Choices

- Chilled water to the floor?
- Raised floor vs flat
- Under-floor or overhead
- AC or DC
- Compute density

Standards & Resources

- Uptime Institute (uptimeinstitute.org)
- Datacenter dynamics presentation (LBL/PGE) -**http://preview.tinyurl.com/2vlopz**
- **BICSI** datacenter standards & design documents ($$)
  - www.bicsi.org
  - ITSIM – Information transport systems installation manual
    - (nee) Telecommunications cabling installation manual (TCIM)
    - Standards, codes, tables, architecture, planning, termination, test
    - 888 pages
- TIA 942
  - Telecommunications infrastructure standard for datacenters
  - Fire, room layout, environment, change control, safety, security, IT infrastructure
- NFPA –70 National Electric Code
- http://hardware.slashdot.org/article.pl?sid=06/09/26/2039213
  - Google calls for simplified power supplies
- ASHRAE – Best Practices for Datacom Facility Energy Efficiency

Cooling

Cooling Math

- 1 BTU raises the temperature of 1 pound of water by 1 degree Fahrenheit and corresponds to 0.293W.
- 1 ton of cooling is 12000 BTU/hr, the energy produced by melting one 'short' ton of ice in 24 hours.
- 1 KWH = 3413 BTU/hr = 1.341 horse**power**
- 1 Therm = 100,000 BTU
- 1 MMBTU = 1,000,000 BTU
- 1 watt = 1 Joule/sec
- 1 ton cooling = 3.515 KW
  - 1 8KW rack requires 2.28 tons of cooling
- **1 40 ton CRAH**

Wet vs dry

- Wet (CRAH)
  - \+ Centralized chilled water is efficient
  - – Chilled water pipes take up space and may leak
  - – Humidity control can be problematic
  - – Makeup water
- Dry (CRAC)
  - \+ can be placed almost anywhere (no major structural pipes)
  - – Less efficient
  - +/- inherent dehumidification (from waste heat)

CRAC (Dry)

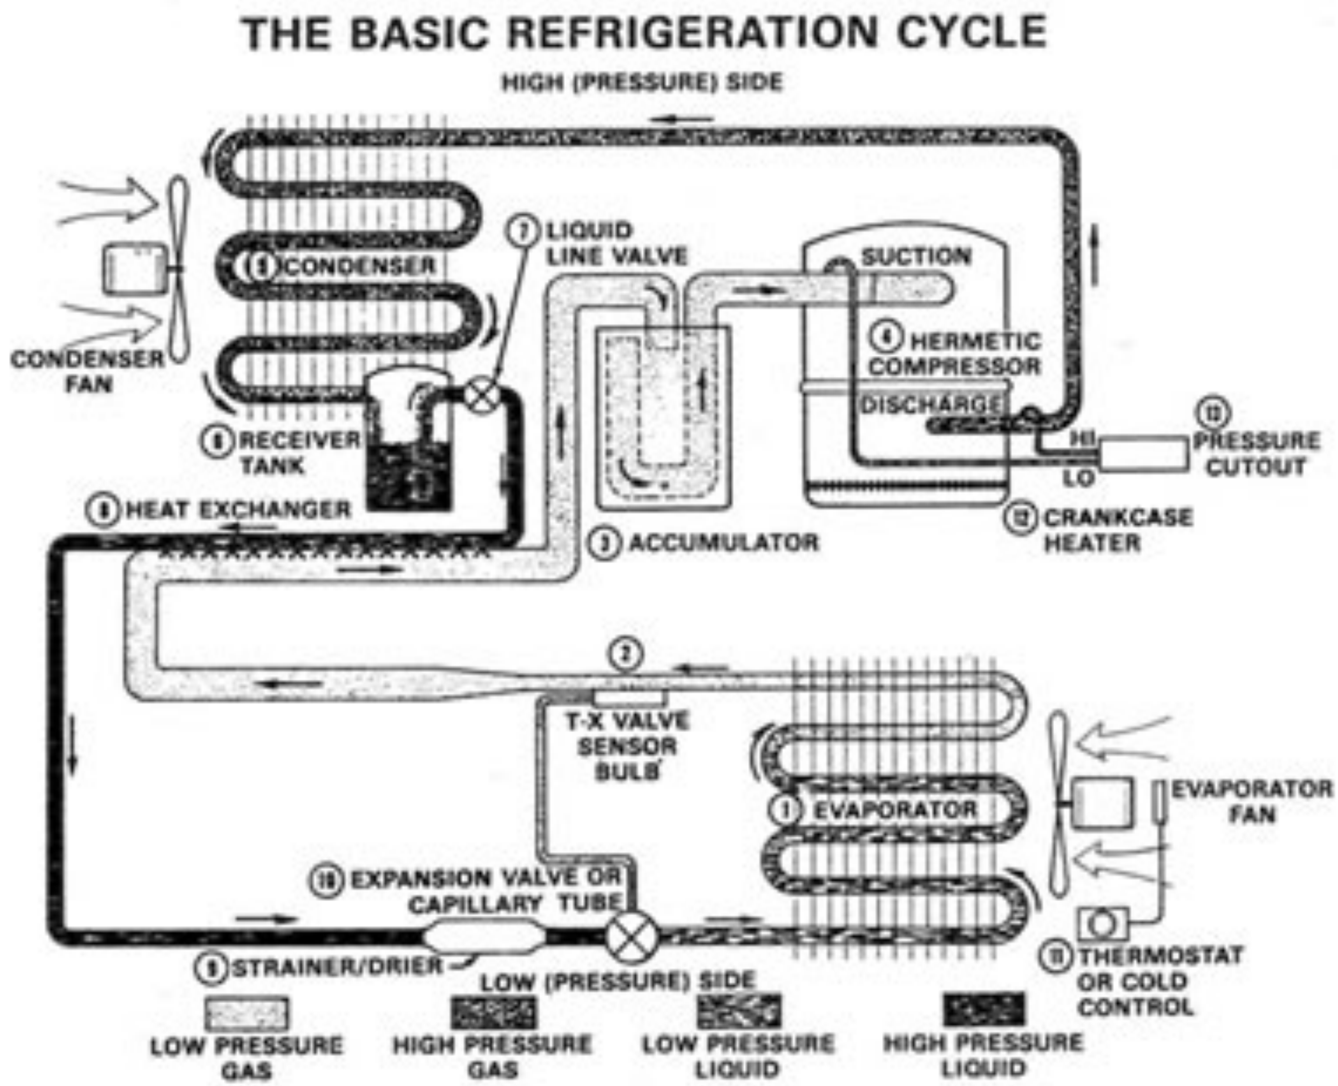
THE BASIC REFRIGERATION CYCLE
HIGH (PRESSURE) SIDE
CONDENSER FAN
5 CONDENSER
7 LIQUID LINE VALVE
SUCTION
4 HERMETIC COMPRESSOR
DISCHARGE
13 PRESSURE CUTOUT
HI
LO
6 RECEIVER TANK
8 HEAT EXCHANGER
3 ACCUMULATOR
12 CRANKCASE HEATER
2
T-X VALVE SENSOR BULB
1 EVAPORATOR
EVAPORATOR FAN
10 EXPANSION VALVE OR CAPILLARY TUBE
9 STRAINER/DRIER
LOW (PRESSURE) SIDE
11 THERMOSTAT OR COLD CONTROL
LOW PRESSURE GAS
HIGH PRESSURE GAS
LOW PRESSURE LIQUID
HIGH PRESSURE LIQUID

Chilled water (Wet) cooling

- Cooling provided by chilled water system
  - BTU/hr = 8.345 lbs/gal *60 minutes/hr * **GPM * DT**
  - DT = delta-T difference in inlet and outlet temperatures
  - GPM = gallon flow per minute
  - 8.345 = weight of 1 gallon of water
- Chilled water cooling capacity can be increased by increasing flow rate or by increasing the delta-T temperature change
- Water has approximately 1000 times the cooling capacity of air

