

Grid Heating: Transforming Cooling Constraints Into Thermal Benefits

P. Brenner¹, D. Thain², D. Latimer¹

¹ University of Notre Dame	² University of Notre Dame
Center for Research Computing	Dept. of Computer Science and Engineering
Notre Dame, IN 46556 USA	Notre Dame, IN 46556 USA
pbrenne1@nd.edu	dthain@nd.edu

Abstract

From 2006 to 2011 the U.S. national energy consumption for powering and cooling IT servers is estimated to grow from a cost of 4.5 to 7.4 billion dollars as reported by a recent EPA study [1] which included current efficiency improvement trends. With growing national concern for energy efficiency and environmental stewardship, current power utilization trends in HPC and data centers cannot continue to scale with computational demands. We introduce a new grid heating framework to promote the efficient growth and sustainment of commercial, academic, and government computation capabilities. We remove cooling expenditures while providing dynamic distributed heating benefits. Experimental results demonstrating temperature control and a municipal grid heating deployment are presented in support of our analytical estimates. We then discuss two example grid heating frameworks in the context of existing distributed computational infrastructures. Additional grid heating challenges and opportunities are reviewed in regards to development, implementation, and deployment.

1 Introduction

It is February and a bright biochemist realizes that she will not be able to complete testing of her new molecular sampling algorithm prior to the conference because of insufficient cooling. As she looks at the icicles outside the window of her drafty office she considers with agitation her multi-node computation cluster inaccessible because of HVAC equipment failure; she

knows something in the world of computation has gone awry. This biochemist is not alone. With the boom in commercially available commodity processors; computational capacity has scaled with the number of nodes. Unfortunately, power and cooling demands followed in near linear step. As an unintended consequence, the HPC and data center infrastructure has developed such that we are now paying premiums to cool servers instead of capitalizing on the available thermal benefit.

Today commercial data centers, academic computation facilities, and national labs recognize that facility operational and infrastructure upgrade costs rival that of the new capital computer equipment [2]. In a recent report based on historical server sales [3] Koomey shows that in 2005 the United States spent 2.7 billion dollars on electrical power to operate and cool IT and HPC servers. From 2000 to 2005 the yearly growth in electrical consumption averaged 15%. A 2007 EPA report [1] yields an even higher estimate with 4.5 billion dollars spent in 2006 and a forecasted 2011 national IT electric energy expenditure of 7.4 billion. The computational technology industry has responded with innovations in processor, hardware, and systems software to reduce the growth of the oppressive power requirements. However, the continued demand for increased capacity by high performance computing (HPC) will drive increases in system scale on par with reductions in hardware energy consumption.

In parallel to these economic IT energy concerns, recognition of public demand for environmental stewardship and sustainability has come center stage. The general public, environmental advocacy groups, and an organization's employees are demanding proactive steps toward conservation and green processes. In response, the technology industry has launched collaboratives such as The Green Grid and multiple individual corporate initiatives such as Sun Microsystems' ECO Responsibility Initiative and IBM's Project Big Green. Corporations, universities, and government labs are trying to reduce their environmental footprint and manage their energy consumption in a world economy where the environmentally sound production of energy is under intense scrutiny. One further example of this new environmental emphasis in HPC is demonstrated by the emergence of a Green500 [4] to complement the Top500 [5] list of Supercomputers.

To meet the continued growth in both the number of HPC users [6] and capability demands of those users [5], new energy focused design paradigms are required. We propose the grid heating (GH) paradigm. GH recognizes that despite evolving lower power architectural technologies, demands for increased capability will drive up power consumption toward economic limits on par with capital equipment costs. At the core of GH is the requirement for efficient re-utilization of this expended electrical energy as thermal energy. In contrast to the design of a single facility around centralized compute infrastructure, GH capitalizes on grid and virtualization technologies to distribute compute infrastructure in-line with existing municipal and industrial thermal requirements.

In this work we first describe grid heating and estimate the economic and environmental benefits of a heating grid. We then present baseline experimental results demonstrating the capability to dynamically control temperature through commodity compute nodes and mature grid software. Our current grid heating deployment to a municipal greenhouse is then reviewed. Two example GH models are proposed to further outline the paradigm. GH is then placed in the context of existing grid infrastructures and related work. Finally we conclude with a discussion of grid heating potentials and challenges.

2 Grid Heating

Grid heating's fundamental core is the recognition that computational infrastructure can be strategically grid distributed in-line with municipal facilities and industrial processes requiring the thermal byproduct that computation delivers. Frameworks built upon this core reduce or remove cooling requirements and realize cost sharing on primary utility expenditures. This contrasts with traditional data center models where a single facility is designed around a centralized compute infrastructure. With similar motives for energy efficiency and environmental stewardship multiple organizations have made great strides in the optimization of traditional centralized data centers; such as the High-Performance Buildings for High-Tech Industries Team at LBNL [7], the ASHRAE Technical Committee 9.9 for Mission Critical Facilities, Technology Spaces, and Electronic Equipment [8], and the Uptime Institute [9]. The discoveries and improvements found for the traditional data center model are of direct benefit to coarse grained grid heating frameworks and of crucial need where the grid heating model is not well matched for an organization.

Individual data centers have re-utilized the thermal energy to the benefit of their own facility, however, to utilize all of the thermal energy effectively and remove the cooling requirement a grid distributed approach is necessary. There are some parallels with industrial co-generation [10] concepts where both electrical and thermal energy are utilized as products of power generation. However in grid heating the products are computation and heat; centering the discussion around computation, energy transformation, and energy transfer. Grid heating models recognize that the transformation and or transportation of the waste heat quickly reduces efficiency and therefore targets the distribution and scale of each heating grid node to match the geographic and thermal requirements of the target heat sink.

