Energy Aware Grid: Global Workload Placement based on Energy Efficiency

Chandrakant Patel, Ratnesh Sharma, Cullen Bash, Sven Graupner
Internet Systems and Storage Laboratory
HP Laboratories Palo Alto
HPL-2002-329
November 21st , 2002*

E-mail: {chandrakant.patel, ratnesh.sharma, cullen.bash, sven.graupner}@hp.com

Grid Computing, energy-efficiency, workload placement, cooling, data center, utility computing

The concept of Grid, based on coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations, is emerging as the new paradigm in distributed and pervasive computing for scientific as well as commercial applications. We assume a global network of data centers housing an aggregation of computing, networking and storage hardware. However, increased compaction of such devices in data centers has created thermal and energy management issues that inhibit sustainability of such a global infrastructure. In this paper, we propose the blueprint of Energy Aware Grid that will provide a global utility infrastructure explicitly incorporating energy efficiency and thermal management among data centers. Designed around an energy-aware co-allocator, workload placement decisions will be made across the Grid, based on data center energy efficiency coefficients. The coefficient, evaluated by the data center's resource allocation manager, is a complex function of the data center thermal management infrastructure and the seasonal and diurnal variations. A detailed procedure for implementation of a test case is provided with an estimate of energy savings to justify the economic viability of such a proposition.

* Internal Accession Date Only
Approved for External Publication

Energy Aware Grid: Global Workload Placement based on Energy Efficiency

Chandrakant Patel, Ratnesh Sharma, Cullen Bash, Sven Graupner
Hewlett-Packard Laboratories
1501 Page Mill Road, Palo Alto, CA 94304
email: {chandrakant.patel, ratnesh.sharma, cullen.bash, sven.graupner}@hp.com

Abstract.

The concept of Grid, based on coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations, is emerging as the new paradigm in distributed and pervasive computing for scientific as well as commercial applications. We assume a global network of data centers housing an aggregation of computing, networking and storage hardware. However, increased compaction of such devices in data centers has created thermal and energy management issues that inhibit sustainability of such a global infrastructure. In this paper, we propose the blueprint of Energy Aware Grid that will provide a global utility infrastructure explicitly incorporating energy efficiency and thermal management among data centers. Designed around an energy-aware co-allocator, workload placement decisions will be made across the Grid, based on data center energy efficiency coefficients. The coefficient, evaluated by the data center's resource allocation manager, is a complex function of the data center thermal management infrastructure and the seasonal and diurnal variations. A detailed procedure for implementation of a test case is provided with an estimate of energy savings to justify the economic viability of such a proposition.

1. Introduction.

Computing will be pervasive, and enablers of pervasive computing will be interconnected data centers housing computing, networking and storage hardware [1]. The expansion of computational grids will lead to an infrastructure of geographically dispersed locations (data centers) where computational resources are offered. The provision of services across such large, distributed and heterogeneous platforms requires a distributed resource management mechanism with capabilities beyond traditional centralized management.

Grids will immensely expand choices for placing computational workloads, choices that can be utilized for increasing energy efficiency by delegating workloads such that resources in less energy-efficient data centers can be released and powered down (along with air conditioning and cooling) saving energy. Figure 1 shows the heat dissipated by the hardware and the energy used to deliver the cooling. A data center, with 1000 racks, approximately 25,000 square feet, would require 10 MW of power for the computing infrastructure. At this heat dissipation, an additional 5 MW would be required to remove the dissipated heat. At $100/MWh, the cooling alone would cost $4 million per annum.

