Thermal Design in the Open Compute Datacenter

Eitan Frachtenberg, Dan Lee, Marco Magarelli, Veerendra Mulay and Jay Park

Facebook
1601 Willow Road
Menlo Park, CA, 94025
{etc,leed,marcom,vmulay}@fb.com

***Abstract*—The advent of Web-based services and cloud computing has instigated an explosive growth in demand for datacenters. Traditionally, Internet companies would lease datacenter space and servers from vendors that often emphasize flexibility over efficiency. But as these companies grew larger, they sought to reduce acquisition and operation costs by building their own datacenters. Facebook reached this stage earlier in 2011 when it unveiled its first customized datacenter in Prineville, Oregon.**

In designing this datacenter, Facebook took a blank-slate approach where all aspects were rethought for maximum efficiency. Although the resulting datacenter is optimized for Facebook's workload, it is general enough to be appeal to a wide variety of applications. This paper describes our choices and innovations in the thermal design of the datacenter building, which employs 100% outside-air economization. The efficiency of this design is manifest in an average infrastructure energy use reduction of 86% compared to leased space, and an overall energy use reduction of 29%. This reduction in turn translates to a power usage efficiency of 1.08, measured over the summer of 2011.

KEY WORDS: Datacenter Cooling, Air Economization, Thermal Design

I. Introduction

Facebook is a social utility serving over 800 million active users worldwide, of which more than half log in on any given day. This scale requires a substantial amount of computer processing power, electronic data storage space, and internet connection bandwidth. Even though the power efficiency of server chips has grown 16× in the past decade, datacenters are consuming increasingly more power, because the user count and products are growing at an even faster rate [1]. Facebook historically has relied solely on datacenter-space collocation providers. Leasing datacenter space and servers afforded Facebook the flexibility it required as a small company, but as its scale grew, so did the associated costs of acquiring and operating leased datacenters and IT equipment. Our analysis showed that we could do better if we designed our own computing facilities. Furthermore, we could share this design under the Open Compute project [2] for open exchange of ideas in this rapidly growing, but mostly proprietary, field.

In 2009, we embarked on a project to build a more cost- and energy-efficient datacenter. Starting from a blank slate, we began looking at all possible combinations of infrastructure and new IT equipment designs. By controlling the building design, software applications and server hardware, we were aiming to achieve these design goals:

- Maximize energy efficiency.
- Minimize initial capital expenses (*capex*).
- Minimize on-going operation expenses and the cost of fully burdened power (*opex*).

We identified three areas in which we could obtain significant improvements in both capex and opex: server design, power distribution, and thermal management. This paper focuses on the third area, the design and optimization of the datacenter cooling scheme, which is a major factor in the the datacenter's high efficiency. (A previous publication describes in detail the server design [3]).

Although capex is certainly a crucial factor in the design of every element of the datacenter, it is hard to overestimate the effect of energy and thermal efficiencies. The cost of fully burdened power has been the focus of many studies [4], [5], [6], [7], [8], and consequently engendered many creative solutions [9], [10], [11], [12]. In this large solution space, we focused on mechanical solutions that leverage a favorable climate location, eliminate the need for wasteful chiller plants, and increase the efficiency of air flow throughout the building. Our design is certainly not the first to propose evaporative datacenter cooling [13], but to the best of our knowledge, it is the largest and most efficient implementation of these proposals.

The key thermal design innovations in this paper are:

- 100% outside air cooling with humidification as needed, providing substantial air conditioning operational cost savings of at least 86% compared to equivalent leased, compressor-cooled datacenter space.
- Temperature and humidity are controlled to take advantage of the full ASHRAE[1] TC 9.9 recommended range.
- The cooling air flows inside the datacenter through room-sized passageways, thus eliminating ducting, reducing airflow restriction and lowering the required fan energy compared to typical datacenters.
- Eliminate the need for chillers, cooling towers and associated capital equipment costs.
- Heat from the IT equipment that would otherwise be wasted is used for office space heating, as well mixed with outside air as necessary to keep the IT equipment within its specifications.

We organize these innovations and others, as well as the design details, by organizing them into functional areas as follows. The next section starts by describing the operational envelope

[1]American Society of Heating, Refrigerating and Air-Conditioning Engineers (www.ashrae.org)

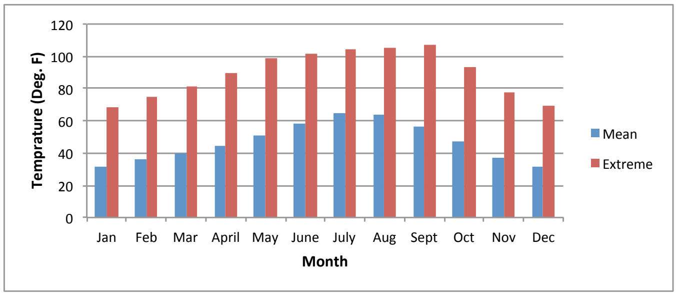

Figure 1. Dry bulb temperatures recorded at Prineville

for which the datacenter was designed, in the context of the climate conditions in the location where it was built. Then, Sec. III shows the complete path the air takes through the cooling system, starting from outside the building, getting cooled and humidified as appropriate, absorbing heat from the IT equipment, and ending back in the atmosphere. We tie these two sections together in Sec. IV, which explains how we modify the intake air characteristics inside the building to match external conditions as required to maintain the operational envelope. Section V then elaborates on the water cycle through the building. Bringing all the pieces back together, Sec. VII discusses the combined efficiency that results from these cooling solutions, as well as areas for further research or improvement.

