regions for Dissemination and Knowledge Deployment (DEN4DEK). [Online]. Available: <http://www.den4dek.org>
- [58] P. Delcloque and A. Bramoullé, "DISSEMINATE, an initial implementation proposal: a new point of departue in call for the 'year 01'?" *ReCALL*, vol. 13, pp. 277–292, 2001.
- [59] E. Newcomer and G. Lomow, *Understanding SOA with web services*. Addison-Wesley, 2005.
- [60] J. Elkington, *Cannibals with forks: the triple bottom line of 21st century business*. New Society Publishers, 1998.
- [61] J. Williams and L. Curtis, "Green: The new computing coat of arms?" *IT PROFESSIONAL*, pp. 12–16, 2008.
- [62] T. G. Grid, "About the green grid," *The Green Grid*, Tech. Rep., 2009. [Online]. Available: <http://www.thegreengrid.org/about-the-green-grid>
- [63] J. Harris, *Green Computing and Green It Best Practices on Regulations and Industry*. Lulu.com, 2008.
- [64] J. West and J. Dedrick, "Proprietary vs. open standards in the network era: An examination of the linux phenomenon," in *System Sciences*. IEEE Press, 2001, p. 10.
- [65] Y. Benkler, "Sharing nicely: on shareable goods and the emergence of sharing as a modality of economic production," *The Yale Law Journal*, vol. 114, no. 2, pp. 273–359, 2004.
- [66] J. Grand, F. Thornton, A. Yarusso, and R. Baer, *Game Console Hacking: Have Fun While Voiding Your Warranty*. Syngress Press, 2004.
- [67] T. Greco, *Money: Understanding and creating alternatives to legal tender*. Chelsea Green, 2001.
- [68] A. Doteuchi, "Community currency and NPOs- A model for solving social issues in the 21st century," *Social Development Research Group, NLI Research*, 2002. [Online]. Available: <http://www.nli-research.co.jp/english/socioeconomics/2002/li0204a.pdf>
- [69] B. Lietaer, "Complementary currencies in japan today: History, originality and relevance," *International Journal of Community Currency Research*, vol. 8, pp. 1–23, 2004.
- [70] Q. Tang, "Economics of web service provisioning: Optimal market structure and intermediary strategies," Ph.D. dissertation, University of Florida, 2004.
- [71] G. Modi, "Service oriented architecture & web 2.0," *Guru Tegh Bahadur Institute of Technology*, Tech. Rep., 2007. [Online]. Available: http://www.gsmodi.com/files/SOA_Web2_Report.pdf
- [72] B. Violino. (2007) How to navigate a sea of SOA standards. [Online]. Available: http://www.cio.com/article/104007/How_to_Navigate_a_Sea_of_SOA_Standards
- [73] R. Schollmeier, "A definition of peer-to-peer networking for the classification of peer-to-peer architectures and applications," in *International Conference on Peer-to-Peer Computing*. IEEE Press, 2002, pp. 101–102.
- [74] M. Bishop, *Computer Security*. Addison-Wesley, 2004.
- [75] I. Craig, *Virtual machines*. Springer, 2006.
- [76] D. Geer, "Chip makers turn to multicore processors," *IEEE Computer*, vol. 38, no. 5, pp. 11–13, 2005.
- [77] M. Zelkowitz, *Advances in Computers: Architectural Issues*. Academic Press, 2007.
- [78] B. Posey, "Multi-core processors: Their implication for Windows," *TechTarget*, Tech. Rep., 2007. [Online]. Available: http://searchwindowserver.techtarget.com/tip/0,289483,sid68_gci1248527,00.html
- [79] P. Mockapetris and K. Dunlap, "Development of the domain name system," *Computer Communication Review*, vol. 18, no. 4, pp. 123–133, 1988.
- [80] M. Maury and D. Kleiner, "E-commerce, ethical commerce?" *Journal of Business Ethics*, vol. 36, no. 1, pp. 21–31, 2002.
- [81] G. Lyon, "Exposing the many reasons not to trust godaddy with your domain names," *NoDaddy.com*, Tech. Rep., 2009. [Online]. Available: <http://nodaddy.com/>
- [82] M. Mueller, *Ruling the root: Internet governance and the taming of cyberspace*. MIT press, 2002.
- [83] (2006) Chinese walls. [Online]. Available: http://www.economist.com/business/displaystory.cfm?story_id=5582257
- [84] D. Goodin, "DNS patch averts doomsday scenario," *The Register*, Tech. Rep., 2008. [Online]. Available: http://www.theregister.co.uk/2008/08/06/kaminsky_black_hat/

- [85] S. Chacon, "About Git," GitHub, Tech. Rep., 2009. [Online]. Available: <http://git-scm.com/about>
- [86] J. Risson and T. Moors, "Survey of research towards robust peer-to-peer networks: Search methods," *Computer Networks*, vol. 50, pp. 3485–3521, 2006.
- [87] A. Razavi, S. Moschyiannis, and P. Krause, "A scale-free business network for digital ecosystems," in *IEEE Conf. on Digital Ecosystems and Technologies*, 2008.
- [88] T. Haerder and A. Reuter, "Principles of transaction-oriented database recovery," *ACM Computing Surveys*, vol. 15, no. 4, pp. 287–317, 1983.
- [89] A. Fox, S. Gribble, Y. Chawathe, E. Brewer, and P. Gauthier, "Cluster-based scalable network services," *ACM SIGOPS Operating Systems Review*, vol. 31, no. 5, pp. 78–91, 1997.
- [90] W. Vogels, "Eventually consistent," *ACM Queue*, vol. 6, 2008.