CRAH

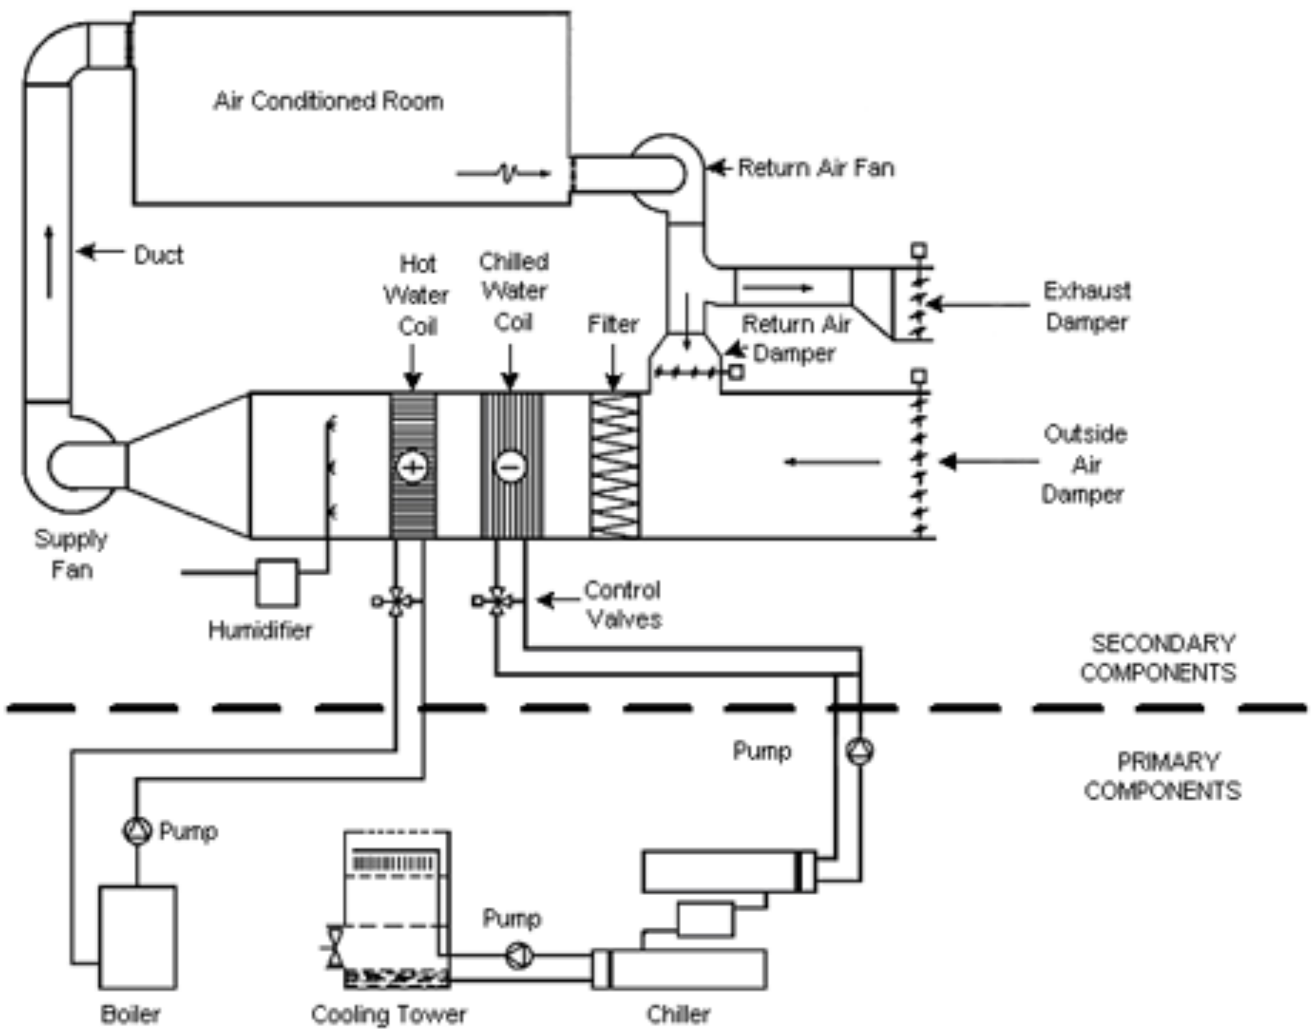

Air Conditioned Room
Return Air Fan
Duct
Hot
Water
Coil
Chilled
Water
Coil
Filter
Return Air
Damper
Exhaust
Damper
Outside
Air
Damper
Supply
Fan
Humidifier
Control
Valves
SECONDARY
COMPONENTS
PRIMARY
COMPONENTS
Pump
Pump
Pump
Boiler
Cooling Tower
Chiller

Air cooling capacity

- For Airflow, compute CFM needed to cool a unit generating W watts with
  - CFM = 3.2*1000*W/deltaT
    - deltaT is in Fahrenheit.
    - (Use 1.76 instead of 3.2 for deltaT in Celsius).
  - To cool a row of 15 cabinets at 20KW per cabinet and 20 degrees F difference between hot and cold aisle requires:
    - 3.2 * 1000 * 20 * 15 / 20 = 48000 CFM
      - (approximate duct size 4′x7′)
- Alternate (also F):
  - CFM = 3.4*VA/1.08/deltaT

Lack of hot/cold isolation leads to…

- Cold air incoming to return
  - Already cooled once and close to dew point
    - Wasted cooling capacity removing water
      - That must be re-added using humidification
      - That could have been used for cooling warm air
  - Units not operating at full capacity/efficiency
- Lower setpoint won't effect temperature and may cause additional condensation

Competing factors

- Highest Delta-T = best cooling efficiency
  - Human comfort
- What about side-to-side cooled devices?
- Most cables are in the hot aisle
  - Pass throughs? Hassles
  - Some vendors have all cables, power and other on front of machine
- Chimney designs limit air flow
  - Forced air = electrical cost
  - But have very good delta-T

Delta-T

- System temperature change
- Cold water and hot air = efficient cooling
  - E.g. air temperature near a large body of water
- Hotter hot aisle and colder water = higher delta-T and more efficient energy profile.
- Mixing hot and cold air results in extra work, low delta-T, and extra energy use.

Humidification

- 1 Static spark = 10,000+ volts

  "One pound of normal data center air (air can be weighed) occupies about 13.6 cubic feet of space and contains about 1/7th an ounce of water."
- Plant Steam
- Steam canister
- Infrared
- Ultrasonic

Humidifier matrix – (APC humidification table 3)			
Humidifier	**Capex**	**Opex**	**Maintenance**
Steam canister	Low	High	Low
Infrared	Low	High	Low
Ultrasonic	High	Low	High

Static electrical buildup: (Source: APC humidification PDF Pg 4)		
Action	**Static building at 80% RH**	**Static building at 20% RH**
Walking across ungrounded raised floor tile	250 Volts	12,000 volts
Walking across synthetic carpet	1500 Volts	35,000 Volts

Humidity guidelines for IT equipment: (source APC humidification Table 2)			
	Allowable RH	Recommended RH	Max dew point
Wiring closets	20-80%	40-60%	70°F
Computer rooms and data centers	20-80%	40-60%	63°F

