Stanford's "4th Generation" District Energy System

Combined Heat & Cooling Opens Path to Sustainability

Joseph C. Stagner, P.E., Executive Director,
Sustainability & Energy Management, Stanford University

The Stanford Central Energy Facility uses renewable electricity as a primary fuel source to heat and cool the university.

Stanford University transformed the district power, heating, and cooling system serving its buildings from natural gas to sustainable, renewable electricity, greatly increasing system efficiency and reducing cost in the process.

In 1987 Stanford University took a giant step forward in efficiency, economics, and environmental stewardship by installing a 50 megawatt natural gas fired combined heat and power (aka CHP or cogeneration) plant to provide electricity, steam, and chilled water for the campus. Three decades later the Cardinal Cogeneration plant has been retired and the Stanford Energy System Innovations (SESI) project has taken the university into the 21st century with an even more efficient system that immediately reduces campus greenhouse gas emissions by 68% and water use by 18%, and is expected to save hundreds of millions of dollars for the university over the next three decades compared to cogeneration. Shifting from gas cogeneration to grid electricity may be opposite of current trends, but heat recovery and renewable power are the path to sustainability for Stanford.

Hot and cold water thermal energy storage and advanced 'model predictive control' assure high efficiency

HEAT RECOVERY

The cornerstone of SESI is recovering waste heat from the campus district chilled water system to meet building heating and hot water needs. This opportunity was discovered in 2008 upon review of cogeneration energy production data as the university embarked upon a renewed effort to improve the sustainability of its energy systems.

Considering that a chilled water system is really just a system for collecting and disposing of unwanted heat, the opportunity to capture and use that waste heat to serve the campus rather than discard it to the atmosphere with evaporative cooling towers was investigated to see if centralized heat pumping could play a part in Stanford's district energy system. With cooling occurring primarily in summer and heating in winter the opportunity for heat recovery was assumed to be modest until Stanford engineers compared the simultaneous delivery of heating and cooling from the cogeneration plant over all hours of the year. This revealed a 71% overlap of heating and cooling produced by the cogeneration facility and illuminated a potential new opportunity for efficiency and cost savings.

Figure 1 shows the typical daily heating and cooling overlap at Stanford for the different seasons throughout the year. In summer, warmer days yield to cooler nights at this geographical location leading to a surplus of available heat part of the day and a deficit at other times. In winter, there is a constant deficit of available heat, and spring and fall are transition periods between the two extremes. Nevertheless, at all times of the year there is some level of chilled water need for conditioning air in the buildings such as for humidity control or for cooling of computers or high energy laboratory equipment, presenting an opportunity for year-round heat recovery.

Viewed on an annual basis the thermal overlap and corresponding opportunity for heat recovery totals 71%, with 88% of campus heating and hot water needs met by recovering 54% of the waste heat from the chilled water system as shown in Figure 2. As the figure shows, in summer 100% of the heating and hot water needs of the campus can be met through recovered waste heat, while part of the heating needs in winter can be met with recovered waste heat.

Figure 1- Daily Heating & Cooling Overlap (Stanford University)

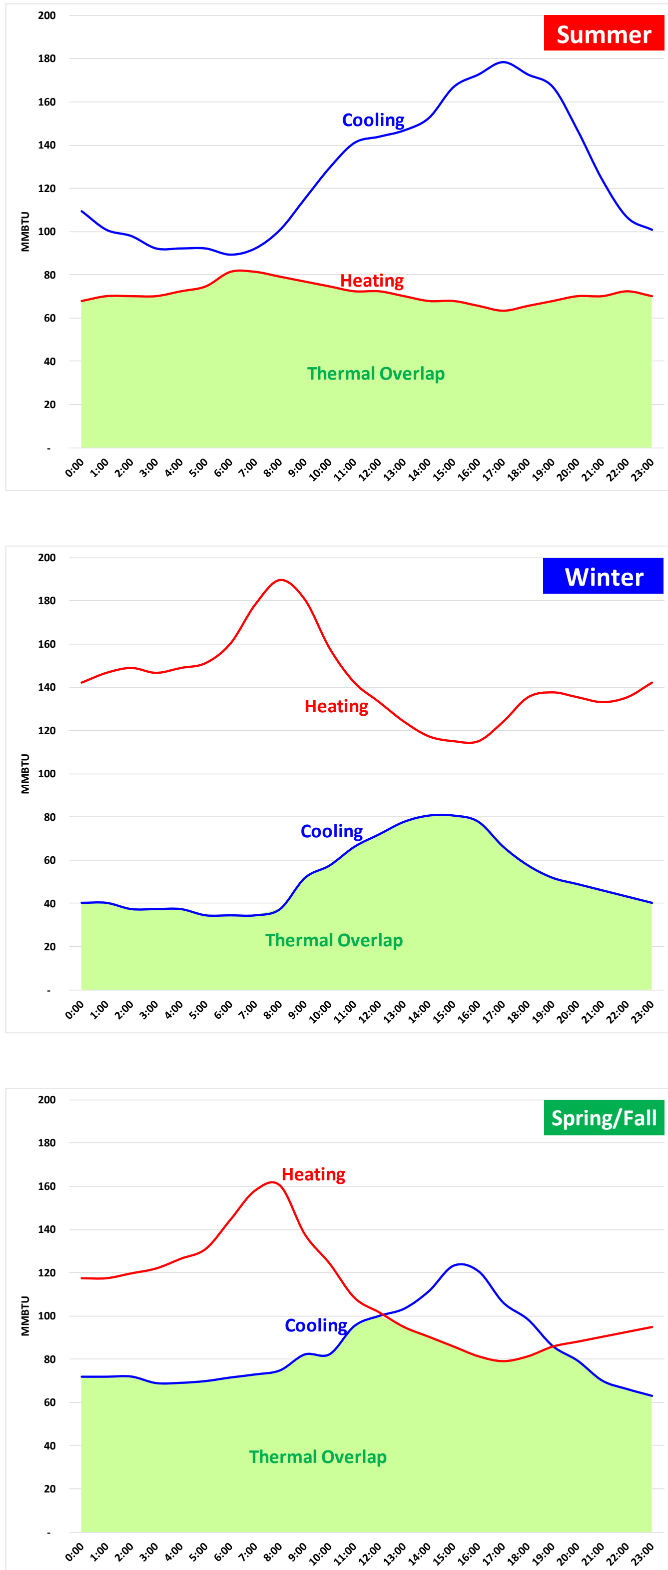

Note that the annual heat recovery potential shown is for the load profiles at Stanford University and with the size of the hot and cold thermal energy storage provided. In Stanford's system the provision of thermal energy storage allows 20% more heat recovery over the course of a year than could otherwise be achieved from only direct real time heat recovery via heap pump without thermal energy storage. Additional thermal energy storage could increase the total annual heat recovery even further, however the benefits of higher amounts of thermal storage diminish exponentially beyond optimal amounts. Conversely, heating needs not met through heat recovery might have to be met by burning fossil fuels in conventional hot water generators at greater financial and environmental expense unless supplemental heat recovery resources are also incorporated into the system such as ground or water source heat exchange. All of these things must be considered in the overall system design, but it is very clear that an optimal amount of thermal energy storage provides huge benefits including greatly increasing heat recovery amounts; reducing required capacity of heat pumps and conventional chillers; and shifting of electricity consumption by the system to more optimal hours in a time-of-use electricity pricing regime.