II. Operational Envelope

We based our design on the 50-year weather data collected at Redmond, OR, which is the closest weather station to Prineville. The maximum dry bulb (DB) temperature recorded in this data was $105.6°F$ (Fig. 1), whereas the maximum wet bulb (WB) temperature recorded was $70.3°F$. The winter extreme condition was recorded as $-30.8°F$ dry bulb temperature at 50% relative humidity (RH). This climate is advantageous for outside air and evaporative cooling; the coincident wet bulb temperatures are low when the dry bulb temperatures tend to be high, allowing free cooling most of the year and efficient use of evaporation when needed. In fact, the fogging system in our design is used more often to raise humidity compared to being used to control server air inlet temperatures.

Note that the actual weather in Prineville is generally quite stable and predictable over time. For example, Figs. 2 and 3 show that the Prineville temperatures measured in November of 2009 are very close to those measured in November of 2011. These charts also demonstrate how typical weather conditions in Prineville require little effort investment in cooling and moisture control. On the other hand, heating the air is attainable for free by mixing outside air with recirculated air heated by the IT equipment.

As indicated on the chart in Fig. 4, the system is designed to handle both these extreme conditions. In fact, the dry bulb temperature considered for summer design limits was $110°F$ instead of $105.6°F$. The supply air temperature in the data hall is controlled between $65°F$ and $80°F$. The moisture content is maintained between $41.9°F$ dew point temperature at the lower

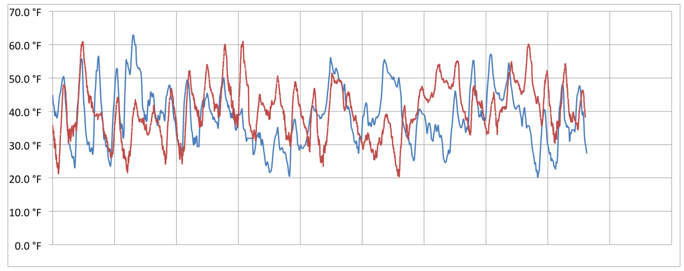

Figure 2. Dry bulb temperatures recorded at Prineville in 2009 (in blue) and measured at the datacenter in 2011 (red), during the month of November, in hourly bins.

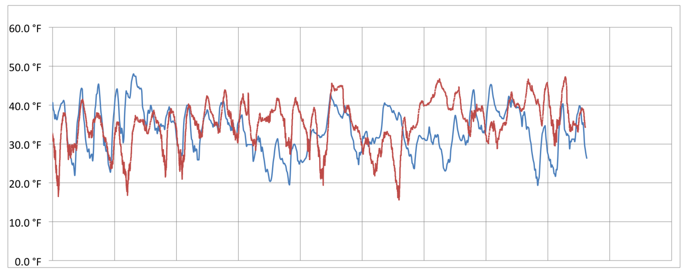

Figure 3. Wet bulb temperatures recorded at Prineville in 2009 (in blue) and measured at the datacenter in 2011 (red), during the month of November, in hourly bins.

end (to prevent condensation) and $65\%\ RH$ at the higher end (we believe that a dew-point limit on the upper end is too conservative).

Table I compares this operational envelope with ASHRAE's recommended operational envelopes. We can see that the operating environment of the Prineville data center is similar to ASHRAE's 2008 recommendations [14], except that the high end moisture level is not limited by the dew point temperature, because we felt it imposed unnecessary restrictions. It is worth noting that equipment manufacturer requirements are typically less stringent than ASHRAE's [13].

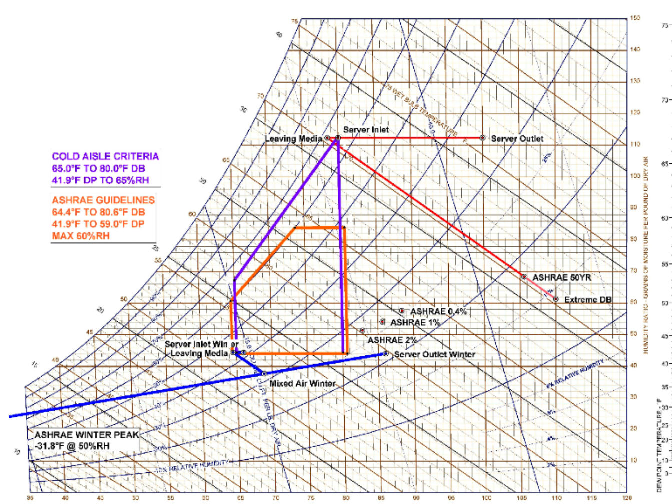

Figure 4. Psychrometric chart of KRDM's 50-year weather conditions

Table I
PRINEVILLE DATACENTER OPERATIONAL ENVELOPE COMPARISON TO ASHRAE'S GUIDELINES.

Criterion	ASHRAE 2004	ASHRAE 2008	Prineville design
Temperature: low end	$68°F$ $(20°C)$	$64.4°F$ $(18°C)$	$65°F(18.3°C)$
Temperature: high end	$77°F$ $(25°C)$	$80.6°F$ $(27°C)$	$80°F$ $(26.7°C)$
Moisture: low end	40% RH	$41.9°FDP$ $(5.5°C)$	$41.9°FDP$ $(5.5°C)$
Moisture: high end	55% RH	60% RH and $59°FDP$ $(15°C)$	65% RH

Figure 5. Rendering of airflow through the datacenter.