ASHRAE PSYCHROMETRIC CHART

Winter 40°F 50% RH

Summer 85°F 80% RH

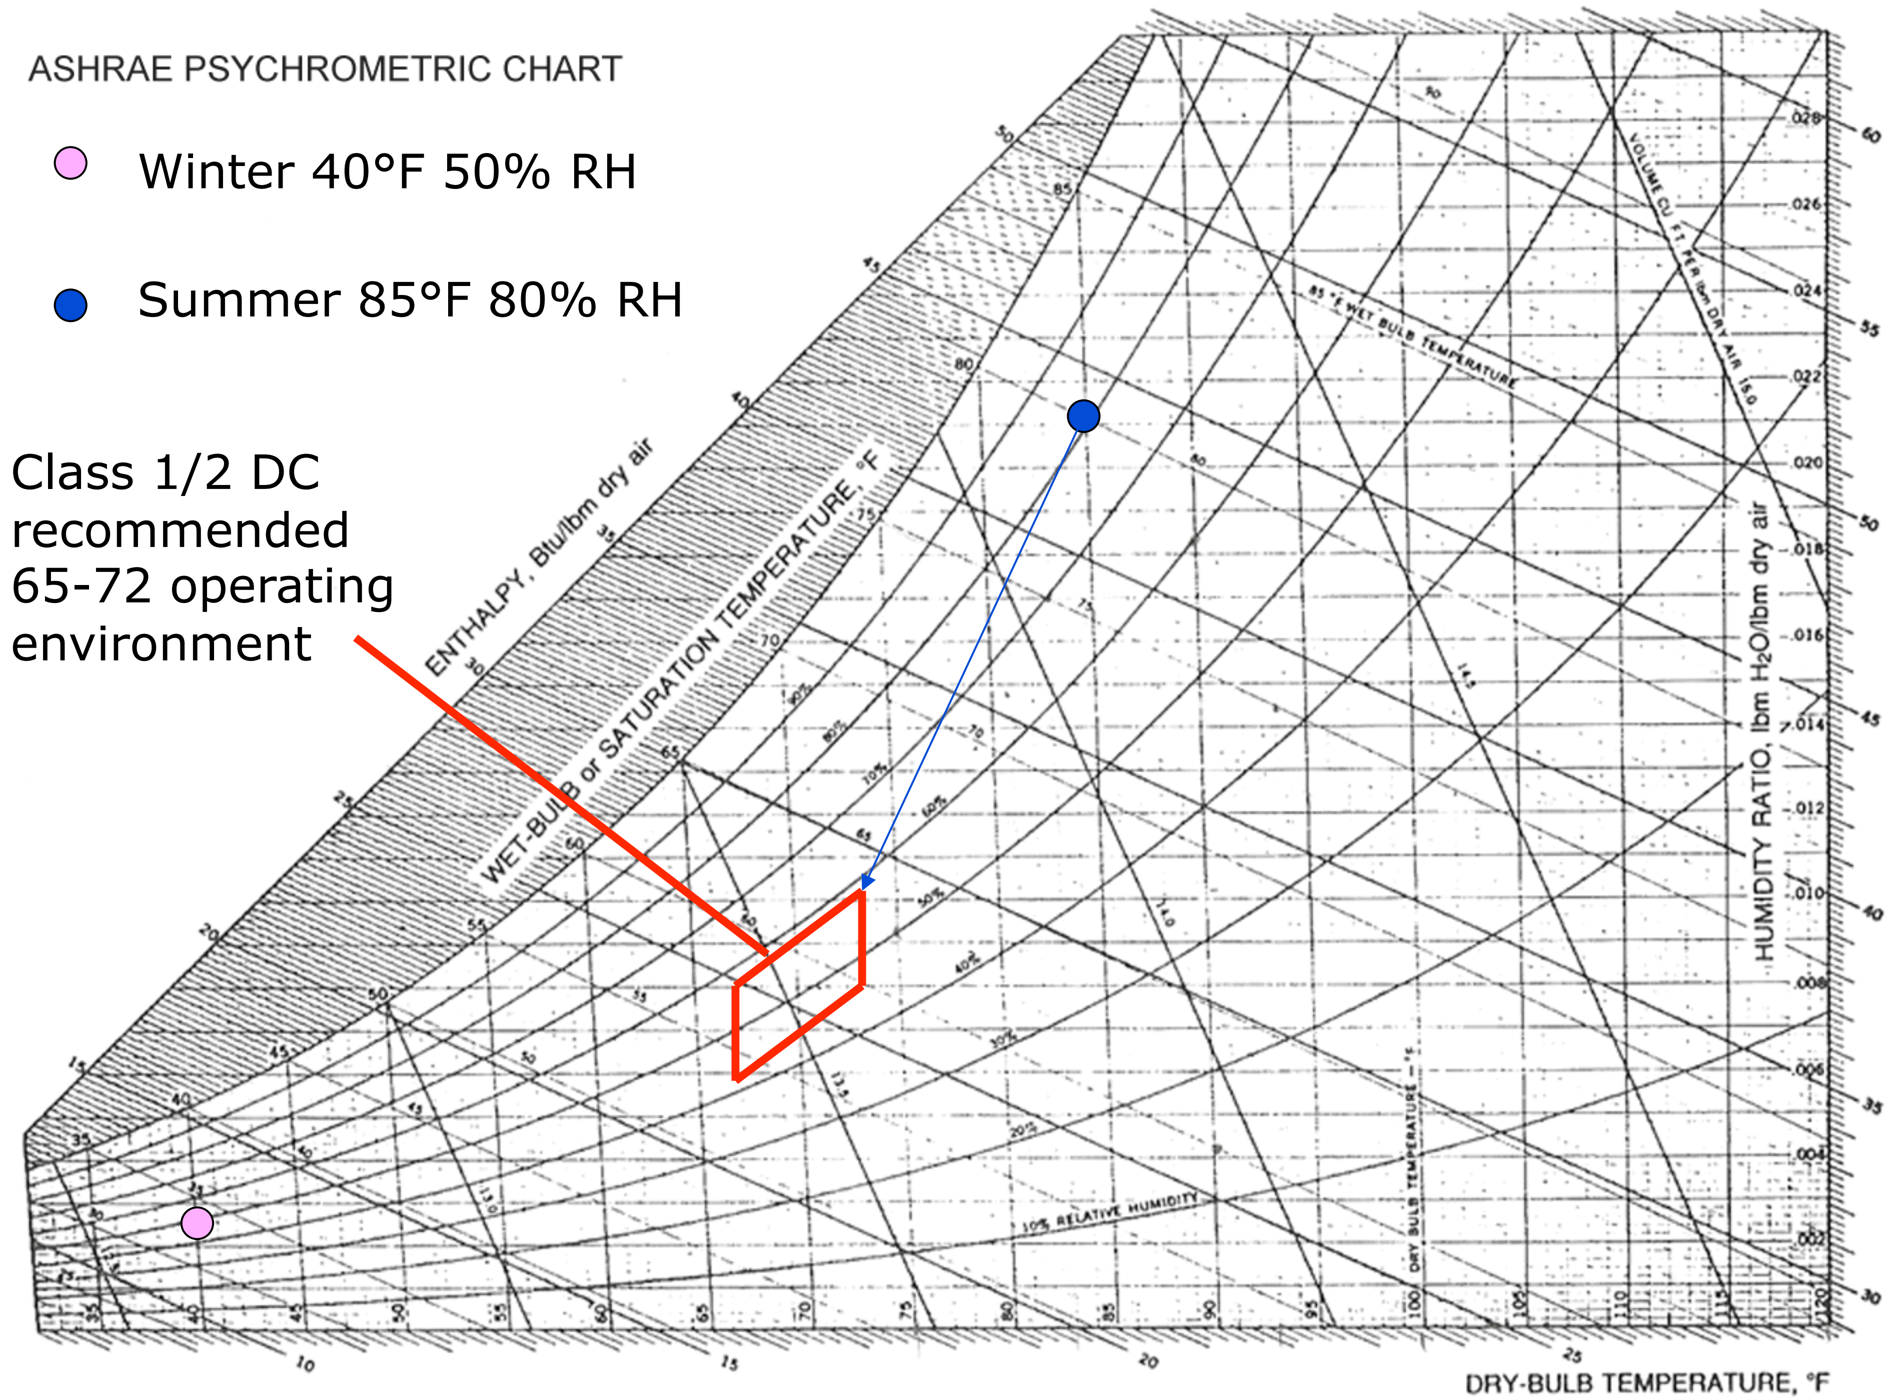

Economizers

- Air side
  - Pump waste heat directly outside (much $$ savings potential)
  - Humidification of cold air
  - Mechanical failure
- Water side
  - Heat exchanger
  - A fan and a pump
  - Water treatment of outdoor loop
    - Maintenance when outside temperature is very cold?
  - DC Lifetime savings, but high up-front costs

