

Automated Synthesis of Sustainable Data Centers

Tom Christian, Yuan Chen, Rocky Shih, Ratnesh Sharma,
Christopher Hoover, Manish Marwah, Amip Shah, Daniel Gmach

Sustainable IT Ecosystem Laboratory
Hewlett-Packard Laboratories
Palo Alto, CA USA
first.last@hp.com

Abstract—Next generation data centers must be designed to meet Service Level Agreements (SLAs) for application performance while reducing costs and environmental impact. Traditional design approaches are manually intensive and must integrate thousands of components at multiple granularities, often with conflicting goals. We propose an Automated Data Center Synthesizer to design Sustainable Data Centers that meet SLA goals, minimize carbon emissions and embedded exergy, are optimally efficient and deliver significantly reduced Total Cost of Ownership (TCO). The paper concludes with a use case study that employs the synthesizer process flow to design an optimal data center to deliver a set of services for a hypothetical city using state of the art sustainable technologies.

Index Terms—carbon footprint, data center, embedded exergy, synthesis, sustainable, Total Cost of Ownership.

I. INTRODUCTION

THE increased demand for Information Technology (IT) Services has resulted in world-wide proliferation of data centers with racks of densely packed blade servers. IT currently accounts for 2% of the worldwide carbon footprint, with a 50% increase expected by 2020 [1]. A large fraction of that growth is expected to come from data centers. Data center owners and operators face regulatory and social pressure to reduce the environmental impact of their facilities, and it's becoming increasingly important to reduce TCO since service costs are still higher than the market can bear in emerging economies, leaving large communities without access to the services most of us take for granted. TCO remains the most critical driver for corporate CIOs.

Data center design is a complicated laborious process involving many trades: architects design the building, mechanical engineers the cooling infrastructure, electrical engineers the power distribution and IT professionals the computing, networking and storage infrastructure. The requirements passed from trade to trade and across design stages are minimal, often on paper. This design approach leaves much of the design space unexplored and can miss better designs, often leading to inefficient or sub-optimal solutions. Design automation offers an alternative that can produce better designs and reduce non-recurring engineering expenses, capital expenses and operating expenses.

II. DATA CENTER SYNTHESIZER

We propose a Data Center Synthesizer to automate and optimize data center design to facilitate meeting sustainability, SLA and TCO goals. The objective is to capture best practices, incorporate sustainability criteria into the design goals and add helpful automation to a process that is now largely manual. An outline of the Synthesizer process and data flow is shown in Figure 1.

A. Services Analysis

Data center design begins with a description of the services to be provided and any constraints on the design. The application sizing step generates equipment lists for services, including computing, storage, and networking requirements. Multi-tier applications represent a wide class of applications in today's data centers. We have developed an SLA-driven approach that can transform Service Level Objectives (SLOs) for multi-tier applications into the computing resources required to deliver those services [2][3].

Our approach combines performance modeling with performance profiling to create models that automatically generate IT equipment requirements for multi-tier applications from high level SLO requirements (Figure 2). It undertakes the sizing task in a systematic way following a series of steps. First, we benchmark each application on a base platform and generate a detailed resource and performance profile for the application on that platform; i.e., the resource demand of each request and service rate. The profiles are obtained by performing statistical regression analysis on historical data [4].

Once we have obtained the profile data for a base platform, we apply cross-platform models [7] to derive the performance and resource profiles for other platforms or equipment. The results are stored in a profile repository. Next, we size the applications. Given SLOs — such as the number of users, response time and throughput requirements — and equipment specifications, we use a queueing performance model and a constraint solver to determine the application's resource requirements; e.g., how many servers of what type are required at each tier of the application to meet the SLOs. It also generates meta-information such as cost, performance, power consumption, etc.