- [91] M. Rappa, "The utility business model and the future of computing services," *IBM Systems Journal*, vol. 43, no. 1, pp. 32–42, 2004.
- [92] H. Attiya and J. Welch, *Distributed computing: fundamentals, simulations, and advanced topics*. Wiley-Interscience, 2004.
- [93] C. Mellor, "Ssd and hdd capacity goes on embiggening," The Register, Tech. Rep., 2009. [Online]. Available: http://www.theregister.co.uk/2009/01/09/ssd_and_hdd_capacity_increases/
- [94] M. Daley, "Software bloat," MattsComputerTrends.com, Tech. Rep., 2009. [Online]. Available: <http://www.mattscomputertrends.com/softwarebloat.html>
- [95] P. Yianilos and S. Sobti, "The evolving field of distributed storage," *IEEE Internet Computing*, vol. 5, pp. 35–39, 2001.
- [96] H. Garcia-Molina, J. Ullman, and J. Widom, *Database Systems: The complete book*. Prentice Hall, 2008.
- [97] "Broadband growth and policies in OECD countries," Organisation for Economic Co-operation and Development, Tech. Rep., 2008.
- [98] B. Cohen, "Incentives build robustness in BitTorrent," in *Workshop on Economics of Peer-to-Peer Systems*, vol. 6, 2003.
- [99] S. Adler, "The Slashdot effect: An analysis of three internet publications," *Linux Gazette*, vol. 38, 1999.
- [100] D. Turner and K. Ross, "A lightweight currency paradigm for the p2p resource market," in *International Conference on Electronic Commerce Research*, 2004.
- [101] M. Papazoglou, "Service-oriented computing: concepts, characteristics and directions," in *International Conference on Web Information Systems Engineering*, T. Catarci, M. Mecella, J. Mylopoulos, and M. Orłowska, Eds. IEEE Press, 2003, pp. 3–12.
- [102] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Symposium on Theory of computing*. ACM, 2009, pp. 169–178.
- [103] M. Singh and M. Huhns, *Service-Oriented Computing: Semantics, Processes, Agents*. Wiley, 2005.
- [104] A. Marinos and P. Krause, "What, not how: A generative approach to service composition," in *Digital Ecosystems and Technologies Conference*. IEEE Press, 2009.
- [105] —, "Using sbvr, rest and relational databases to develop information systems native to the digital ecosystem," in *Digital Ecosystems and Technologies Conference*. IEEE Press, 2009.
- [106] S. Arseanrapoj, "Mercurial: Source control management system," Mercurial, Tech. Rep., 2009. [Online]. Available: <http://mercurial.selenic.com/wiki/>
- [107] Tigris, "Subversion: open source version control system," Tigris.org, Tech. Rep., 2009. [Online]. Available: <http://subversion.tigris.org/>
- [108] D. Price, "Cvs - concurrent versions system," nongnu.org, Tech. Rep., 2009. [Online]. Available: <http://www.nongnu.org/cvs/>
- [109] Heebie Blog, "Wikipedia Fundraising: The real truth," Heebie Intuitive Design, Tech. Rep., 2009. [Online]. Available: <http://blog.heebie.co.uk/wikipedia-fundraising-real-truth>
- [110] A. Modine, "Wales' personal begging earns last \$2m," The Register, Tech. Rep., 2009. [Online]. Available: http://www.theregister.co.uk/2009/01/02/wikipedia_fundraising_2m_jan_2/
- [111] W. Roelf, "Wikipedia founder mulls revenue options," Reuters, Tech. Rep., 2007. [Online]. Available: <http://www.reuters.com/article/internetNews/idUSL1964587420070420>
- [112] H. Leslie, "Wikipedia to run out of money?" Digital-Lifestyles, Tech. Rep., 2007. [Online]. Available: <http://digital-lifestyles.info/2007/02/12/wikipedia-to-run-out-of-money/>
- [113] C. Metz, "Google accused of avoiding youtube revenues," *The Register*, 2009. [Online]. Available: http://www.theregister.co.uk/2009/06/18/google_youtube_loses/
- [114] D. Silversmith, "Google losing up to \$1.65m a day on youtube," *Internet Evolution*, 2009. [Online]. Available: http://www.internetevolution.com/author.asp?section_id=715&doc_id=175123&
- [115] M. Arrington, "Google relies on akamai to stream youtube live; 700,000 concurrent viewers," *TechCrunch*, 2008. [Online]. Available: <http://www.techcrunch.com/2008/11/22/google-relies-on-akamai-to-stream-youtube-live-700000-concurrent-viewers/>
- [116] "Open cloud manifesto," OpenCloudManifesto.org, Tech. Rep., 2009. [Online]. Available: <http://www.opencloudmanifesto.org/Open%20Cloud%20Manifesto.pdf>
- [117] C. Metz, "What's an open cloud? the manifesto's not telling," *The Register*, 2009. [Online]. Available: http://www.theregister.co.uk/2009/03/31/amazon_on_cloud_manifesto/
- [118] C. Hewitt, "ORGs for scalable, robust, privacy-friendly client Cloud Computing," *IEEE Internet Computing*, vol. 12, no. 5, 2008.
- [119] A. Avram, "Architecting for green computing," *InfoQ.com*, 2008. [Online]. Available: <http://www.infoq.com/news/2008/12/Architecture-Green-Computing>
- [120] A. Weiss, "Computing in the clouds," *netWorker*, vol. 11, pp. 16–25, ACM Press, 2007.
- [121] G. Briscoe and A. Marinos, "Digital ecosystems in the clouds: Towards community cloud computing," in *Digital Ecosystems and Technologies Conference*. IEEE Press, 2009. [Online]. Available: <http://arxiv.org/abs/0903.0694>