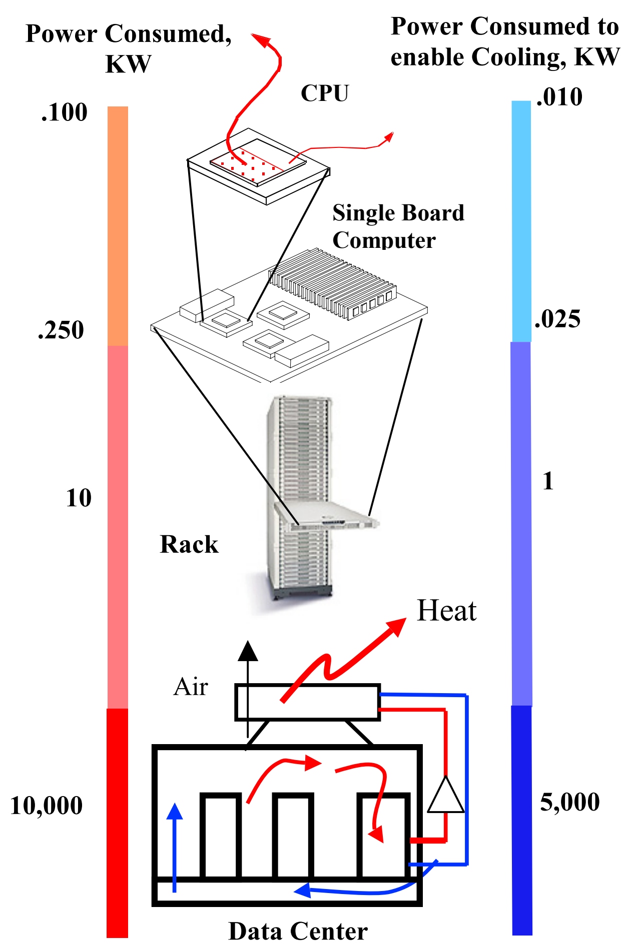

Figure 1. Power used by physical plant.

Furthermore, continuing resources commoditization has widened the spread between average and peak utilization of resources. Studies show [2] an average utilization of a resource infrastructure of 20% with the remaining 80% provisioned for rare peak loads and increased availability. In average, 80% of the resources are not utilized, yet consuming power, generating heat, absorbing cooling capacity – another potential for saving energy by partially powering down spare resource capacity, yet keeping resources available when they are needed.

Data center thermo-mechanical architecture.

Figure 2 indicates a typical data center cooling infrastructure with raised floor plenum. A modern data center is a heterogeneous mix of hardware and software delivering a multiplicity of services. It typically, has rows of racks packed with multiple air-cooled hardware supplied with cold air from the air conditioning units (see Fig. 2). The hot exhaust from the hardware is returned to the air conditioning units. The heat extracted by the air conditioning unit is transferred to the external environs of the data center. At the high compaction possible with today's slim servers and the ensuing high heat density deployment, the airflow rates are very high in the data center. The heat dissipation by a server rack, when fully loaded, stands at 10 KW. Similarly, networking and storage components aggregate to 5 to 10 KW at rack level. In addition to this type of physical configuration, there exist "bladed" servers – multiple single board computers packaged at even denser pitch in a single enclosure.

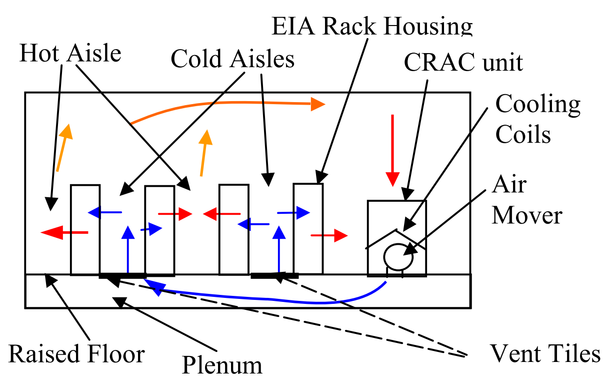

Figure 2. Typical under floor air cooling data center configuration with room return.

From the representation in Fig.1, one can see that energy to avail the cooling in chips and systems is approximately 10 % of the total heat dissipated. However, for data center it is 50% of the total heat dissipated. The difference stems from the additional thermodynamic work done in cooling the air. The systems are cooled by fans that move the air across the hot components (flow work), whereas the data center has the work associated with reducing the temperature of the return air by means of reverse power cycle. The work used to achieve the temperature reduction, and the work associated with moving the air in the data center and the condenser, add up to approximately 50% (based on use of direct expansion air conditioning units) of the total heat load in a data center (see Fig. 3).

In Fig.3, the evaporator provides cold air to the data center while the condenser dissipates heat to the ambient. Although several variations of this cycle are in existence, they all have the same principle of operation [4].

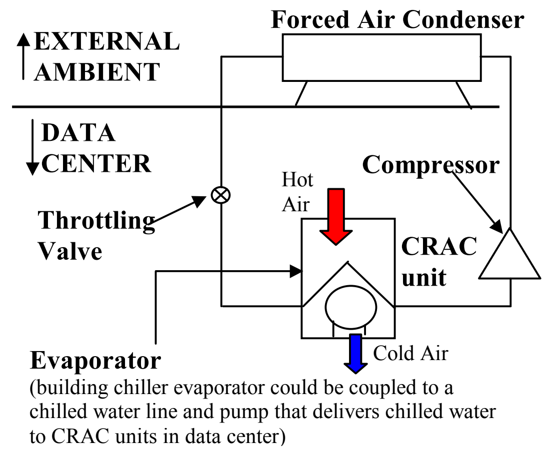

Figure 3. Data center cooling system using direct expansion Vapor compression refrigeration cycle.

