Waste Heat Recovery Using Coupled 2-Phase Cooling & Heat-Pump Driven Absorption Refrigeration

R. Khalid[1], S. G. Schon[2], A. Ortega[1], A. P. Wemhoff [1]
[1]Department of Mechanical Engineering, [2]QuantaCool Corp.
Villanova University,
800 Lancaster Avenue
Villanova, PA, USA, 19085
Email: aaron.wemhoff@villanova.edu

ABSTRACT

Data center waste heat recovery is an energy efficient and economically viable option when the data center is near other facilities. This study explores boosting the vapor from a data center's two-phase server waste heat using a novel vapor re-compression system. The boosted waste heat contains sufficient thermal energy to drive an absorption refrigeration (AR) chiller to obtain a stream of cold fluid. Alternatively, the chiller can be bypassed to use the boosted heat directly in the neighboring facility. These approaches are modeled to enable the estimation of energy savings and economic benefits under different cooling and heating loads. These calculations also indicate situations where the compressor should be by-passed, and the server exhaust simply sent to the condenser for rejection to the ambient. The analysis is focused on the waste heat recovery and re-use potential of a 1 MW data center with a mid-size office building in the vicinity. The paper focuses on two key issues: 1) the overall efficiency of the waste heat system, including neighboring sites, when boosted with the novel re-compression system versus a turn-key heat pump coupled with warm-water cooling and 2) a comparison of the relative capital and operating costs of the two systems. The results show that the novel recompression system has a lower cost and a higher thermodynamic efficiency than a warm-water cooling system coupled to a conventional heat pump, and the use of the novel waste heat recovery system is viable in areas with high electricity cost.

KEY WORDS: data center, vapor recompression, co-location, energy efficiency, thermosyphon based, passive design

NOMENCLATURE

ABS.	absorption
AR	absorption refrigeration
BRS	boost and recovery system
COP	coefficient of performance
DC	data center
DX	direct expansion
FoM	figure of merit
HX	heat exchanger
$\dot{Q}$	heat transfer rate, Watt
$\dot{W}$	rate of doing work, Watt

Greek symbols

ε	effectiveness (of heat exchanger)

Subscripts

evap	evaporator
mp	microprocessor

INTRODUCTION

Data centers can be viewed as giant warehouses or small cabinet enclosures that ultimately discharge heat into the ambient environment. The heat generated at the server level must be cooled by a system that typically uses electricity, sometimes in conjunction with free air cooling. This scheme presents two disadvantages in terms of the electric energy consumed by the cooling system and the discharge of heat to the ambient without any intermediate use.

Although waste heat recovery is widely utilized in high temperature applications such as engine gas exhaust [1, 2] or gas turbine exhaust [3], this concept is seldom heard of in the context of data centers due to the low temperature, and hence low quality, of the server exhaust. The two elements required in waste heat to make its capture and utilization useful are sufficient quantity and high quality. Quantity refers to the amount of waste heat that is available from a process or the amount of heat that can be successfully recovered while quality outlines the temperature at which that waste heat is available. Thermodynamically, they represent the waste heat recovery process from a first and second law perspective - without sufficient quality, the quantity does little good and vice versa. Data center waste heat, although in abundant quantity depending on the scale of the data center, is generally of low quality and hence does not present itself for further use – as a result, it is generally discarded or used only for less energy intensive applications such as water desalination [4, 5], drying biomass [5] or anaerobic digestion [6].

The growing use of liquid cooling allows waste heat recovery to become a promising option for data centers, since liquid cooling allows higher coolant temperatures than air because of its superior heat transfer. Facebook just recently announced a new data center in Singapore that is going to feature liquid cooling [7]. Similarly, Google's TPU's are cooled via liquid cooling [8]. Thus, this growing popularity means that liquid cooling is starting to feature in more mainstream applications and is no longer a niche for applications such as high-performance computing. Liquid cooling traditionally uses water as the cooling medium due to its ease of availability, usage safety and favorable thermal properties such as high heat capacity. Hence, liquid cooling lends itself more favorably to waste heat capture and boosting than air cooling due to the higher temperatures achieved due to water's higher thermal heat capacity.

Several studies look at low-grade waste heat recovery in data centers. Ebrahimi, Jones and Fleischer [9], in their paper on a review of data center cooling technology, did a thorough review of the existing methods for low-grade waste heat recovery. Based on the eight methods reviewed, they recommended absorption refrigeration and power generation via an organic Rankine cycle as the most promising techniques for low-grade waste heat recovery in data centers.

Further expanding on their initial assessment, Ebrahimi et al. [10] did a thorough thermo-economic study of data center waste heat recovery using absorption refrigeration. They developed thermodynamic models for absorption refrigeration using ammonia/water and water/lithium bromide while validating them against two previously published experimental studies. Their results showed higher COP values for water/LiBr for the range of temperatures common in data center waste heat. In addition, the water/LiBr cycle is less costly and less complicated to implement due to fewer parts, with crystallization of the LiBr as the only major issue. The authors rectified this by recommending the maximum concentration of LiBr in the strong absorber solution to be capped at 65%.

Further, Ebrahimi et al. used the novel on-chip liquid cooling method proposed by Marcinichen et al. [11, 12] and coupled that with an AR cycle by replacing the condenser of the on-chip cooler with the generator of the AR machine. Thus, the AR generator is effectively provided by heat from the microprocessor chip on the servers, not only cooling them but also utilizing the condenser 'exhaust' as 'fuel' for the AR generator. They redefined the COP of this coupled system as COP = Q_{evap}/Q_{mp}, where Q_{evap} is the cooling effect generated or heat load removed by the evaporator of the AR machine and Q_{mp} is the heat generated by the microprocessor of the server cooled by liquid cooling.