Grid heating deployments must also address the physical and practical considerations of operating a computational infrastructure. A sample (non-inclusive) list of such considerations must include basic hardware operational requirements of temperature, humidity, and air particulate. Example practical factors include suitable bandwidth/data locality, security, sysadmin access, reliability/redundancy, and acoustics. Each of these factors is influenced by the chosen granularity of the heating grid. We provide examples in the discussion of our Greenhouse deployment and a subsequent review of two example grid heating frameworks.

2.1 Economic Estimates

To estimate the primary economic and environmental benefits of the grid heating paradigm we focus on quantifying the power utilized in normal system operation and the cooling equipment necessary to permit said operation. The following general assumption is made: the HPC community averages high server utilization on the order of 90+% annually. Note we do not here consider fine grained CPU utilization. For large parallel jobs CPU utilization is known to vary greatly based on variables such as the parallel programming paradigm and interconnect technology. To quantify potential grid heating savings

we start with a 1 megawatt server load and then calculate the savings if the same computational capability were deployed as a heating grid.

The US Department of Energy (DOE) Energy Information Administration [11] reports the following national average price of electric per kilowatt-hour (KWh) for residential, commercial, and industrial usage respectively: 0.0945, 0.0867, and 0.0573 dollars. Using the commercial figure, the corresponding cost for a megawatt of server power to a 24/7 data center or HPC facility over the course of one year would be \$760,000. We feel this is a reasonable estimate as numerous HPC and data center facilities are in or near large urban areas where the cost per KWh may be higher than the national average. Note this is the cost for server power alone, not inclusive of overhead/additional utility costs.

2.1.1 Cooling Cost Savings

For cooling costs we simply provide an estimate for the additional energy consumption as a percentage of that consumed by the computational equipment. We conservatively do not factor in the costs of cooling equipment capital expenditures and maintenance as we do not provide a comparative (although likely much smaller) analysis of capital costs to integrate computation resources into a grid heating framework. For simplicity we consider the predominant air cooled configurations, with recognition that a more thorough discussion would include currently developing and deployed water and spray cooled technologies. The total efficiency of air cooled strategies varies with scale from small space computer room air conditioning (CRAC) units to large scale air handling unit (AHU), chiller, and cooling tower configurations. To simplify the analysis we reference the works by Koomey [3], Brill [2], and Malone [12] which argue for a SI-EER (or “Power Usage Effectiveness”) of 2 (total facility power required = hardware requirement * SI-EER). The SI-EER includes all overhead power requirements but Brill breaks out optimal values for cooling at 0.3 for the chillers and 0.1 for the blowers. We therefore estimate a year round cooling cost savings of 0.4 MW or \$304,000.

2.1.2 Grid Heating Savings

In a grid heating paradigm the thermal byproduct is delivered directly as a facility or industrial process heating benefit. While the heating requirement of traditional facilities varies with geographic location and season, many industrial process heat requires are reasonably constant and independent of these factors. Further the scale of each component of a computational heating grid is sized to provide optimal heat load while minimizing cooling costs over an annual evaluation. An example is given in the experimental section. For the industrial application we argue grid heating removes the year round cooling requirement (\$304,000) and the thermal byproduct is of direct benefit throughout the year (\$760,000) leading to a power savings of \$1,364,000 or 100% of server and cooling load. For a seasonally effected deployment the figure is estimated at a 50% power savings on server and cooling load. This does not include savings that would be associated with removing the capital and maintenance costs of HVAC equipment, nor the associated additional electrical infrastructure to provide 24/7

power to the HVAC equipment. Further modifications to the thermal biproduct value would be required for comparisons with steam, oil, and natural gas heating infrastructures.

3 Experimentation

3.1 Temperature Control for a Grid of Fine Granularity

Given the basic observation that the electrical energy consumed by modern computational infrastructure is predominantly transformed into thermal energy (we acknowledge minor transformations into mechanical work by the fans, hard drive, etc.); a thermal energy surplus on the order of the input electrical energy will be locally available. We accepted this basic observation and selected an experiment to evaluate the more challenging effective utilization and control of this thermal source. We constructed a dynamically tunable grid heating framework of commodity components and successfully demonstrated the temperature control of a large multiuser workspace.

The experimental goal was to provide thermal control of a multiuser workspace through dynamic resource utilization and job scheduling. *Note that the goal of this test set was to study thermal control, recognition of preferred 24x7x365 operation is addressed via configurations described in later sections.* The workspace selected was a functional system administration office measuring 14x18x9 in units of feet. The workspace had HVAC controls independent of the larger facility and these were disabled during the tests. Tests were run during the evenings and or weekends as not to introduce variance from the occupying personnel. In the room we constructed a small grid heating cluster consisting of 6 dual processor Sun V60 rack mounted nodes which provided us with twelve 3GHz Xeon CPUs. The nodes were networked and configured with Red Hat Enterprise Linux 4. On each node a Condor [13] client was installed, reporting to the Notre Dame campus Condor pool. A temperature probe with accuracy greater than 0.1°F was placed to take room temperature readings every 5 minutes.

The computational/thermal production of the grid heating cluster was exercised through the calculation of molecular dynamics trajectories for biomolecular conformational analysis as discussed in previous work [14, 15]. To isolate the experiments from any variations due to characteristics of other Condor jobs, the grid heating cluster was configured to accept only these long running (greater than 1 month) molecular dynamics simulations. There were sufficient jobs in the Condor queue such that whenever a processor on the grid heating cluster was available a suitably matching job from the Condor queue would start running. Dynamic thermal control was achieved by the varied execution of these jobs through reconfiguration of Condor machine attribute files at regular intervals in response to the current room temperature and designated target temperatures.