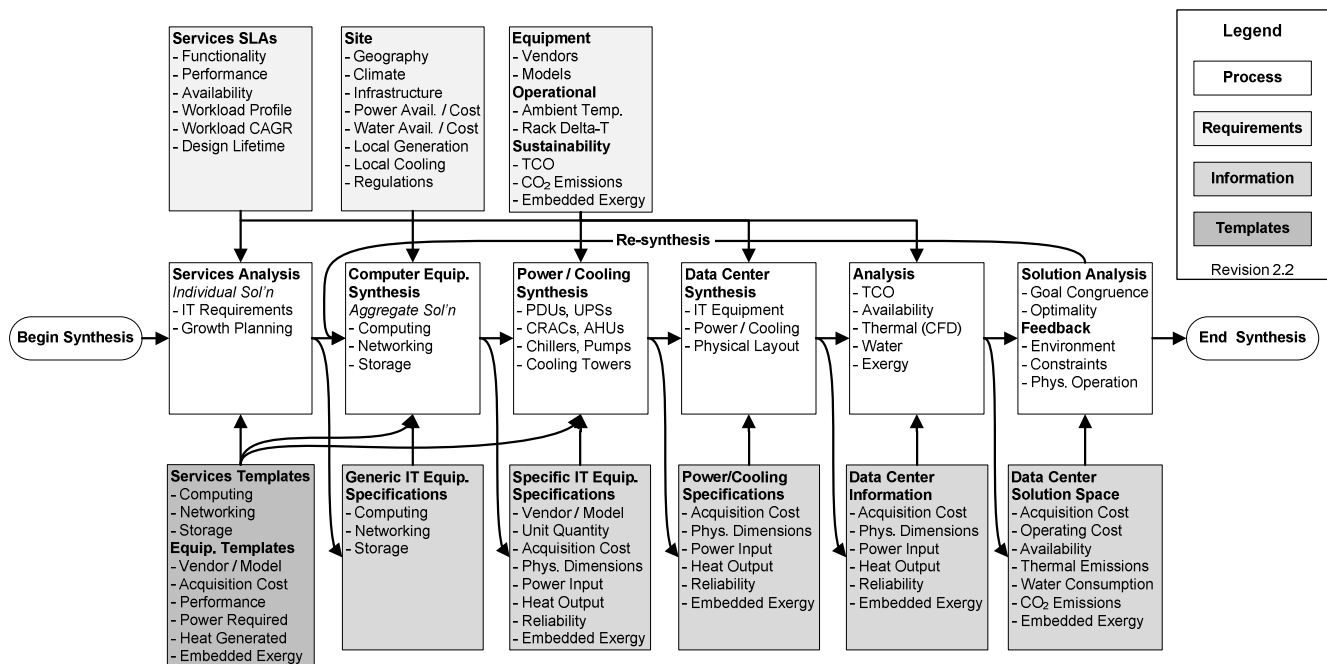


Figure 1 - Data Center Synthesizer Process and Data Flow

The key component of our sizing approach is an analytical performance model for multi-tier applications, where each tier is modeled as a multi-station queueing center [4, 8]. We combine the performance model with the application resource profile and formalize the sizing problem as a constraint satisfaction problem. A constraint solver is currently under development. We consider different platform specifications as separate initial conditions and generate multiple solutions. We compare the different solutions and select an optimal one based on the desired constraints and criteria in the computer equipment synthesis step. An advantage of our approach is that once an application profile is created, it can be used repeatedly to perform sizing for different SLOs, platforms or equipment. Further, if the workload changes we need only re-solve the constraint satisfaction problem with new parameters.

The salient features of our sizing approach can be summarized as follows:

- **SLA-driven:** Unlike traditional sizing tools, our approach takes into account SLOs specified in Service Level Agreements (SLAs). Ensuring adequate computing performance to meet SLAs is a critical requirement in data center design.
- **Generic:** Traditional sizing tools are application-specific and different applications require different sizing tools. Our approach is sufficiently general to be applied to a wide class of commonly used multi-tier applications with any number of tiers and transactions. Moreover, our model can deal with dynamically changing workloads in both request volume and transaction mix.
- **Practical:** Our approach simplifies the modeling process. We do benchmarking once on a base platform. If the workload or response time changes, we re-solve the constraint satisfaction problem with new parameters. Further, our approach is non-intrusive in the sense that it requires no instrumentation; the data used in our approach

are readily available from standard system and application monitoring.