They further identified a figure of merit (FoM) as the ratio of Q_{mp} to that of Q_{evap}. This FoM effectively defines the number of individual liquid cooled servers or racks that are needed to produce enough cooling to cool a server or rack through the combined system. Using data from two US-based vendors of AR machines, they identified that a given data center can be entirely cooled using these combined AR machines without any CRAC/CRAH input. As such, they identified the payback time for these AR machines to be as short as 4-5 months for a 25-ton AR chiller from the vendor Yazaki Energy Systems Inc.

Kim et. al. [13], theoretically studied the feasibility of using an absorption based miniature heat pump for electronics cooling. Water/LiBr pair was used as the working fluid in their analysis. The properties of LiBr solution were taken from Yuan and Herold [14]. Their novel design incorporated a dual-channel micro-evaporator that significantly reduces pressure drop for microchannel two-phase flow as compared to single channel micro-evaporators. A combined condenser/desorber component interacting thermally and separated by a hydrophobic membrane is also featured in their miniature heat pump. The membrane allows for mass exchange as water evaporates from the weak solution in the desorber and moves across the membrane to the condenser where it is cooled to a subcooled liquid state. The condenser is air cooled using offset fin strips to enhance heat transfer. A miniature pump is featured to increase the pressure of the weak solution leaving the absorber to that of the desorber. The pump's displaced volume and power consumption are small and hence negligible. Their entire system envelopes a volume 150mm x 150mm x 100mm. Through their work, they demonstrate the feasibility of such a miniaturized system for electronics cooling, with the ability to remove 100W of heat.

Haywood et. al. [15] experimentally studied the thermodynamic feasibility of using server waste heat in a data center to run an absorption driven chiller. They identified IBM's 'zero-emission' data center [16] project as having the same goal but a different approach as to how to reuse the waste heat to drive the cooling process. The authors' water-fired chiller used water/LiBr as the working pair. They identified several factors affecting the performance of the absorption chiller including the cooling water inlet temperature, heat medium flowrate and temperature of the heat input to the desorber, which determines the resultant cooling capacity at the evaporator. Using performance curves comparing COP and the generator and evaporator heat load to the generator temp, they show that there is a trade-off between maximum efficiency (highest COP) and maximum cooling (highest evaporator heat load). In general, their results showed that the hotter the heat medium, the lower the cooling water temperature or the higher the heat medium flowrate, the larger the heat load that could be removed at the evaporator.

The authors argued that when the heat input to the desorber is essentially free, as in the case of waste heat which is a by-product of server operation, then cooling capacity, not efficiency, is the main consideration. The main challenge is (1) capturing enough sufficient temperature waste heat and (2) transporting that heat from the CPUs on the servers to the desorber of the AR chiller. Thus, both quality and quantity of waste heat need to be considered.

The paper also introduces the concepts of Highest Heat Fraction (HHF) and Capture Fraction (CF) which point to the highest heat that can be generated from a server and the amount of that heat which can be captured, respectively. They identified 70% contribution of a server heat dissipation coming from CPUs. Further, a capture fraction of 85% is possible with liquid cooling using cold plates with water running through them. Since water is the refrigerant in the AR chiller, water-water heat exchange is convenient and more efficient as compared to air to water heat exchange. The authors identified the resistance between the server and the cold plate as the biggest bottleneck in heat capture and recommended incorporating a high-conductivity thermal interface material (TIM) between the two surfaces.

Lastly, the authors utilized solar thermal energy to supplement server waste heat. In this case, PUE[1] values of below unity are possible and facility energy consumption can be negative, since the definition of PUE does not account for energy recovery. Thus, they recommend using Energy Reuse Effectiveness, ERE, as a metric that can potentially enable a data center to be a net producer rather than a consumer of energy.

Liquid coolants, especially water, pose their own hazards near electronics but their harm can be mitigated by switching to relatively inert coolants such as Freon (R-22, R-134a etc.).

[1] Defined as the ratio of total facility energy consumption to IT equipment energy consumption

Common commercially available refrigerants present their own issues such as low boiling points and the subsequent high vapor pressures needed for reliable server operation, which further poses a health and safety issue for the data center operators in case of a leakage. The required high vapor pressures result in thicker piping with larger diameters and similarly larger equipment that can sustain those pressures, thus increasing the CAPEX of the boost and recovery system.

Low-grade waste heat capture and recovery from data centers has been significantly explored in the past. However, boosting this low-grade heat to make it more thermo-economically feasible is an area that lacks in the open literature. Hence, this paper explores (i) the use of lower-pressure refrigerants such as R-245fa and R-1233zd flowing through two-phase cold plates for server cooling, (ii) the concept of boosting the vapor from the server exhaust to make it high quality, and (iii) using the boosted vapor to supplement the heating and cooling requirements of a co-located facility.

DATA CENTER LAYOUT

Figure 1 shows the configuration of cooling a single rack using refrigerant flowing through two-phase cold plates mounted in series or parallel on each server. Refrigerant exits the cold plates with a quality of about 60% and gets collected in a liquid-vapor separator. Within the separator, liquid settles at the bottom while the vapor rises above the liquid due to its lower density. The liquid refrigerant then flows down using gravity from the separator into the cold plates, gathers heat while cooling the servers and the cycle repeats. Conversely, the vapor in each separator rises due to buoyancy and collects in the manifold above each row.