Figure 2- Annual Heat Recovery Potential (Stanford University)

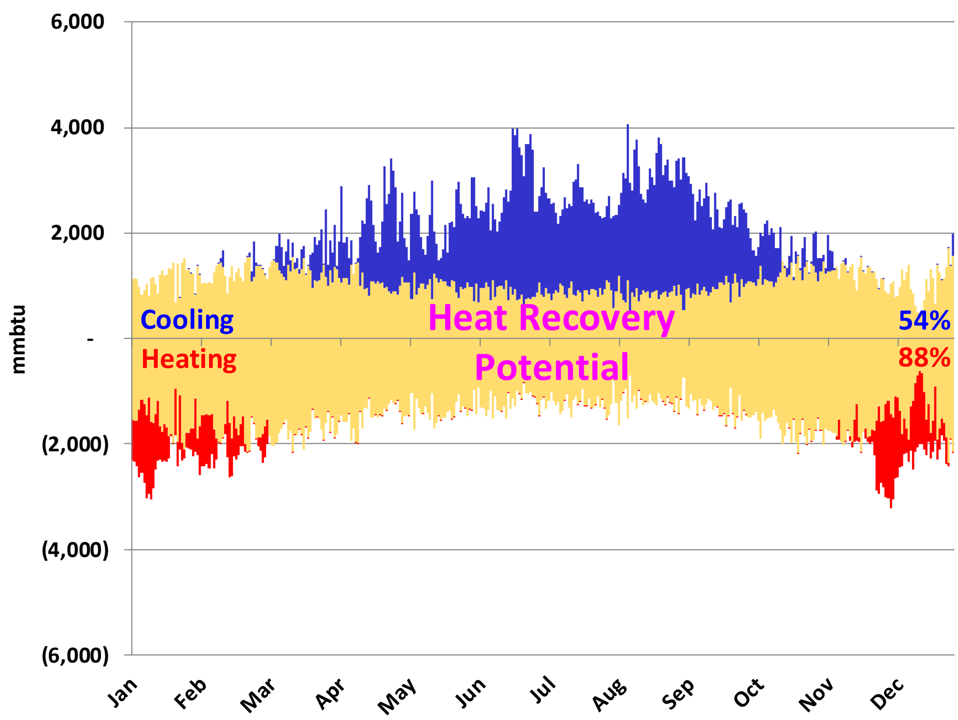

WHY SO MUCH THERMAL OVERLAP?

The astounding amount of annual heating and cooling overlap revealed through analysis of the energy system production data was met with skepticism and had to be verified with extra diligence. Since hourly heating and cooling production data from the campus cogeneration plant (which supplied all the energy to the campus district energy system) wasn't routinely measured and recorded, detailed efforts were required to assemble this information. It took about one half full time engineering assistant about 6 months to assemble it over the summer of 2009. There was concern that errors could have occurred in putting this data together this way using equipment run time and power consumption logs and other such techniques. In order to verify the dataset on which the heat recovery potential was based, several audits were conducted, including reconciling monthly totals from the hourly load tables to the cogeneration monthly production reports and corresponding invoices to the university for energy supply. Estimates of total campus loads were also prepared by extrapolating hourly energy consumption logs for the limited number of building energy management systems to the total size and makeup of the university building fleet served by the district energy system and allowing for nominal lien loss estimates between the cogeneration plant and building meters. All these audits affirmed the likely accuracy of the detailed hourly heating and cooling load data produced 'long hand' from cogeneration plant equipment logs. However even if accurate, questions arose as to whether the year in question was a representative weather and thermal load sample. To address this question hourly load data was subsequently assembled for several additional years and further verified the basic hourly load data set, with an understanding that this could vary with weather from year to year but also noting that even when it did that did not materially change the basic conclusion that the heat recovery potential for the Stanford campus was always in the upper 60% to upper 70% range and therefore could be the cornerstone to a new type of district energy system centered on heat recovery.

It became clear that the hourly heat overlap data was real, and the next question became "why was this happening"? Examination of the reasons for this thermal overlap revealed that while domestic hot water needs comprised a small amount of year-round heat load it was large building 'reheat' process loads for the high energy intensity research buildings on the Stanford campus that were the reason for this amount of coincident heating and cooling. Building reheat occurs in modern buildings when incoming air has to first be cooled to remove humidity, then reheated up to the final temperature desired in the building spaces. While building mechanical HVAC engineers have long known and understood this and often incorporate 'heat wheels' or economizers to take advantage of it at a single building scale, those that design and operate central energy systems to serve groups of buildings often may be unaware of this. Or they do not realize the scale of heat recovery potential this may represent because of over simplistic models engrained in us by how we operate our homes such as "heating occurs in winter and cooling in summer" so likely little to no overlap. Or because they don't know what if anything could be done in a central district energy system to take advantage of thermal overlap or if doing so would be economical. Still other district energy system operators, when presented this information and challenged to examine the heat recovery potential in their systems, assumed that there could be little or no overlap in their systems if they were in a climate of far greater heating and cooling extremes than Stanford's mild coastal climate. However, the results revealed otherwise for those that dug deeper.

To further test the phenomenon, in 2010 and 2011 hourly heating and cooling load data was obtained by engineers operating several other university district energy systems in very different climates and the amount of thermal overlap was examined to determine if similarities existed that could confirm what was happening at Stanford. It was very difficult to obtain such data because the universities examined, like Stanford, were not actively monitoring the hourly heating and cooling production from their district energy systems. Nevertheless, data was obtained from another major California university located in a much different climate than Stanford, as well as from two Midwest universities, a small liberal arts college and a major research university in New England. This small sampling showed a 55% thermal overlap at the other California university while the two Midwest and one New England universities had about 30% (liberal arts college) and 50% (large research universities) thermal overlaps in a much different climate than either of the California universities.

Through discussions with building HVAC engineers it was further determined that the building reheat process would always occur and that even aggressive building efficiency improvement programs would likely just reduce the overall magnitude of heating and cooling but probably not change the relative thermal overlap in the entire district energy system.

Given these checks it was determined that the phenomenon of significant concurrent heating and cooling in district energy systems was a widespread occurrence resulting from the needs for the building reheat process as well as domestic hot water production; that this would likely continue into the future indefinitely regardless of overall building efficiency; and that the data was therefore reliable for Stanford's future energy planning.

INITIAL FEASIBILITY

With such a large opportunity for heat recovery identified the next steps were to determine if heat pumps were available at this scale, and would it be economical to base a district energy system around heat recovery.

Stanford engineers spoke with chiller and refrigeration equipment manufacturers and found that heat recovery chillers (HRCs), i.e. water to water heat pumps, could be made in sizes over 1,000 tons from which a central energy system could be based. However, a key variable in HRC efficiency would be the temperature of hot water required. Leaving condenser water temperatures of less than 100F for typical chillers would not be sufficient for a district hot water system, and the higher that temperature needed to be, the more electricity that would be required by the heat pumps to produce it. The good news was that whatever that electricity usage had to be, it would be converted to machine heat mostly transferred to the leaving condenser water and therefore add to the useful hot water produced for a heat recovery system, rather than adding to the waste heat burden on evaporative cooling towers as with the existing conventional chilling process at the cogeneration plant.

To determine the likely range of temperatures required for a district hot water system at Stanford, engineers examined current building circulating hot water (hydronic) system operating temperatures in winter (when heating loads are highest) and determined that most could operate at 165F, meaning that a district hot water supply temperature of 170F would be sufficient when accounting for loss across the heat exchanger interface between the district and building circulating hot water systems. Unlike the previous district steam to building hot water system interface which required a steam-to-hot water heat exchanger in each building, it was noted that the university could forego heat exchangers between district and building hot water circulating systems in order to reduce cost and increase efficiency somewhat, and indeed that is what Ball State University chose to do in implementing a district hot water system at their campus in Indiana. However, concern about mismatched building versus district energy piping materials and resulting potentially destructive water chemistry, coupled with concern over risks of building hot water leaks compromising the entire district hot water system, compelled Stanford to include district hot water to building hot water heat exchangers in each building in the final system design.