- **Incorporation of Sustainability Metrics:** We include sustainability metrics in the input and output of our sizing approach, which enables us to select and optimize the solutions in terms of sustainability goals in next steps.

B. Computer Equipment Synthesis

The application sizing step produces requirements for computing, storage and networking equipment. The computer equipment synthesis step compares the requirements with the available makes and models of equipment and selects the most appropriate solution. This is accomplished by comparing the output of the services analysis step with the equipment templates, each of which contains complete specifications for a particular piece of equipment. Attributes include vendor, model, relative performance, acquisition cost, power consumption, heat generation, embedded exergy and embedded carbon emissions. Performance is the first requirement for selection: if the equipment is not able to meet the service SLAs, it will probably not be an acceptable solution; equipment that is close to meeting the SLAs may be retained for customer consideration if it has other desirable attributes, such as much lower TCO or sustainability impact. It is also important for the selected equipment to accommodate planned workload growth without major disruption of the data center.

Low acquisition cost is always an important selection criterion, but a higher price can sometimes be offset by lower operating costs over the design lifetime of the data center. In data centers, equipment operating costs derive mainly from power consumption. Hence, a low-power alternative that offers acceptable performance may deliver better TCO than the higher-performing alternative. It is also the case that virtually all power going into a data center must eventually be

removed as heat. Power consumption to remove that heat can range from 10-30% of the power consumed by the computing equipment for current best-case practices and may be as high as 100% in legacy data centers.

Computer equipment requirements may include direct selection by vendor and model, either in support of the

specified services or as an adjunct to the equipment that will be selected for those services. In the future, legacy equipment will be incorporated into the process using this capability since a majority of data centers will not have the luxury of a greenfield design.

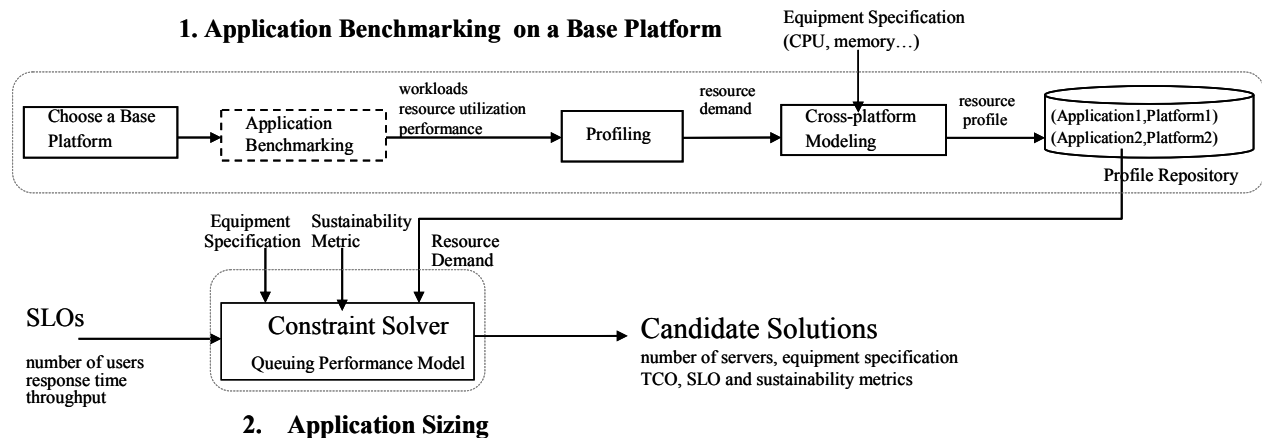


Figure 2 - Sizing of Multi-tier Applications

C. Power / Cooling Synthesis

Power distribution and cooling equipment are selected to meet the operating characteristics of the computing, storage and networking equipment. The rack inlet and outlet air temperatures, required airflow and planned ambient operating temperature guide selection of the cooling equipment. Using workload and growth profiles from the services specifications combined with knowledge of real-time performance management capabilities for software service instances and data center hardware, the Synthesizer can specify a minimal solution that meets requirements and allows for economical future expansion.