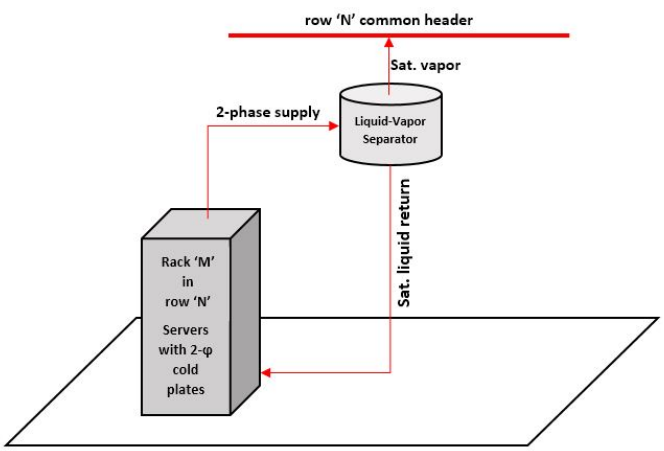

Fig. 1 Cooling a Single Rack Using Two-Phase Refrigerant and Cold Plates

From thereon, the combined vapor from all the racks in each row gets collected in the manifold above each row and gets sent to a common header of the boost and recovery system (BRS). Figure 2 details the complete plumbing setup for the BRS, indicating the row manifolds and common header for the entire system. The figure represents a 1MW data center split into 50 racks of 20kW each. Rack power densities in data centers vary between 5 kW on the low side to about 30 kW for high-end enterprise data centers [17]. Thus, the 20kW option chosen for this study represents the median density of most enterprise racks in operation.

The combined vapor in the common header flows through a regenerative heat exchanger before being boosted through a compressor. Large centrifugal compressors can handle the vapor flowrate from a 1 MW and larger data center. The footprint of such a compressor along with the motor and a variable speed drive would be about 12' x 10'. Legacy data centers can install the compressor assembly outside the data center airspace while future data centers can make space allowance the size of the compressor footprint within the data center. Fig. 2 shows the compressor along with other components of the BRS placed in one corner of the data center airspace.

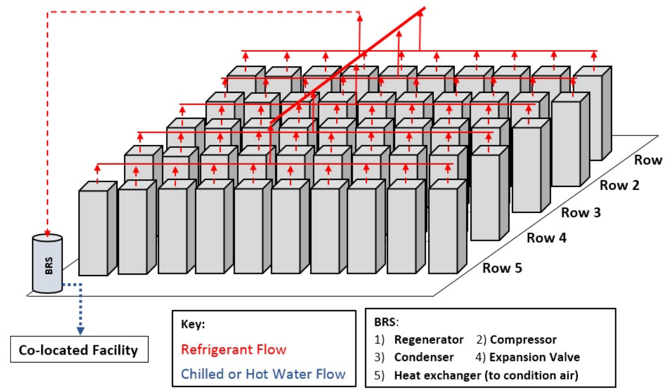

Fig. 2 Data Center Floor Plan showing 50 racks of 20kW each arranged in five rows

The details of the BRS are depicted in Fig. 3 using a 2-D schematic. The compressor boosts the temperature and pressure of the incoming vapor while itself consuming electric power. Based on seasonal demand, the boosted vapor then has sufficient quality and quantity to either drive an absorption chiller to produce chilled water for cooling a co-located facility through a separate heat exchanger or directly heat the water or air in the heat exchanger to supplement the heating requirements of that facility. In off-demand times, the regenerator and compressor of the BRS can simply be by-passed and the vapor from the common header sent to an overhead air-cooled condenser to reject the data center heat to the ambient.

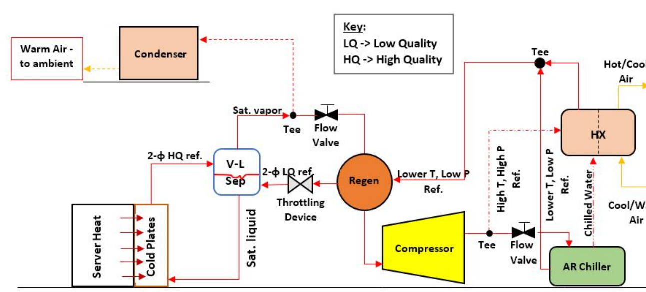

Fig. 3 Overall Schematic of the Boost and Recovery System showing the Use of the Boosted and Non-Boosted Refrigerant

To complete the loop, liquid refrigerant from either the absorption chiller or neighboring facility's heat exchanger flows through the regenerator in case of boosting or from the elevated (rooftop) condenser to be throttled through an expansion valve to a low pressure two-phase mixture that gets collected in a common, larger separator. From there, liquid refrigerant can be drained to each rack's individual separator to cool the servers while the vapor gets sent to the common header to be sent back to the compressor.

WASTE HEAT RECOVERY AND BOOSTING SYSTEMS

This paper looks at two separate systems for data center waste heat capture and recovery. System I is typically used for waste heat capture and has been reported in the literature. This system includes a pumped water loop to cool servers where the water loop is coupled to a heat pump loop running a refrigerant. The other option is a novel system that has been proposed for this study. This system features a single refrigerant loop using gravity driven two-phase refrigerant to cool the servers with the vapor refrigerant being compressed in a vapor recompression system.

Boost and Recovery System I

Figure 4 shows System I for waste heat recovery and boosting. This system consists of warm water-based server cooling coupled to a heat pump using an intermediate water chiller.

