

Gerard Briscoe and Alexandros Marinos

Digital ecosystems in the clouds: towards community cloud computing

Book section

Original citation:

Briscoe, Gerard and Marinos, Alexandros (2009) Digital ecosystems in the clouds: towards community cloud computing. In: IEEE, (corp. ed.) 2009 3rd IEEE International Conference on Digital Ecosystems and Technologies (DEST 2009). Institute of Electrical and Electronics Engineers (IEEE), New York, USA, pp. 103-108.

© 2009 [IEEE](#)

This version available at: <http://eprints.lse.ac.uk/26664/>

Available in LSE Research Online: January 2010

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of London School of Economics and Political Science's products or services.

Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

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.

This document is the author's submitted version of the book section. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Digital Ecosystems in the Clouds: Towards Community Cloud Computing

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

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

Abstract— Cloud Computing is rising fast, with its data centres growing at an unprecedented rate. However, this has come with concerns of privacy, efficiency at the expense of resilience, and environmental sustainability, because of the dependence on Cloud vendors such as Google, Amazon, and Microsoft. Community Cloud Computing makes use of the principles of Digital Ecosystems to provide a paradigm for Clouds in the community, offering an alternative architecture for the use cases of Cloud Computing. It is more technically challenging to deal with issues of distributed computing, such as latency, differential resource management, and additional security requirements. 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 Clouds, Community Cloud Computing, Digital Ecosystems, 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, including the lack of information privacy [2]. Also, the data centres required for Cloud Computing are growing exponentially [3], creating an ever-increasing *carbon footprint*, raising environmental concerns [4], [5].

The social paradigms and technologies of Digital Ecosystems, including the community ownership of digital infrastructure, can remedy these concerns. So, Cloud Computing combined with the principles of Digital Ecosystems provides a compelling socio-technical conceptualisation for sustainable distributed computing, utilising the spare resources of networked personal computers to provide the facilities of a virtual *data centre* to form collectively a Community Cloud.

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 how it is depicted in computer network diagrams to abstract the complex infrastructure it conceals [6]. It can be seen as a commercial evolution of the academia-oriented Grid Computing [7], succeeding where Utility Computing struggled [8], [9]. It is being promoted as the cutting edge of scalable web application development [2], in which




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. Green symbolises resource consumption, and yellow resource provision. The role of coordinator for resource provision is designated by red, and is centrally controlled.

dynamically scalable and often virtualised resources are provided as a service over the Internet [10], [1], [11], [12], with users having no knowledge of, expertise in, or control over the technology infrastructure of the Cloud supporting them [13]. It currently has significant momentum in two extremes of the web development industry [2], [1]: the consumer web technology incumbents who have resource surpluses in their vast *data centres*, and various consumers and start-ups that do not have access to such computational resources. Cloud Computing conceptually incorporates software-as-a-service (SaaS) [14], Web 2.0 [15] 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. Green symbolises resource consumption, and yellow resource provision. The role of coordinator for resource provision is designated by red, and is centrally controlled. From the figure, it can be seen that coordination and resource provision are centrally controlled, even if the central node is implemented as a distributed grid, which is the usual incarnation of a data centre. Providers, who are also the controllers, are usually companies with other web activities that require large computing resources, and in their efforts to scale their primary businesses they 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.