Apart from inherent characteristics of the cooling system, the complex thermo-fluid behavior in each data center creates non-uniform temperature and airflow patterns. Unlike a modern home, where defined room boundaries and low heat loads can be adequately monitored and controlled by a single temperature sensor, the data center has no well defined boundaries. In fact, the state of the art data center lacks an adequate control mechanism to dissipate high heat loads.

Prior research conducted by Patel *et al.*[6] has shown the potential methods and benefits of apportioning cooling resources based on the dynamic needs of the data center. In this paper, we propose an energy-aware co-allocator that redistributes computing within the global network of data centers. We also examine the methods for evaluating energy efficiency and thermal management parameters applicable to any data center cooling infrastructure.

2. Methodology.

Exploiting energy-saving potential requires awareness of the resource management system of properties of the underlying physical world. This section describes how this can be achieved in the Globus resource management architecture [3]. Figure 4 shows the Globus resource management architecture with extensions for taking energy-related information into account.

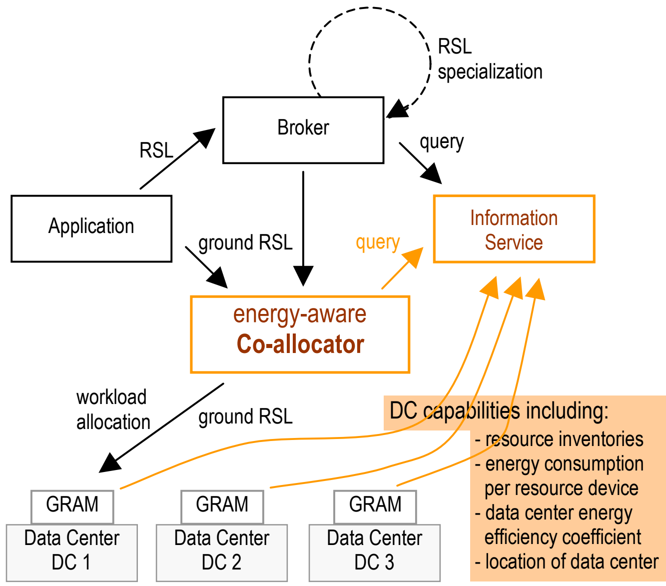

Figure 4. Grid resource management architecture [3].

Application's workload specification.

The application specifies resource needs in form of a RSL (Resource Specification Language) specification. The process of resolving higher-ordered RSL specifications into elementary, ground resource specifications through a broker infrastructure (specialization) remains the same as shown in [3]. A ground RSL specification consists of a list of physical resources (machines, storage units, devices) needed to perform a computation. Ground RSL specifications are understood by GRAMs (Globus Resource Allocation Managers). GRAMs in our scenario represent entire data centers managing their resources offered to the grid. Availability provided, a GRAM allocates RSL ground resources from its resource pool for a scheduled time period and assigns them to a computation.

Information about Data Centers.

The grid uses an information service (GIS) where information about resources and their availability is registered. The co-allocator uses this information for making a decision about which GRAM will be assigned a given workload. Data centers need to register their capabilities and properties in the information service such that the co-allocator has them available. An energy-aware co-allocator assumes the following information about a data center:

- the offered *resource types* in a data center (types of machines, devices, operating systems, etc.),
- the (static) *capacity* of those resources (total number of resource instances),
- the schedule of current *allocations* and future *reservations* of resources,
- the *energy efficiency coefficient* (introduced later) of the data center where resources are located, and
- the *location* of the data center in a network coordinate system.

Resource types include all static descriptions (attributes, properties) of a physical resource such as a machine, a storage unit, or other devices exposed to the grid as resources. Properties include the energy consumption per device that will later allow estimating the "energy consumption" of a workload with a given set of resource devices and an estimated duration.

The energy efficiency coefficient of a data center represents the energy cost when placing a workload in a particular data center. This cost depends on the factors explained later and may vary seasonally and/or depending on the location of the data center.