Practical CRAH usage

- Disable humidification on all but 1 units
  - Or use purpose-specific humidifier
  - Or make sure *n* units are calibrated ***precisely*** for relative input settings
- Have most units’ humidity setpoint well above actual humidity. Adjust one down only when humidity is needed, then adjust up.
  - Prevent a humidity war
- Keep an eye out for changing conditions
  - Open doors, new equipment, failing fans, all change return air
  - Automate
- Increasing server room cold temperature 1°F can yield a 1% decrease in operating costs.
- Disable reheat

Backdraft and Setpoint relativity

```
> ~/bin/lieberts
liebert2:
System is  on, at 10% capacity, Cool  on Heat off hum off dehum off
Setpoint: 78F 45%        Actual: 77F 34%

liebert3:
System is off, at  0% capacity, Cool off Heat off hum off dehum off
Setpoint: 81F 43%        Actual: 53F 24%

liebert14:
System is  on, at 100% capacity, Cool  on Heat off hum off dehum
off
Setpoint: 77F 45%        Actual: 81F 29%
```

liebert3 has no damper

Current cooling options

Liebert XDK
(17-25/25-35KW)
(Aka Knuerr CoolTherm)

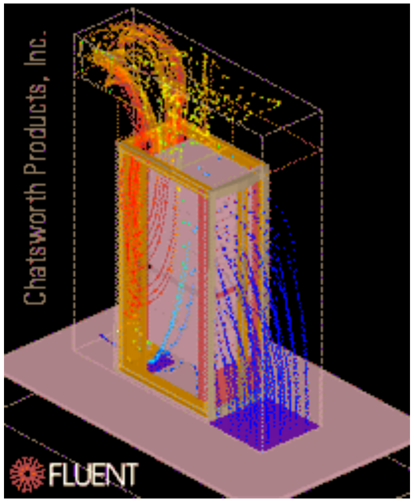

Chatsworth passive
(4-8KW)

Sanmina-sci ecobay
(25KW) - EOL

APC rear door (to 25KW non-redundant)

+ comfortable environment
+/- medium density
- Matched blower pressurization
+/- moderate electrical usage
+ high delta-T

More cabinet stuff

IBM rear door
(coolBlue)

Rittal LCP (30)
(for 42U™ racks)

+ density
+ noise reduction
- cost
- Water near computers

HP (freaking huge)
Freon loop!

Supplemental cooling

RC – to 30KW

RP 0-70KW

Inline SC – up to 7KW

APC inline cooling accessories

Hot aisle containment

Liebert - refrigerant/water heat exchangers
XDV - to 10KW
XDO – to 20KW

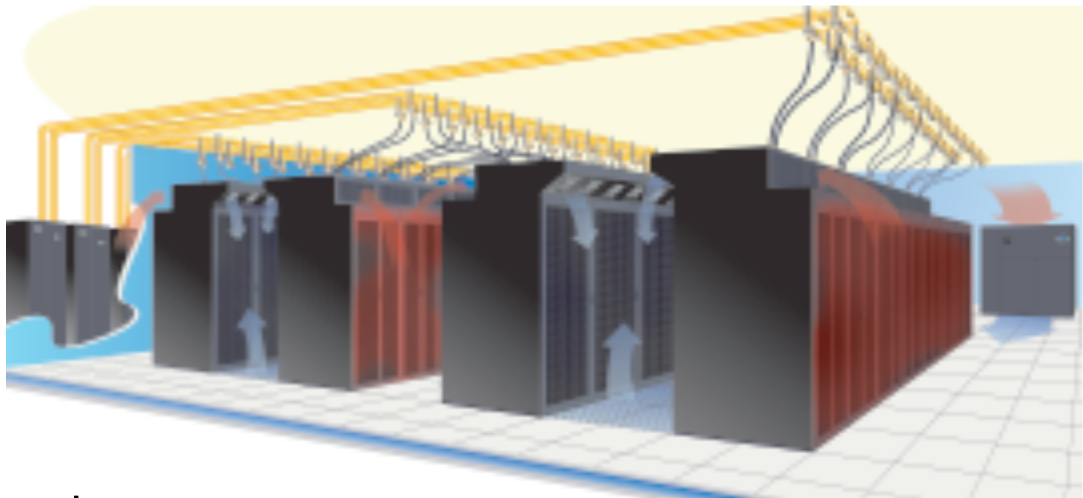

Peripheral refrigerator
H/X

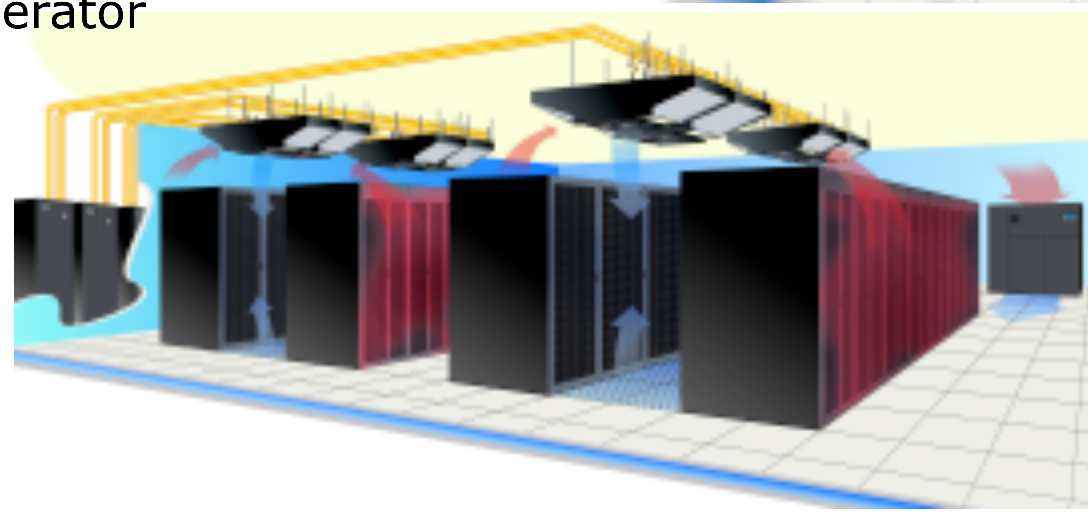

Flooring

- Rolling vs Static loading
- Cement/epoxy
  - + high weight load
    - Usually – consult your building engineer, especially for multi-floor
  - – Bad for chilled water
    - Possible high density implications
- Raised floor
  - High load =~ $25/sqft
  - + place for piping of chilled water
  - +/- cabling?
  - – Mixing power and water may lead to "unhappiness"
  - + forced air distribution
    - – high floor = high expense but builtin plenum)
  - – What about your elevator? (if necessary)
  - – Pushing heavy stuff up ramps is not fun – lifts?
  - – Zinc whiskers

Fire suppression

- Pre-action
  - + Air filled pipes
  - + Smoke detection
  - + Heat detection at sprinkler head
  - – Cleanup
  - – Interior fires?
- Under floor detection?
- Dry agent fire suppression (Aero-K, CO2, etc.)
  - + Interior fires extinguished
  - + Minimal downtime; no dry-out period
  - – Room sealing
  - – Corrosion potential
  - – **Asphyxiation is a career limiting experience**

wiring

- Under floor? Overhead?
  - Tracing
  - Distance limitations
  - Rats nests (Snake Tray!)
- BICSI standards
  - (Building Industry Consulting Service International)
- TIA standards

Cable management

- Wire ducts
- Ladders
- Fiber ducting
  - Bend radius
- Vertical vs horizontal








Power Busway (Starline)

Power

- Higher voltage = better efficiency
- Fewer conversions = better efficiency
- Power factor correction
- 3 phase power (polyphase)
- DC vs AC