Effective use of the alternative power and cooling solutions offered by the local environment can yield significant improvement in TCO and carbon emissions. Climate and wind characteristics may be entered as characteristic zones that enable decisions about alternative solutions without the need for an extensive database of local attributes, or these values may be entered directly.

Some local knowledge is essential in this step. For example, the cost and availability of power must be specified; there are certain locations in the world where data centers must run using on-site generators because reliable grid power simply doesn't exist. Similarly, the availability or lack of abundant water may suggest or rule out certain cooling solutions. Regulations may also play a role as demand for sustainable solutions increases.

Supply side infrastructure includes heat and power sources that generate electrical power and chilled water (or cold air) for operation of datacenter. The essential inputs to synthesis are:

- Data center power and cooling demand profiles with desired availability levels.
- Utility electricity and natural gas tariffs.

- Capital, operating and maintenance, and fuel costs of the various alternatives.
- Available energy sources.
- Available distributed energy conversion technologies, such as fuel cells, photovoltaic, etc.
- Rate of carbon emissions from the macro-grid and from on-site power generation of power and heat.
- Carbon tax rates or allowable emissions.
- Thermodynamic parameters governing the use of combined heat and power technologies and energy storage.
- Waste heat utilization strategy (e.g., use of waste heat for desalination, building heating, digesters, etc.).

The objective of the power/cooling synthesis step is to minimize the total cost of ownership by maximizing the utilization of available energy to deliver power and remove heat at the desired availability and emissions levels. A typical window for such optimization would be a year.

Outputs of this process step include:

- Choice of technologies and installed capacities.
- Operating schedules and policies for control and management.
- Total cost of ownership.
- Total carbon emissions.
- Coefficient of Performance (COP) of the ensemble, which is the ratio of the change in heat to the supplied power.

Certain design outcomes can help in site selection based on geographical conditions and local regulations. Others can gate the level of availability of critical IT services. The scope of the supply side infrastructure can also be explored to include societal impacts like job creation and revenue generation for the data center facility.

D. Data Center Synthesis

The next step is projection of a candidate data center solution. This step results in a complete model and physical layout of the data center, incorporating all computing, networking and storage devices and the power and cooling infrastructure required to support them. Not included in the design is the building itself, although the principal characteristics of the computer room will be included.

Data center layout will be guided by design best practices and will incorporate an automated connection to existing modeling tools for further analysis.

E. Data Center Analysis

The candidate data center is then analyzed to determine its operating characteristics. The analysis includes sustainability criteria, performance, availability, thermal characteristics and TCO. The sustainability criteria include power consumption for computing and cooling, operational and embedded carbon emissions and heat load. The sustainability analysis also considers embedded exergy, which accounts for the loss in exergy (usable energy) for all components and manufacturing processes associated with each piece of equipment, from raw materials extraction to finished product.

Some characteristics can be calculated or measured directly, such as power consumption. Knowledge about the operating policies, such as real-time management of computing loads to minimize power consumption, will be incorporated into these calculations. Other characteristics, such as availability, must be inferred from aggregated equipment reliability attributes and known industry best practices.

This step includes a complete thermal analysis of the operation of the data center using existing Computational Fluid Dynamics (CFD) analysis tools, such as *6SigmaRoom* from Future Facilities, Limited [5]. Maximizing the efficiency of a data center typically involves running at a higher ambient temperature than has been the practice in the past and that necessitates a thorough and accurate understanding of the thermal operating characteristics and the impact on reliability.