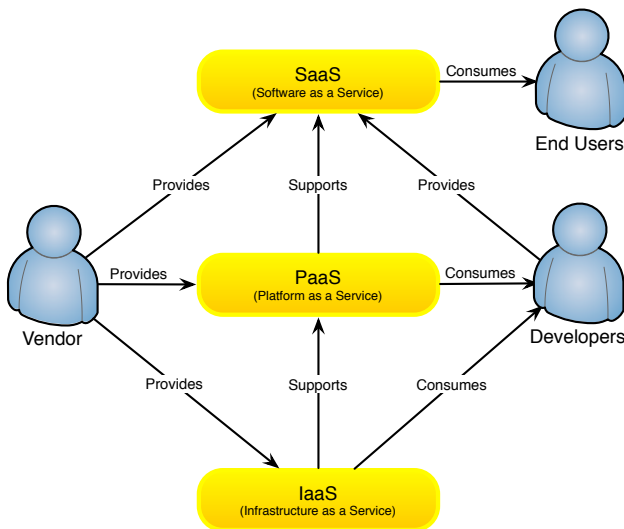


Figure 2. Abstractions of Cloud Computing: There is a significant *buzz* [16] around Cloud Computing, but there is little clarity about which offerings actually qualify and their interrelation. The key to resolving this confusion is by realising that the various offerings fall into different levels of abstraction, aimed at different market segments.

A. Layers of Abstraction

There is a significant *buzz* [16] around Cloud Computing, but there is little clarity about which offerings actually qualify and their interrelation. The key to resolving this confusion is by realising 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)* [17]: At the most basic level of the Cloud Computing offerings, there are providers such as Amazon [18] and Mosso [19], 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)* [20]: One level of abstraction above, services like Google App Engine [21] 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* instead of *relational databases*.

3) *Software as a Service (SaaS)* [14]: At the consumer-facing level are the most popular examples of Cloud Computing, with well-defined applications offering 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 take on multiple roles, with *vendors* developing services for the *end users*, or *developers* utilising the services of others for their own. Yet, within a Cloud the role of provider, and therefore controller, can only be occupied by a single entity, the *vendor*.

B. Concerns

The Cloud Computing model is not without concerns, as others have noted [22], [2], and we consider the following as primary:

1) *Economics of Failure*: The uptime 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 these failures do occur it has a cascade effect, taking down organisations depending on the Cloud. This was illustrated by the Amazon (S3) Cloud outage [23], which took with it several other dependent businesses. So, failures are now system-wide, instead of being partial or localised. Therefore, the efficiencies gained from centralising infrastructure for Cloud Computing will increasingly be 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 [24], but its failure [25] highlighted 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. 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 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* [3] of the data centres required for Cloud Computing. With the industry expected to exceed the airline industry by 2020 [5], this raises sustainability concerns [4]. The industry is being motivated to address the problem by legislation [5], [26], the operational limit of power grids (not being able to power any more servers) [27], and the potential financial benefits [28], [5]. Their primary solution is the use of *virtualisation* to maximise resource utilisation [29], but the problem remains [30], [31].

While these issues are endemic to current Cloud Computing, they are not flaws in the Cloud conceptualisation, but of the vendor provision and implementation of Clouds

[21], [18]. There are attempts to address some of these concerns, such as avoiding vendor lock-in through a portability layer between different vendor Clouds, called a meta-Cloud or a Cloud-of-Clouds [32]. While this would avoid vendor lock-in to an extent, it will not alleviate the other concerns raised. Also, there is an open source implementation of the Amazon (EC2) Cloud [18], called Eucalyptus [33], which allows for a data centre to execute code compatible with Amazon's Cloud, providing *private internal* Clouds. This would avoid vendor lock-in and provide information privacy, but only for those with their own data centres, and so is not really Cloud Computing (which by definition is to avoid needing one's own data centre [1]). So, vendor Clouds remain synonymous with Cloud Computing [10], [1], [11], [12], while our response is an alternative model for the Cloud conceptualisation infused with the principles of Digital Ecosystems.

III. DIGITAL ECOSYSTEMS

Digital Ecosystems are distributed adaptive open socio-technical systems, with properties of self-organisation, scalability and sustainability, inspired by natural ecosystems [34], and are emerging as a novel approach to the catalysis of sustainable regional development driven by SMEs [35]. Digital Ecosystems aim to help local economic actors become active players in globalisation [36], valorising their local culture and vocations, and enabling them to interact and create value networks [37] at the global level. Increasingly this approach, dubbed *glocalisation*, is being considered a successful strategy of globalisation that preserves regional growth and identity [38], and has been embraced by the mayors and decision-makers of thousands of municipalities [39]. The community focused on the deployment of Digital Ecosystems, REgions for Digital Ecosystems Network (REDEN) [35], is supported by projects such as the Digital Ecosystems Network of regions for (4) DissEmination and Knowledge Deployment (DEN4DEK) [40], a thematic network that aims to share experiences and disseminate all the necessary knowledge that will allow regions to plan an effective deployment 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.