III. AIR FLOW

To keep the IT equipment within the specifications of Table I, the cooling scheme is designed as a "one-pass" system, employing outside air economization year-round. An evaporative cooling and humidification (EC/H) system is used to maintain the temperature and humidity of supply air within the operational envelope. A schematic rendering of the airflow through the building is shown in Fig. 5. The left side of the building in the figure (where the air inlets are) faces west, which coincides with the northwesterly prevailing winds in the summer (Fig. 6). Incidentally, the backup generators are located on the opposite side (east of the building), from where the winds are least likely to blow exhaust gases into the datacenter air inlets.

The data center is a two-storied building. The first floor holds the data hall and office space, along with the receiving yard and storage area. There is a large plenum above the data hall for hot return air. The second floor is a built-up mechanical penthouse that holds the air handling equipment line-ups. These line-ups are divided into the intake corridor, the filter room, the EC/H room, the fan-wall room, the supply corridor, and finally, the exhaust corridor. Fig. 10 depicts the building's different steps through which the air flows (see also Figs. 5,11,12 to see how these parts fit together in the building). The photos show the following areas in the air's path:

(a) The exterior wall of the penthouse consists of vertical static louvers. These louvers have "S" shaped cross-section, which helps keep precipitation from getting into the corridor, and facilitates water drainage through drain lines to drain pans.

(b) Outside air enters the intake corridor.

(c) The outside air is then introduced into the filter room, which acts as a mixing chamber. On one side of this chamber are motorized dampers for outside air (top) as well as return air (bottom). Depending on temperature and humidity of the outside air, these dampers modulate to vary the proportion of outside air and return air. This mixed air then exits the mixing chamber through a filter wall. The filter wall consists of a 2" pleated pre-filter followed by a MERV 13 filter.

(d) After passing through both filters, the mixed air enters into the EC/H room. The EC/H system uses high pressure pumps and atomizing nozzles to spray a fine mist into the mixed air stream. Multiple cooling stages are modulated based on the temperature and humidity of the supply air.

(e) The sprayed air then passes through a mist eliminator media which arrests any water molecules that are not evaporated, thereby preventing a moisture carry-over. The extra water is collected in drain pans and returned to the water loop for further processing and recirculation (Sec. V).

(f) After the mist eliminator, the air goes through a fan-wall: an array of plug fans assembled in matrix form. It is this array that does the actual pulling of air throughout the preceding sections, all the way from outside the building. On typical weather days, it is also the only moving-parts element of the building's cooling system.

(g) The supply corridor contains shafts that open into the data hall below, through which the cold air naturally descends.

(h) In the data hall, the cabinets are laid out in hot aisle/cold aisle arrangement. Server fans, aided by a pressure differential, pull an afflux of cold air over the motherboards, and exhaust the air after it picked up heat into contained hot aisles. The return air then raises from the hot aisles to the return air plenum. This containment isolates the supply air from the air exiting the IT equipment, thus avoiding the mixing of the two air streams, as well as the recirculation of hot air and the bypass of the supply air.

(i) From the return air plenum, the hot return air is introduced into the mixing chamber if the outside air conditions dictate it. The modulating dampers determine the quantity of the return air for mixing and the remainder of the hot return air is rejected to the atmosphere via relief fans in the exhaust corridor. In typical operation, these fans remain idle.

Note the absence of air ducts: the entire building serves as the air ducting systems [15]. Nevertheless, impedance to air circulation and hot/cold air mixing is still kept to a minimum, as shown in Fig. 11.

IV. SEQUENCE OF OPERATION

The cold aisle temperature and humidity is maintained within the operational envelope by mixing outside air with return air, and then misting it as necessary. The proportion of mixing and amount of misting required depends on outside air conditions, which are determined in real time by integrating data from temperature and humidity sensors strewn throughout the datacenter.

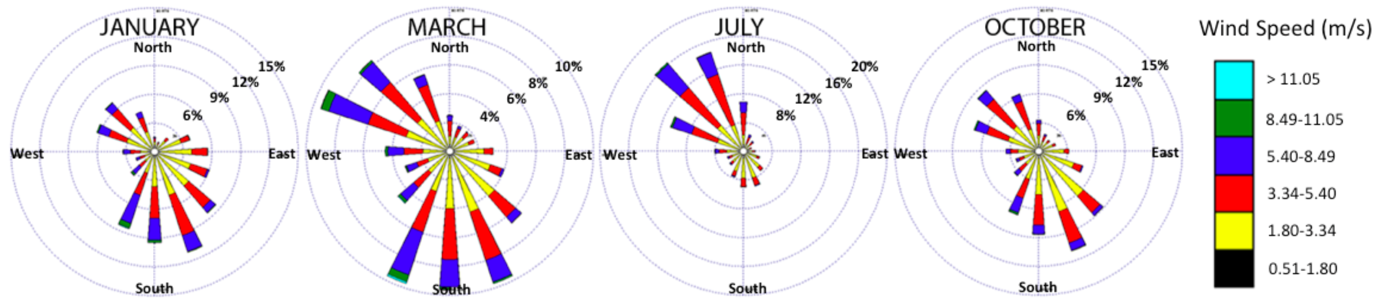

Figure 6. Wind data for KRDM (the nearest FAA weather reporting station at the Redomond, OR airport). Each compass rose represents a circular histogram of wind direction in hourly bins, for that month. During the hottest month, July, the prevailing winds come from the northwest.