The output of the data center analysis step is a data point in n-space that represents the combined attributes of the candidate data center.

F. Solution Analysis

Using the result of the previous step, the candidate solution is compared with customer-specified goals for the data center and with other solutions from previous iterations. Provided that the solution meets the SLA objectives for services – or is reasonably close to those goals – it becomes a candidate for consideration as the recommended solution.

Arriving at an optimal or near-optimal synthesis requires solution of a constraint optimization problem [13], with hard constraints such as performance against the SLA and multiple soft constraints such as carbon emissions goals. Regular analytic techniques won't work because not every constraint can be modeled as an algebraic function. One of the future challenges in automating the process will be finding a way to characterize, evaluate and optimize multiple constraints with

no convenient closed-form representation.

G. Re-Synthesis

If the data center design does not meet key customer requirements, or if there are other potential design solutions that might represent improvements, a re-synthesis process uses the results to modify the specifications and generate another candidate data center. In a future version of the Synthesizer, iteration will continue in a fully-automated loop until an optimal solution is reached. The objective is to evaluate multiple solutions and fill as much of the n-space as possible to enable both objective and subjective evaluation of the range of possibilities, an exercise made possible by the re-synthesis step.

The Data Center Synthesizer is one element of the HP Labs Sustainable Data Center project. It will be accompanied by two other modules that will provide real-time management of software service instances and of the power, cooling and computing hardware infrastructure. One deliverable of the data center design solution will be the initial operating policies for real-time management of the data center. During actual operation, data will be collected to evaluate the results of operational decisions and, ultimately, those results will be fed back to the Synthesizer to further refine its assumptions about real-time management capabilities and recommended initial operating policies. This step is also forms a re-synthesis loop, but with a longer period than for the initial design evaluation process.

III. USE CASE VALIDATION

We have evaluated our approach through a use case that involves design of a hypothetical data center for a small city. We are required to build a data center that will support:

- E-mail services for 40,000 city residents,
- Human Resource Management for 1500 businesses,
- Real-time infrastructure monitoring for 10,500 sensors spread throughout the city to track temperature, water, power, transit, weather and air quality, and
- A high-performance, cloud computing environment with approximately 2000 cores for the local university.

Our hypothetical city has a seaside location and a desert climate; saltwater is available in abundance, but fresh water is scarce. Sunshine is plentiful all year long and rainfall is infrequent, an obvious opportunity for photovoltaic power generation. There is a predictable prevailing breeze that is reasonably consistent in direction, but it lacks the velocity required for wind power generation. The mean daily low/high temperatures range from 12/24°C in the winter months to 29/42°C in the summer, making it possible to cool a data center with outside air during the evening for at least part of the year.

Municipal grid power is generated by natural gas and may be purchased for \$0.09/KWH. The CO₂ emission rate for grid power is estimated to be 0.6 kg/KWH; the CO₂ emission rate for solar generated power it is considered to be negligible in this analysis.

A. Data Center Design

Starting with these requirements, we generated several data center designs following the process and data flow described in Figure 1, but using manual processes since the Synthesizer is not yet operational. We integrated the output of application configurators [6], UPS/PDU sizing tools and Future Facilities 6SigmaRoom CFD software, calculated the operational energy cost and carbon emissions and performed reliability analyses. A detailed description of the optimal data center follows in Table 1.

