Leveraging Cloud Computing and High Performance Computing Advances for Next-generation Architecture, Urban Design and Construction Projects

Francesco Iorio[1], Jane L. Snowdon[2]

[1] Autodesk Research
210 King Street East
Toronto, ON, Canada, M5A1J7
francesco.iorio@autodesk.com

[2] IBM T. J. Watson Research Center
1101 Kitchawan Road, Route 134
Yorktown Heights, NY, USA, 10598
snowdonj@us.ibm.com

Keywords: Cloud Computing, High Performance Computing, Architecture, Urban Design, Construction, Simulation, Optimization.

Abstract

Architecture and urban design projects are constantly breaking barriers of scale and complexity and continuously seek improved efficiency, sustainability, building energy performance, and cost-effectiveness. Simulation and large-scale data processing are now fundamental elements of this process. Recent advances in algorithms and computational power offer the means to address the complex dynamics of an integrated whole building system. However, scalability is a significant barrier to the realization of whole building systems tools for design, control and optimization.

This position paper presents a set of techniques such as fast design parameter-space exploration, large-scale high-accuracy simulation, and integrated multi-disciplinary optimization for semi- or fully-automated designs. These techniques are extremely computing intensive, and have traditionally only been available to the research community. But, once enabled by advances in cloud computing and high performance computing, these techniques can facilitate the interactive design process resulting in improved outcomes and reduced development cycle times.

1. INTRODUCTION

According to the World Business Council for Sustainable Development (2009) and the U. S. Department of Energy (DOE) (U.S. DOE, Energy Information Administration 2008a), buildings account for 40% of the world's total energy consumption and, in 2005, and nine gigatons of global carbon dioxide (CO_2) emissions, well above the transportation and industry sectors. In the United States alone, commercial and residential buildings account for 38% of all CO_2 emissions and 72% of electricity consumption (U.S. DOE, Energy Information Administration 2008b; U.S. DOE, Environmental Information Administration 2008c). Furthermore, buildings use 13.6% of all potable water, or 15 trillion gallons per year, and 40% of raw materials globally, which is equivalent to 3 billion tons annually (Hutson et al. 2004, Lenssen and Roodman 1995). Much of the energy consumption by commercial buildings is spent on lighting (26%), followed by heating and cooling (13% and 14%, respectively) (U.S. DOE 2007). Investing in energy efficient light bulbs and insulation materials and in automated shading has proven to reduce the energy demands on cooling and lighting (Lee et al. 2007). However, incremental improvements achieved by implementing individual energy efficient technologies alone are not sufficient to successful meet the challenging objectives set forth by the Intergovernmental Panel on Climate Change (IPCC) and other directives issued by cities, for example PlaNYC in New York City (IPCC 2007, PLANYC 2007).

Reducing energy consumption by buildings and improving energy efficiency in buildings are two approaches to help achieve the goals. Frost and Sullivan (2009) cites the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) stating only 11% of building cost is construction, while 50% of the lifecycle cost is in the operation of the building.

Nobel Prize laureate and U. S. Secretary of Energy Dr. Steven Chu stated: "We need … computer design tools for commercial and residential buildings that enable reductions in energy consumption of up to 80 percent with investments that will pay for themselves in less than 10 years." (Chu 2009). This grand challenge will require a system integration approach to building design, aided by computer tools with embedded energy analysis, computer monitoring, and real-time control of building systems.

A smarter building can be defined as integrating major building systems on a common network where information and functionality between systems is shared to improve energy efficiency, operational effectiveness, and occupant satisfaction (IBM, 2010). A smarter building is a complex system of systems that span heating and air conditioning, lighting, security, access control, entertainment, people movers, water, and monitoring, control and maintenance systems. Ideally, these systems would have well managed and integrated physical and digital infrastructures that make the building safe, comfortable, and functional for its occupants and sustainable for the environment. Through the use of sensors, digital smart IP enabled meters and submeters, digital controls, and analysis tools to automatically monitor and control services for its users, buildings can be studied in much greater detail to investigate whole systems solutions. Since the average lifetime of a commercial building is fifty years, it is crucial to focus on the design, construction, and operation of both new buildings and retrofits.