Multiple test scenarios were designed and evaluated to capture the grid heating clusters ability to meet and oscillate about temperature targets. In Figure 1 we present one such test which captures the key capabilities and demonstrates the achievement of our experimental goal. The test scenario covered 48hrs starting with all machines powered off and a room

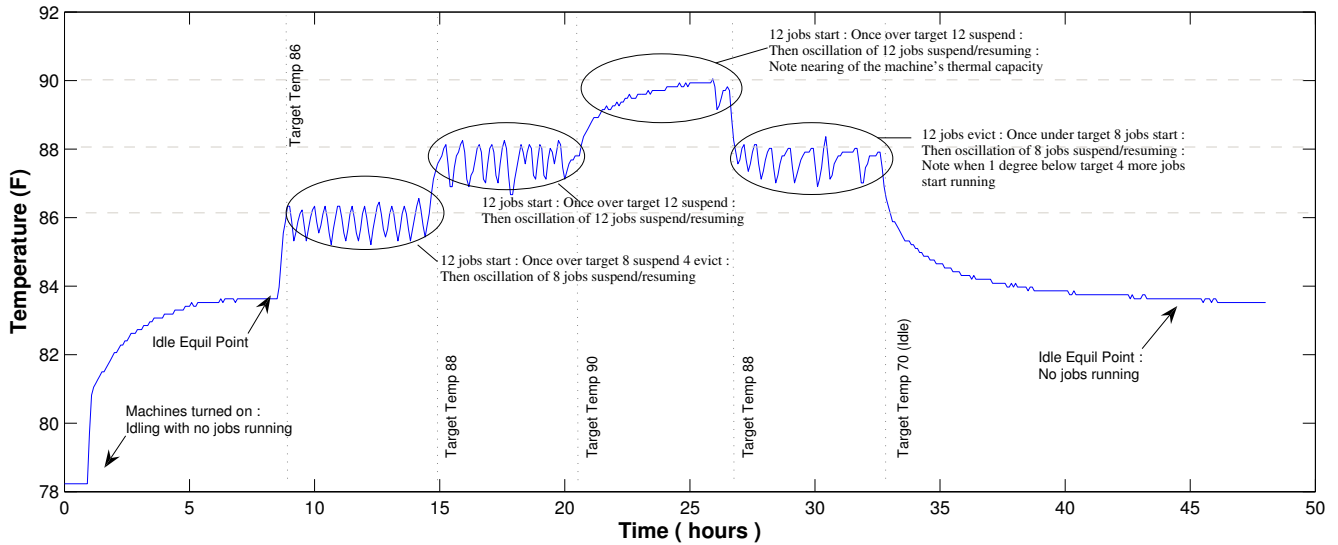


Figure 1. Dynamic Temperature Control

temperature just above 78°F. Simply powering the six machines on raised the room temperature nearly 6°F. We note that for our test equipment the idling power consumption (thermal production) is a large portion of total CPU saturated power consumption; however as idling behavior is tunable (standby, hibernate, etc...), machine design dependent (component selection such as power supply), and less frequent in HPC we omit further discussion and focus on the thermal contribution and behavior under computational load.

Once the room had equilibrated to a stable temperature with machines idling, a test script was run to adjust target room temperatures and prompt the grid heating cluster nodes to dynamically accept jobs. As shown in Figure 1 the target temperature was updated every 6 hours from 86° to 88° to 90° to 88° and then down to 70°F which effectively put the machines back in the idle condition. The condor machines were reconfigured (attribute file updated) every 5 minutes with the new room temperature and target temperature. Machines would run jobs until the target temperature was reached and then either suspend or evict. The oscillations in the figure demonstrate the thermal variance in the room as jobs moved from suspend to resume status. Note the graph also reflects the properties of the machines maximum thermal capacity. When the first temperature target is set to 86°F the 12 processors accept jobs and almost immediately reach the target whereas it takes slightly longer to reach the first 88°F target and much longer to reach the 90°F.

One minor optimization was implemented to demonstrate the concept of job eviction in addition to suspension. This reduced the degree of oscillation about the target. Jobs were evicted instead of suspended when the total thermal capacity of the processors was much greater than the specified target, and then prevented from restarting until the room temperature was greater than 1 degree below the target. This is one simple example shows how the oscillations may be reduced and job execution may be optimized, further evaluation of thermal tuning, machine utilization, and job progress optimizations are

left to future and peer works. Given these successful test results for our first implementation we next offer two example frameworks for utilizing this thermal benefit.

3.2 Greenhouse Deployment

With the successful demonstration of dynamic temperature control we proceeded with our first grid heating deployment at the City of South Bend Botanical Conservatories and Greenhouse (SBGH). The SBGH is a traditional glass greenhouse facility with multiple interconnected structures of which the most recently constructed was in the 1970s. Municipal funds are not available for capital upgrades and rising annual natural gas heating costs grew to over \$115,000 in 2006. Primary heat is provided via natural gas boilers that provide steam heat to the majority of the building. The desert botanical collection, where our deployment is located, is found in a dome structure directly heated by multiple natural gas heaters (and mobile propane heaters during the coldest months).