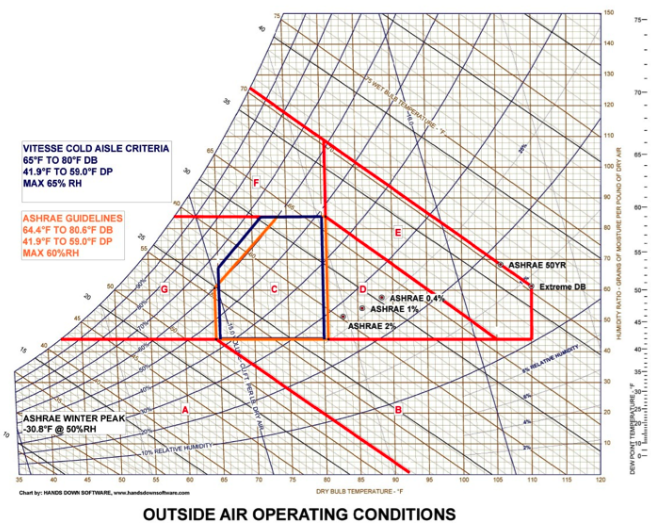

Figure 7. Psychrometric chart of distinct regions of operation

There are eight distinct operational regions as shown in Fig. 7, which cover all possible outside air conditions. The sequence in which the air handling line-ups respond while in those regions, is as follows.

Region A ($< 52^oFDB$ and $< 41.9^oFDP$): When outside air conditions lie within this region, the target supply air dry bulb temperature is $65°F$. The outside and return air dampers modulate to mix both airstreams. If required, the EC/H system stages on to provide the necessary humidification for maintaining wet bulb temperature of the supply air at $54°F$ and the dew point temperature at $42°F$.

Region B ($> 52^oFDB$ and $< 41.9^oFDP$): This region calls for 100% outside air. The return air dampers are completely closed and the outside air dampers are fully open. EC/H stages on to provide the required humidification or cooling. The supply air dry bulb temperature is maintained between $65°F$ and $80°F$ while the dew point temperature is maintained at $42°F$.

Region C ($> 65^oFDB$ and $> 41.9^oFDP$ and $< 80^oFDB$ and $< 59.0^oFDP$ and $< 65\%RH$): In this region too the return air dampers are completely closed and the outside air dampers are fully open. 100% outside air is admitted. The EC/H system remains off, since no evaporative cooling or humidification is required. The outside air is delivered into the data hall "as is" (after filtration).

Region D ($> 80^oFDB$ and $> 41.9^oFDP$ and $< 65.76^oFWB$): The economizer is at 100% in this region as well, meaning that outside air is not mixed with return air. EC/H stages on to provide required humidification or cooling. The supply air dry bulb temperature is maintained at $80°F$ while dew point temperature is maintained between $42°F$ and $59°F$.

Region E ($> 80^oFDB$ and $> 41.9^oFDP$ and $> 65.76^oFWB$): Once more, the dampers modulate to bring in 100% outside air. EC/H stages on to provide the required humidification or cooling. The supply air dry bulb temperature is maintained at $80°F$ while dew point temperature is kept above $59°F$.

Region F ($< 80^oFDB$ and $> 59.0^oFDP$ and $> 65.76^oFWB$): In this region, the dampers modulate to mix outside air with return air to increase cold aisle temperature as necessary for reducing cold aisle relative humidity to a 65% maximum. The supply air temperature is maintained between $65°F$ and $80°F$. The dew point temperature is kept above $59°F$. The direct evaporation system is bypassed, since no evaporative cooling or humidification is required.

Region G ($> 65^oFDB$ and $< 59.0^oFDP$ and > 65 and $< 63^oFWB$ or $> 65^oFDB$ and $> 41.9^oFDP$ and $> 65\%RH$ and $< 59.0^oFDP$ or $< 65.76^oFDB$: Again, the dampers modulate to mix outside air with return air to increase cold aisle temperature as necessary for reducing cold aisle relative humidity to a 65% maximum. The supply air temperature is maintained above 65^oF. and the dew point temperature is kept below $59°F$. The direct evaporation system is bypassed, since no evaporative cooling or humidification is required.

Region H (Unacceptable OA conditions): When outside air is inadmissible to the datacenter (such as excessive smoke or dust particulates in the air), the external dampers are shut.

V. Water Cycle

The direct EC/H system described in Sec. III requires water that is treated by a reverse osmosis (RO) process. The purpose of the RO system is to remove impurities from the water that could otherwise clog the misting nozzles, since the orifices on

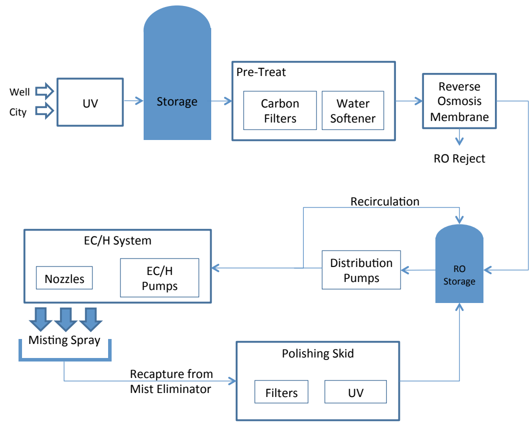

Figure 8. Water flow through the datacenter.

the misting nozzles are on the micron scale, and the datacenter has several thousand misting nozzles.

The datacenter has two RO plants, each serving half of the building. Each plant processes the water as follows (Fig. 8). The water is stored in an outdoor, above ground, storage tank with enough capacity to support 48 hours of operation during peak 50 year BIN weather conditions. There are two sources of water for the tank: the primary supply is from an on-site well, and the secondary is from the Prineville municipal water system. Within the storage system, the water is intermittently circulated through an ultraviolet (UV) filter for disinfection and to ensure that the water does not stagnate for an extended period of time. Next, the water is treated in the RO filter room. Two parallel $50\,PSI$ booster pumps (with a third pump for redundancy) pull water from the tank and through three sets of carbon filters and water softeners, to purify and remove minerals from the water. Three RO pump skids then receive the water and pump the water through the RO membranes at $50\,PSI$, to further remove large molecules and ions. We found that the RO process in Prineville produces $\approx$ 67% purified water from the input well water.