Furthermore, while a handful of buildings were being operated with hydronic temperatures as high as 190F, a review of building HVAC system designs indicated that those temperatures could be reduced to 165F either through simple operational changes or with modest modifications of equipment. An actual winter test program was implemented and confirmed those assumptions. A similar review of summer conditions indicated that all buildings could operate with a district hot water supply temperature of 160F or perhaps as low as 150F. A return hot water temperature of 130F to 140F was estimated given these supply temperatures for winter and summer.

With a tentative design criteria of 150F to 170F leaving condenser water temperature for the hot water supply system, 130F to 140F return hot water, and nominal chilled water supply and return temperatures of 42F and 56F that already existed with the campus chilled water system respectively, preliminary estimates from manufacturers were an efficiency of about 1.5KW/ton for the HRCs. This was used to examine the basic economic feasibility of heat recovery.

With each ton of cooling processed (12,000 btu) the HRC would produce about 17,100 btu of hot water with the added heat from conversion of the 1.5KW of electricity input (~5,100 btu). At an electricity cost of $.10/KWH and natural gas cost of $5.50 per million btu the cost to process one ton-hour of cooling and produce 17,100 btu of heat would be $.15 via a combined heating and cooling process. The cost to produce the same amount of heating and cooling with new high efficiency chillers at .5KW/ton and gas boilers at 85% HHV (higher heating value) efficiency would be $.05 for electricity and $.11 for natural gas for a total of $.16. While it was recognized that economics would have to be tested over a broader range of potential long term future gas and electricity prices, and with capital and operating & maintenance costs included, heat recovery appeared more economical than using conventional boilers and chillers for heating and cooling on a raw energy basis, but would it be more efficient and sustainable?

SHP v CHC: EFFICIENCY

To compare the efficiency and sustainability of a conventional separate heat & power (SHP) system employing gas heating and electric chilling powered by grid electricity to a combined heat & cooling (CHC) system powered primarily by grid electricity, a common fuel of natural gas was selected.

To assure objectivity in the analysis, all new equipment was assumed in each case, with on-site gas heating at 85% efficiency on an HHV basis; electric chiller + cooling tower average efficiency of .5KW/ton; and heat recovery chiller average efficiency of 1.5KW/ton. Data from the California Energy Commission showed that in 2011 the average efficiency of the existing California grid gas power plant fleet was 48% HHV, while new gas grid plants were averaging 51% and progression to 55% was expected. Given that all-new equipment was being modeled in all cases to assure objectivity, an average new gas grid power plant efficiency of 51% was assumed for the SHP and CHC options.

With these factors, the total amount of natural gas required to process one ton-hour of cooling and produce 17,100 btu of heat would be (.5 KWH * 3,413 btu/KWH ÷ .51 = 3,346 btu) + (17,100 btu ÷ .85 = 20,118 btu) = 23,464 btu of gas for the SHP option. For CHC the total gas required would be 1.5 KWH * 3,413 btu/KWH ÷ .51 = 10,038 btu. From a sustainability perspective CHC would use 57% less natural gas than SHP and produce 57% less greenhouse gas emissions. Significant water savings from curtailment of the use of evaporative cooling to reject waste heat would also result with the CHC option. The next step was to compare SHP and CHC to a modern gas cogeneration option in terms of efficiency and sustainability.

SHP v CHC v CHP: EFFICIENCY

To compare the efficiency of a new SHP or CHC system to a new Combined Heat & Power (CHP) system the full trigeneration (power, heat, and cooling) loads of the university would have to be modeled to assure a complete comparison. Stanford's trigeneration loads were close to a three-way balance when the modeling was performed and were forecast to remain that way for the long term, so for simplicity, an even three way balance was assumed for general initial comparisons. Natural gas was again chosen as the common fuel, with the factors outlined earlier for heat recovery chiller, gas heater, and electric chiller efficiencies. A range of cogeneration plant prime movers (heat and power) including gas turbines and internal combustion engines