The design consists of liquid water at about 35°C circulating through water blocks mounted on servers to cool them. The heated water exiting the water blocks at 38°C flows through an intermediate heat exchanger coupling the water loop to a refrigerant based heat pump loop. The water exits the HX at a lower temperature and is pumped back to the water blocks. The refrigerant exits the HX as a vapor and is preheated through a regenerator before being drawn into a compressor, where it gets boosted to superheated state at high pressure.

This higher temperature, higher pressure refrigerant has now more quantity and quality of heat than the water exiting the water blocks. It can now be used to drive an absorption chiller to produce chilled water at its evaporator or simply run through a condenser to heat water for further use.

The heated refrigerant has more exergy than the cooling water exiting the water blocks. The only significant power input is that to the compressor, since the heat generated by the servers to heat the water is a by-product of an existing process and hence can be considered 'free'. Thus, this type of system ideally lends itself in the use of absorption refrigeration cycles since the heat required to drive the generator is coming from a source deemed to be free and undesirable. However, since absorption chillers require the incoming heat source driving the generator to be at least 75°C to have a significant COP, the boosting or compressor part is unavoidable.

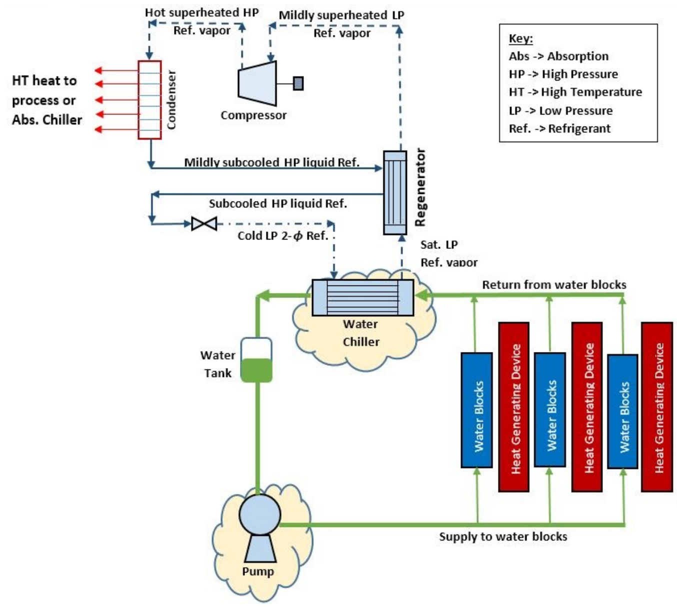

Fig. 4 System I for Recovering and Upgrading Waste Heat - Warm Water Cooling with Indirect Heat Pump

Boost and Recovery System II

The second system, System II, proposed for this study is a novel vapor recompression system utilizing two-phase refrigerant based cooling coupled to a heat pump. Additionally, the refrigerant is gravity-fed to two-phase cold plates mounted on servers using a thermosyphon-based design rather than an active one employing a pump. Figure 5 details this system.

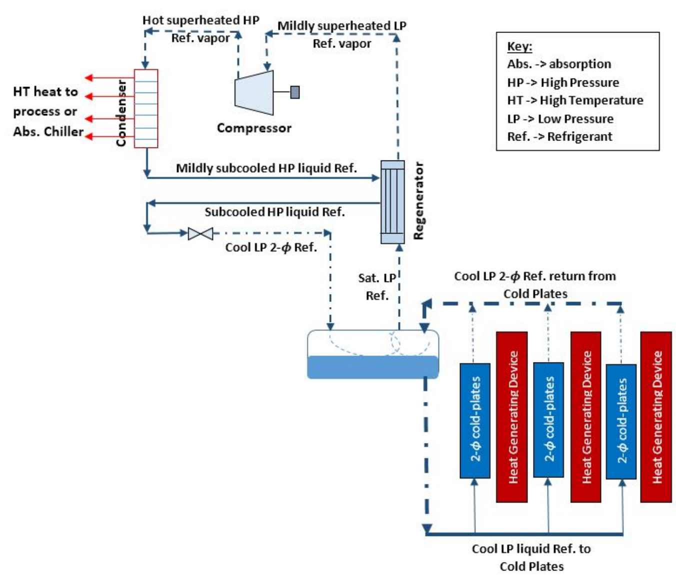

Fig. 5 System II for Recovering and Upgrading Waste Heat – Thermosyphon based two-phase cooling with Direct Vapor Recompression

Saturated liquid refrigerant at about 38°C (similar to System I's water block exit temperature) enters the cold plates while absorbing heat dissipated from the servers. The refrigerant exits as a two-phase liquid-vapor mixture as separate streams from each cold plate and gets collected in the liquid-vapor separator. From there, liquid refrigerant is drained back into the cold plates using gravity, while the vapor is pre-heated through a regenerator before going into the compressor to be superheated at a higher pressure.

This superheated refrigerant vapor at a higher temperature and pressure possesses greater energy and exergy than the two-phase mixture leaving the cold plates. Hence, it can also be used to produce electric power, cooling or heating as in the System I case. However, this process is more efficient than the primary method as explained in the next section.

Comparison of System I and II

System II is more efficient than System I because of two reasons:

1. The proposed system eliminates the intermediate heat exchanger (labeled "water chiller" in Fig. 4). Hence, this reduces exergy losses due to heat transfer from the water loop to the refrigerant loop.
2. The proposed system uses a passive design and eliminates the water circulation pump. Hence, it saves the pump power.