Servers		
E-mail	22	ProLiant BL465c
HRM	16	ProLiant BL460c
Real-time	4	ProLiant BL465c
HPC Cloud	240	ProLiant DL165G5p
Storage		
Shared	9	EVA4400/8D, 380TB
Networking		
Shared	7	ProCurve 2900-48G
Shared	1	ProCurve 3500yl-48G
Power and Cooling		
Shared	1	PDU
Shared	1	UPS
CRAC	2	Chilled water AHU
DX	1	Tower, Chiller
Layout		
Room size: L 28 ft. x W 26 ft. x H 10 ft.		
Floor plenum: 24 in. (no obstructions)		
Ceiling plenum: 36 in. (no obstructions)		
9 @ 2 KW storage racks		
2 @ 9 KW E-mail/HRM/RT server racks		
6 @ 15 KW HPC cloud racks		
1 @ empty rack (for future expansion)		
Total heat load: 126 KW		

Table 1 - Optimal Data Center Characteristics

The data center layout was generated manually for this use case then analyzed using Future Facilities 6SigmaRoom. The data center layout is shown in Figure 3 and the thermal map from CFD analysis is shown in Figure 4. The relatively slow performance of CFD analysis software presents an issue when attempting fast, iterative solutions. We are working with CFD vendors and other researchers to explore fast, alternative mechanisms that could indicate quickly if a solution is likely to be acceptable or not. Once a small number of candidate solutions has been identified, they could be subjected to a detailed CFD analysis to verify the "ballpark" results.

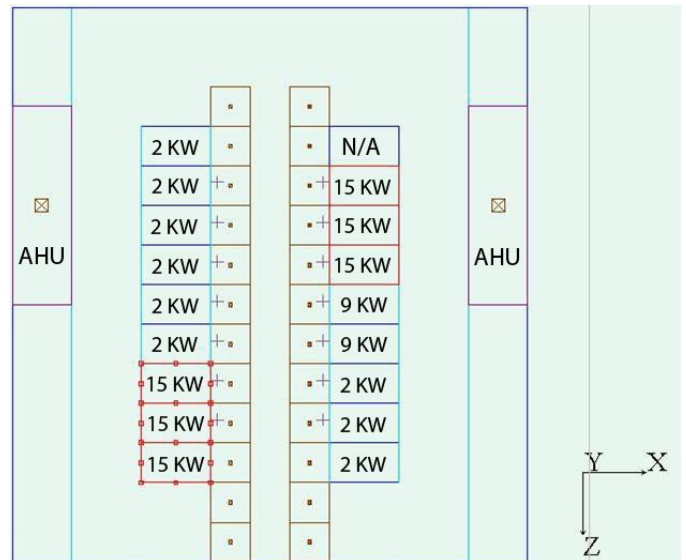


Figure 3 - Data Center Floor Layout




Figure 4 - Data Center Thermal Analysis

B. Supply-side Infrastructure

The design proposes a combination of photovoltaic and grid power for the data center. Because of the abundant sunshine, photovoltaic was an obvious choice. A 600KW photovoltaic array is the primary source of power for the data center, supported by grid power at night and when conditions are not adequate for solar generation. Excess power generated during the day will be sold back to the city to offset costs, and it provides a generous margin for future growth.

We considered, but did not utilize, incineration of municipal solid waste as an alternative energy source.

C. Demand-side Infrastructure

Although our city is in a desert environment, we determined that it is possible to use outside air for up to 70% of the cooling requirements during the winter months and up to 30% during the rest of the year. This results in reduced power

consumption and operating costs for cooling. Perhaps more importantly, it results in reduced consumption of fresh water. Sea water can be used in cooling towers, but not without complications. The balance of cooling requirements are met with a conventional chiller and cooling tower installation.

We considered, but did not utilize, heat rejection to the sea and adsorption cooling as additional possibilities. The current design also makes no assumptions about alternative use of waste heat; that is a possible enhancement for the future.

IV. CONVENTIONAL APPROACHES

It is typically the case that data center designers use limited static information about loading and operating conditions when specifying computing hardware and the supporting infrastructure. To ensure adequate capacity, power and cooling requirements are calculated assuming all equipment is in operation and under nominal load. SLAs for supported services are not considered for purposes of optimization, and it is most often the case that data center designers don't plan for or consider the possibility of real-time load management or facilities management to optimize resource utilization.