Figure 3. Relative Efficiency of SHP, CHP, and CHC

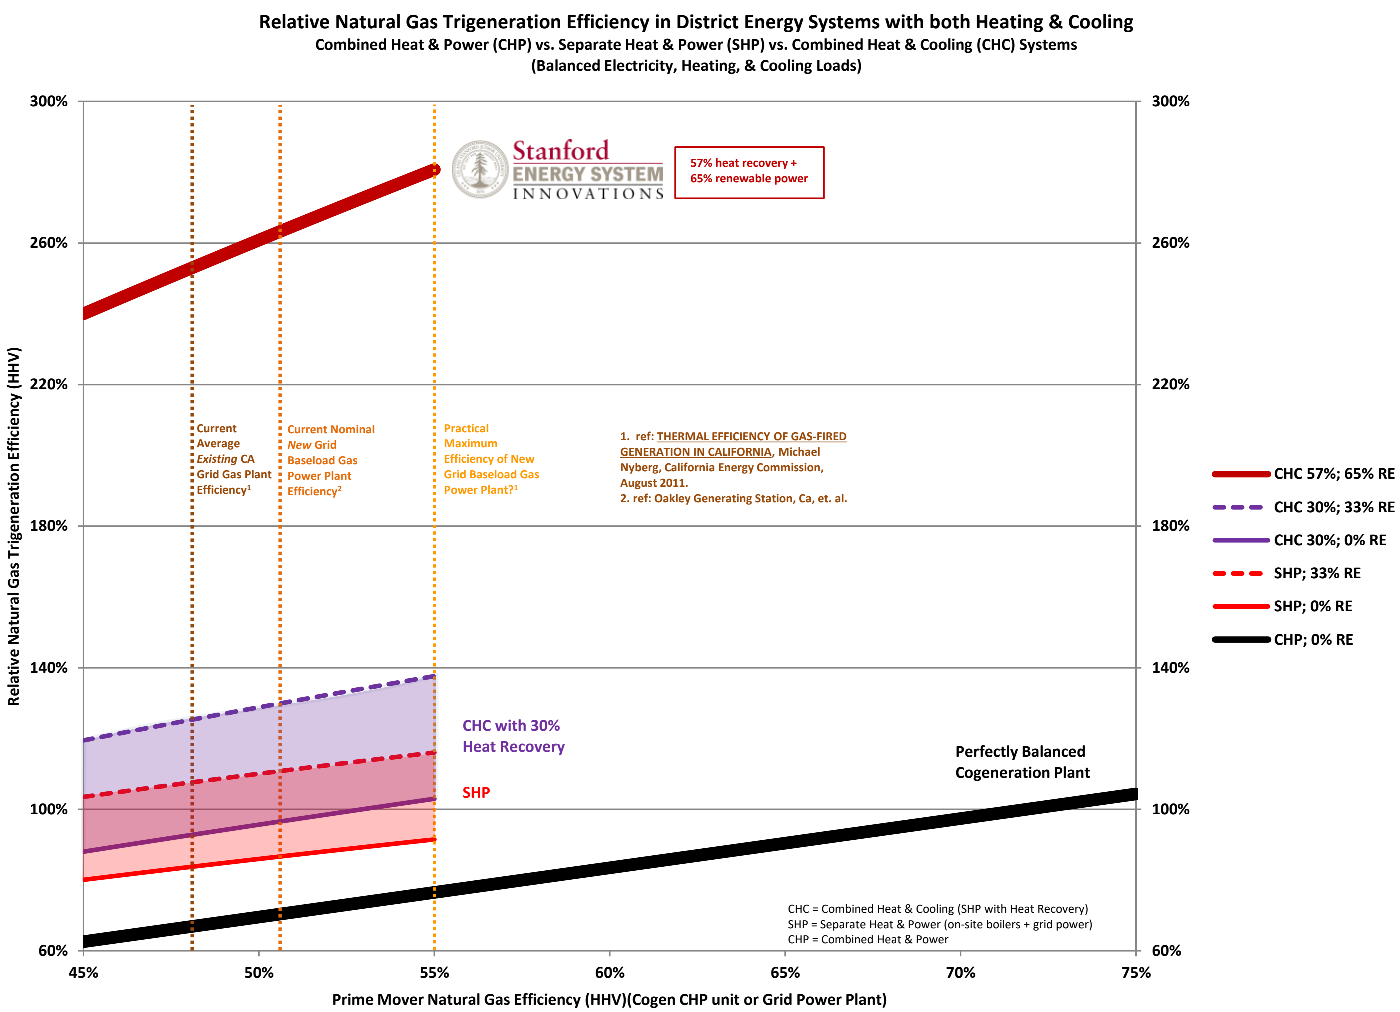

and efficiencies ranging from 45% to 75% HHV were modeled for the CHP option. For chilling in the CHP option an optimum arrangement using 100% electric chilling was used since this was more efficient than incorporating steam powered chillers into the scheme.

Figure 3 shows the relative natural gas efficiency of SHP, CHC, and CHP with balanced electricity, heating, and cooling loads with a range of 0% to 33% renewable electricity and heat recovery respectively in the SHP and CHC options. An upper bound of 33% for renewable electricity content and heat recovery was selected to reflect the California Renewable Portfolio Standard requirement then in effect of at least 33% renewables on the grid by 2020 and the typical minimum heat recovery potential in university district energy systems based on the small initial sampling set from four universities, respectively. Also shown on the chart is the CHC option Stanford ultimately selected, which incorporates 65% renewable electricity generation and 75% heat recovery. As with a Coefficient of Performance (COP) for thermal equipment that exceeds 1 the overall efficiency of a trigeneration system as portrayed in the analysis can exceed 100% since the energy required to reject heat via the cooling process is typically much less than the quantity of heat rejected. Presenting the relative efficiency of trigeneration systems in this way, using a common fuel such as natural gas, was determined to be a useful way to compare the overall efficiency and sustainability of the different energy system options Stanford was contemplating. This comparison indicated that with a small percentage of renewable electricity and/or heat recovery incorporated into the system an SHP or CHC based system would use less total natural gas and result in less greenhouse gas emissions than even the best new gas cogeneration plant suitable for Stanford’s energy loads.

HOT WATER v STEAM

It was determined that a central energy facility based around heat recovery was more efficient and sustainable than gas fired CHP, but that it would require converting the campus district heating system from steam to hot water. An examination of the existing steam distribution system at Stanford revealed an average 14% line loss across the entire system, while discussions with district energy hot water system operators in St. Paul, Minnesota and in northern Europe indicated line losses of 4% for low temperature hot water. It was noted that this efficiency gain was not included in the analysis represented by Figure 3 and that a conversion to hot water would increase overall system efficiency and sustainability even more for all three options. From these discussions, it was further determined that such a conversion, while perhaps the biggest challenge to making such a transformation of the district energy system, would also yield considerable operation and maintenance savings given the relative simplicity of a hot water distribution system compared to a steam distribution and condensate return system which incorporates condensate traps to remove water from the steam lines, pressure reducing fittings to manage different steam pressures at different points in the system, and thermal expansion devices to accommodate the wider range of temperature changes seen in steam systems at different operating conditions, amongst others. Furthermore, given the age of the existing steam system a conversion to hot water would also erase a great deal of future capital renewal costs anticipated to be required within the not too distant future. Reports of steam to hot water conversions by numerous city district heating systems in Europe for these same reasons further bolstered the case for making such a change regardless of which energy production option was chosen.

From an efficiency and sustainability standpoint it was clear that SHP and CHP with a small percentage of renewable electricity generation in the mix were superior to CHP powered by natural gas, and that the significant challenge of converting from steam to hot water was not a deal breaker and indeed may be a worthy project on a standalone basis. From a sustainability perspective, it was noted that the CHP option could in theory be augmented with carbon capture and sequestration at some point in the future, or that it could be powered by sustainable biogas, either of which would allow it to compete with SHP and CHC in sustainability terms. However, these possibilities were set aside for further consideration pending an economic comparison of the options. If a base CHP system using natural gas could compete with SHP and CHC economically then further studies on the potential application of CCS or sustainable biogas would also be performed.

SHP v CHC v CHP: ECONOMICS

To compare SHP, CHC, and CHP economics objectively, assumptions on long term electricity and natural gas prices would have to be well informed and established in advance of preparing detailed energy models for the different options. To develop estimates of both energy sources Stanford used long term forecasts from several consultants; the federal Energy Information Administration; faculty engaged in energy policy & economics; staff expertise; and a review of the past five decades of trends in energy pricing. Sensitivity ranges of plus or minus 20% the baseline gas and electricity prices forecasts were also applied for testing bounds in the subsequent energy and economic models that would be developed in detail for all SHP, CHC, and CHP options. At the time of formulating this criteria, the impact of gas fracking on forward looking gas prices was known and included. What was not included was an assumption that Stanford could achieve Direct Access to California electricity markets in order to manage the selection, cost, and renewable content of the electricity it would purchase off the grid to power the campus and a mostly electric central energy facility in the SHP and CHC schemes. Instead the assumption was that the university would continue to purchase electricity from the local investor owned utility (PG&E) indefinitely. As discussed later this assumption proved to be conservative because the university did achieve Direct Access and substantially reduced its cost of electricity while increasing its renewable content.

In addition to objective energy cost forecasts, an objective analysis of the efficiency of the three options was required so that bias for or against any of the technologies would not influence the outcome. To achieve this, two different staff engineering managers with consultant support developed detailed energy models for the respective gas CHP and electric SHP/CHC options with a goal to devise the most efficient system possible.

For the CHP option combined cycle gas turbine systems with either back pressure or condensing steam turbines were investigated, along with state of the art new reciprocating internal combustion systems. The CHP team further developed a hybrid cogeneration/heat recovery system that employed an optimally sized reciprocating internal combustion engine scheme and up to 20% of heat recovery from the chilled water system via heat pumps.

To develop models for the most optimal CHC system the engineering team soon realized that an energy modeling program would be needed to show how the combination of heat pumps, chillers, gas hot water generators, and both hot water and cold water thermal energy storage could best be sized and operated in the presence of a dynamic real-time electricity pricing market such as California's wholesale power markets operated by the California Independent System Operator (CAISO). Such software could not be found on the market so Stanford utility department engineers developed a program called the Central Energy Plant Optimization Model (CEPOM) to perform this modeling.

CEPOM

To determine if a particular set of equipment and hot and cold thermal storage tanks would meet the forecasted heating and cooling loads of the university, one approach would be to design the system around the maximum respective heating and cooling days expected each year, for if the system could meet those then it could meet lesser demands as well. However, to assure this to be the case the system would have to be designed to assure 100% thermal energy storage recharge within the 24 hour period of those maximum load days as anything less would open the system to shortfalls if successive similar high load days occurred and gradually depleted the tanks. This then could result in oversizing the equipment fleet or thermal storage tanks by eliminating the effective diversity factor of the varying load profiles of successive days, even high load days.