Along with ground specifications, the application may provide further information such as a resource use profile and the expected duration of the computation. Further information may also include hints for preferred hosting locations of the computation or deadlines when the computation is expected to start or to finish.

Co-allocator: Data Center selection process.

For allocating a workload specified in ground RSL, the co-allocator will choose one (or potentially more) GRAMs as destination(s) for the workload. The co-allocator uses a selection process over a sequence of steps for identifying the workload's final destination(s). Initially, the list contains all data centers that have registered capabilities with the information service. The co-allocator obtains this list from the information service and excludes data centers whose capabilities do not meet criteria in any of the following steps.

Steps include necessary conditions (1-4) and policy conditions. Necessary conditions are:

1. Functional: the appropriate types of resources (types of machines, operating systems, software) must be available that can perform the workload (static information). GRAMs not offering the entire set of resource types are excluded.
2. Quantitative: sufficient amounts (instances) of resources must be available to meet the demand specified in the ground RSL (static information).
3. Schedule: sufficient amounts of resource instances must be available for allocation within the time frame the application has determined for the workload (dynamic information referring to the GRAM's schedule of resource allocations and reservations).

4. Constraints: restrictions, if provided by the application, that may prevent allocating a workload in a certain data center and hence must be obeyed (static information).

If the set of data centers that have passed all necessary conditions is empty, the workload cannot be allocated under the given specifications. If it is not empty, a set of data center candidates exists among which policies can be applied to determine the final destination, one of the policies being energy efficiency.

3. Generating an Energy Policy.

In the following section, we outline the differentiators that provide the basis for generating energy efficiency coefficient χ for each data center used in the Energy-Aware Co-Allocator.

Low condenser temperature.

The efficiency (η) of a classic Carnot power cycle[4] is a function of temperature of heat addition and temperature of heat rejection.

$$\eta = 1 - \frac{T_{heatrejection}}{T_{heataddition}} \qquad (1)$$

As the temperature of heat addition rises, the efficiency of conversion of heat to work increases. An identical effect occurs if the temperature of heat rejection to the environment drops. As mentioned before (see Figure 3), heat extraction system in a data center is based on a variation of reverse power cycle (also known as vapor compression cycle) with heat addition in the evaporator and heat rejection in the condenser at constant pressure and temperature. Figure 5 shows a typical pressure (P)-enthalpy (h) diagram for a vapor compression cycle using refrigerant R134a; with heat addition in the evaporator (C-D), work input at the compressor (D-A) and heat rejection at the condenser (A-B). Processes C-D and A-B occur at constant temperatures referred, hitherto, as evaporator temperature (T_{evap}) and condenser temperature (T_{cond}), respectively.

Heat is extracted from the data center at the evaporator, denoted by Q_{evap}. The coefficient of performance (COP) of the cooling system is the ratio of desired output (i.e. Q_{evap}) over the work input (i.e. W_c). Heat rejected at the condenser Q_{cond} is sum of compressor work W_c and evaporator heat addition Q_{evap}. Lower condenser temperature improves coefficient of performance of cooling system by reducing the required compressor work to provide the same amount of cooling (i.e. Q_{evap}).

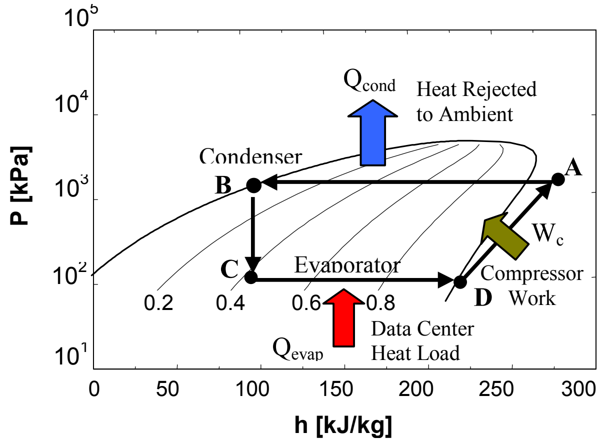

Figure 5: Typical Vapor compression cycle for heat rejection from data centers using R134a refrigerant.

This is indicated in a COP versus condenser temperature plot in Fig. 6. The COP results are based on an evaporator temperature of 10C and a compressor isentropic efficiency of 60%. Since heat can only be rejected to the ambient surroundings over a negative temperature gradient [4], the ambient temperature gates the temperature of heat rejection to the external environment (*i.e.* condenser temperatures). Consequently, ambient temperatures place a theoretical limit on the maximum efficiency of data center cooling system.