In a traditional market-based economy, made up of sellers and buyers, the parties exchange property [41], while in a new network-based economy, made up of servers and clients, the parties share access to services and experiences [41]. Digital Ecosystems aim to support network-based economies reliant on next-generation ICT that will extend the Service-Oriented Architecture (SOA) concept [42] with the automatic combining of available and applicable services in a scalable architecture, to meet business user requests for applications that facilitate business processes. Scalable resource provision has yet to be considered in Digital Ecosystems research. Without it Digital Ecosystems risk being subsumed into vendor Clouds at the infrastructure level, while striving for decentralisation at the service level, which would clearly be incompatible with its principles. So, the realisation of the Digital Ecosystems vision requires a form of Cloud




Figure 3. Community Cloud: Created from shaping the underutilised resources of user machines, with nodes potentially taking on all roles, *consumer*, *producer*, and most importantly *coordinator*. Green symbolises resource consumption, yellow resource provision, and red resource coordination and administration.

Computing – but within the principle of community-based infrastructure, where individual users share ownership [43].

IV. COMMUNITY CLOUD

Community Cloud Computing arises from concerns over the control of vendors in Cloud Computing and the observation that analogous concerns drive Digital Ecosystems research. It aspires to combine the principles of Digital Ecosystems with the use cases of Cloud Computing. Replacing vendor Clouds by shaping the underutilised resources of user machines to form a Community Cloud, with nodes potentially taking on all roles, *consumer*, *producer*, and most importantly *coordinator*, as shown in Figure 3.

A. Conceptualisation

The conceptualisation of the Community Cloud draws from Cloud Computing [1], Digital Ecosystems [34] and Green Computing [44]. It's a paradigm for Cloud Computing in the *community*, without dependence on Cloud vendors, such as Google, Amazon, or Microsoft.

1) *Openness*: Removing the 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 [45] that has emerged in code, standards and data, but has not until now been expressed in the realm of hosted services.

2) *Community*: The Community Cloud is as much a social structure as a technology paradigm [46], because of the community ownership of the infrastructure. This community ownership carries 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) *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 will 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 it will do so gracefully, non-destructively, and with minimal downtime, as the unaffected nodes compensate for the failure.

4) *Convenience and Control*: The Community Cloud, unlike vendor Clouds, has no inherent conflict between convenience and control, because its community ownership provides for democratic distributed control.

5) *Environmental Sustainability*: The Community Cloud will have a significantly smaller *carbon footprint* than vendor Clouds, because making use of underutilised user machines will require much 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.

B. Architecture

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, this poses challenges on many different levels, but many are aligned with currently active research topics in Digital Ecosystems.

1) *Core Infrastructure*: Before proceeding to the resource exchange and service composition, the nodes will be deployed as isolated *virtual machines*, forming a fully distributed peer-to-peer network, providing support for distributed identity and coordination.

a) *Virtual Machines (VMs)*: Executing arbitrary code in the machine of a resource-providing user will require a *sandbox* for the guest code, a VM to protect the host. The role of the VM is to make system resources *safely* available to the Cloud in a measurable manner. Fortunately, feasibility has been shown with heavyweight VMs such as the Java Virtual Machine and Common Language Runtime, and with the lightweight JavaScript VMs present in most modern web browsers. Furthermore, the age [47] of *multi-core processors* has resulted in unused and underutilised cores being commonplace in modern personal computers [48], which lend themselves well to the deployment and background execution of Community Cloud facing VMs.