Such single day plant sizing models also do not provide an accurate estimate of the total gas and electricity energy required across all 8,760 hours of the year to operate any particular system for use in calculating annual energy costs, greenhouse gas emissions, and water use. Furthermore, the presence of thermal energy storage, changing hourly outside air temperatures (cooling tower efficiency) and varying market hourly electricity pricing allows efficiency and cost savings optimization opportunities through 'model predictive control' that also cannot be determined or implemented with single day plant sizing models.

Finally, in an SHP or CHC scheme the central energy facility represents a large portion of the overall district energy system electrical footprint on the grid, with CHC accounting for about one third of the total while building electricity use (air handlers, lighting, refrigeration, computers, and other miscellaneous plug loads) accounts for two thirds at Stanford. Knowing this it was determined that optimization of the district thermal energy system electrical loads and the essentially non-dispatchable background building electricity loads together would yield the most economic and efficient overall energy system for the university. Accordingly, variables for hourly market electricity prices as well as peak monthly transmission demand charges were added to the model.

CEPOM was developed from 2009 to 2010 by Stanford utility department engineers. When completed and successfully tested for multiple full year system operations simulations the program was then used via iteration to optimize the design of a system to meet Stanford's forecast energy loads in 2015 since this was the first year of anticipated operation due to the impending expiration of the cogeneration plant contract between the university and the cogen system owner in March of 2015.

Future energy load forecasts for the university from 2015 through 2050 were prepared based on the known capital construction plans of the university over the first ten years of the study period and based on average annual growth rates thereafter established by the university long range planning process. CEPOM was then used to model optimum CHC system design in 5 year increments from 2020 through 2050 to determine the most optimal long term configuration for the system to inform initial system site election and central energy facility design, including thermal storage tank sizing.

Figure 4 shows a sample CEPOM model output for a given scenario, providing the optimal hourly dispatch of energy system equipment at the central energy facility, the thermal loads being served, and the changing levels of the hot and cold thermal energy storage tanks over time. The sample is for a selected two-day period however the CEPOM model output provides the same data for all 8760 hours in a year and then sums up resulting annual system operating data.

In addition to CEPOM for modeling CHC and SHP options Stanford university utilities engineers also developed a separate model for cogeneration options, including combined cycle and internal combustion engine based schemes. Eventually these separate models could be merged into one single district energy system model that could be used for virtually any potential system configuration to show the relative efficiencies and costs of each option in one place.

Figure 4- Central Energy Plant Optimization Model (CEPOM)