The University of Notre Dame has partnered with the City of South Bend on a three phase grid heating deployment to provide direct thermal load to the SBGH and year round operations of our computational equipment. Phase 1 has been completed successfully, and entailed the initial deployment of sufficient computational infrastructure to enable remote submission of scientific simulations from the Notre Dame campus to the SBGH. Primary phase 1 components included a traditional compute rack located within the desert collection dome structure. The rack was configured with traditional 1U compute nodes of similar to that used in the thermal control tests. Network interconnectivity is provided by a local ISP contracted through the City of South Bend. Successful completion was marked by the facilitation of molecular dynamics simulations performed by the LCLS research group to further their research in accelerated trajectory generation [16].

We are now in phase two of the deployment with a focus on scaling and securing the compute infrastructure. In this regard we give an overview of the practical considerations addressed, starting with the physical factors. Operational temperatures and humidity for the hardware should be within ASHRAE specifications [8]. We evaluated the seasonal fluctuations of South Bend temperatures by reviewing information from the National Climatic Data Center [17] as shown in Figure 2. We observed that the average highs do not exceed the maximum allowable inlet face temperature and therefore predict the suitability of year round operations utilizing outside air. For the small number of hours annually which are above 90°F we will dynamically idle the machines as shown previously. Heat evacuation will be handled by an existing ventilation fan at the top of the dome structure and aided by natural draft. Humidity in the facility is currently controlled year round for the plant life. Proper mixing of inlet and inside air during the cold low humidity winter months will be managed by existing outside air louvers near the compute rack front faces. We lowered the priority to investigate particulate load based on LBNL findings [18] indicating outside air particulate loads well below EPA and manufacturer guidelines; however, readings relative to the local vicinity are planned.

Phase two entails improvements with respect to many of the practical factors for deploying computational infrastructure.

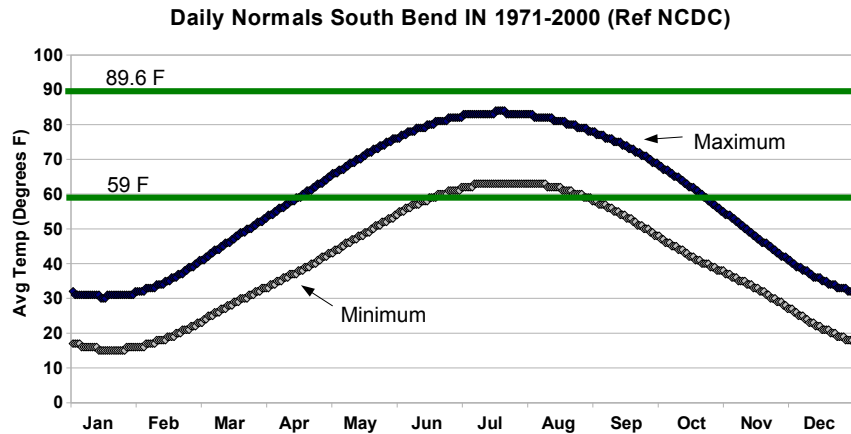


Figure 2. Daily Temperature Normals with ASHRAE Allowable Max/Min Temperatures

Internet security is currently enabled through router and OS firewalls in addition to traditional administrator passwords. Physical security is currently provided in multiple parts: monitored access to this public facility by SBGH staff that manage paid admissions, locking the rack itself, and mechanically securing the power cord from the rack to the outlet receptacle. Phase 2 will entail an ORM process to consider upgrades such as placing power and network cabling in rigid conduit. System administrators have regular access to the facility which is located about ten minutes from the main campus and of equal distance to our off campus remote HPC facility. Regarding reliability and redundancy, there are currently single points of failure in power and cooling (ceiling fan), phase 2 includes a cost benefit evaluation of upgrading to a Tier 1 or 2 configuration as designated by the Uptime Institute. As the deployment is internal to the Greenhouse (as opposed to an enclosure placed outside of the Greenhouse), acoustics are an important consideration. Currently the audible levels of the 1U nodes are on par with the multiple natural gas heaters and ceiling fan, however, a reduction in waste noise is preferred. Custom devices discussed in the next section are being designed to address this concern. During phase 2 the bandwidth will remain under the provision of the contracted ISP which limits us to compute bound simulations with small data transfer requirements (15MB/s), phase 3 includes a connection to the St. Joe Valley Metronet fiber-optic network shared by South Bend and Notre Dame.

4 Grid Heating Frameworks

Following our argument for the advantageous utilization of thermal energy and the experimental validation of our initial implementation it is helpful to provide general frameworks for example deployments. We briefly outline two potential frameworks to highlight the diversity in configurations. The first, GH Clusters, focuses on a tightly interconnected multi-

node scalable configuration to meet facility heating requirements while the second, GH Appliances, focuses on a widely distributed fine grained network of nodes for many small localized thermal benefits. These examples are chosen to represent practical implementations with minimal change to existing infrastructures. The identification of more exotic frameworks with novel hardware, software, and network configurations to yield greater economic/environmental impact are left to peer and future works.

4.1 GH Clusters for Facilities and Processes

A GH Cluster framework is based upon a scalable tightly interconnected multi-node configuration similar to the traditional rack configurations. A full 40U rack, partial rack, or multiple racks could be selected to best match the annual heating requirement of a particular facility. The GH Cluster could be a stand alone enclosure collocated with an industrial heat sink or facility. Waste heat transferred via air or liquid media would be directly utilized. Examples include air heating of a greenhouse, hot water preheat for a hospital, or primary heat for a water treatment plant. Fully integrated modular enclosures to deliver the heating fluid are commercially available, such as Sun MicroSystem's Black Box and Rackable System's Ice Cube.