Provisions for data center availability are left to the judgment of the designer. There are well-regarded recommendations for component redundancy to achieve various tiers of availability, such as those recommended by the Uptime Institute [9], but conventional guidelines such as these have been challenged by providers like Google, which provisions with low-cost commodity hardware and uses their software stack to achieve high availability [10]. Lacking the proper resources to generate an optimal solution, the usual result is over-provisioning to ensure adequate performance and availability. That solution will achieve the performance and availability goals, but could result in excessive TCO.

In the past, little or no thought was given to the environmental footprint of the data center. The first LEED-certified [11] data center was designed by EYP Mission Critical Facilities (now a division of the Hewlett-Packard Company) for Fannie Mae's Urbana Technology Center (UTC) located in Maryland, in 2005 [12]. Attention to sustainability issues, most often seen as a focus on reduced power consumption, is now much more common.

V. CONCLUSIONS

The Sustainable Data Center is a vision that is taking shape through several integrated projects at Hewlett-Packard Laboratories. The Data Center Synthesizer, an element of the Sustainable Data Center, is intended to facilitate the design of a data center that is optimum for a specific set of services, with particular emphasis on improved sustainability of the solution and reduced TCO.

While some of the objectives proposed for the Synthesizer can now be achieved through manual operations, the proposed approach will ultimately enable a holistic data center design that is appropriately provisioned to the user's needs, meets performance SLAs and has minimal environmental impact. The planned automation will incorporate design best practices

and enable multiple design alternatives to be explored and evaluated in much less time than is now required. In addition, knowledge of operational performance characteristics and real-time management capabilities will enable accurate estimates of TCO and associated design trade-offs. Our future work will focus on specification and characterization of the elements involved in each process step, methods to analyze the relative merits of each candidate solution, and automation of the iterative loop.

REFERENCES

- [1] The Climate Group, "SMART 2020: Enabling the low carbon economy in the information age", 2008.
- [2] G. J. Janakiraman, J. R. Santos, Y. Turner, "Automated multi-tier system design for service availability," HP Labs Technical Report, HPL-2003-109.
- [3] G. J. Janakiraman, J. Renato Santos, Y. Turner, "Automated system design for availability", HP Labs Technical Report, HPL-2004-54.
- [4] Y. Chen, S. Iyer, A. Sahai, D. Milojicic, "A Systematic and Practical Approach to Generating Policies from Service Level Objectives". Proceedings of the 11th IFIP/IEEE International Symposium on Integrated Network Management (IM 2009), 2009. (To appear)
- [5] Further information about Future Facilities Ltd. Software is available at <http://www.futurefacilities.com/software/6SigmaRoom.htm>.
- [6] Hewlett-Packard Company has solution configurators and application sizers for many products and software services. They are available at: <http://egui.houston.hp.com/eGlue/eco/begin.do>.
- [7] C. Stewart, T. Kelly, A. Zhang, K. Shen, "A Dollar from 15 Cents: Cross-Platform Management for Internet Services", Proceedings of USENIX Annual Tech 2008.
- [8] Y. Chen, S. Iyer, X. Liu, D. Milojicic, A. Sahai, "Translating Service Level Objectives to Lower Level Policies for Multi-tier Services", Journal of Cluster Computing (CC), Special Issue, 2008.
- [9] W. Pitt Turner IV, John H. Seader, Vince Renaud, Kenneth G. Brill, "Tier Classifications Define Site Infrastructure," Uptime Institute white paper, 2008.
- [10] Luiz André Barroso, Jeffrey Dean, Urs Hölzle, "Web Search for a Planet: The Google Cluster Architecture," IEEE Micro, Volume 23, Issue 2, March-April 2003. pages 22-28.
- [11] U.S. Green Building Council, <http://www.usgbc.org/>.
- [12] William J. Kosik, "Current Thinking in Data Center Energy and Environmental Topics," <http://www.sustainablefacility.com/>, October 12, 2007.
- [13] Rina Dechter, Constraint Processing, Morgan Kaufmann Publishers, 2003.