An integrated design process (IDP) for new buildings and building retrofits that incorporates energy simulation, lighting analysis, computational fluid dynamics and digital information about a building's structure, occupancy, and plumbing, electrical and mechanical systems and costs, will enable architects, builders and contractors, facilities managers, and building owner/operators to achieve optimized building performance through the selection of the best materials, windows, equipment, subsystems and processes. These virtual models and tools, combined with cloud and high performance computing (HPC), will improve productivity and allow rapid assessment of design alternatives (Lunzer and Hornbæk, 2008) to optimize the energy performance of a building. A shift in culture among building stakeholders is needed to adopt collaboration, interoperability, and the effective and active use of operational building data as the way of the future.

Cloud computing is an emerging style of computing in which applications, data and IT resources are provided as services to users over the web. Cloud computing is a way of managing large numbers of highly virtualized resources such that, from a management perspective, they resemble a single large scalable resource, which in turn can be used to deliver services. Several cloud computing paradigms exist, for example, Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS) (Zhu et al. 2009). This is a disruptive computing and business paradigm, poised to change the traditional organizations' IT infrastructure management and business practices, and to dramatically improve productivity and cost effectiveness of software solutions. As an infrastructure paradigm, cloud computing is extremely well suited to efficiently handle traditional "web" workloads, whereas its adoption in different areas of computing is not without technical and strategic challenges that, without solution, inhibits acceptance and adoption.

We present an analysis of the benefits and challenges of using cloud computing and HPC in the context of the design and operation of next-generation architecture and urban design projects, and suggest some potential solutions.

2. ARCHITECTURE, URBAN DESIGN PROJECTS AND SOFTWARE APPLICATIONS

Architecture and construction projects require a vast array of software solutions that range from Product Lifecycle Management (PLM) to Computer Aided Design (CAD) systems, logistics, and energy usage simulation to environmental impact analysis. These software tools improve the productivity of designers, architects, engineers, builders, managers and facility workers who need to operate in the best way to successfully complete the assigned projects within budget and time constraints.

Simulation has been traditionally used in architectural and engineering projects to validate designs created by individual domain experts, and is usually performed on reduced-complexity system models in order to contain the overall analysis time and cost. The United States Green Building Council (USGBC) encourages the use of simulation in its Leadership in Energy and Environmental Design (LEED) certification process.

Today's building energy simulators are not able to capture the multi-scale dynamics of a whole building system and the majority of models provide results relevant only to steady state conditions. More recently though, the requirements for additional efficiency and cost reduction in large-scale projects, combined with the availability of more sophisticated tools to evaluate and predict the resulting project's performance (e.g. a building's total energy consumption) as accurately as possible, are improving the collaborative nature of design. With these more sophisticated tools come increases in the amount of data produced, which ultimately requires increases in compute power to process various simulation-driven optimization strategies.

The PLM range of software applications is relatively easy to migrate to the cloud environment, as their fundamental operations are of a transactional nature and large, distributed databases have been commonly used in cloud-based web applications for some time. These characteristics facilitate the economic justification of a cloud-based business model. CAD and more general design applications, on the other hand, especially when applied to large-scale projects, present significant challenges that must be addressed to allow the full exploitation of cloud computing economies of scale.

Simulation is also a discipline, which, despite the vast amount of expertise accumulated over recent decades, has not been fully exploited as a process optimization tool, mostly due to the inherently high costs associated with dedicated HPC resources. Fully addressing these challenges opens up tremendous new opportunities that can revolutionize the way architecture and urban design projects are planned and executed.

Below, computing advances are outlined where scalable resources can have a significant impact on the development of solutions in support of the grand challenges of this domain.

2.1. Intelligent Application Execution Prediction

Design software is characterized by a workflow composed of alternating stages: a user performs an operation (action stage) and then reviews the results of the operation (review stage). When large amounts of computing resources are available, computing power can be dedicated to predict a user's behavior in order to reduce the perceived application latency. Implicitly, the more action-review workflow cycles that can be performed during the design process, the better the outcome. Predictive methods can greatly reduce the time taken for each cycle thereby allowing more cycles to take place in the same amount of time.