The GH Cluster could also be integrated into the existing HVAC infrastructure as a complementary or primary furnace. This would require no modification to the facility's internal ductwork or sensor network. Proper duct transition sections would be required to modulate airflow through the GH cluster and into the distribution system. To reduce pressure drop nodes may be spaced as to provide sufficiently unrestricted flow. Vertical or horizontal form factors can be selected as to best meet the existing HVAC configuration. A high level schematic with the GH Cluster as a stand alone enclosure is shown in Figure 3.

The deployment of GH Clusters would have direct benefit to institutional and commercial facilities as well as individual homes. For example a university campus could heat each dorm while providing desired scientific computation, a corporation could integrate its data centers into the HVAC infrastructure, and individual home owners could receive heating bill deductions correlated to the use of their GH Cluster as part of a national computational grid. The GH Cluster framework is well suited to meet many of the associated implementation challenges. During warm months the excess thermal load could be diverted from the distribution ductwork to an outside exhaust or heat pump. Suitable air conditioning for electronics (particulate removal and humidity adjustment) may already be in place up stream of the GH Cluster. Because the computational equipment for the facility is consolidated, routing a fast dedicated network connection is not cost prohibitive. The tight interconnection of the nodes allows parallel computations requiring low latency communication (MPI jobs for example) to run efficiently. Further, by coupling the GHC with an existing furnace infrastructure, costly job evictions shown in the baseline experiment could be eliminated as the furnace modulates to provide *target temp - GHC max output*.

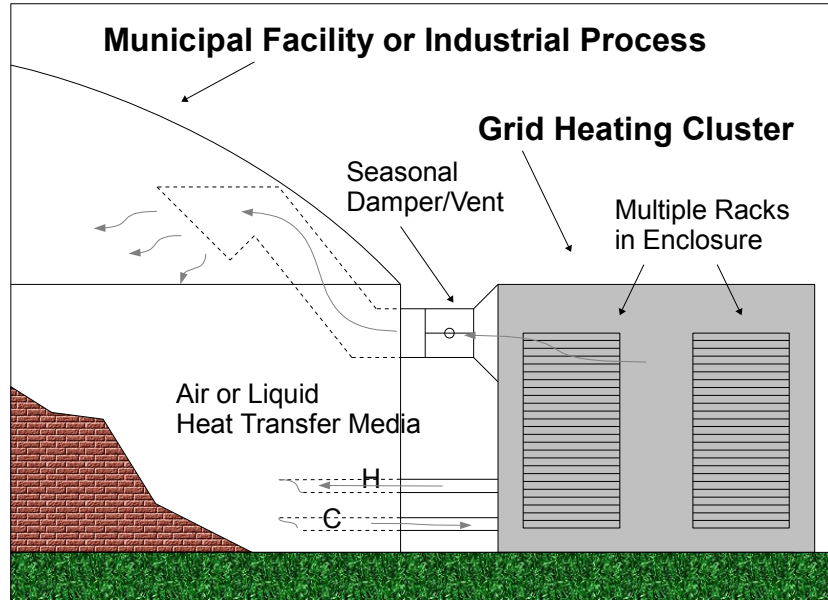


Figure 3. GH Cluster Schematic

4.2 GH Appliances for Spaces

The second GH framework we introduce is that of the GH Appliance. A GH Appliance configuration is based upon the wide distribution of independent mobile computational resources so placed as to capitalize on their thermal energy product. Multiple appliances may be utilized in near proximity however the model is based upon their independent function regardless of possible collective deployments. While traditional workstations may be a suitable form of GHA, we focus on a model GHA with a minimal set of hardware designed to maximize computational capability. For example consider a commodity wireless enabled mother board, multi-core processor, RAM, power supply, and flash based storage. This device could be simply plugged into an electric power source and dynamically provide computation by accepting jobs based on the owner specified set of thermal conditions. The computational input directives and output data files would be wirelessly transmitted to the governing grid engine. We have completed the first prototype GHA with additional authors [19] and are currently working to integrate it into the Greenhouse deployment.

The deployment of GH Appliances would meet multiple proximity based thermal requirements such as individual work and residential spaces, remote environmental control of electro/mechanical equipment, and localized heating in large or open air facilities such as stadiums, greenhouses, museums etc... Figure 4 provides an abstract visual understanding of the appliance and some potential deployments. Grid heating devices could replace space heaters in a cubical based office environment. They could also perform dynamic heating for thermally exposed or remote electro/mechanical equipment such as mechanical pumps, or locally heat sensitive sections of a large centrally modulated greenhouse. These example