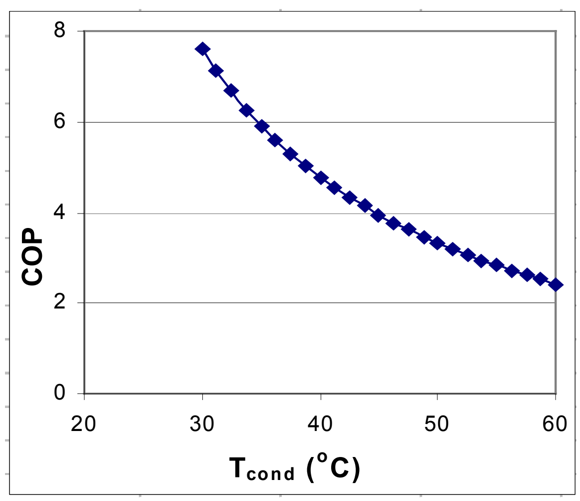

Figure 6. COP variation with temperature of heat rejection (T_{cond}).

Workload placement in data centers located in regions with higher ambient temperatures can increase the energy consumption per unit workload.

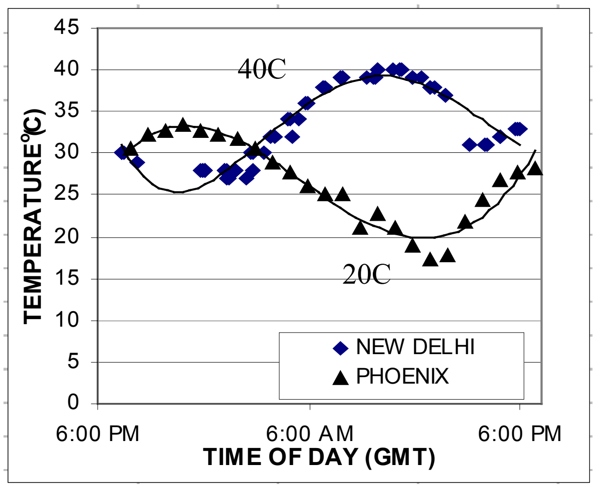

Figure 7. Diurnal temperature variations in New Delhi and Phoenix.

To elaborate; Figure 7 shows the diurnal temperature variations in two cities on a typical summer day in May, 2002. Afternoon temperatures in New Delhi reach a maximum of 40C when the night temperatures in Phoenix dip below 20C. Assuming that the condenser temperature is 10C higher than the ambient temperature at this time of operation, data centers in New Delhi and Phoenix would have condenser temperatures of 50 and 30C, respectively. From the COP curve in Figure 6, we ascertain that the COPs for these operating conditions to be 3.32 and 7.61, respectively. This result clearly indicates that the workload placement in New Delhi would be 56% more energy intensive than that in Phoenix at that time of the day. On an Energy-Aware Grid, workload placement should be carried out based on lowest available heat rejection temperature.

Note that the concept of COP is developed above utilizing a vapor-compression refrigeration cycle as an example. This concept, however, is valid in general for any type of cooling infrastructure and is affected by the type of infrastructure deployed along with additional factors discussed subsequently.

Relative humidity (RH).

Depending on the ambient humidity, cooling of data center supply air may also involve, inadvertently, condensing moisture from the air. For example, cooling air at 30C 50%RH to 15C 98%RH involves condensation of 3 grams of moisture for every kilogram of air. Thus, about 30% of the actual cooling capacity is wasted in extraction of latent heat during condensation. The condensation process leads to latent heat load on the cooling system, not accounted for by the sensible cooling capacity provided by the system. Such an extraneous load reduces the effective sensible cooling obtainable from the cooling system. Typically, outside air makes up 10% of total recirculation volume flow rate of air in a data center. Therefore, the effect of relative humidity of ambient air on data center cooling performance is an order of magnitude lower than that of condenser temperature. However, higher ambient relative humidity is a potential disadvantage because it negates the use of less energy-intensive methods like evaporative cooling for data center cooling systems. Relative humidity levels for the cases discussed before are shown in Fig. 8. At the time of interest, almost identical relative humidity levels prevail in the two cities, thus, indicating negligible effect of humidity.