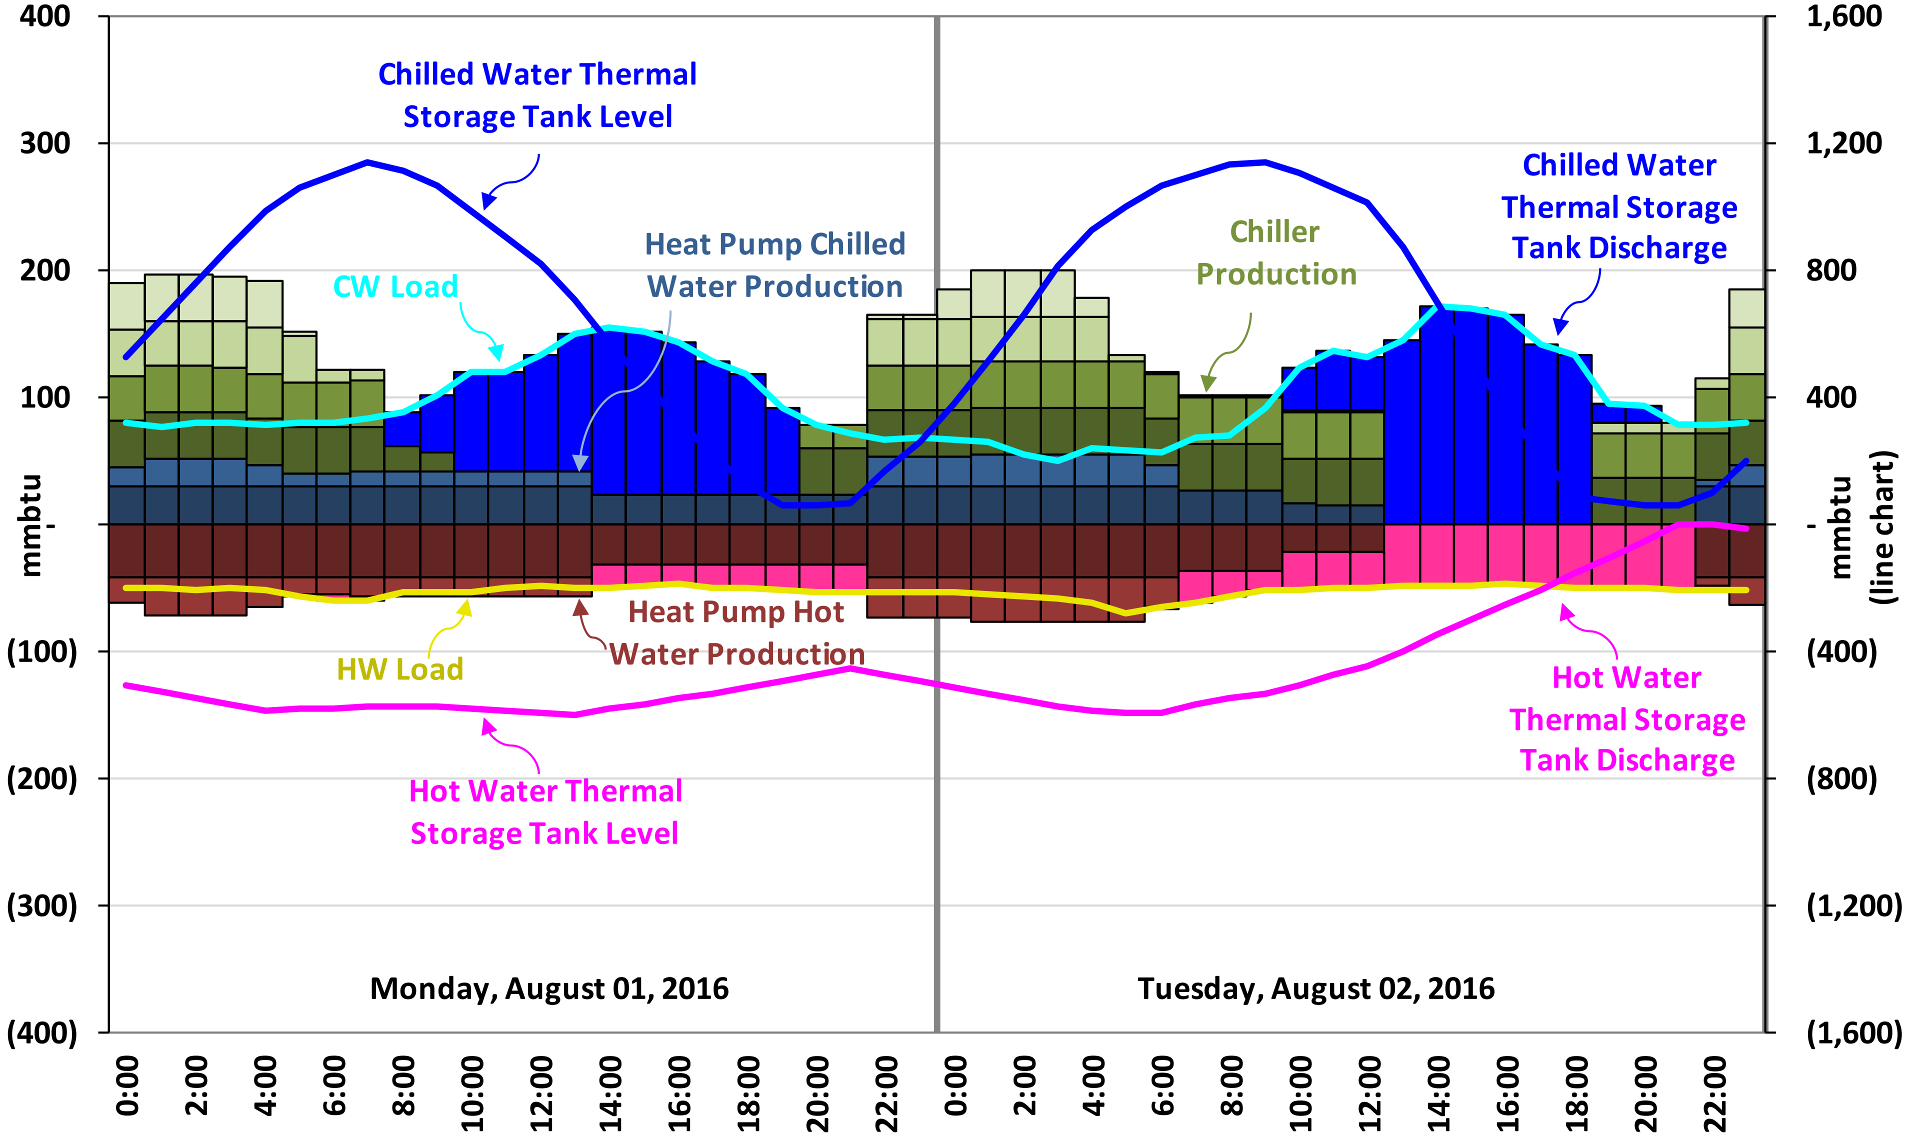

COST COMPARISONS

With CEPOM providing the energy model for CHC and SHC, and the cogeneration model providing the energy model for CHP options, in 2011 the university developed an overall Energy & Cost Model (ECM) for campus energy supply by which the economics of the different options could be compared. The ECM provided energy, O&M, and capital cost modeling for every year between 2015 and 2050 for all options as well as greenhouse gas emissions and water use. A planning horizon of 2050 was selected to represent the average anticipated life span of the energy system components being considered. Present Value Cost (PVC) analysis was performed to provide the relative economics for all the options in 2011 dollars. Inflation, escalation, and discount rates were developed by university finance and administration authorities with peer review by university business school faculty and external consultants. The overall ECM was also peer reviewed by faculty and external consultants.

Figure 5 shows the economic and sustainability comparisons of the options developed for providing energy to the university over the planning period. Options were grouped into those that retained steam as the method for delivering heat to the buildings versus installing a new hot water distribution system for this purpose; and then by those options that were primarily natural gas fueled versus those which would be primarily powered by imported grid electricity.

As the figure shows, all options except the Business as Usual (BAU) third party cogeneration arrangement then in place appeared similar in overall long term cost, given the reliability of long term market gas and electricity price projections. The BAU option was significantly higher in cost because it was the only option utilizing a third-party ownership and operating arrangement which would come with accompanying higher return on capital investment and ongoing annual overhead and profit costs than all the other options, which would be university owned and operated.

Figure 5. Comparison of Energy Supply Options

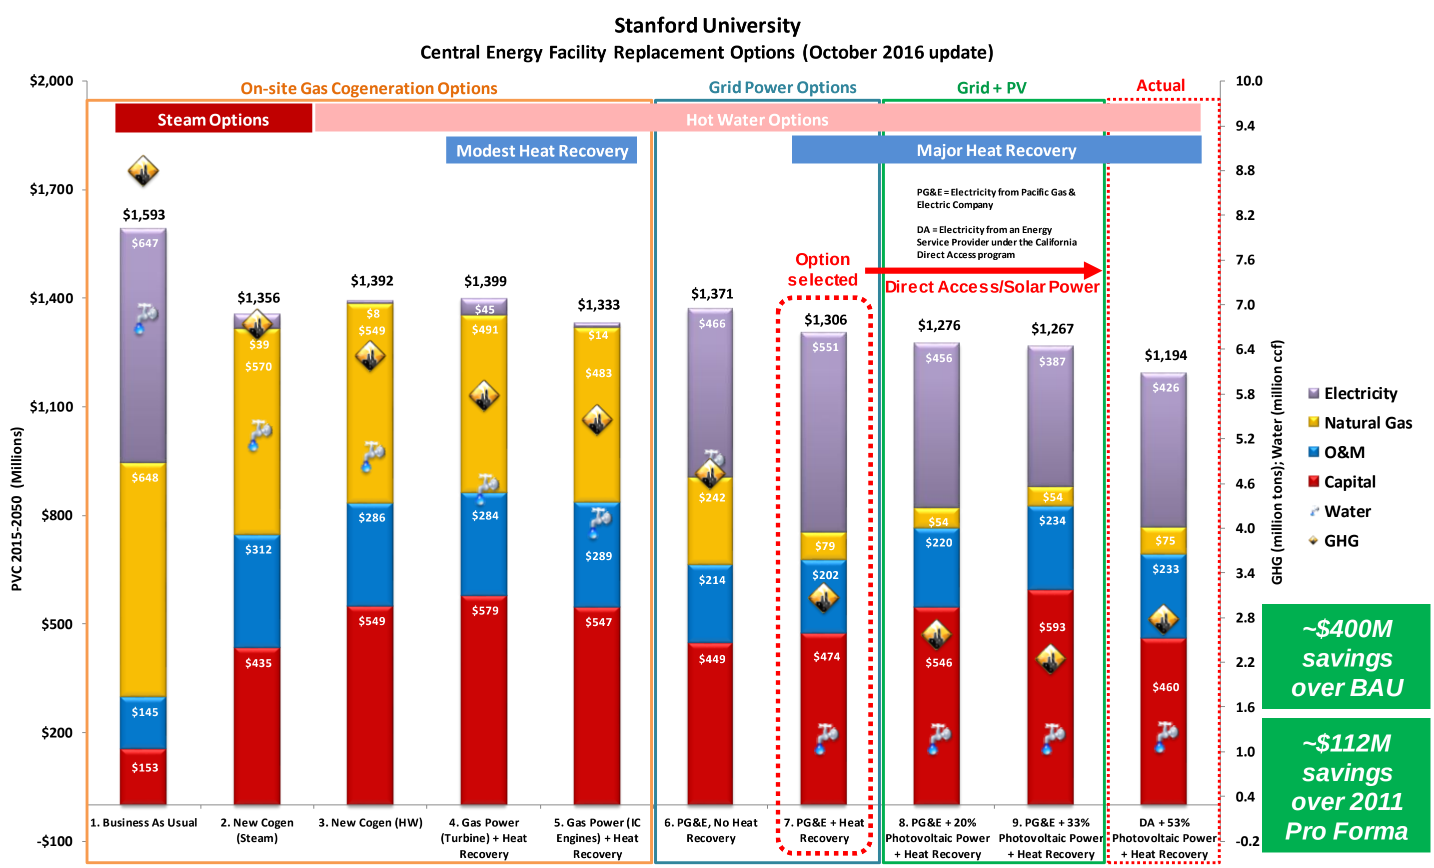

OTHER CONSIDERATIONS

In addition to the overall economics of the options other key considerations included sustainability, future adaptability, ease and time required for implementation, ease of university ownership and operation, system reliability and resiliency, and relative initial capital outlay of the options.

Sustainability considerations included initial and long-term greenhouse gas and other air pollutant emissions, pathways to reduce those emissions, and water consumption.

Future adaptability considerations addressed the ability of the scheme to incorporate technological advances in energy production over time.