In the review stages, while the user is evaluating the effects of the last operation or set of operations, the application will analyze the user's workflow and attempt to predict which set of operations the user is most likely to perform next. Then it will allocate computing resources to actually perform in parallel all those operations, speculating that one of them will actually be selected by the user in the next action stage (Igarashi and Hughes, 2001). If the speculation is successful, the latency of the selected operation can be significantly reduced by overlapping it with the time the user takes to review his/her progress.

2.2. Fast Parameter-space Exploration

Design, architecture and engineering are disciplines that require the evaluation of numerous options and the expertise of individual domain experts provides guidance towards the choices that must be made at all stages of a project. Despite the skill level of a project's individual contributors, there is a limit to the variety of options and combinations of components an architect or an engineer can evaluate in a reasonable amount of time during each design cycle on a large project.

It is necessary to develop effective methodologies and tools to improve the speed of design space exploration for optimal configurations that reduce energy consumption, improve overall efficiency, and/or the total cost of the project itself in a robust way so that variations in the different uses of buildings, energy loads, and degradations in equipment operation can be considered in a timely manner. Access to large amounts of computing power would lead to the ability to evaluate hundreds or thousands of different design variations and their effects on the overall project. This would provide an opportunity for more informed decision-making and for further confirmation of a choice made by a subject matter expert.

2.3. Large-scale Simulation

Next-generation architecture projects face increasing budget constraints, demands for greater efficiency and reductions in environmental impact. As the multitude of variables inherent in construction projects increases and the need for more sophisticated simulation models for design and construction grows, for example to better understand the real non-linear behavior of materials, a much higher level of detail in the simulations are required in order to have the most accurate possible analyses. In the past, greatly simplified models have been considered sufficient for estimating energy consumption. While the potentially double-digit error ranges that result may be acceptable for buildings expecting typical energy consumption, these simplifications are not viable for a low-energy building where a difference of 10% would be unacceptable for a theoretically net-zero energy building.

Complex and very detailed analysis models require simulation to a scale only available to large private, academic, and government research laboratories. Recent advances in computational power, computer software, and clever mathematical algorithms and heuristics offer industries access to such tools across a building's life cycle providing the means to address the complex dynamics of a whole integrated building system. Burns et al. (2010) discuss the needs, the current state and gaps, and the research approach for simulation-based design.

2.4. Multi-disciplinary Optimization for Semi-automatic Design

Multi-disciplinary optimization can be defined as the field of engineering that uses optimization methods to solve design problems incorporating a number of disciplines. The importance of sound software design practices is often underestimated when developing large-scale integrated simulations. Typically each discipline develops or uses software modules and well-defined interfaces between these modules are lacking.

As a further refinement to fast parameter-space exploration, numerical analysis principles can be applied to find optimal project configurations based on soft and hard constraints such as energy consumption, carbon emissions, among others. Architecture designs can be defined as completely parametric problems as illustrated in Shea, Aish and Gourtovaia (2005). The concept can be extended to define not only shapes but also other characteristics of architecture projects, such as the choice of construction materials for every component, placement of fixtures, shape and size of windows, mechanical systems integration strategies, and so on.

When combined in a single problem domain, the number of variables easily overwhelms the capabilities of an individual simulation. While analysis of the results of individual simulations and fast parameter space exploration can provide strong guidance, multi-disciplinary optimization combines analytical methods with domain-specific heuristics to provide a very integrated decision-making tool. This technique can result in dramatic improvements in overall cost reduction and efficiency, and if applied to a vast majority of future architecture and engineering projects, it can substantially reduce the environmental impact of next-generation urban planning, but not without substantial increases in the complexity of the system.

3. CHALLENGES

3.1. CAD User Experience

Interactivity and complex 3D visualization are both fundamental characteristics of modern CAD systems. The main source of potential compromises in using a cloud-based interactive computing environment is the latency introduced by the communication network, which in most circumstances is the Internet. The very low response latency typical of CAD applications running on a user's workstation is hard to replicate while running in the cloud, although the large amount of computing power required for increasingly important operations like complex simulation and multi-disciplinary optimization is not available in a desktop system.