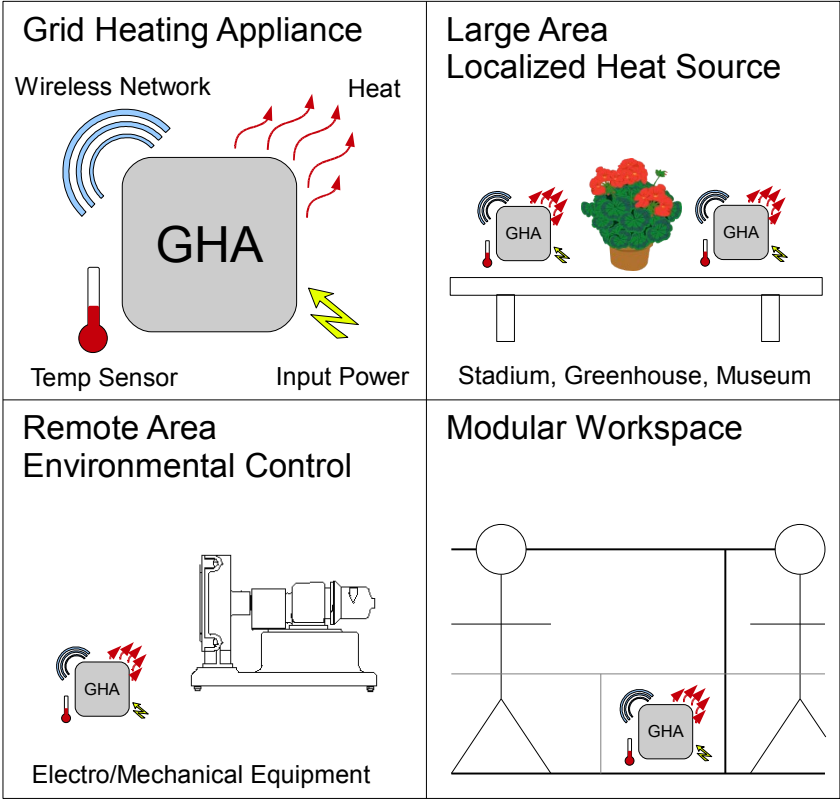


Figure 4. GH Appliance Schematic

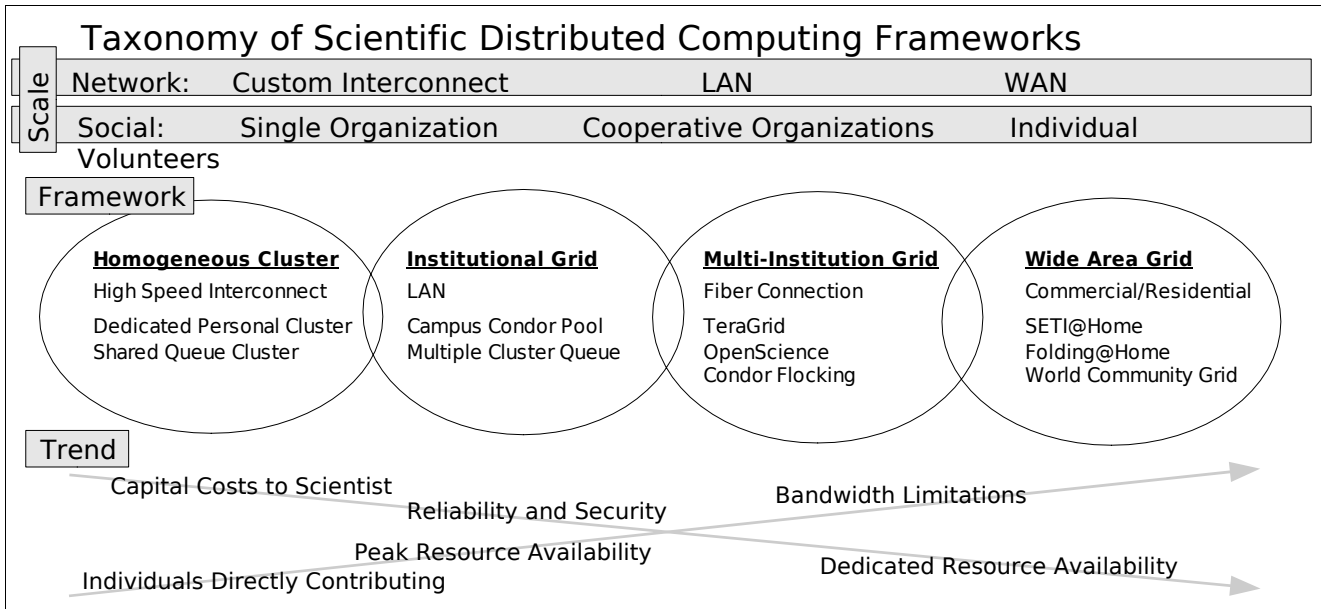


Figure 5. Taxonomy of scientific distributed system frameworks

deployments are seed concepts for many additional application tailored configurations to be implemented and tested as part of peer and future works.

5 Discussion and Related Work

For either of the two fore mentioned GH frameworks to be considered viable their integration with existing HPC and scientific distributed computing frameworks must be considered. In Figure 5 we provide a general taxonomy of the dominant frameworks currently available for scientific computation and indicate trends in economy and efficiency. We focus on the properties of the three grid frameworks, with additional comments on migrating the single location cluster resources into a grid heating configuration.

Wide area server-client computational grid models such as SETI@home [20], BOINC [21], FightAIDS@home [22], and Folding@home [23] have taken scientific computing into universities, offices, and households around the world. Folding@home alone has employed over 1,000,000 unique CPUs since October 2000. Traditionally the software enabling these models has been highly application specific however endeavors such as the BOINC project and World Community Grid are working to provide an application agnostic model. All of the models rely on the volunteer resources of individual users and companies distributed worldwide. The systems often scale from thousands to hundreds of thousands of autonomous volunteered resources however this leads to limitations in reliability and bandwidth. This existing model maps quite well to the grid heating appliances framework. Jobs would be prioritized on the target temperature of the corresponding spaces. This correlates with the dominant workstation utilization pattern where machines are often idle through the cool night hours.

In contrast to the wide-area individual ‘volunteer’ framework, major application agnostic projects such as TeraGrid [24] and OpenScience [25] have worked to unite large universities and national labs providing general purpose access to both high performance computation clusters and individual institutionally owned workstations. The framework is tied together with general purpose computational resource management tools such as Condor [13] and Globus [26], which provide secure and structured access to the geographically and semantically (architecture, operating system, ownership) differentiated distributed systems. The institutional management of these resources tends toward higher security and reliability. The two GH frameworks could be used jointly in this construct. For example utilizing grid heating appliances as part of the multi institution condor pool and grid heating clusters accessible through an OpenScience or TeraGrid portal.