b) *P2P Networking*: At this most fundamental level, nodes will have to be interconnected to form a peer-to-peer network. It will have to be specifically engineered to provide high resilience while avoiding single points of control and failure, which would make a decentralised super-peer based control mechanism [49] insufficient. A completely distributed peer-to-peer network is required, immune to super-peer failure.

c) *Distributed Identity/Trust*: The performing of tasks beyond the networking requires nodes to identify each other and keep historical context. This identification must be performed in a fully distributed manner, which has implications as most identity schemes are based on an identity provider controlling provision. Additionally, trust should be tracked as a multi-dimensional variable, including considerations such as uptime, performance characteristics, and reputation.

2) *Resource layer*: As the networking infrastructure is now in place, we can discuss the first consumer-facing uses of the *virtual data centre* of the Community Cloud. At its core, Cloud Computing is about using resources from the Cloud. The Community Cloud will offer the usage experience of Cloud Computing on the platform-as-a-service (PaaS) layer and above. *Utility Computing* [50]

scenarios, such as access to raw storage and computation, will be made available at the PaaS layer. Access to these abstract resources for service deployment will then provide the software-as-a-service (SaaS) layer.

a) *Distributed Computation*: The field has a long history of successful incarnations in its centrally controlled form. However, Community Cloud Computing will need to take inspiration from Grid Computing [51] to provide distributed coordination of the computational capabilities that nodes offer to the Community Cloud.

b) *Distributed Persistence*: The Community Cloud will naturally require storage on its participating nodes, taking advantage of the ever-increasing surplus on most personal computers [52]. 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 more relaxed. Such varying requirements call for a combination of approaches, including distributed storage [53], distributed databases [54] and key-value stores [55]. Information privacy in the Community Cloud will be provided by the encryption of user information when on remote nodes, only being unencrypted when cached on the user's node, allowing for the secure and distributed storage of information.

c) *Bandwidth Management*: The Community Cloud will probably require more bandwidth than vendor Clouds, but can take advantage of the ever-increasing bandwidth and deployment of broadband [56]. Also, peer-to-peer protocols such as BitTorrent have made the distribution of information over networks much less bandwidth-intensive for providers, accomplished by using the downloading peers as repeaters of the information they receive. Community Cloud Computing will have to adopt such approaches to ensure efficient use of available network bandwidth, to avoid fluctuations or sudden rises in demand (e.g. the Slashdot effect) burdening parts of the network.

d) *Community Currency*: An important theme in the Community Cloud is the notion of nodes being contributors as well as consumers, which will require a *community currency*¹ (redeemable against resources in the community) to reward users for offering valued resources. It will also allow for traditional Cloud vendors to participate, by offering their resources to the Community Cloud to gather considerable *community currency*, which they can then monetise against participants who cannot contribute as much as they consume (i.e. running a *community currency* deficit). To avoid predicting or hard-coding the relative cost of resources (storage, computation, bandwidth), their prices can fluctuate based on market demand.

3) *Service Layer*: Cloud Computing can be said to represent a move from service-oriented [42] to service-driven architectures, making services explicitly dependent on other providers instead of building on self-sufficient

¹ In economics, a *community currency* is a medium (currency) for exchanging goods and services within a community, that is not backed by a central authority (e.g. national government) [57], and which need not be restricted geographically despite sometimes being called a local currency [58].

resource locations. Community Cloud Computing makes this more explicit, breaking down the stand-alone service paradigm, with any service by default being composed of resources contributed by multiple participants. The following sections describe the core infrastructural services that the Community Cloud needs to provide.