The purified RO water is then pumped into two RO storage tanks, sized for one hour of operation in the event of an RO pump skid failure at full load. Two parallel distribution pumps at $45\,PSI$, (with again a redundant third pump) circulate the RO water through another UV filter and up, to the E/CH system pump skids in the mechanical penthouse. The EC/H pump skids increase the water pressure from $45\,PSI$ to $1000\,PSI$ through the misting nozzles. Approximately 85% of the misted water is evaporated into the air stream, with the remaining 15% recaptured in the mist eliminator, to minimize water carry-over. This recaptured water is then brought back to the RO room, treated via a polishing skid with a UV and micron filter, and finally piped back into the RO storage tanks for reuse. (The intent of the RO water reclaim via the mist eliminator is to purify the RO water that has been potentially contaminated by the penthouse air stream.) The air that has now been conditioned via the EC/H system is then supplied into the datacenter via fan arrays and dry wall shafts.

Based on operational data collected thus far, we estimate the Prineville datacenter water usage effectiveness (WUE) at $0.31 Liter/KWh$ (we will publish more accurate measurements after a full year of operation). Since very little has been disclosed so far on the WUE metric for other datacenters [16], we cannot put this number in perspective, but we believe it to compare favorably with estimates for typical datacenters [17].

VI. Structural Design

For maximum efficiency, even the building design is deeply integrated with the thermal design. Fig. 12 shows a mechanical cross-section of the building plan through the cold aisle. It depicts many of the cooling aspects already discussed: the air mixing chamber, the fan array, the air filter, the EC/H system, the supply and return air shafts, the water drain from the mist eliminator back to the RO tanks, and relief fans. It also shows numerous examples of how the building structure as a whole supports the task of cooling the data hall, while minimizing cost and energy use, as discussed next.

One of the first things to note is that the entire building is used for the movement of cooling air, as mentioned in the introduction. The facility does not use ductwork to cool the servers in the data hall. Instead, they receive cooled supply air from fourteen vertical shafts located every 16 feet down the center spine of the main aisle. Effectively, this makes the data hall a pressurized plenum. The equipment's hot exhaust air is routed through a ceiling return plenum to either the exhaust fans or partially recirculated to preheat outside intake air. The server cabinets are arranged in rows in cold aisle and hot aisle configuration. The contained hot aisles extend up to a return air plenum and a static pressure differential is used to minimize server fan speeds and increase overall efficiency.

Another notable feature in the figure is that the beams that support the upper floors are all intentionally aligned behind battery cabinets, not IT equipment. The columns stand in the hot aisle, and positioning them behind working servers would have interfered with the server cooling airflow. Battery cabinets, on the other hand, produce minimal heat and airflow, providing an excellent location for the columns.

On the left (west) side of Fig. 12 we can see the office space, which is structurally separate from the data hall, and not carved out of it as in normal datacenters. This increases thermal efficiency both because we do not need to expend as much energy to cool office space as we do server space, and because we avoid impeding air flow in the data hall with extraneous offices. It also allows lighting the office space with more sunglight than is normally accessible or desired in the data hall.

Finally, the building also supports rainwater collection for irrigation and toilets through a collection tank in the central yard. The facility also provides a $100kW$ solar array to support noncritical IT equipment. Although not directly related to the thermal design, this design features do affect the overall environmental footprint of the datacenter, and further reduce its water and power consumption.

VII. Discussion

Many of the ideas discussed in this paper had originated previously in different forms. For example, using evaporative cooling has been suggested and even implemented before [13]. Another example is the use of 100% outside-air economization, even when outside temperatures are hot [18]. In Prineville's case, this idea makes a lot of sense, because even the warmest summer days are still cooler than the server's outlet air, so recirculating and chilling air is far less efficient. What makes this datacenter unique, perhaps, is its low overhead, in terms of water usage, cost, and power usage efficiency (PUE).

These efficiencies are intertwined. The main reason for the lower water usage (WUE of $0.31 Liter/KWh$, as discussed in Sec. V) is that very little power is used for cooling in the EC/H system, improving PUE and reducing the amount of required cooling, and consequently, water evaporation [17]. Cost is also tightly linked to efficiency. Obviously, opex is directly affected by the amount of overall power used in the datacenter, so lower PUE translates to lower opex. But capex too is affected by efficiency, sometimes indirectly. Lower PUE means less overall power, which requires less power distribution equipment to provision and lower backup power requirements, reducing cost. In some cases, higher efficiency components and a distributed backup power solution can actually be obtained with cheaper and fewer components too [3]). And fewer components, in turn, also increase reliability, further reducing cost and increasing efficiency.

Since so much hinges on efficiency, it is important to measure it accurately. To quantify power waste, we continually measure power output of the datacenter in two points: at the utility connection outside the building, before the transformation down to $480VAC$, and again at the output of the reactor power panel at each server row, right before the IT equipment. The former number divided by the latter yields the PUE metric. This metric is not ideal, and can be controversial [19]. For example, the energy spent on server fans counts as "useful" energy spent in the IT equipment, so one could artificially lower PUE by shifting the cooling burden from building fans to the less efficient server fans.[2] Nevertheless, it is a reasonable approximation of datacenter power efficiency, and it quantifies the "waste" power that does not end up in the IT equipment. Our continuous PUE samples throughout the summer of 2011 varied in the range 1.06–1.1, and averaged 1.083 at full load (Fig. 9). A waste of only 8% on cooling energy and power distribution compares favorably with the industry's best [13], [20], especially when taking the seasonally warm temperatures into account. We expect the full-year PUE to average around 1.07. Compare it, for example, to a best-practice leased datacenter with a PUE of 1.5 [21]. Our design uses 86% less energy on overhead ($\frac{0.5-0.07}{0.5} = 86\%$) and 29% overall less energy in the datacenter ($\frac{1.5-1.07}{1.5} = 29\%$).