3.2. Data Transfer and Data Security

Organizations operating in the engineering sector generate large amounts of security sensitive data such as project plans, CAD drawings, and project schedules. Transferring security sensitive data to a public cloud is then a serious trust issue for some organizations, who cannot afford the risk of having data stolen or compromised. Transferring the data to the cloud for on-demand processing can be impractical and cumbersome, especially in the presence of large datasets, which is common for large engineering projects; some form of incremental data transfer is then required. Buecker et al. (2009) provide technical guidance on cloud security.

Faced with exponential data growth and increasing regulatory requirements, organizations want to protect their data and applications at reduced costs. Services exist today that enable security-rich managed protection of critical data on-site or off-site, help reduce the total cost of ownership with cloud-based data backup solutions, and ease management of various industry regulation requirements.

3.3. Simulation and Algorithm Scalability in the Cloud

The scale and complexity of the problems that future simulation algorithms will be required to solve, combined with the fundamental architecture of cloud systems, does not suit the traditional HPC model which relies on a much higher level of predictability from the underlying hardware and infrastructure. Standard cloud nodes are not well suited to run classical HPC algorithms, so new variations of algorithms are needed to make better use of the commodity hardware based cloud configurations. Only recently cloud providers are beginning to offer on-demand services for high performance computing.

3.4. Fault Tolerance

Moving an entire organization's business on to cloud systems can have dramatic consequences in case of faults like network outages. The benefit and the appeal of using public or semi-private computing clouds would be greatly diminished if fault-tolerance were not appropriately handled and if the solutions were not properly demonstrated to organizations. Extreme flexibility in the underlying infrastructure needs to be factored in when designing solutions that must be able to accommodate sudden faults and/or sudden increases in computing demands.

4. POTENTIAL SOLUTIONS

4.1. CAD User Experience

In order to avoid disrupting the user experience, a careful evaluation of applications and user behavior is necessary. CAD applications exhibit a peculiar kind of workflow that alternates phases of intense interaction to perform simple operations and review, to phases of intense computation, necessary to perform complex operations. Systems like "Snowbird/Snowflock" described in Lagar-Cavilla et al. (2007) are capable of detecting and reacting to this alternation performing incremental, transparent, live virtual machine migration from a user's workstation to the cloud and vice-versa. This capability provides an extremely useful new paradigm to extend application computing scalability using vast cloud-based resources, while minimizing changes in the application's user interface modules and general behavior.

4.2. Data Transfer and Data Security

To overcome data transfer latency to the cloud, two strategies offer some potential solutions:

- I/O operations originating from a users' workstation and directed to an organization's central storage system such as a Storage Area Network can be intercepted by dedicated appliances, which can batch, compress and encrypt a number of popular storage protocols such as iSCSI, and forwarded to the cloud's storage system. In case of network fault the operations are queued and transmitted after the network is restored.
- User's workstations are partially virtualized, with "shadow" virtual machines running in the cloud to replicate all the user's operations, including storage accesses.

Any of these strategies allow for a persistent copy of the relevant data to be maintained and updated in the cloud, ready for use by a number of scalable computing workloads.

Data security can also be augmented by using partitioning schemes, such that a complete dataset is never completely available in a single system, similar to distributed, peer-to-peer content distribution systems; this strategy is particularly well suited for large datasets representing models on which simulation is performed, as domain decomposition is then necessary to distribute the calculation over a large amount of computing nodes.

The application of distributed data exchange security models based on virtualization like Trusted Virtual Domains (TVD) in Linwood Griffin et al. (2005) and Bussani et al. (2005) would also be a further step towards the comprehensive system required by organizations to entrust data to a cloud provider.

4.3. Simulation and Algorithm Scalability in the Cloud

Due to the extreme commoditization of the architecture of cloud systems, there are several unique characteristics that must be accounted for to avoid hard scalability limitations. Still, several lessons can be learned from past research in distributed computing to leverage the very dynamic nature of these systems.

First, workloads must be created having a single virtual machine as the fundamental computing entity, relying on as little external information and communication as possible.