Local institutional grids such as a Condor pool or unified cluster queue are not complicated by the requirements of multiple organizations or the social maintenance of a globally distributed volunteer pool. Additionally, the network speed and reliability is often more robust allowing for enhanced generation and management of large datasets. A single organization’s Condor pool however is not free from design requirements to handle resource failure and job eviction. Further, A centralized queue to access multiple campus clusters will likely not abstract all of the heterogeneity between resources. Most importantly the maximum number of available resources is limited to those funded and maintained by the single institution. Here again the two GH frameworks could be used jointly in this construct with the specific benefit of adding computational capacity while directly reducing the single institutions operating costs. GH appliances could be distributed to members requesting local working temperatures above the optimal facility heating target and GH clusters would provide primary air and water heating to facilities.

It is also important to note evolving commercial applications in grid, utility, and cloud computing [27, 28, 29] that would directly benefit from a grid heating configuration. As the computational infrastructure configuration and locality is obscured from the end user the flexibility to distribute and configure grows, allowing for additional economic and environmental optimizations. Along these lines the growing acceptance of virtualization [30, 31, 32] in commercial applications will also allow greater flexibility in the design and deployment of infrastructure.

While these grid frameworks are evolving, a large body of work has studied the problem of managing energy, heat, and load in large centralized HPC and data centers. In addition to those works cited earlier in this paper Schmidt et al [33, 34] provide a good overview of the mechanical issues of cooling units, heat sinks, fans, and so forth. Given an adequate mechanical infrastructure, several server management techniques can also be applied to reduce energy costs. For example, inactive servers can be shut down, or load migrated as more energy efficient hardware becomes idle/available. Chase [35] and Bradley [36] describe techniques for balancing performance, cost, and energy in this situation. However, blind server selection can lead to hot spots, which reduce component lifetime and increase energy costs. To avoid hot spots, it is necessary to map the relation between components and heat [37] and then shape loads so as to evenly distribute the heat [38]. Further large institutions such as the University of Illinois and NCSA are taking a holistic look at their campus utilities infrastructure to

efficiently operate their data center. Despite the efficiency benefits of these new techniques, Patel et al. [39] report that a typical data center still consumes about as much energy for cooling as it does for productive work. The advances in efficient computational infrastructure and a GH framework will serve in tandem to provide greater computational capability while reducing economic and environmental costs.

6 Conclusions

We have introduced the grid heating computational infrastructure to promote the efficient growth and sustainment of computational capability while providing for greater economic viability and environmental stewardship. Grid heating removes the overhead costs of thermal cooling and capitalizes on the available heat product. Our theoretical analysis shows from 50 to 100% savings in server and cooling energy consumption. Baseline experiments successfully demonstrated the capability to dynamically provide fine grained (1°F) temperature control with a grid heating framework. Phase one of our greenhouse deployment has been completed successfully, prompting a second phase to expand scale and capability. The grid heating cluster and grid heating appliance models were presented to provide diverse examples of practical implementations requiring minimal modifications to existing thermal infrastructures. The analysis, experimentation, and implementations were then discussed in relation to existing computation grid configurations and research to reduce HPC and data center power consumption.

This work lays the foundation for numerous exciting peer and future research projects focused on the development of unique grid heating frameworks. A few example development challenges and opportunities include the management of wireless appliances through grid engines, security in non-traditional IT locations/utilization, dynamic seasonal modulation, interfaces with existing HVAC infrastructure and thermally optimal hardware. Implementations of increasingly larger scale will provide more accurate energy savings and reclamation estimates. Grid Heating holds the promise to fundamentally change the deployment of computational infrastructures and the potential to evolve individual and institutional facility heating mechanisms.

7 Acknowledgments

This work was partially supported by a Federal NWICG grant. Computational resources were made available by the Center for Research Computing and the Department of Computer Science and Engineering at the University of Notre Dame. We thank our many municipal collaborators at the City of South Bend Botanical Garden and Greenhouse. The authors are appreciative of insightful conversations with Miron Livny, Christopher Sweet, Ed Bensman, Rich Sudlow, and Curt Freeland.

References

- [1] U.S. Environmental Protection Agency: Report to congress on server and data center energy efficiency public law 109-431. Technical report, U.S. E.P.A (Aug 2007)
- [2] Brill, K.G.: Data center energy efficiency and productivity. Technical report, The Uptime Institute (2007)
- [3] Koomey, J.G.: Estimating total power consumption by servers in the u.s. and the world. Technical report, Lawrence Berkley National Laboratory (Feb 2007)
- [4] Sharma, S., Hsu, C.H., chun Feng, W.: Making a case for a green500 list. In: IEEE International Parallel and Distributed Processing Symposium (IPDPS 2006)/ Workshop on High Performance - Power Aware Computing. (2006)
- [5] Dongarra, J.J., Meuer, H.W., Strohmaier, E.: 29th top500 Supercomputer Sites. Technical report, Top500.org (June 2007)
- [6] Tabor Research: Intersect360(tm) Market Advisory Service: Buyer Views. Technical report (August 2007)
- [7] Mills, E., Shamshoian, G., Blazek, M., Naughton, P., Seese, R., Tschudi, W., Sartor, D.: The business case for energy management in high-tech industries. *Journal of Energy Efficiency* **1** (2007) 1–16
- [8] 9.9, T.: High Density Data Centers - Case Studies and Best Practices. ASHRAE, Atlanta, GA (2008)
- [9] Brill, K.G.: Special report: Energy efficiency strategies survey results. Technical report, The Uptime Institute (April 2008)
- [10] Mise, A., Nishi, H., Sato, H.: Energy assessment of cogeneration and air conditioning system on Keio University Shonan Fujisawa campus. In: Proceedings of the 2006 IEEE International Conference on Industrial Informatics, IEEE (2006) 387–392
- [11] Energy Information Administration: Average retail price of electricity to ultimate customers by end-use sector. Technical report, U.S. D.O.E (October 2006)
- [12] Malone, C., Belady, C.: Metrics to characterize data center & it equipment energy use. In: Proceedings of the Digital Power Forum. (September 2006)
- [13] Litzkow, M., Livny, M., Mutka, M.: Condor - A hunter of idle workstations. In: Proc. International Conference of Distributed Computing Systems. (1988)