The main lesson here is that co-designing all aspects of the datacenter together, as opposed to piecemeal from commodity components, can yield higher efficiencies in all of its aspects.

[2] In fact, our server design achieves the opposite effect with high-efficiency fans and lower overall power usage compared to commodity servers [3].

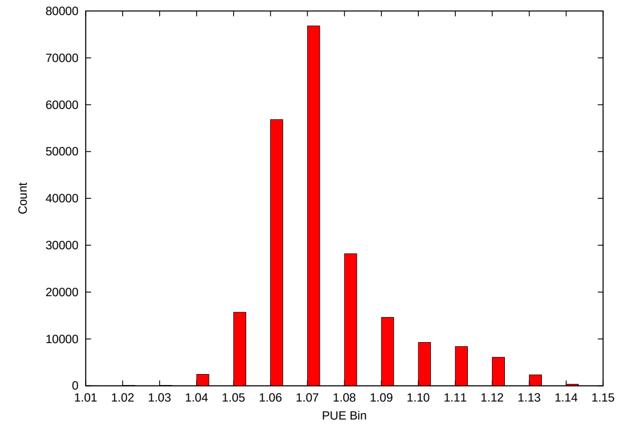

Figure 9. PUE measurements in hourly bins over the summer of 2011

The overall efficiency is higher than the sum of its parts, because each part can make assumptions about the others. For example, the EC/H system can assume that the custom-made servers have better airflow than commodity servers [3], and thus expend less energy on pressure and temperature differentials.

Still, these features represent only the first iteration of our datacenter design. Although they have resulted in high efficiency, we have learned some lessons along the way. One challenge we encountered, as an example, was keeping our air handler lineups from "fighting" with each other as they dealt with the rapid changes in the temperature and humidity of the outside air between day and night (e.g., if outside air dampers of one lineup are at 70%, the adjacent lineups would have their outside air dampers at 20-30$\%$). This alternate modulation, or fighting, could lead to stratification of the air streams. This issue was fixed by modifying the proportional-integrative derivative (PID) control loop. Other areas for future research also include improved server design and power distribution.

VIII. Conclusions

This paper presented the thermal design of Facebook's first custom-made datacenter. Some of the most interesting properties of this design are [15]: A) 100% Outside-air economization with full-wall, low-resistance filters; B) The entire building serves for air movement, with no ducting and with large plenum areas; C) Mist-based, no-process evaporative cooling; and D) Large, efficient impellers with variable frequency drive.

We have also described the permissive environmental envelope in the datacenter design, the operational sequence to remain within this envelope based on outside air conditions, and the cycle of water used in the evaporative cooling.

Throughout the first six months of operation, this design has proven to use significantly less energy than our colocation datacenters, while at the same time costing less to equip and operate. It has also been awarded the US Green Building Council's LEED Gold Award, as well as Engineering News-Record's Best Green Building Award for 2011. But this design is only a first iteration. We are currently building two additional datacenters in Forest City, NC, USA and Lulea, Sweden, expanding on this design as a basis, and incorporating additional improvements that are still in the planning phase. Once these new datacenters are operational, we plan to measure and report their efficiency and modifications, as well as contribute the design back to the Open Compute Project. Our hope is that by opening up the design for comments and contribution through the Open Compute Project, we will be able to move rapidly toward further improvements.

REFERENCES

[1] C. L. Belady, "In the data center, power and cooling costs more than the IT equipment it supports," *Electronics Cooling*, vol. 23, Feb. 2007, www.electronics-cooling.com/2007/02/.

[2] Facebook, "The Open Compute server architecture specifications," www.opencompute.org, Apr. 2011.

[3] E. Frachtenberg, A. Heydari, H. Li, A. Michael, J. Na, A. Nisbet, and P. Sarti, "High-efficiency server design," in *Proceedings of the 24th IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, Seattle, WA, Nov. 2011.

[4] J. Hamilton, "Cost of power in large-scale data centers," perspectives.mvdirona.com/2008/11/28/CostOfPowerInLargeScaleDataCenters.aspx, Nov. 2008.

[5] ——, "Internet scale storage," in *ACM Special Interest Group on the Management of Data Annual Conference (SIGMOD'11)*. Athens, Greece: ACM, Jun. 2011, slides available at mvdirona.com/jrh/TalksAndPapers/JamesHamilton_Sigmod2011Keynote.pdf.

[6] J. G. Koomey, "Estimating total power consumption by servers in the US and the world," LBNL, Tech. Rep., Feb. 2007, citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.87.5562&rep=rep1&type=pdf.

[7] J. Koomey, "Growth in data center electricity use 2005 to 2010," Analytics Press, Oakland, CA, Tech. Rep., Aug. 2011, www.analyticspress.com/datacenters.html.

[8] A. Qureshi, "Plugging into energy market diversity," in *Proceedings of the Seventh ACM Workshop on Hot Topics in Networks (HotNets)*, Oct. 2008, citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.161.8738.

[9] X. Fan, W. dietrich Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *Proceedings of the 34th International Symposium on Computer Architecture (ISCA)*, San Diego, CA, Jun. 2007, citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.84.9052.