Workload placement in an Energy-Aware Grid should avoid the potential disadvantages associated with high ambient humidity conditions. Data center located in regions with low seasonal humidity and ambient temperature can directly utilize outside air to provide cooling. As explained in architecture section, direct use of outside air increases cooling efficiency by reducing energy consumed for thermodynamic work.

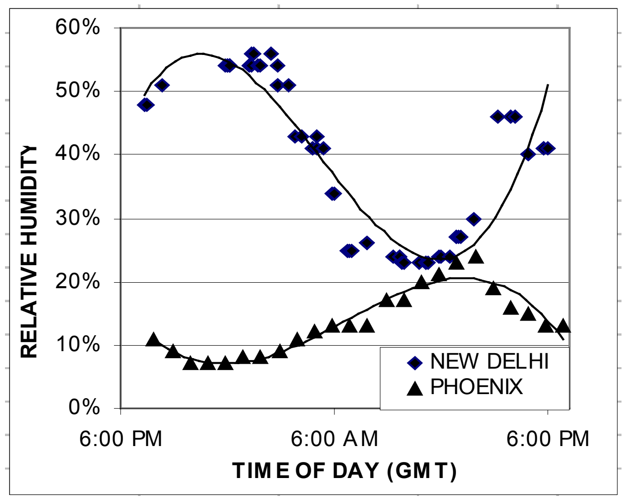

Figure 8. Diurnal variations of Relative Humidity in New Delhi and Phoenix.

Cooling load.

Coefficient of performance of cooling systems also varies with load. For chilled water systems, COP can deteriorate by as much as 20% if the load drops to 50% of rated capacity. This is primarily, due to the effect of sizing of individual components. Workload placements across data centers in an Energy-Aware Grid should strive to maintain optimum load levels for highest possible coefficient of performance.

Ground as a heat sink.

Use of ground as a sink for rejecting heat can provide higher COP at a slightly higher initial cost. The diurnal temperature variation is barely observable below a depth of 1 m while the seasonal temperature wave is not observable below 20 m. Providing heat rejection systems at different subsurface locations can provide a stable and low heat rejection temperature all year round. As indicated before, a drop of 20C in condenser temperatures can reduce energy consumption by more than 50%. Figure 9 shows a typical ground coupled heat rejection system.

Heat from the condenser is rejected to the ground instead of the air, by circulating glycol or water through underground piping. Since the heat transfer occurs across liquid and solid earth crust as opposed to liquid and air for conventional air-side condensers, lower temperature gradients can be sustained. Similar benefits can be obtained by providing ground coupled heat rejection systems near water bodies. Such systems have high energy efficiency ratios (EER=BTU/hr per Watt of electricity consumed) of 12~15 all year round. Energy-Aware Grid should be aware of efficiency of these systems for prospective workload placement during adverse ambient conditions.

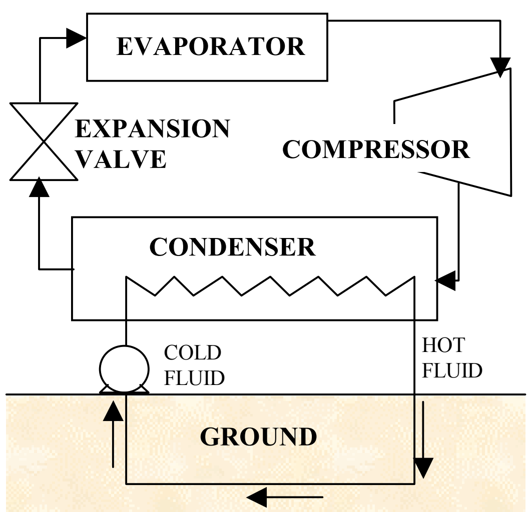

Figure 9. Typical vapor compression system with ground coupled heat rejection system.

Local Thermal Management.

Energy efficiency alone may be a sufficient condition for workload placement across the grid, but its certainly not the necessary condition. Prior studies [5] have shown that data centers with mal-adjusted rack and CRAC unit layouts can have hotspots even at reasonable operating efficiencies. Workload placement motivated by energy efficiency should also be managed to ensure that the data center temperatures are well below the redlining values. The ease and flexibility of addressing thermal management issues is a direct function of the data center infrastructure and will vary across the data centers participating in the Grid.