Ease and time required for implementation considered the logistical impacts that the project might impose on an operating research and teaching university.

Ease of university ownership and operation addressed whether the system was technologically simple enough for a university to own and operate without the need for expensive third-party support.

System reliability and resiliency considered the relative reliability of each system to support research, health care and other critical operations and how resilient the system would be in the face of emergencies such as a loss of offsite power or gas deliveries, earthquakes, floods, or other challenges.

Relative initial capital outlay of the different options was important to consider insofar as how much university debt capacity would be required to implement the system versus other campus capital needs.

SYSTEM SELECTION & IMPLEMENTATION

Given similar estimated long-term costs Stanford selected the grid electricity powered CHC option because it provided the most long-term flexibility in energy choices; would result in immediate 50% greenhouse gas reductions and 18% water savings; would open a pathway to full energy system sustainability if 100% renewable electricity supply were used to power the system; it was simple enough for the university to own and operate; presented an average initial capital outlay of the options and did not challenge the university's debt capacity with respect to other capital needs at the time.

It was also considered the most reliable and resilient of the options in that it provided three fuel options for providing continuous heating supply to the campus (electricity, natural gas, or liquid fuel) as well as providing a day's chilled water supply via thermal energy storage and emergency generator powered discharge pumps. Although the university had previously had on-site electricity generation and several of the options could also provide that, past operational records of the previous cogeneration plant indicated that as a complex machine it caused as many electrical and thermal service outages to the campus as it had shielded the university from. Furthermore, with the steam distribution based options when steam pressure was lost via equipment outages it took several hours to rebuild system pressure in order to supply all locations with adequate heating service and with hot water that would not be the case. Finally, past operational records also showed that loss of offsite power supply via grid generation and transmission service was rare at about once every 5 to 10 years for several hours duration at most. Given this high grid reliability and the fact that all critical building and infrastructure electrical loads were already backed up by local emergency generators the loss of offsite power was not seen as a major risk and certainly not one that would dictate the need for an on-site electricity generating function such as gas cogeneration at the expense of efficiency, economics, and sustainability.

These findings were peer reviewed by numerous external consultants as well as Stanford faculty knowledgeable in these fields and even a six-member subcommittee of the Board of Trustees appointed to oversee the analysis and development of energy supply options for the university. Customer and community meetings revealed no opposition to the analysis or recommendations except one which requested a more detailed analysis of the relative exergy (second law of thermodynamics) of each of the options. That analysis was performed to the satisfaction of the energy faculty member that requested it and he then provided his support for the finding and recommendations.

The new design was named the Stanford Energy System Innovations (SESI) and was approved by University Trustees in December 2011. Detailed system design commenced in January 2012 using equipment selections provided by CEPOM. Construction began in October 2012 and was completed in March 2015.

In late 2012, through a limited state regulatory lottery process, Stanford gained Direct Access authority over purchase of its electricity supply instead of having to use local investor owned utility (Pacific Gas & Electric Company) power. The university then entered into long term Power Purchase Agreements (PPAs) to have new on- and off-campus solar power generation plants installed to provide 53% of the university's electricity. This increased SESI GHG reduction from 50% to 68% and reduced estimated long-term cost by $112 million. The stacked bars on the extreme right of Figure 5 show the economic and sustainability improvements over the approved 2011 SESI pro forma resulting from the university gaining Direct Access.

HOW THE SYSTEM WORKS

SESI included construction of a new Combined Heating & Cooling Central Energy Facility focusing on heat recovery; a new high voltage substation; installation of 22 miles of new hot water distribution piping from the new CEF to campus buildings; and conversion of 155 building primary heat exchangers from steam to hot water including replacing steam meters with hot water meters for energy metering and billing to each facility. Each building was equipped with a new heat exchanger interface between the district hot water supply loop and the building hydronic system in order to minimize the risks of system leaks and incompatibility between building hydronic system components of various ages and materials and the new hot water supply loop and central energy facility equipment.

The new high voltage substation was previously planned as the old substation had outlived its capacity and useful life. Its configuration was the same for all options and its cost was included in the present value cost calculations for all options as well.

Figure 6 shows the Central Energy Facility and the five major components of heat recovery chillers; conventional chillers & cooling towers; gas hot water generators; hot water thermal storage; and cold water thermal energy storage. The relative percentages of waste heat recovery and heat supplied from heat recovery as shown in Figure 6 are different from those of the original 2011 pro forma shown in Figure 2 as