System II is simpler in design and is a single loop system (refrigerant only), while System I is a dual loop system (refrigerant and water). Coupling the two loops using a heat exchanger adds latency to the system for varying server loads, so the primary system is slower to respond to changes in server utilization than the proposed system.

Independent simulations were run in ASPEN Plus for the two systems to determine the compressor power requirement for a 1000 kW (1MW) heat load at the servers. The results are shown in Figs. 6 and 7 below. In both cases, the refrigerant (R-245fa) is boosted to 110 psig and this boosted waste heat is used to either directly drive an absorption chiller to produce chilled water in the 7-12°C range or hot water at 70°C through a counterflow heat exchanger.

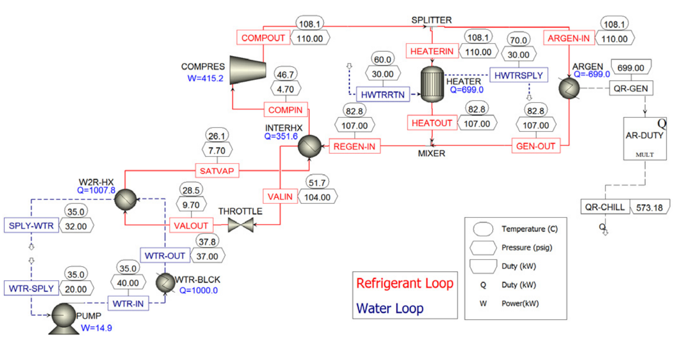

Fig. 6 Boost and Recovery System I Simulated in ASPEN Plus Showing Combined Heating and Cooling Modes

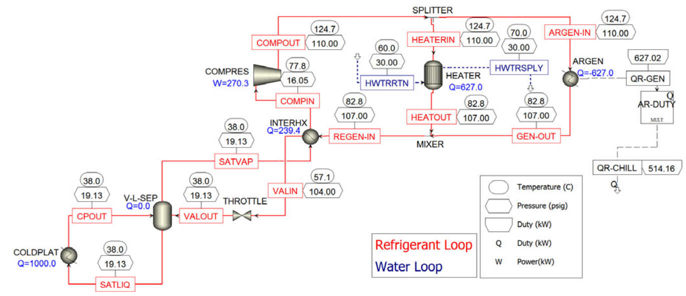

Fig. 7 Boost and Recovery System II Simulated in ASPEN Plus Showing Combined Heating and Cooling Modes

For each system, ASPEN Plus determines the compressor power requirement along with the heat imparted to the chiller's generator and/or the water heater (For the purpose of illustration, this example assumes half the vapor is used for heating and half cooling). These are tabulated in Table 1 below.

Table 1. Simulated Values for Systems I and II in Combined Heating and Cooling Mode for a 1 MW Data Center

Component	System I (kW)	System II (kW)
Compressor	415	270
AR Generator	699	627
Water Heater	699	627

For the same condensing temperature of the chiller and the supply hot water temperature (82.8°C and 70°C respectively), Table 1 shows that the compressor work requirement for System II is 35% lower than System I. In terms of system efficiency, we define a metric called recovery COP, calculated as follows for the heating mode:

$$COP_{rec,heating} = \dot{Q}_{water\text{-}heater}/\dot{W}_{comp} \qquad (1)$$

i.e. the ratio of the rate at which heat is imparted to the water heater heat exchanger to the compressor power input. Similarly, a cooling mode recovery COP can be defined as follows:

$$COP_{rec,cooling} = COP_{chiller} \cdot \dot{Q}_{AR\text{-}gen} / \dot{W}_{comp} \qquad (2)$$

Assuming the same heat load imparted to the water heater and the chiller's generator (as shown in figures 6 and 7) as well as the same chiller COP value, equations (1) and (2) dictate that the recovery COP is only dependent on the compressor power input and as such, is inversely proportional to it. Since System II consumes 35% less compressor power, its COP is correspondingly higher. Thus, System II is more efficient than System I in both heating and cooling mode.

Further, for waste-heat driven absorption chillers, the generator/desorber condensing temperatures are typically less than 100°C. Literature shows the COPs of such chillers to be below unity [9, 12]. Considering the same heat load at the water heater and AR generator again, a chiller COP of less than unity dictates that the recovery COP for cooling be lower than that for heating as per equation (2).

SUPPLEMENTING ENERGY REQUIREMENTS OF A CO-LOCATED FACILITY

The goal of boosting the refrigerant is to produce hot or chilled water by either running the refrigerant through a water-cooled heat exchanger or running it through the regenerator of an absorption chiller respectively. The hot or cold water can then be supplied to a co-located facility. This mitigates the cost of fuel used by the facility to heat or cool water that is used by its HVAC system.

Cooling Load Analysis

Consider a 1 MW data center and a boost and recovery system comprising of a compressor, water cooled heat exchanger (water-heater) and a LiBr based absorption chiller. Based on seasonal demand, the water-heater may be by-passed and all the exhaust from the compressor sent to the absorption chiller to maximize cooling load and produce chilled water only. In this case, additional simulation results show that the cooling capacity generated at the chiller's evaporator is simply twice that shown in figures 6 and 7. Hence those values can be used for further calculations.