Based on well-established CFD studies in data center modeling [6], [7], we propose the use of data center-level thermal multiplier to account for the ability of the data center infrastructure to cope with new workload placement.

$$\xi = 1/SHI = \left(\overline{T}_{rackoutlet} - T_{ref}\right)/\left(\overline{T}_{rackinlet} - T_{ref}\right) \qquad (2)$$

where T_{ref} denotes the air supply temperature to the data center. The numerator and denominator represent the temperature difference at the rack outlet and inlet, respectively. SHI denotes effect of hot air infiltration at the inlet to server or rack [7]. When the rack inlet temperature rises, systems become more vulnerable to failure and reliability problems. Higher ξ indicates a greater potential for such vulnerability.

Energy Efficiency Coefficient.

Data center energy efficiency coefficient is a composite indicator of energy efficiency and thermal management of a data center. The efficiency coefficient of i^{th} data center is given by

$$\chi_i = \xi_i \cdot COP_i \qquad (3)$$

High ξ indicates better thermal management with reduced flow work. High COP indicates the conversion efficiency of data center cooling system in generating cold air. At low ξ, placing new workloads runs the risk of disturbing the thermal management of the data center. Above a certain threshold ξ, the energy efficiency coefficient χ provides a basis for distributing new workloads.

4. Energy-aware Workload Distribution in a Grid.

The co-allocator has a global picture of all data centers with their energy efficiency coefficients χ_i. Within the scope of the remaining subset of candidate data centers, the co-allocator can choose those data center with the highest performance index at the time of the placement:

$$WPI = \max(\chi_i)\forall i \qquad (4)$$

However, depending on the quantum of data or application migration needed for workload placement, factors like long distances and different time zones may outweigh the energy efficiency policy. There is a need to weigh the energy efficiency function based on localization. Several weighing schemes based on inverse-square law formulation or logarithmic decay can be used to capture the effect of locality. Equation (4), can thus be modified to

$$WPI = \max\left(f(\chi_i \cdot \omega_i)\right)\forall i \qquad (5)$$

where ω_i is a locality factor inversely proportional to the distance between the i^{th} data center and the source of workload placement. Distance is related to a network coordinate system that reflects hops and available bandwidths. ω_i represents a counter weight for energy-driven placement decisions that involve long distances between the source of workload and destination data center. (The amount of data to be transferred for realizing a placement is currently not considered.)

Since the co-allocator knows the energy efficiency coefficients of all data center, its energy-efficiency policy can discharge load from data centers with low coefficients by disregarding them for further placement decisions. This will cause those data centers becoming less loaded with increasing opportunity for powering down resources and consequently reducing their energy consumption.

The energy efficient workload placement policy also observes daily or seasonal fluctuations of data center energy coefficients. As Figures 3 and 4 show, the daily temperature charts of the two data centers located in New Delhi and Phoenix have impact on the energy efficiency coefficient of those data centers. When necessary conditions and resource schedules permit, shorter workloads (less than 12 hour) should preferably be placed in data centers over night in respective time zones having lower outside temperature such that air conditioning in the "hot" area can be reduced by not utilizing resources there and giving the data center the opportunity for powering down resources and air conditioning during these times.

Besides choosing locations in "cold time zones" for placing workloads, another dimension is in time. When the deadlines of the workload permit, performing the workload may be delayed to the night hours achieving a similar effect of discharging load from day periods.

Workload Placement Example.

Consider a co-allocator charged with placing a particular workload from a digital animation company. The workload consists primarily of an animation rendering application that requires the exclusive use of 64 processors for 8 continuous hours starting at 6:00 AM GMT of a given date.

The co-allocator initiates the 3-step process indicated below to identify an appropriate data center:

Step 1. Conduct a search for the data centers that meet the set of necessary conditions listed above;

Step 2. Determine the physical location of the requested resources in each of the data centers identified in Step 1, and obtain ξ for those locations from the GRAM of each data center. Eliminate all data centers with low ξ (ξ<4);

Step 3. If the remaining set of data centers is > 1, use WPI to determine the final placement of the load.

Upon completion of Step 1 above, it is determined that three placement scenarios exist that meet the necessary workload placement conditions. Two of the options exist within a single data center in New Delhi that is particularly lightly loaded during the time in which the resources are required. The other set of resources are housed in a data center in Phoenix.

Table I lists the potential resources to be used within each of the qualifying data centers along with the local ξ, COP, and χ for each of the data centers.