- [14] Brenner, P., Wozniak, J.M., Thain, D., Striegel, A., Peng, J.W., Izaguirre, J.A.: Biomolecular Path Sampling Enabled by the PINS Distributed Simulation Framework. In: Proc. Sixth IEEE Workshop on High Perf. Computing in Biology HiCOMB 2007. (March 2007)
- [15] Wozniak, J.M., Brenner, P., Thain, D., Striegel, A., Izaguirre, J.A.: Generosity and gluttony in gems: Grid enabled molecular simulations. In: IEEE International Symposium on High Performance Distributed Computing Proceedings. (2005)
- [16] Sweet, C., Petrone, P., Pande, V., Izaguirre, J.: Normal mode partitioning of Langevin dynamics of biomolecules. *J. Chem. Phys.* (2008) to appear.
- [17] Center, N.C.D.: Daily normals of temperature, precipitation, and heating and cooling degree days for south bend, in: *Climatology of the United States NO.84 (1971-2000)*
- [18] Tschudi, W.: 24x7 energy efficiency presentation. Technical report, Lawrence Berkley National Laboratory (Feb 2007)
- [19] Kelly, M., Lammie, M., Brenner, P.: Grid heating appliances: Development and deployment. In Progress (2008)
- [20] Anderson, D.P., Cobb, J., Korpela, E., Lebofsky, M., Werthimer, D.: SETI@home: An experiment in public-resource computing. *Communications of the ACM* **45**(11) (2002)
- [21] Anderson, D.P.: Boinc: A system for public-resource computing and storage. In: 5th IEEE/ACM International Workshop on Grid Computing, Pittsburg, USA, IEEE/ACM (2004) 4–10
- [22] Chang, M.W., Lindstrom, W., Olson, A.J., Belew, R.K.: Analysis of hiv wild-type and mutant structures via in silico docking against diverse ligand libraries. *Journal of Chemical Information and Modeling* (2007)
- [23] Shirts, M., Pande, V.S.: Screen savers of the world unite! *Science* **290** (2000) 1903–1904
- [24] Catlett, C.: Teragrid <http://www.teragrid.org>.
- [25] Livny, M.: Open Science Grid NSF annual report (2007) <http://www.opensciencegrid.org>.
- [26] Foster, I., Kesselman, C.: Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications* **11** (1997)
- [27] Eilam, T., Appleby, K., Breh, J., Breiter, G., Daur, H., Fakhouri, S.A., Hunt, G.D.H., Lu, T., Miller, S.D., Mummert, L.B., Pershing, J.A., Wagner, H.: Using a utility computing framework to develop utility systems. *IBM Systems Journal* **43** (2004) 97–120

- [28] Raghavan, B., Vishwanath, K., Ramabhadran, S., Yocum, K., Snoeren, A.C.: Cloud control with distributed rate limiting. In: SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, New York, NY, USA, ACM (2007) 337–348
- [29] Mitchell, W.: Data center in a BOX. *Scientific American* **297** (Aug 2007) 90–93
- [30] Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., Warfield, A.: Xen and the art of virtualization. In: SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles, New York, NY, USA, ACM (2003) 164–177
- [31] Figueiredo, R., Dinda, P., Fortes, J.: Resource virtualization renaissance. *Computer* **38** (2005) 28–31
- [32] Uhlig, R., Neiger, G., Rodgers, D., Santoni, A., Martins, F., Anderson, A., Bennett, S., Kagi, A., Leung, F., Smith, L.: Intel virtualization technology. *Computer* **38** (2005) 48–56
- [33] Schmidt, R., Cruz, E., Iyengar, M.: Challenges of data center thermal management. *IBM Journal of Research and Development* **49** (July 2005)
- [34] Schmidt, R., Iyengar, M.: Best practices for data center thermal and energy management - review of literature. *ASHRAE Transactions* **113** (Jan 2007)
- [35] Chase, J., Anderson, D., Thakar, P., Vahdat, A., Doyle, R.: Managing energy and server resources in hosting centers. In: Symposium on Operating Systems Principles. (2001)
- [36] Bradley, D., Harper, R., Hunter, S.: Workload based power management for parallel computer systems. *IBM Journal of Research and Development* **47** (2003) 703–718
- [37] Moore, J., Chase, J., Ranganathan, P.: Weatherman: Automated, online, and predictive thermal mapping and management for data centers. In: IEEE International Conference on Autonomic Computing. (june 2006)
- [38] Moore, J., Chase, J., Ranganathan, P.: Making scheduling “cool”: Temperature-aware workload placement in data centers. In: USENIX Annual Technical Conference. (April 2005)
- [39] Patel, C.D., Bash, C.E., Sharma, R., Beitelmal, M.: Smart cooling of data centers. In: Proceedings of IPACK. (July 2003)