High voltage power distribution

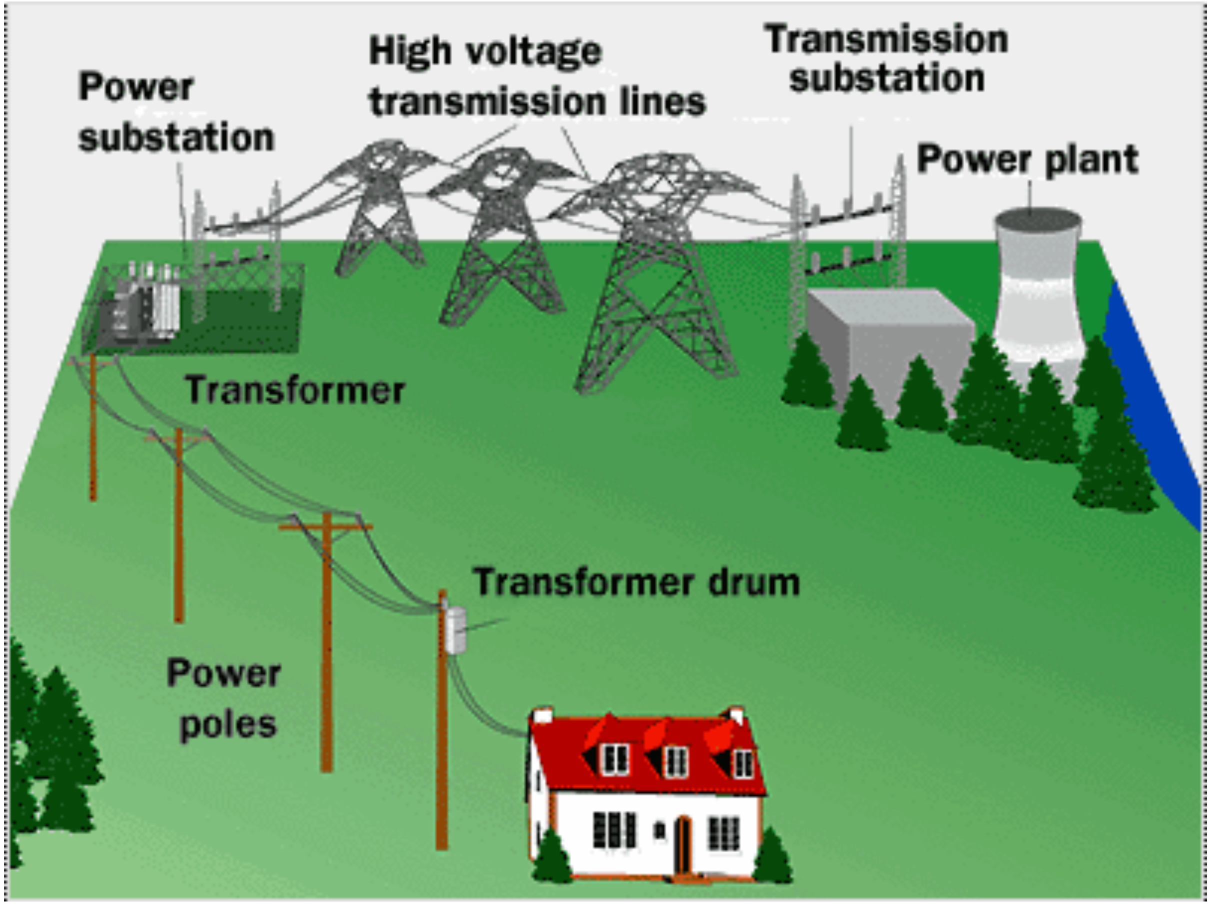

Power supply

- 1MW =~ 1000 homes
- 1 utility generator =~ 4MW
- 1 medium-large datacenter =~ 10MW
- Transmission -> 23KV 3ph -> 12.5KV 3ph-> 480V 3ph -> UPS -> 480V 1-3 ph -> 208V 3ph -> 120V single phase -> 12V/5V ATX
- Each conversion uses 1-2+%
  - The PS in your computer can be as much as 35% efficiency drop. (old ones suck worse)
    - Redundant PSUs are worse still
  - All loss is dissipated as heat which you have to remove, using more electricity

Transformers

- Varying electrical fields induce magnetic field B in ferromagnetic material
- Current entering through primary coil creates magnetic flux which creates current in secondary coil
- Secondary current is proportional to ratio of number of input winds to output winds

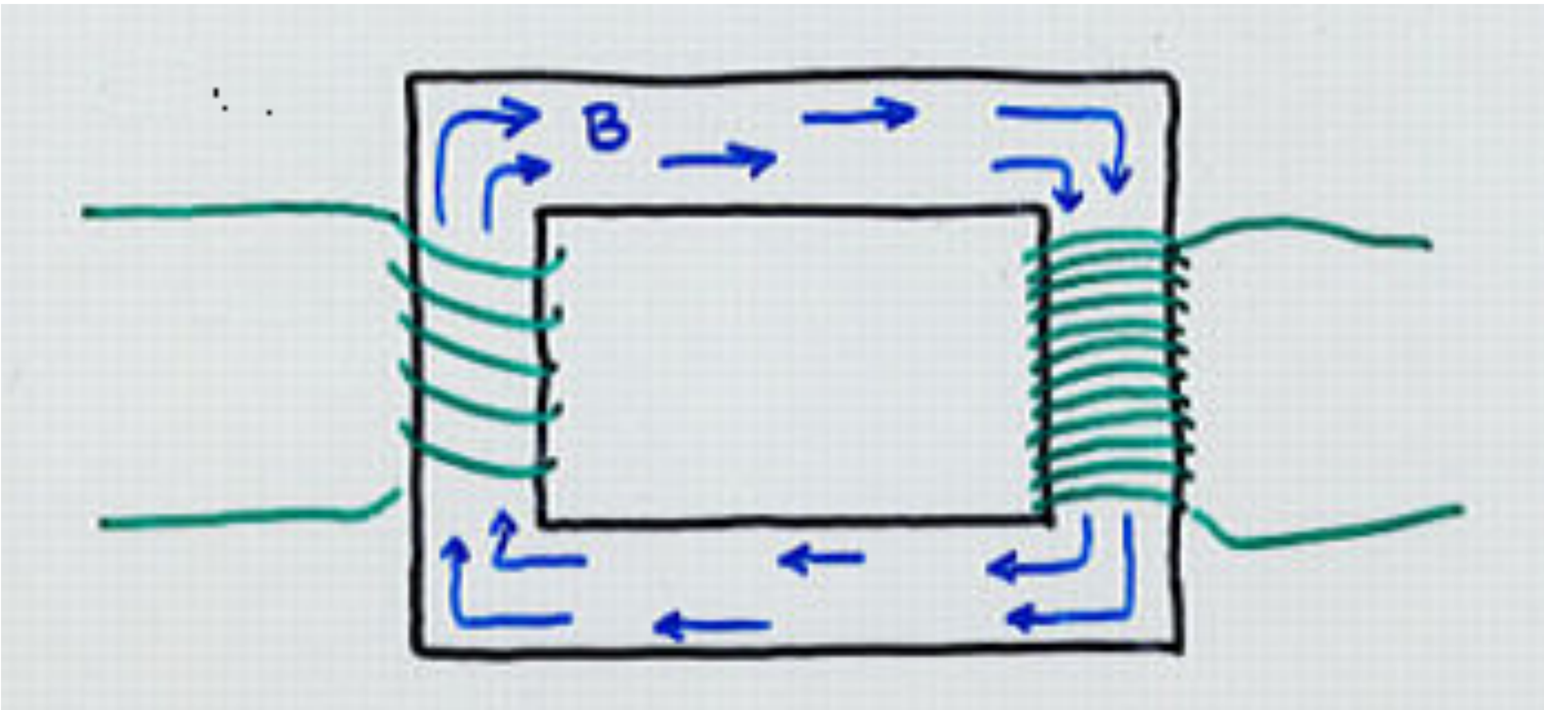

More on transformers

- Keep transformer loads below 80%.
- Overloaded transformers experience core saturation which distorts waveforms.
- Distorted waveforms cause excess heating in the loads.