a) Distributed Service Repository: The repository of the Community Cloud must provide persistence, as with traditional service repositories [59], for the pointers to services and the semantic descriptions of services. To support the absence of principal (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) Remote Service Execution: When a service is needed to fulfil a request, but is not currently instantiated on a suitable node, a copy should be retrieved from the repository and instantiated as needed. This allows for flexible responsiveness and resilience to unpredictable traffic spikes. Nodes are naturally interested in executing services as their purpose is to gather *community currency* for their users. A developer should note the resource cost of a service in its description, allowing for pre-execution resource budgeting by nodes, and post-execution *community currency* payments by consumers. It is in a developer's own interest to mark resource costs correctly, because over-budgeting will burden their users and under-budgeting will cause premature service termination. Additionally, developers could add a subsidy to promote their services. Remote service execution must be secured against potentially compromised nodes, because while unable to access a complete traffic log of the services they execute, they could potentially access the business logic of the services they execute. Otherwise, we would be replacing the vendor *looking in* problem, with an anyone *looking in* problem.

V. WIKIPEDIA IN THE COMMUNITY CLOUD

Wikipedia suffers from an ever-increasing demand for resources and bandwidth, without a stable revenue source for support [60]. Their current funding model requires a continuous influx of monetary donations for the maintenance and expansion of their infrastructure [61], the alternative being contentious advertising revenues [60], which has caused a long-standing conflict within their community [62]. While it would provide a more scalable funding model, the fear is it would compromise the public's trust in the content [63]. Alternatively, the Community Cloud could provide a self-sustaining scalable resource provision model without risk of compromising 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 Community Cloud Computing, it would be distributed throughout the Community Cloud alongside other services, which in this context can be as simple as a webpage or as complex as necessary. Participants in the Community Cloud will have a node on their machine, which when active accumulates *community*

currency by providing resources to fulfil service requests from other nodes. These service requests can be as simple as instantiating a simple HTML page or executing a server-side script, the core operations of Wikipedia. Participants can then use their amassed *community currency* to interact with Wikipedia, performing a search or retrieving a page. More complicated tasks, such as editing a Wikipedia page, will require an update to the distributed storage of the Community Cloud, achieved by transmitting the new data through its network of nodes, most likely resulting in an eventual consistency model [64].

We have discussed Wikipedia in the Community Cloud, but the latter is not limited to not-for-profit organisations, being just as beneficial to for-profit businesses. Its organisational model for resource provision moves the cost of service provision to the user base, effectively creating a micro-payment scheme, which dramatically lowers the barrier of entry for innovative start-ups.

VI. CONCLUSIONS

We have presented a socio-technical conceptualisation for sustainable distributed computing, the Community Cloud. The Community Cloud is an alternative to Cloud Computing, created from blending its usage scenarios with the principles of Digital Ecosystems. Community Cloud Computing 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 platform they wish to use. 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.

ACKNOWLEDGEMENTS

We would like to thank for helpful discussions, Georgios Exarchakos, Paul Krause, and Nick Antonopoulos of the university of Surrey, and Jo Stanley of the University of Cambridge. This work was supported by the EU-funded Open Philosophies for Associative Autopoietic Digital Ecosystems (OPAALS) Network of Excellence (NoE), Contract No. FP6/IST-034824.

REFERENCES

- [1] M. Haynie, "Enterprise cloud services: Deriving business value from Cloud Computing," Micro Focus, Tech. Rep., 2009.
- [2] 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, Tech. Rep., 2009. [Online]. Available: <http://d1smfj0g31qzek.cloudfront.net/abovetheclouds.pdf>
- [3] 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>
- [4] P. Mckenna, "Can we stop the internet destroying our planet?" *New Scientist*, vol. 197, no. 2637, pp. 20–21, 2008.
- [5] 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
- [6] J. Scanlon and B. Wieners. (1999) The internet cloud. [Online]. Available: <http://www.thestandard.com/article/0,1902,5466,00.html>
- [7] 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.