Large centrifugal compressors for such systems typically have isentropic efficiencies of around 78% with drive train efficiencies around 94%, yielding an overall efficiency of about 73%. Similarly, LiBr based absorption chillers operating at a desorber temperature between 75 – 85°C have a COP value of about 0.82. The cooling rate at the chiller's evaporator can be calculated as:

$$\dot{Q}_{evap} = COP_{chiller} \cdot \dot{Q}_{AR\text{-}gen} \tag{3}$$

This system can produce chilled water between 7 – 12°C which can further be used to meet the cooling demands of a neighboring facility by cooling the supply air to about 15°C. The recovery COP calculated in Table 2 is based on the cooling rate produced at the chiller's evaporator.

Table 2. Comparison of Cooling Produced From a 1 MW Data Center Using Systems I and II

Parameter	System I	System II
$\dot{Q}_{chiller,\ evap}$ (kW)	1150	1030
COP_{rec}	2.76	3.80

Table 2 shows that both Systems I and II are thermodynamically feasible and that System II has a COP slightly higher than a typical residential air conditioner, which is about 3.0 [18]). Furthermore, as mentioned before, the relative COPs of Systems I and II depend only on the compressor input. However, the individual COPs rely on the absorption chiller COP which in turn is influenced by the generator source and condensing temperature. Previous studies show this dependence to be minor for the range of temperatures involved.

Now consider the cooling load of a typical medium-sized office building in various locations across the US [19]. The locations chosen are listed in Table 3 along with their respective ASHRAE climate zones. The BRS output is considered the same across all four locations but the office building cooling load varies. The cooling requirements for this building can be adequately met by the above given boost and recovery system, as evidenced in Table 4 below. Also, it is evident that warmer climates benefit from higher cooling annual savings compared to colder climates.

Table 3. Chosen Locations across the US for Data Center Activity with Colocation

S. No.	Place	ASHRAE Climate Zone
Location 1	Miami, FL	1A
Location 2	Phoenix, AZ	2B
Location 3	San Francisco, CA	3C
Location 4	Boulder, CO	5B

Table 4. Comparison of Cooling Load Requirement of a Typical Medium-Sized Office Building in Various Locations

Parameter	Loc. 1	Loc. 2	Loc. 3	Loc. 4
DX system sized capacity (kW)	424	429	347	357
System I buildings served	2.70	2.67	3.30	3.21
System II buildings served	2.43	2.40	2.96	2.88
Annual cooling electricity usage per building (kWh) * 10^3	235	202	38.9	0.50
Average annual electricity rate ($/kWh)	0.086	0.097	0.146	0.037
System I annual savings ($ 000)	54.6	52.4	18.8	0.059
System II annual savings ($ 000)	49.0	47.0	16.8	0.053

Heating Load Analysis

Now consider the same BRS operating in a 1 MW data center. In heating mode, such as during the winter season, the boosted refrigerant by-passes the absorption chiller and all the flow from the compressor's exhaust is directed to the water-heater heat exchanger. Additional simulations show that such a scenario would result in twice the heat rate imparted to the water heater than that shown in figures 6 and 7.

The hot water exiting the water-heater at 70°C can then be supplied to a neighboring facility to either 1) heat air to about 35°C in the building's HVAC system to meet heating load requirements, in contrast to the cooling mode, or 2) used directly to supplement the building's hot water requirements. Table 5 shows the recovery COP obtained for both Systems I and II operating in heating mode.

Table 5. Comparison of Heating Produced From a 1 MW Data Center Using Systems I and II

Parameter	System I	System II
$\dot{Q}_{water\text{-}heater}$ (kW)	1400	1250
COP_{rec}	3.37	4.64

Table 6 compares the performance of Systems I and II in meeting the electric heating requirement of the same medium sized office building considered in the cooling mode section across the same four locations. From the results of Table 6, Location 3 (i.e. San Francisco) derives the most benefit from supplemental heating due to the high electricity cost.

Table 6. Comparison of Heating Load Requirement of a Typical Medium-Sized Office Building in Various Locations

Parameter	Loc. 1	Loc. 2	Loc. 3	Loc. 4
Heating system sized capacity (kW)	33.9	37.1	39.5	104.3
System I buildings served	41.2	37.7	35.4	13.4
System II buildings served	37.0	33.8	31.7	12.0
Annual heating electricity usage per building (kWh) * 10^3	4.10	39.9	81.7	89.0
Average annual electricity rate ($/kWh)	0.086	0.097	0.146	0.037
System I annual savings ($ 000)	14.5	146	422	44.1
System II annual savings ($ 000)	13.0	131	379	39.6

It should be pointed out that for System I, the compressor boosts the refrigerant to a pressure and temperature of 110 psig and 108.1°C. However, liquid refrigerant remains in the compressor at those discharge conditions which can lead to mechanical damage in the long run.

A sensitivity analysis on System I indicates that the refrigerant needs to boosted to at least 140 psig for no liquid to exist in the compressor. Figure 8 shows the compressor discharge and generator condensing temperatures versus the compressor discharge pressure, indicating the pressure values at which liquid exists at the compressor outlet. Figure 8 also indicates the minimum condensing temperature for the AR generator to be 92°C for ensuring no liquid entrainment in the compressor; this is higher than that for System II which can safely function at lower condensing temperatures of around 82°C which leads to lower compressor discharge temperatures and hence lesser power consumption. This again proves that System II outperforms System I.