[10] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and limitations of tapping into stored energy for datacenters," in *Proceedings of the 38th International Symposium on Computer Architecture (ISCA)*, ser. (ISCA'11). New York, NY, USA: ACM, 2011, pp. 341–352, www.cse.psu.edu/~sgovinda/papers/isca11.pdf.

[11] A. Krioukov, P. Mohan, S. Alspaugh, L. Keys, D. Culler, and R. Katz, "NapSAC: design and implementation of a power-proportional web cluster," *ACM SIGCOMM Computer Communication Review*, vol. 41, pp. 102–108, 2011, ccr.sigcomm.org/drupal/files/p102-green12t-krioukovA.pdf.

[12] K. Lim, P. Ranganathan, J. Chang, C. Patel, T. Mudge, and S. Reinhardt, "Understanding and designing new server architectures for emerging warehouse-computing environments," *ACM SIGARCH Computuer Architecture News*, vol. 36, no. 3, pp. 315–326, 2008, portal.acm.org/citation.cfm?id=1382148.

[13] EcoCooling, "2.5 MW data centre is kept cool by evaporative cooling," www.modbs.co.uk/news/fullstory.php/aid/8384/2.5_MW_data_centre_is_kept__cool_by_eva, Oct. 2010.

[14] M. Hydeman, "Implications of current thermal guidelines for data center energy use," *ASHRAE journal*, pp. 30–41, Aug. 2010.

[15] J. Hamilton, "Open Compute mechanical system design," perspectives.mvdirona.com/2011/04/09/OpenComputeMechanicalSystemDesign.aspx, Apr. 2011.

[16] R. Miller, "Water usage: The next disclosure challenge?" www.datacenterknowledge.com/archives/2010/04/07/water-usage-the-next-disclosure-challenge/, Apr. 2010.

[17] R. Cockrell, "Quneching the thirst of power hungry data centers," *Data Center Journal*, May 2009, www.write-angle.com/wp-content/uploads/2011/03/DCJarticle050809FINAL.pdf.

[18] M. K. Patterson, "The effect of data center temperature on energy efficiency," in *Proceedings of the 11th Intersociety Conference on Thermal and Thermomecachnical Phenomena in Electronic Systems (ITHERM)*, Orlando, FL, May 2008, pp. 1167–2274, www.ecoinfo.cnrs.fr/IMG/pdf/Michael_K_Patterson_-_The_effect_of_Data_Center_Temperature_on_Energy_Efficiency.pdf.

[19] R. Miller, "Microsoft eliminates server fans, despite PUE hit," www.datacenterknowledge.com/archives/2011/01/31/microsoft-eliminates-server-fans-despite-pue-hit/, Jan. 2011.

[20] Google, "Quarterly PUE benchmark data," www.google.com/corporate/datacenter/efficiency-measurements.html, Jul. 2011.

[21] "Report to congress on server and data center energy efficiency," Environmental Protection Agency, Tech. Rep., Aug. 2007.

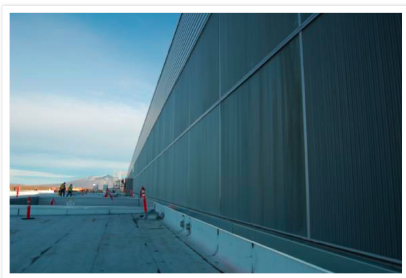

(a) vertical louvers for outside air

(b) outside air intake corridor

(c) mixing and filter room

(d) EC/H room

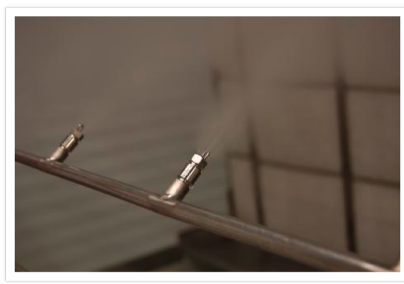

(e) spray nozzles closeup

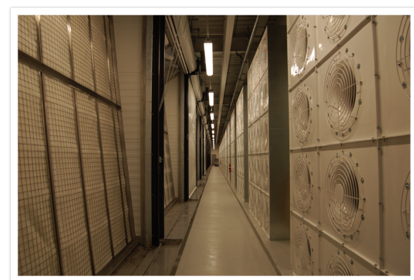

(f) mist eliminator and fan wall array

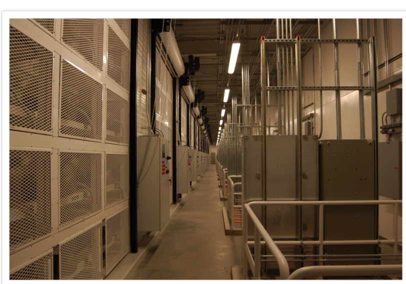

(g) supply shaft for cold air down to data hall

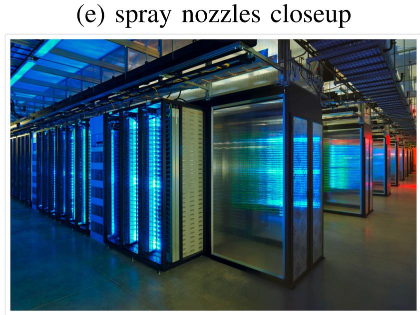

(h) data hall with hot-aisle containment

(i) relief fans

Figure 10. Photos of the datacenter along the path of the airflow.