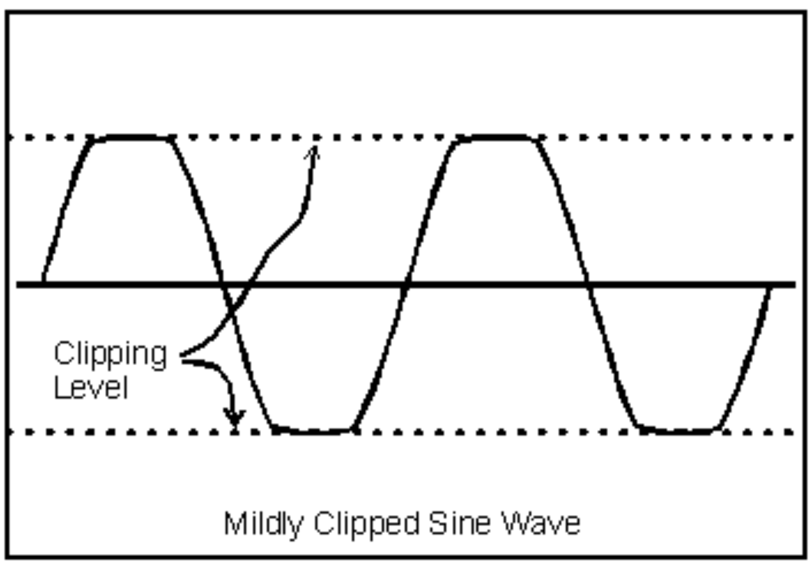

Power factor

- The ratio of real power to apparent power.
  - Real power – capacity for performing work in a unit of time
  - Apparent power – current multiplied by voltage – can appear to be higher than real power because of inefficiencies, distortions and loading effects
  - A typical modern PC has a powerfactor of 90% or higher. Higher is better.
- Also measured as VA (UPS requirements) vs Watts (heat dissipation and power utilization)

3 phase power

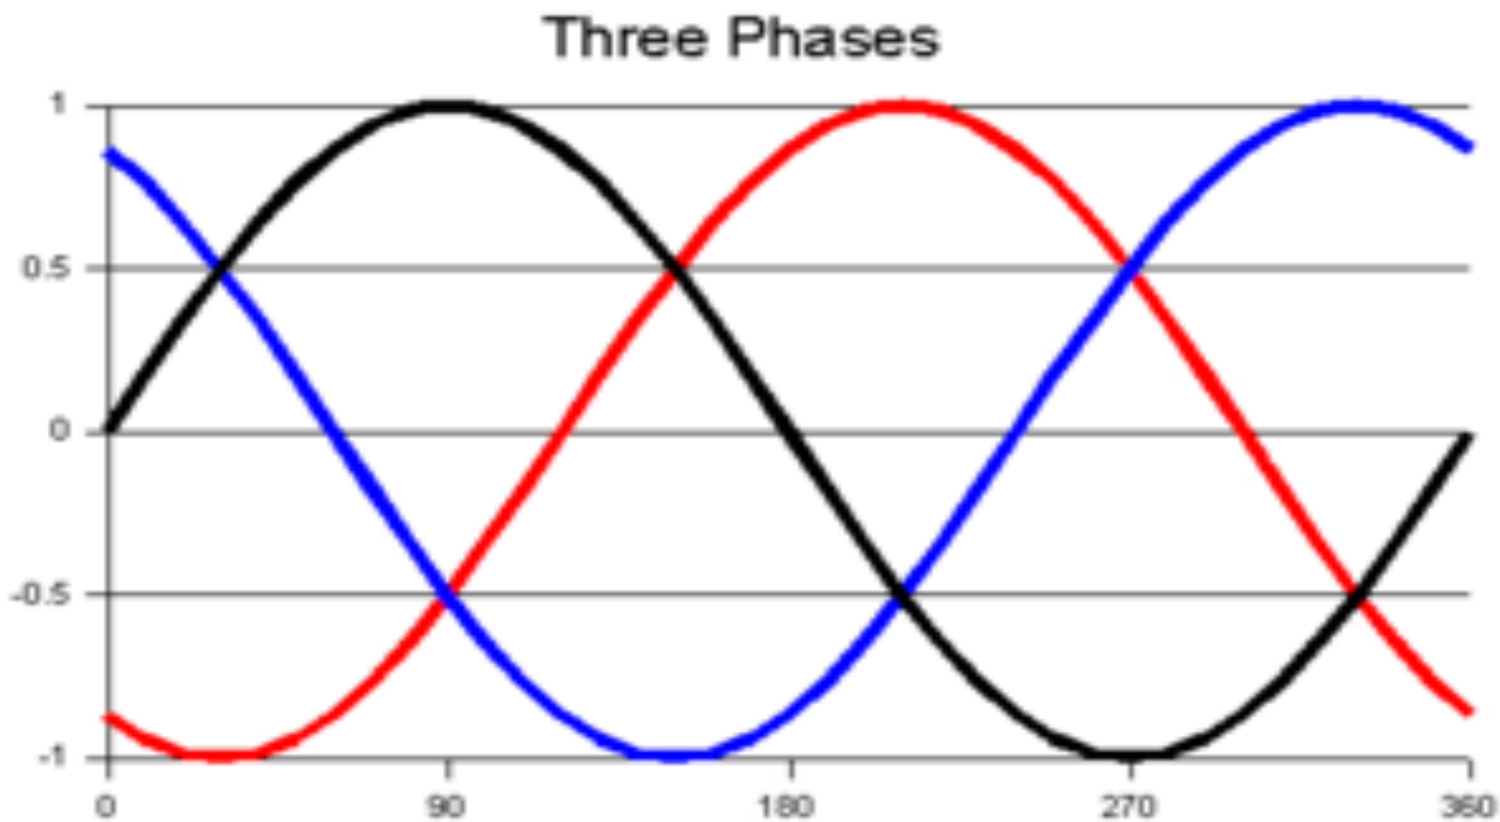

3 phase power

- 3-phase power systems can provide 173% more power than a single-phase system.
  - Smaller conductor
- 3-phase power allows heavy duty industrial equipment to operate more smoothly and efficiently.

3 phase generator/motor

- Simple design
- High starting torque
- High efficiency
- pumps, fans, blowers, compressors, conveyor drives
- More compact
- Less expensive
- Less vibration
- More durability