- [8] T. Foremski. (2006) Sun services CTO says utility computing acceptance is slow going. [Online]. Available: <http://blogs.zdnet.com/Foremski/?p=33>
- [9] A. Orłowski. (2005) The Cell chip - how will MS and Intel face the music? [Online]. Available: http://www.theregister.co.uk/2005/02/03/cell_analysis_part_two/
- [10] G. Gruman and E. Knorr. (2008) What Cloud Computing really means. [Online]. Available: http://www.infoworld.com/article/08/04/07/15FE-cloud-computing-reality_1.html
- [11] Gartner. (2008) Cloud Computing will be as influential as e-business. [Online]. Available: <http://www.gartner.com/it/page.jsp?id=707508>
- [12] P. Gaw. (2008) What's the difference between Cloud Computing and SaaS? [Online]. Available: <http://blog.fortiva.com/fortivablog/2008/05/what-is-the-dif.html>
- [13] K. Danielson. (2008) Distinguishing Cloud Computing from Utility Computing. [Online]. Available: http://www.ebizq.net/blogs/saasweek/2008/03/distinguishing_cloud_computing/
- [14] M. Turner, D. Budgen, and P. Brereton, "Turning software into a service," *Computer*, vol. 36, no. 10, pp. 38–44, 2003.
- [15] 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.oreilynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>
- [16] B. Worthen. (2008) Overuse clouds buzz term's meaning. [Online]. Available: <http://blogs.wsj.com/biztech/2008/09/23/overuse-clouds-buzz-terms-meaning/>
- [17] A. Newman, A. Steinberg, and J. Thomas, *Enterprise 2.0 Implementation*. McGraw-Hill Osborne Media, 2008.
- [18] Amazon. (2009) Amazon Elastic Compute Cloud (EC2). [Online]. Available: <http://aws.amazon.com/ec2/>
- [19] Mosso. (2009) Deploy and scale websites, servers and storage in minutes. [Online]. Available: <http://www.mosso.com/>
- [20] 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.
- [21] Google. (2009) Google App Engine: Run your web apps on Google's infrastructure. [Online]. Available: <http://code.google.com/appengine/>
- [22] B. Johnson. (2008) Cloud Computing is a trap, warns GNU founder Richard Stallman. [Online]. Available: <http://www.guardian.co.uk/technology/2008/sep/29/cloud.computing.richard.stallman>
- [23] A. Modine. (2008) Web startups crumble under Amazon S3 outage. [Online]. Available: http://www.theregister.co.uk/2008/02/15/amazon_s3_outage_feb_2008/
- [24] J. Montgomery. (2008) Google Apps sees 99.9% uptime, proves "cloud reliability". [Online]. Available: <http://tech.bloge.com/Structure:0/2008/11/02/google-apps-sees-999-uptime-proves-cloud-reliability/>
- [25] J. Perez. (2007) Google Apps customers miffed over downtime. [Online]. Available: http://www.pcworld.com/businesscenter/article/130234/google_apps_customers_miffed_over_downtime.html
- [26] "EPA report to congress on server and data center energy efficiency," US Congress, Tech. Rep., 2007.
- [27] R. Miller. (2006) NSA maxes out Baltimore power grid. [Online]. Available: <http://www.datacenterknowledge.com/archives/2006/08/06/nsa-maxes-out-baltimore-power-grid/>
- [28] K. McIsaac, "The data centre goes green, the CFO saves money," Intelligent Business Research Services, Tech. Rep., 2007.
- [29] R. Talaber, T. Brey, and L. Lamers, "Using virtualization to improve data center efficiency," The Green Grid, Tech. Rep., 2009.
- [30] K. Brill, "The invisible crisis in the data center: The economic meltdown of Moore's law," Uptime Institute, Tech. Rep., 2007.
- [31] J. Brodtkin. (2008) Gartner in 'green' data centre warning. [Online]. Available: <http://www.techworld.com/green-it/news/index.cfm?newsid=106292>
- [32] C. Metz. (2009) The Meta Cloud - flying data centers enter fourth dimension. [Online]. Available: http://www.theregister.co.uk/2009/02/24/the_meta_cloud/
- [33] 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.