Figure 6. Stanford SESI Central Energy Facility

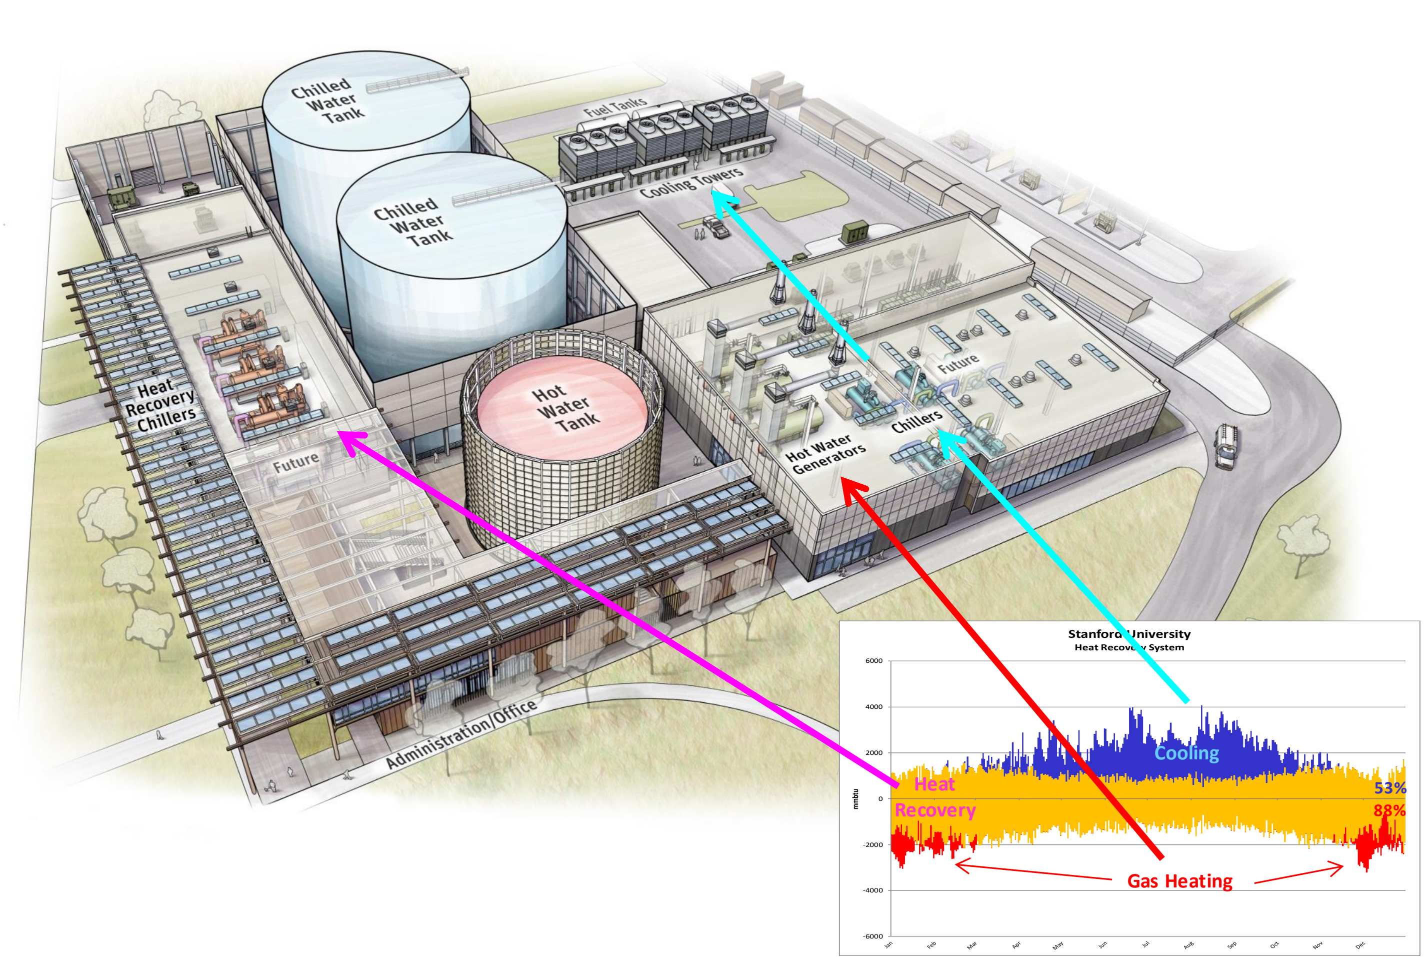

they are for different years and based on different campus thermal load forecasts which change over time.

ACTUAL PERFORMANCE SINCE 2015 START UP

An independent audit of first year SESI performance was conducted by an external consulting firm and the results were shared with Stanford leadership and its Board of Trustees in October 2016. The audit revealed that SESI was performing slightly more efficient and about 10% lower cost than the 2011 pro forma approved by the Board. The slightly higher efficiency achieved was attributed to conservative equipment performance specifications provided by manufacturers in order to assure purchase contract compliance, while most of the unexpected cost savings was due to the University gaining Direct Access to state electricity markets and lowering its purchased electricity cost. As of this report SESI continues to perform as expected and is achieving the economic and sustainability results predicted.

LESSONS LEARNED

As a state of the art new district energy system unforeseen challenges were anticipated with SESI. There have been no serious system problems leading to large unexpected costs or system performance issues, however several lessons learned have occurred over the first two and a half years of operation, including:

- The large Heat Recovery Chillers (heat pumps) require a longer start up cycle to achieve specified hot and cold water temperature output than gas hot water generators and regular chillers respectively. This results in partially intemperate (10F to 15F above or below nominal) slugs of hot and cold water being injected into the respective district hot and cold water distribution loops to campus buildings for about 15 minutes after start up. Fortunately, campus building HVAC systems were able to compensate for these temporary short slugs of water outside nominal temperature specifications and degradation of building heating and cooling services did not result. Nevertheless, some minor piping and software changes are being implemented at the CEF to partially mitigate these issues. It is recommended that any future CHC systems consider the longer start up cycle for heat pumps and include provisions for compensating for this.
- Heat Pump, Hot Water Generator, and Hot Water Thermal energy storage hot water discharge temperatures can vary from each other and in some cases the differences can be large enough (20F or so) to result in thermal stratification in main CEF hot water discharge piping. While this did not result in any impact to building HVAC system performance it did cause inaccurate readings of total heat flow and amount of heat in the hot water thermal storage tank. Therefore, the Stanford CEF is being retrofitted with an in-line mixer in the main hot water temperature piping header to prevent this stratification and yield more accurate measurement of hot water energy flows.
- The large Heat Recovery Chillers suffered mechanical failures after two years of operation, requiring several months down time for factory repairs. The root cause of the failures was determined to be inadequate machine internal refrigerant backflow valves and problems with internal machine control logic. These problems occurred because this line of heavy industrial chiller had not been used in the ways designed for SESI and incomplete communication and understanding of intended machine use between the chiller manufacturer and system controls vendor resulted in machine control programming problems. The fix was a redesign of the simple mechanical backflow valves with faster acting pneumatic valves and changes to machine internal control logic, and was performed at no cost to Stanford as a machine warranty issue.
- A year after SESI was designed about 15% of campus heating load intended to be served by the new system was withdrawn because Stanford Hospitals could not convert some of their building systems to hot water in time for system start up due to logistical conflicts associated with construction of a new adult hospital and major addition to the children's hospital. Instead the hospitals constructed a small new unmanned steam plant to sustain steam flow to the facilities which could not be converted. The loss of this major heat load changed the thermal load balance of SESI from a design of 70% of chilled water waste heat recovered to meet 80% of system heating loads to 53% of chilled water waste heat being recovered and meeting 88% of the now lower system heating load. The net effect of the

hospitals failure to convert some of their facilities from steam to hot water is underutilization of the heat recovery chiller equipment capacity, overutilization of the conventional chiller capacity, slightly less greenhouse gas reduction and water savings than could otherwise be achieved, and slightly higher overall system cost. However even with this last-minute change SESI has met or exceeded overall system goals and has proven to be a far superior system in economic and sustainability terms than any other system studied. Eventual conversion of the affected hospital space between 2019 and 2021 when hospital logistics may allow it is now being studied.

Positive lessons learned included the fact that SESI became operational during a several year drought in California and its reduction of campus water use by 18% immediately enabled the university to comply with statewide water conservation requirements of 25% without significant hardship to campus operations. After SESI was approved and construction underway the university also obtained Direct Access to state electricity markets which enabled it to reduce anticipated electricity costs by 20% while immediately reducing greenhouse gas emissions by 68% through long term power purchase agreements to have 53% of the university's power consumption supplied by new solar photovoltaic plants. The subsequent increase in the state's renewable portfolio standard for general market power from 33% by 2020 to 50% by 2030, affecting the 47% of campus electricity being supplied as general market power, further reduced the forecast greenhouse gas emissions from SESI without any action required by the university and reinforced the value of fuel source flexibility through electrification instead of relying primarily on natural gas via cogeneration.

CONCLUSIONS

As an 'electrification + renewable electricity' option, SESI has proven to be the most efficient, economic, and sustainable district energy system possible for serving Stanford University buildings energy needs. Sustainable biogas to support continued carbon-based combustion energy processes is in limited supply, often uneconomical, and is not a scalable solution globally. Likewise, carbon capture and storage technologies are also unavailable or significantly non-competitive in cost to electrification and renewable electricity processes for building power, heating, and cooling. Also, sustainable non-carbon-based chemical/mechanical combustion processes (e.g. hydrogen fuel cells using hydrogen produced with renewable electricity instead of being reformed from natural gas) are less efficient and more costly than renewable electricity options.

For many of these reasons the International Energy Agency, United Nations Environmental Programme, US DOE National Labs, and others have concluded that the path to sustainable building power, heating, and cooling is via electrification and renewable electricity supply. The use of electric powered heat pumps for supplying building heating and hot water enables this electrification.

SESI represents perhaps the first large scale operational example of what the United Nations refers to as a 4th Generation district energy system that employs heat recovery and renewable electricity to achieve a sustainable building energy supply system. SESI was recognized as the Best of the Best Engineering project in the United States in 2015 and the Best Global Green project in the world in 2016 by Engineering News Record. SESI also receive the California Governor's Environmental and Economic Leadership Award (GEELA) in 2015. SESI created new knowledge and tools in how to develop transformational thermal energy microgrids and these are being widely shared with the world. The transformation of Stanford's energy system was accomplished in less than 5 years, allowing Stanford to exceed state, national, and international greenhouse gas emission reduction targets decades early. It is now apparent that conversion of as many processes as possible to electricity and the use of renewable electricity is the easiest, fastest, and most economic path to greenhouse gas emission reduction, not only for building power and thermal loads, but also for transportation and other energy consuming human processes.