Second, new programming models that account for the dynamic and heterogeneous nature, and the massive scale of cloud systems, need to emerge to mitigate algorithm development complexity and to avoid continuous manual refactoring subject to the volatile underlying platforms and networks. For example, taking some hints from numerical optimization as described in Iorio (2009), an effective layering strategy can be highly beneficial.

Third, when the individual computing entity is a pre-configured virtual machine, adaptive, dynamic domain decomposition and computing strategies must also be adopted to leverage the extreme amount of computing potentially available to cloud-based applications.

4.4. Fault Tolerance

The fully virtualized nature of cloud systems allows for a more streamlined management of fault tolerance. Applications and workloads designed to run within a virtual machine are completely agnostic of their location, their individual performance, and the exact topology of the network fabric.

The characteristics of the infrastructure and the constraints posed by it on the applications actually result in a much easier exploitation of multiple fault-tolerance strategies that depend on this strict separation of information. As an example, modern cloud infrastructures can produce full virtual machine snapshots in a very short time, and upon detection of a software or hardware issue can provision a new virtual machine that maintains the same characteristics of the one that failed.

This feature makes the adoption of a "checkpointing" fault-tolerance strategy much easier from the application developer standpoint in a very elegant way. Traditional HPC systems, based on a message passing interface, either do not possess this characteristic, or must be designed and implemented at the application level, which is much more time consuming and error prone.

5. LIVING LAB CASE STUDIES

The discrepancy between the modeled and actual energy performance in completed construction projects is becoming clearer to practitioners and certification bodies. This trend, together with the recognition of the need for widespread implementation of massive energy efficiency improvements in both new and existing buildings, has led to a number of projects that provide testbeds for experimentation called "living labs." For a commercial building, the Digital 210 King project in Toronto (Attar et al. 2010) brings together a detailed Building Information Model (Eastman, 2007) with sensor data collected from the building control system as well as meters and submeters. At a sustainable community level, the Millennium Water Olympic Village project in Vancouver (Bayley et al. 2009) applied a multitude of efficiency measures to great effect achieving a close to net-zero energy community. At an urban level, the highly publicized yet-unrealized Masdar City project in Abu Dhabi (www.masdar.ae) aims to be entirely fueled by renewable resources, primarily from solar energy.

At regional and national levels, the U.S. Department of Energy recently announced the establishment of an Energy Innovation Hub focused on developing technologies to make buildings more energy efficient (U.S. DOE 2010). Led by The Pennsylvania State University, the Energy Innovation Hub is located at the Philadelphia Navy Yard Clean Energy Campus and will function as a living lab. Hub members include nine universities, two U.S. National Laboratories (Lawrence Livermore and Princeton), and five industrial firms including IBM. A highly multidisciplinary approach is underway to take full advantage of the diverse knowledge sets needed to develop a significantly improved outcome. As lead for Integrated Design Processes and Computational Tools, IBM is collaborating with other Hub members including Carnegie Mellon University, Lawrence Livermore National Lab, Penn State, Purdue, United Technologies Research Center, and Virginia Tech on this task. IBM will develop a cloud computing framework and infrastructure to support the interoperability of information and tools that will help ensure a range of small to mid-sized to large firms can take advantage of these technologies. In addition, IBM will implement performance monitoring and diagnostic systems for the integrated control systems.

6. CONCLUSION AND FUTURE OUTLOOK

The cloud environment offers the very compelling opportunity of using computing power on a scale previously unavailable to most users of design and engineering software applications. The main challenges are thus both in defining novel avenues to use that computing power and in ensuring that the user experience remains familiar to widely adopted standards, augmenting the user's capabilities rather than exacting new workflow compromises.

Promising techniques that leverage the already virtualized nature of the cloud environment exist and can be used to exploit the vast opportunities cloud computing has to offer to the engineering community, and will make it a viable option for architecture and engineering organizations to increase the efficiency of their practices and operations.

Moreover the constant reduction in the cost of cloud-based computing resources will dramatically increase the amount of computing power available on demand, thus enabling full democratization of an entire portfolio of powerful tools, with the prospect of both reductions in project costs and environmental impact, and increases in overall efficiency over the entire project lifecycle.