- [34] Wikipedia. Digital Ecosystem. [Online]. Available: http://en.wikipedia.org/wiki/Digital_ecosystem
- [35] L. Rivera León. Regions for Digital Ecosystems Network (REDEN). [Online]. Available: <http://reden.opaals.org/>
- [36] P. Dini, G. Lombardo, R. Mansell, A. Razavi, S. Moschioniannis, 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.
- [37] F. Nachira, A. Nicolai, P. Dini, M. Le Louarn, and L. Rivera León, Eds., *Digital Business Ecosystems*. European Commission, 2007.
- [38] R. Robertson, "Globalisation or glocalisation," *Journal of International Communication*, vol. 1, pp. 33–52, 1994.
- [39] 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
- [40] L. Rivera León. Digital Ecosystems Network of Regions for Dissemination and Knowledge Deployment (DEN4DEK). [Online]. Available: <http://www.den4dek.org>
- [41] 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.
- [42] E. Newcomer and G. Lomow, *Understanding SOA with web services*. Addison-Wesley, 2005.
- [43] 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>
- [44] J. Harris, *Green Computing and Green IT Best Practices*. Lulu.com, 2008.
- [45] 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.
- [46] 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.
- [47] D. Geer, "Chip makers turn to multicore processors," *IEEE Computer*, vol. 38, no. 5, pp. 11–13, 2005.
- [48] B. Posey. (2007) Multi-core processors: Their implication for Windows. [Online]. Available: http://searchwindowserver.techtarget.com/tip/0,289483,sid68_gci1248527,00.html
- [49] J. Risson and T. Moors, "Survey of research towards robust peer-to-peer networks: Search methods," *Computer Networks*, 2006.
- [50] M. Rappa, "The utility business model and the future of computing services," *IBM Systems Journal*, vol. 43, no. 1, pp. 32–42, 2004.
- [51] F. Berman, G. Fox, and A. Hey, *Grid Computing: Making the global infrastructure a reality*. Wiley, 2003.
- [52] M. Daley. (2009) Software bloat. [Online]. Available: <http://www.mattscomputertrends.com/softwarebloat.html>
- [53] P. Yianilos and S. Sobti, "The evolving field of distributed storage," *IEEE Internet Computing*, vol. 5, pp. 35–39, 2001.
- [54] H. Garcia-Molina, J. Ullman, and J. Widom, *Database Systems: The complete book*. Prentice Hall, 2008.
- [55] T. Bain. (2008) Is the relational database doomed? [Online]. Available: http://www.readwriteweb.com/archives/is_the_relational_database_doomed.php
- [56] "Broadband growth and policies in OECD countries," Organisation for Economic Co-operation and Development, Tech. Rep., 2008.
- [57] T. Greco, *Money: Understanding and creating alternatives to legal tender*. Chelsea Green, 2001.
- [58] A. Doteuchi, "Community currency and NPOs- A model for solving social issues in the 21st century," Social Development Research Group, NLI Research, Tech. Rep., 2002. [Online]. Available: <http://www.nli-research.co.jp/english/socioeconomics/2002/li0204a.pdf>
- [59] 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.
- [60] Heebie Blog. (2009) Wikipedia Fundraising: The real truth. [Online]. Available: <http://blog.heebie.co.uk/wikipedia-fundraising-real-truth>
- [61] A. Modine. (2009) Wales' personal begging earns last \$2m. [Online]. Available: http://www.theregister.co.uk/2009/01/02/wikipedia_fundraising_2m_jan_2/
- [62] W. Roelf. (2007) Wikipedia founder mulls revenue options. [Online]. Available: <http://www.reuters.com/article/internetNews/idUSL1964587420070420>
- [63] H. Leslie. (2007) Wikipedia to run out of money? [Online]. Available: <http://digital-lifestyles.info/2007/02/12/wikipedia-to-run-out-of-money/>
- [64] W. Vogels, "Eventually consistent," *ACM Queue*, vol. 6, 2008.
- [65] C. Hewitt, "ORGs for scalable, robust, privacy-friendly client Cloud Computing," *IEEE Internet Computing*, vol. 12, no. 5, 2008.