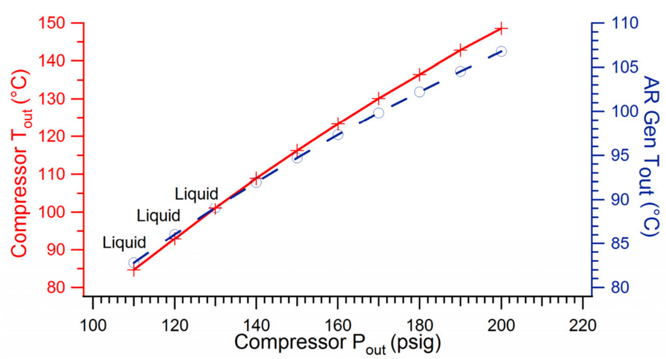

Fig. 8 Variation of Compressor Discharge Temperature and Generator Condensing Temperature with Compressor Discharge Pressure

Choice of Refrigerants for the Boost & Recovery System

Refrigerant 245fa was chosen as the working fluid for the BRS simulations because of its moderate vapor pressure and availability. As outlined in the introduction, R-245fa and similar refrigerants present an excellent choice for electronics cooling and boosting the resultant waste heat due to their favorable thermal properties. Table 7 presents a comparison of the commercially used refrigerant R134a and the proposed refrigerant R245fa with their environmentally friendly counterparts.

Table 7. Refrigerant Comparison for the Boost and Recovery System

Parameter	R-134a	R-245fa	R-1224yd	R-1234yf	R-1233zd
NBP (°C)	-15.0	14.9	14.0	-29.0	18.0
MW (g/mol)	101.5	134.0	148.5	114.0	130.5
ODP	0	0	0	0	0
GWP	1370	1030	1	4	1
Flamm. Limits	None	None	None	6.2–12.3%	None
Availability	Yes	Yes	Limited	Yes	Yes
Relative Cost	Low	Low	High	Medium	Medium

From Table 7, refrigerants 1224yd and 1233zd establish themselves as the most suitable candidates for the boost and recovery system based on their excellent physical properties such as relatively high boiling points and their low global warming potential. The higher boiling point allows these refrigerants to cycle through the two-phase cold plates at relatively low pressures (consider 19 psig from Figures 6 and 7, versus about 125 psig for R-134a at a saturation temperature of about 38°C through the cold plates). This relatively low pressure allows lower compressor and exchanger design pressures and hence reduced capital costs.

Finally, R-1224yd and R-1233zd have significantly lower ozone depletion and global warming potentials than R-134a and hence meet the more stringent environmental protocols enforced in various nations around the world, particularly in Europe. Further, among these two refrigerants, R-1233zd is more readily available and has a lower cost due to its lower fluorine content. Hence it is the preferred refrigerant for the boost and recovery system.

Economic Comparison between System I and System II

Tables 8 and 9 present the technical specifications/costing source and CAPEX comparison respectively between the two systems under consideration.

Table 8. Equipment Technical Specifications and Costing Source for Systems I and II

Component	Specifications/Costing Source
V-L separator	Multi-feed flash drum/ASPEN Plus (System II only)
Regenerator	Counter current Shell & Tube/ASPEN Plus
Water-Ref HX	Counter current Shell & Tube/ASPEN Plus (for System I only)
Water-Heater	Counter current Shell & Tube/ASPEN Plus
Compressor	Single-Stage Elliott Centrifugal Compressor; Frame Size: 10M; 650 hp; 20639 RPM; 78% eff. /Elliott Turbo Group
Pump	Sundyne KSMK radially split, Single or Two-stage centrifugal pump; 23 hp required; max. flow 13209 gpm/Sundyne (for System I only)
Cold Plates	Water based (System I) and 2-phase refrigerant based (System II)/QuantaCool
Abs. Chiller	High pressure (116 psig) steam-driven 330RT Water-Cooled/ [20]

Table 9. Capital Expenditure Comparison between Systems I and II (In Thousands)

Component	System I ($)	System II ($)
V-L separator	N/A	22
Regenerator	33	19
Water-Ref HX	169	N/A
Water-Heater HX	39	31
Compressor	1,000	750
Pump	14	N/A
Cold Plates	Similar Cost	Similar Cost
Absorption Chiller	393	393
Total Cost	**1,950**	**1,430**

Table 9 shows that System II neither requires a water-to-refrigerant chiller nor a pump. The elimination of these two components reduces the initial cost required for constructing System II. A comparison of the total cost for both systems shows that System II requires approximately 27% lower CAPEX than that of System I, making System II the economically preferred choice in addition to the higher thermodynamic efficiency.

In cooling mode alone, System II consumes 270 kW of compressor power while removing 1030 kW (292RT) of heat load at the chiller's evaporator. A 300RT DX-based air-cooled mechanical chiller has a typical maximum EER of 10.3 [21]. This translates into a COP value of 3.02 for the stand-alone chiller. This chiller would hence consume 341 kW of electric power to produce a 292 RT. For comparison, System I consumes 415 kW in cooling mode. Thus,

$$\dot{W}_{comp,Sys\ I} < \dot{W}_{chiller} < \dot{W}_{comp,Sys\ II} \quad (4)$$

Thus, equation (4) shows that System I is not an attractive option compared to a stand-alone chiller, while the power savings realized from implementing System II versus a stand-alone chiller comes out to be:

$$\Delta\dot{W}_{savings} = \dot{W}_{chiller} - \dot{W}_{comp,Sys\ II} = 70\ \text{kW} \quad (5)$$

System II consumes about 21 % less power than a comparable stand-alone chiller.

Summary & Conclusions

This paper introduces two competing systems for the simultaneous thermal management and upgrading of server waste heat from data centers. System I focuses on warm water-based server cooling while System II introduces the concept of a coupled data center two-phase cooling system integrated with direct vapor recompression (2PCVR). The boosted heat is then used to supplement the heating or cooling requirements of a co-located midsize office building. The two systems are analyzed from a thermo-economic perspective to determine their relative merit.