Acknowledgements

Francesco Iorio acknowledges all members of Autodesk Research for the excellent insight, and Prof. Eyal de Lara at University of Toronto for their work on cloud systems and for the discussions on application behavior in cloud computing.

Jane Snowdon gratefully acknowledges Bruce D'Amora and Florence Hudson for their review and useful comments.

References

ATTAR, R., PRABHU, V., GLUECK, M., AND KHAN, A. 2010. 210 KING STREET: A DATASET FOR INTEGRATED PERFORMANCE ASSESSMENT. IN PROCEEDINGS OF THE 2010 SPRING SIMULATION MULTICONFERENCE (SPRINGSIM '10). ACM, NEW YORK, NY, USA, ARTICLE 177 , 4 PAGES.

BAYLEY, R. 2009. THE CHALLENGE SERIES – MILLENNIUM WATER: THE SOUTHEAST FALSE CREEK OLYMPIC VILLAGE.

BUECKER, A., LODEWIJKX, K., MOSS, H., SKAPINETZ, K., WAIDNER, M. 2009. CLOUD SECURITY GUIDANCE: IBM RECOMMENDATIONS FOR THE IMPLEMENTATION OF CLOUD SECURITY. IBM REDPAPER, 2009.

BURNS, J., COLLIS, S., GROSH, J., JACOBSON, C., JOHANSEN, H., MEZIC, I., NARAYANAN, S., WETTER, M. 2010. APPLIED AND COMPUTATIONAL MATHEMATICS CHALLENGES FOR DESIGN, DYNAMICS AND CONTROL OF ENERGY SYSTEMS. DRAFT WHITEPAPER V3. http://www.engineering.ucsb.edu/~mgroup/wiki/images/1/17/WP_Computation_Bulidings_November_2010_Version_3.docx

BUSSANI, A., ET AL. 2005. TRUSTED VIRTUAL DOMAINS: SECURE FOUNDATIONS FOR BUSINESS AND ITS SERVICES. IBM RESEARCH REPORT RC23792.

CHU, S. 2009. STATEMENT OF STEVEN CHU SECRETARY OF ENERGY BEFORE THE COMMITTEE ON SCIENCE AND TECHNOLOGY. U.S. HOUSE OF REPRESENTATIVES, WASHINGTON, D.C., U.S.A. http://www.energy.gov/news/7055.htm

EASTMAN, C., TEICHOLZ, P., SACKS, R. AND LISTON, K. 2007. BIM HANDBOOK: A GUIDE TO BUILDING INFORMATION MODELING FOR OWNERS, MANAGERS, DESIGNERS, ENGINEERS AND CONTRACTORS. JOHN WILEY & SONS, NEW JERSEY,108.

ENERGY EFFICIENCY AND RENEWABLE ENERGY, U.S. DEPARTMENT OF ENERGY, 2007. 2007 BUILDINGS ENERGY DATA BOOK. http://buildingsdatabook.eren.doe.gov/docs/DataBooks/2007_BEDB.pdf

FROST AND SULLIVAN 2009. BRIGHT GREEN BUILDINGS: CONVERGENCE OF GREEN AND INTELLIGENT BUILDINGS, HTTP://WWW.CABA.ORG/BRIGHTGREEN

HUTSON, S.S., BARBER, N.L., KENNY, J.F., LINSEY, K.S., LUMIA, D.S., AND MAUPIN, M.A., 2004. ESTIMATED USE OF WATER IN THE UNITED STATES IN 2000: RESTON, VA., U.S. GEOLOGICAL SURVEY CIRCULAR 1268, 46 P.

IBM 2010. SMARTER PLANET PRIMER. HTTP://SMARTERPLANET.TUMBLR.COM/POST/58347576/A-SMART-PLANET-PRIMER-TO-VIEW-FULL-SCREEN-CLICK

IGARASHI, T. AND HUGHES, J. 2001. A SUGGESTIVE INTERFACE FOR 3D DRAWING. UIST '01. ACM, PP. 173-181.

IORIO, F. 2009. ALGORITHMS RE-ENGINEERING AS A FUNDAMENTAL STEP TOWARDS EXPLOITABLE HYBRID COMPUTING FOR ENGINEERING SIMULATIONS. CECAM WORKSHOP: ALGORITHMIC RE-ENGINEERING FOR MODERN NON-CONVENTIONAL PROCESSING UNITS.