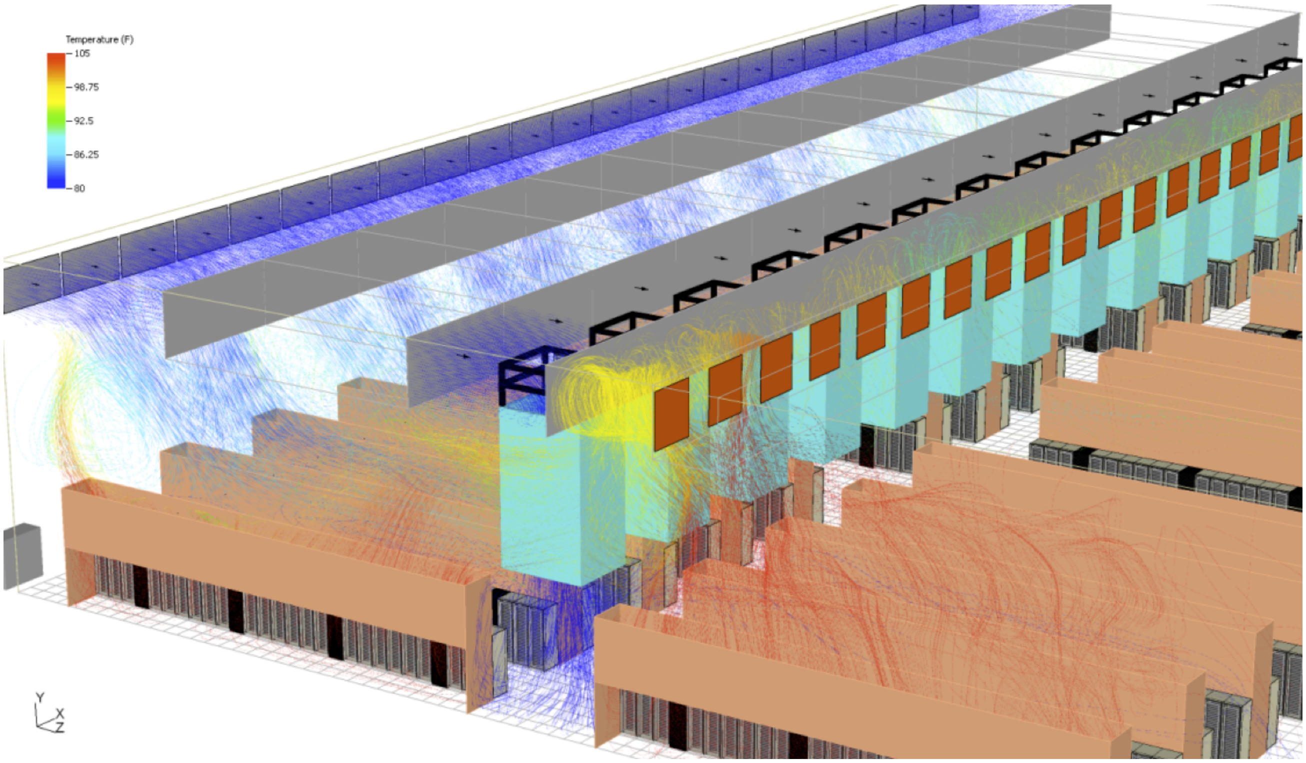

Figure 11. 6SigmaDC simulation of air circulation through the building.

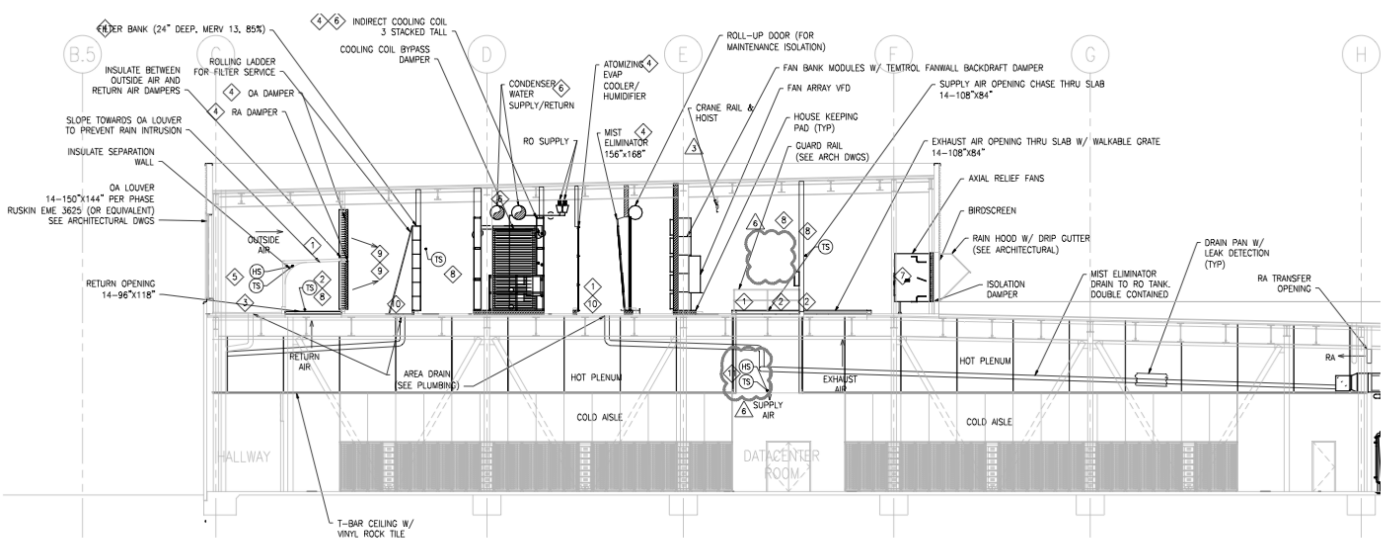

Figure 12. Mechanical section of cold aisle.