	Data Center 1 (Phoenix)	Data Center 2a (New Delhi)	Data Center 2b (New Delhi)
Resource Data	32 2-way 2U systems @ 400 W each	32 2-way 1U systems @ 300 W each	4 16-way Bladed systems @ 2000 W each
ξ	6.7	10	3.3
COP	7.6	3.3	3.3
χ	50.9	33.0	10.9
Power Consumption of Workload	9,600 W	9,600 W	8,000 W
Power Consumption of Cooling Resources	1,263 W	2,909 W	2,424 W
Combined Power Consumption of Workload Placement	10,863 W	12,509 W	10,424 W
Total Energy Consumption	312 MJ	360 MJ	300 MJ

Data center 1, located in Phoenix, has 64 processors housed in 32 1U systems while data center 2 is lightly loaded and contains the required resources in the same 32 1U systems or 4 bladed systems. The bladed systems are particularly attractive because they consume less power per processing unit than the set of 1U systems. However, bladed systems are designed to maximize the number of processors that can be deployed in a rack, increasing rack power densities and subsequently placing a disproportionate burden on a data center's cooling system. This tends to result in higher local inlet air temperatures to racks that contain these systems. In this particular example, a thermal multiplier, ξ, of 3.3 is given by the data center 2 GRAM for the location within the New Delhi data center that houses the bladed systems. Additionally, a heat index of 6.7 is given to the location within the same data center in which the 1U systems reside. Thermal multipliers less than 4 generally result in high inlet air temperatures and reduced system reliability [8]; therefore the bladed option can be removed from the list of possibilities.

Of the remaining two options, the 1U systems in the New Delhi and Phoenix data centers consume identical amounts of energy. However the COP of the data center during the time of day the resources are required is considerably better in the Phoenix data center resulting in an energy efficiency coefficient, χ, significantly higher than that of the New Delhi center. This further results in a reduction in cooling resource power consumption of 56% and a reduction in total energy consumption of 13% over the New Delhi option. Again clearly indicating that the Phoenix data center is the best choice for the given workload.

5. Summary.

The aggregation of compute, networking and storage nodes in a data center leads to high heat dissipation at building level. The compaction available with today's server design makes it possible to reach 3 to 4 times the heat densities over prevailing installations.

In the context of the emerging Grid and delivering computing as a utility, there will be a proliferation of high-density data centers around the globe. The immense energy consumption by these installations will require a fresh approach in employing global compute resources. In this paper, we have presented an energy-aware policy for distributing computational workload in the Grid resource management architecture. The basic notion is introducing a data center energy coefficient that is taken into account as policy by an energy-ware co-allocator for making a placement decision for a given compute workload. This coefficient is determined by the thermal properties of each data centers cooling infrastructure including regional and seasonal variations An example with three data centers located in two different time zones has shown an estimate of energy savings giving sufficient reason for the economic viability of our approach.

With continuing commoditization of resources and the increase in flexibility coming with the Grid, policies such as the presented energy-aware policy are becoming more and more relevant.

6. References.

[1] Friedrich, R., Patel, C.D.: "Towards Planetary Scale Computing - Technical Challenges for Next Generation Internet Computing", *THERMES 2002*, Santa Fe, New Mexico, Jan 2002.

[2] Bruce Maggs: *Global Internet Content Delivery*, Stanford Internet and Distributed Systems Seminar, May 5, 2002. http://www.stanford.edu/class/cs548/.

[3] Ian Foster, Carl Kesselman (Eds.): *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, 1999.

[4] Reynolds, W.C. and Perkins, H.C.: "Engineering thermodynamics", McGraw Hill Inc., Chapter 9, 1977.

[5] Patel, C.D., Bash, C.E., Belady, C., Stahl, L., Sullivan, D.: "Computational fluid dynamics modeling of high compute density data centers to assure system inlet air specifications", *Proceedings of IPACK'01 – The PacificRim/ASME International Electronics Packaging Technical Conference and Exhibition*, Kauai, Hawaii, July 2001.

[6] Patel, C.D., Sharma, R.K, Bash, C.E., Beitelmal, A.: "Thermal Considerations in Cooling Large Scale High Compute Density Data Centers", *ITherm 2002 - Eighth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems,* San Diego, California, May 2002.

[7] Sharma, R.K, Bash. C.E., Patel, C.D., "Dimensionless parameters for evaluation of thermal design and performance of large-scale data centers", 8th ASME/AIAA Joint Thermophysics and Heat Transfer Conf., St. Louis, MO, June 2002.

[8] Bash, C.E., Patel, C.D. and Sharma, R.K., "Efficient Thermal Management of Data Centers – Immediate and Long-Term Research Needs", to be published in *J. HVAC&R Res.*, 2003.