Simplified 4:2 3-phase transformer

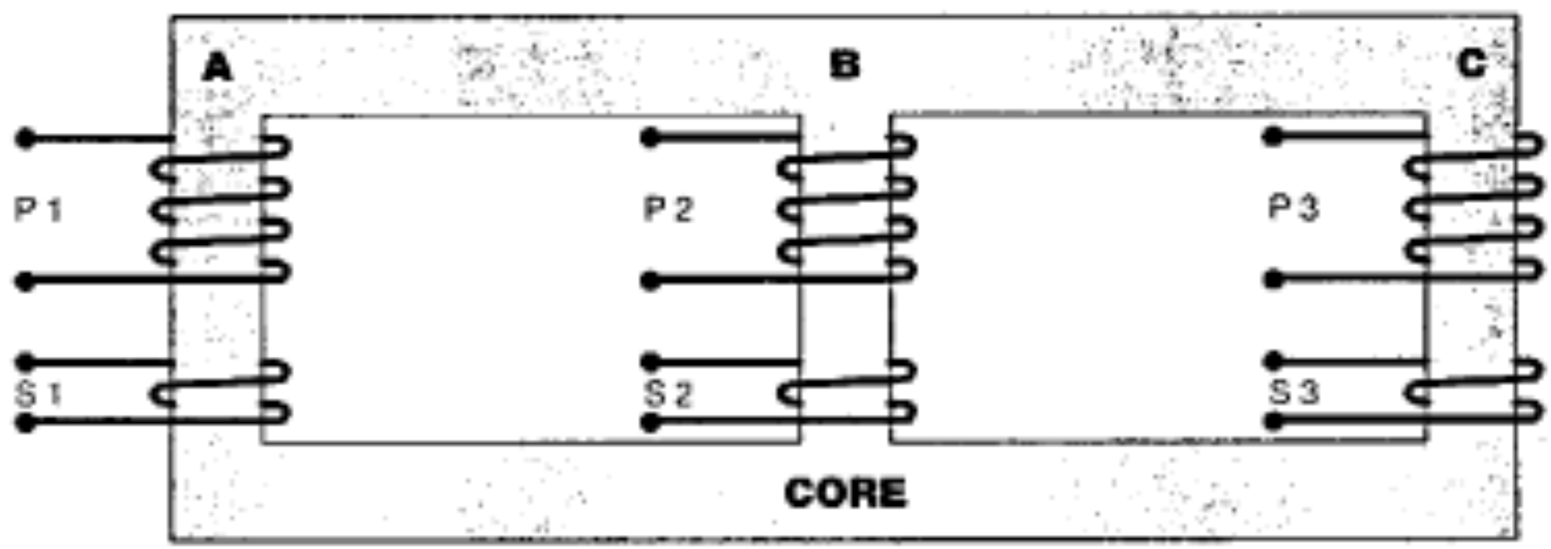

3 phase power

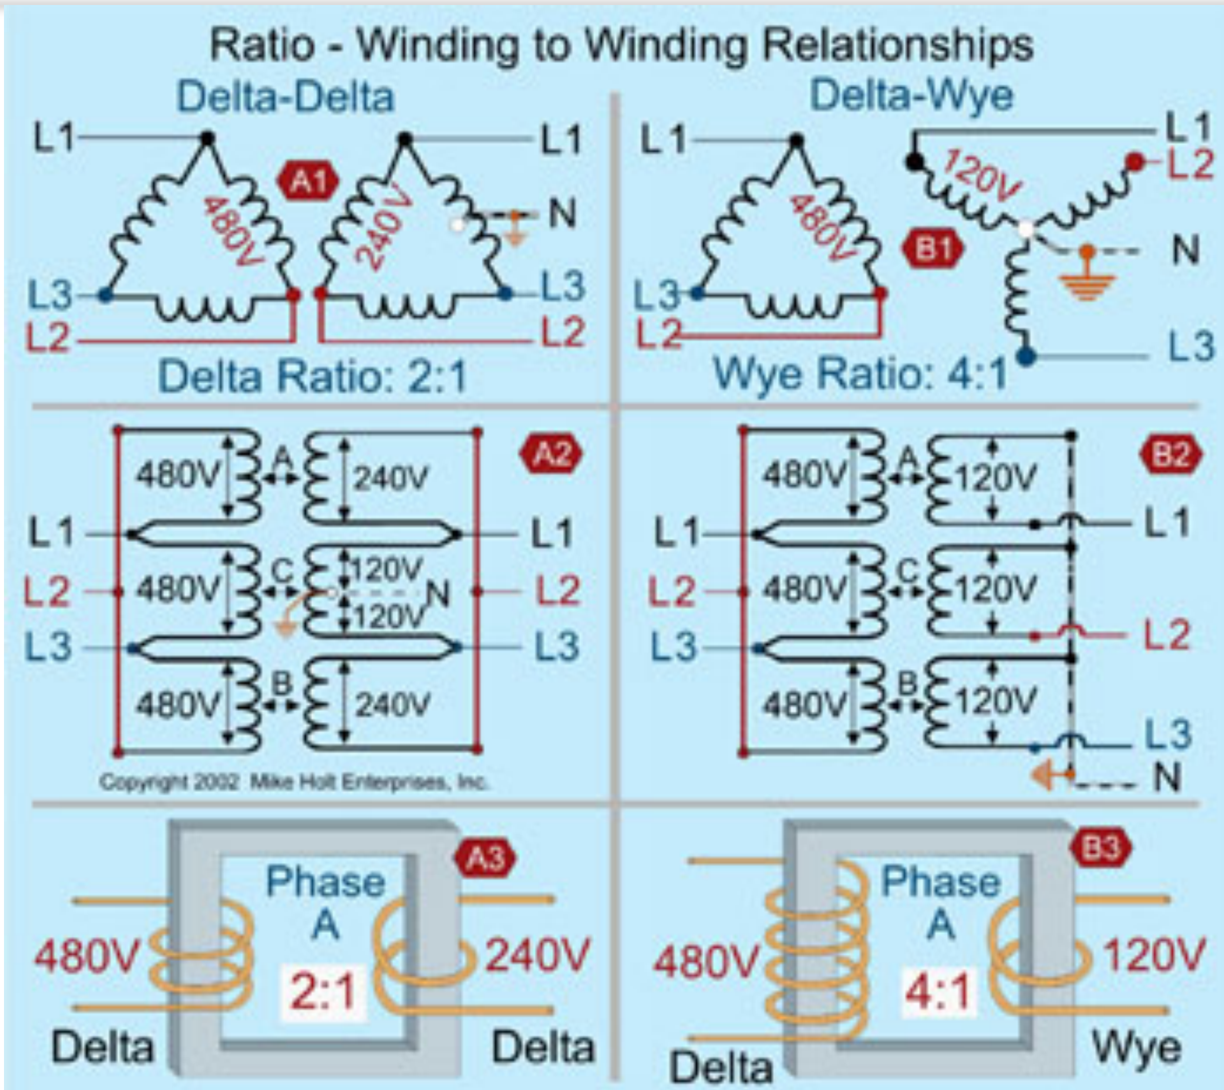

WYE vs Delta

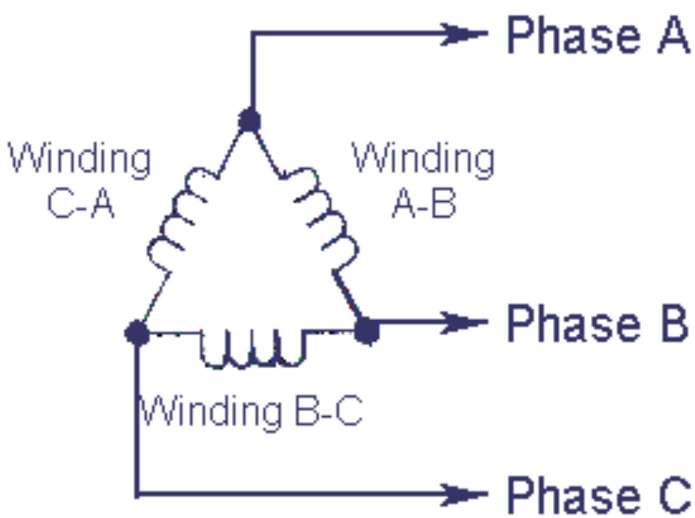

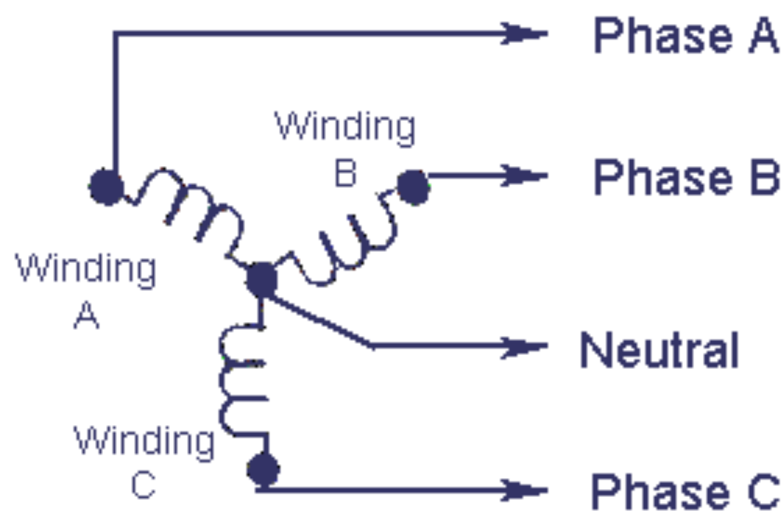

L16-30 (480)