IPCC 2007: METZ, B., DAVIDSON, O.R., BOSCH, P.R., DAVE, R., MEYER, L.A. (EDS). 2007. CLIMATE CHANGE 2007: MITIGATION OF CLIMATE CHANGE. CONTRIBUTION OF WORKING GROUP III TO THE FOURTH ASSESSMENT REPORT OF THE INTERGOVERNMENTAL PANEL ON CLIMATE CHANGE, CAMBRIDGE UNIVERSITY PRESS, 851 PP.

LAGAR-CAVILLA, H. ANDR'ES, TOLIA, N., DE LARA, E., SATYANARAYANAN, M., O'HALLORON, D. 2007. INTERACTIVE RESOURCE-INTENSIVE APPLICATIONS MADE EASY. ACM/IFIP/USENIX 8TH INTERNATIONAL MIDDLEWARE CONFERENCE. PP. 143-163.

LEE, E. S., HUGHES, G. D., CLEAR, R. D., FERNANDES, L. L., KILICCOTE, S., PIETTE, M. A., RUBINSTEIN, F. M., SELKOWITZ, S. E. 2007. DAYLIGHTING THE NEW YORK TIMES HEADQUARTERS BUILDING, FINAL REPORT: COMMISSIONING DAYLIGHTING SYSTEMS AND ESTIMATION OF DEMAND RESPONSE, LAWRENCE BERKELEY NATIONAL LABORATORY, BERKELEY, CA.

LINWOOD GRIFFIN, J., JAEGER, T., PEREZ, R., SAILER, R., VAN DOORN, L., C´ACERES, R. 2005. TRUSTED VIRTUAL DOMAINS: TOWARD SECURE DISTRIBUTED SERVICES. FIRST WORKSHOP ON HOT TOPICS IN SYSTEM DEPENDABILITY.

LUNZER, A. AND HORNBÆK, K. 2008. SUBJUNCTIVE INTERFACES: EXTENDING APPLICATIONS TO SUPPORT PARALLEL SETUP, VIEWING AND CONTROL OF ALTERNATIVE SCENARIOS. ACM TRANSACTIONS ON COMPUTER-HUMAN INTERACTION. 14, 4, ARTICLE 17. 44 PP.

PLANYC: A GREENER, GREATER NEW YORK. 2007. http://www.nyc.gov/html/planyc2030/downloads/pdf/full_report.pdf

ROODMAN, D.M., LENSSEN, N. 1995. A BUILDING REVOLUTION: HOW ECOLOGY AND HEALTH CONCERNS ARE TRANSFORMING

Construction. Worldwatch Paper 124, Worldwatch Institute, Washington, D.C.

Shea, K., Aish, R., Gourtovaia, M. 2005. Towards integrated performance-driven generative design tools. Automation in Construction, 14(2). pp. 253-264.

U.S. Department Of Energy, Energy Information Administration 2008a. Emissions of Greenhouse Gases Report. http://www.eia.doe.gov/oiaf/1605/ggrpt/pdf/0573(2008).pdf

U.S. Department Of Energy, Energy Information Administration 2008b. Assumptions to the Annual Energy Outlook. http://www.eia.doe.gov/oiaf/aeo/assumption/pdf/0554(2008).pdf

U.S. Department Of Energy, Environmental Information Administration 2008c. EIA Annual Energy Outlook 2008. http://www.eia.gov/oiaf/aeo/pdf/0383(2008).pdf

U.S. Department of Energy. 2010. Penn State to Lead Philadelphia-Based Team that will Pioneer New Energy-Efficient Building Designs. http://www.energy.gov/news/9380.htm

World Business Council for Sustainable Development. 2009. Transforming the Market: Energy Efficiency in Buildings. www.wbcsd.org.

Zhu, J., Fang, X., Guo, Z., Niu, M., Cao, F., Yue, S., Liu, Q. 2009. IBM Cloud Computing Powering a Smarter Planet. CloudCom 2009, LNCS 5931, pp.621-625.