Considering combined cooling and recovery, System II is better than System I due to its higher recovery COP in both heating and cooling mode and its relative simplicity. That simplicity translates into lower CAPEX and OPEX for this system. System II is about 30% more efficient in recovering heat than System I, eliminates any water near the servers and other electronics, operates at lower pressures making it safer and less prone to leakages and is a single-loop system. Additionally, the system demonstrates a swift response to dynamic server loads and makes effective use of environmentally friendly refrigerants. It eliminates the use of a pump and intermediate chiller while making the remaining equipment sizes smaller, hence saving about 33% on CAPEX. Lesser equipment also leads to lower OPEX and the single-loop design with lower pressures makes maintenance easier.

Lastly, for System II in cooling mode (when compared to a standalone chiller) shows that such a system is well suited for areas of high urban density with high electricity costs such as parts of the US, Western Europe and East Asia. Hence this study recommends the latter system for data center waste heat boosting and recovery applications.

Future Work

Further work entails a detailed cradle-to-grave exergy analysis for Systems I and II to establish which system is better from a second-law perspective.

Acknowledgments

This material is based upon work supported by the NSF IUCRC Award No. IIP-1738782. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] Y. Yingxiao, Y. Zhaocheng, M. Jiayi and L. Shiyu, "Design and simulation of exhaust gas waste heat recovery system of gasoline engine based on Stirling cycle," 2013 Intl. Conference on Materials for Renewable Energy and Environment, Aug 2014, China

[2] R. Thakar, S. Bhosle and L. Subhash, "Design of heat exchanger for waste heat recovery from exhaust gas of diesel engine," Procedia Manufacturing, vol. 20, pp. 372-376, Feb 2018

[3] D. Bella and A. Francis, "Gas turbine engine exhaust waste heat recovery navy shipboard module development," Supercritical CO_2 Power Cycle Symposium, May 2011, Boulder Colorado

[4] S. Sondur, K. Gross and M. Li, "Data center cooling system integrated with low-temperature desalination and intelligent energy-aware control," Easy Chair Reprint no. 637, Nov 2018

[5] M. Wahlroos, M. Parssinen, S. Rinne, S. Syri and J. Manner, "Future views on waste heat utilization – case of data centers in northern Europe," Renewable and Sustainable Energy Reviews, vol. 82, pp. 1749-1764, 2018

[6] B. Boissevain, "Waste heat utilization in an anaerobic digestion system," MS Thesis, Utah State University, Logan, Utah, 2012

[7] https://www.bbc.co.uk/news/technology-45432723

[8] Google Liquid Cooled AI Chips: https://www.datacenterknowledge .com/google-alphabet/google-brings-liquid-cooling-data-centers-cool-latest-ai-chips

[9] K. Ebrahimi, G. F. Jones and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery

opportunities," Renewable and Sustainable Energy Reviews, vol. 31, pp. 622-638, 2014

[10] K. Ebrahimi, G. F. Jones and A. S. Fleischer, "Thermo-economic analysis of steady state waste heat recovery in data centers using absorption refrigeration," Applied Energy, vol. 139, pp. 384-397, 2015

[11] J. B. Marcinichen, J. A. Olivier and J. R. Thome, "On-chip two-phase cooling of datacenters: cooling systems and energy recovery evaluation," Applied Thermal Engineering, vol. 41, pp. 36-51, 2012

[12] J. B. Marcinichen, J. R. Thome and B. Michel, "Cooling of microprocessors with microevaporation: a novel two-phase cooling cycle," International Journal of Refrigeration, vol. 33, no. 7, pp. 1264-1276, 2010

[13] Y. J. Kim, Y. K. Joshi and A. G. Federov, "An absorption based miniature heat pump system for electronics cooling," Intl. Journal of Refrigeration, vol. 31, pp. 23-33, 2008

[14] Z. Yuan and K. E. Herold, "Thermodynamic properties of aqueous lithium bromide using a multiproperty free energy correlation," International Journal of Heating, Ventilating, Air-Conditioning & Refrigeration, vol. 11, no. 3, pp. 377-393, 2005

[15] A. Haywood, J. Sherbeck, P. Phelan, G. Varsamoplous and S. K. S. Gupta, "Thermodynamic feasibility of harvesting data center waste heat to drive an absorption chiller," Energy Conversion & Management, vol. 58, pp. 26-34, 2012

[16] T. Brunschwiler, B. Smith, E. Ruetsche and B. Michel, "Toward zero-emission data centers through direct reuse of thermal energy," IBM Journal of Research & Development, vol. 53, no. 11, pp. 1-3, 2009

[17] K. Brown, W. Torell and V. Avelar, "Choosing the optimal data center power density," White Paper no. 156, rev. 0, Schneider Electric

[18] D. Cutler, J. Winkler, N. Kruis, C. Christensen and M. Brandemuehl, "Improved Modeling of Residential Air Conditioners and Heat Pumps for Energy Calculations", Technical Report, NREL/TP-5500-56354, National Renewable Energy Lab, January 2013

[19] https://www.energy.gov/eere/buildings/commercial-reference-buildings

[20] Absorption Chillers: https://www.energy.gov/sites/prod/files/2017/06/f35/CHP-Absorption%20Chiller-compliant.pdf

[21] Carrier AQUASNAP 30RB: https://www.carrier.com/commercial/en/us/products/chillers-components/air-cooled-chillers/30rb/