L16-20 (480)

L21-30 (208/120)

L22-30 (480/277)

Delta attachment vs Wye

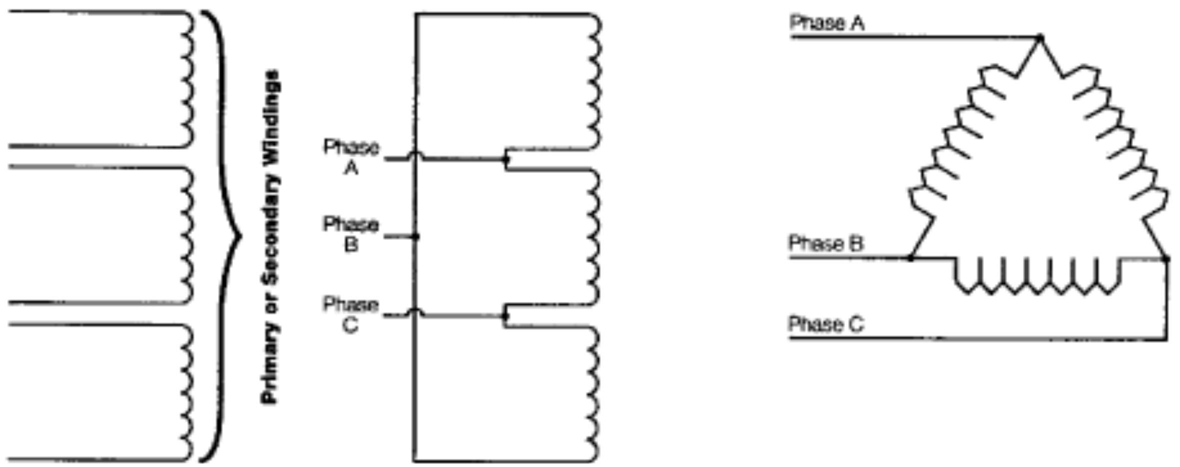

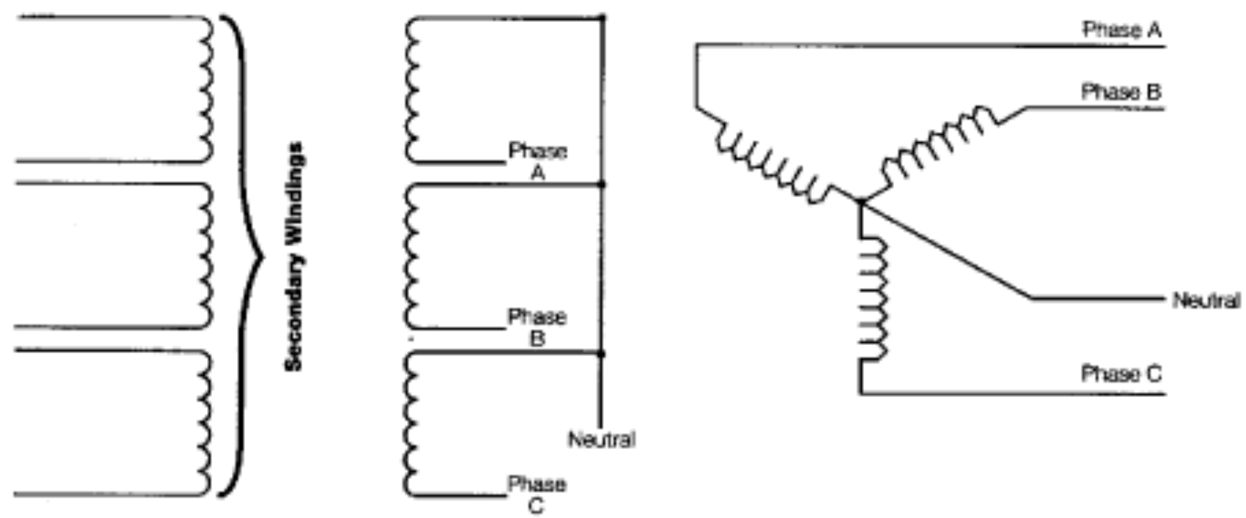

Delta vs WYE

- Delta
  - + Fewer conductors (3-4)
  - – One voltage (e.g. 208)
  - + reduced harmonic potential
  - VPhase=VLine
  - ILine=IPhase×√3.
  - IPhase=ILine÷√3
- Wye
  - + Can run 208 and 120 (or 480 and 277)
  - – More conductors (5)
  - – Switch mode power supplies operating at 120 generate harmonics on the neutral. Harmonics are additive leading to potential overheat and fire hazard.
  - VPhase=VLine÷√3
    VLine=VPhase×√3
  - IPhase=ILine

WYE

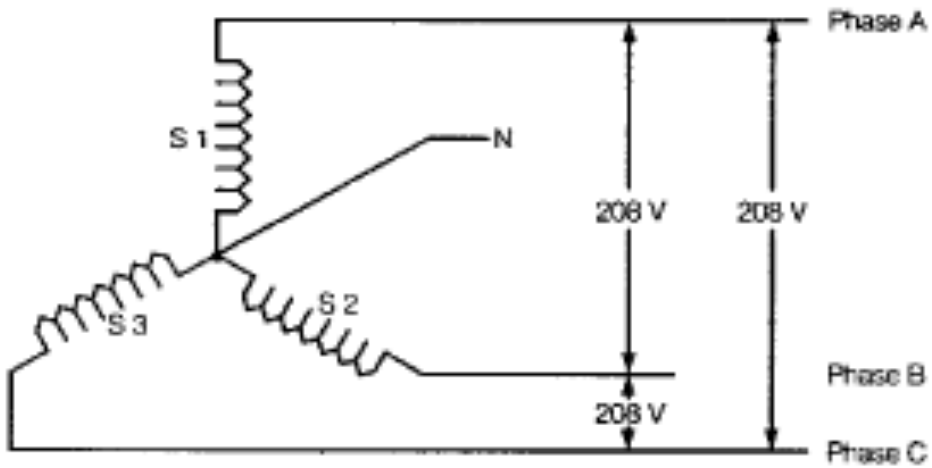
Phase A
S 1
N
208 V
208 V
S 3
S 2
Phase B
208 V
Phase C

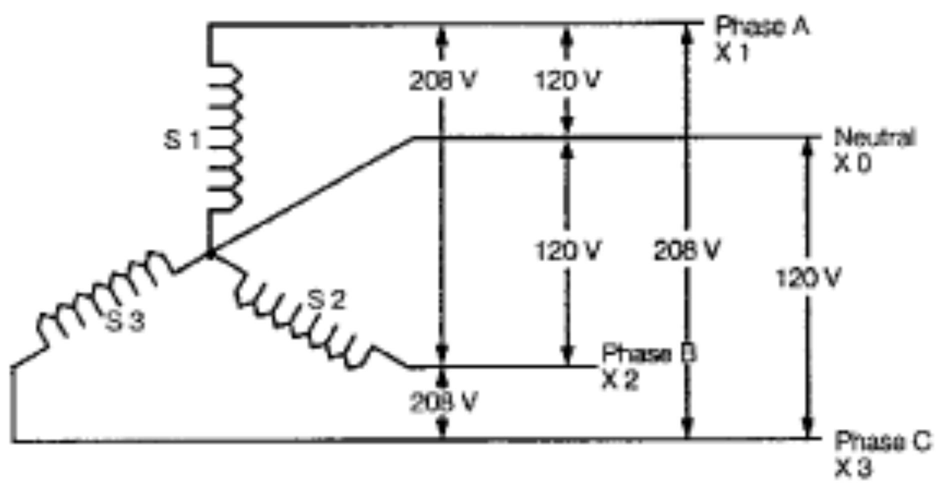
Phase A
X 1
208 V
120 V
S 1
Neutral
X 0
120 V
208 V
120 V
S 3
S 2
Phase B
X 2
208 V
Phase C
X 3

Switched mode power supply (SMPS)

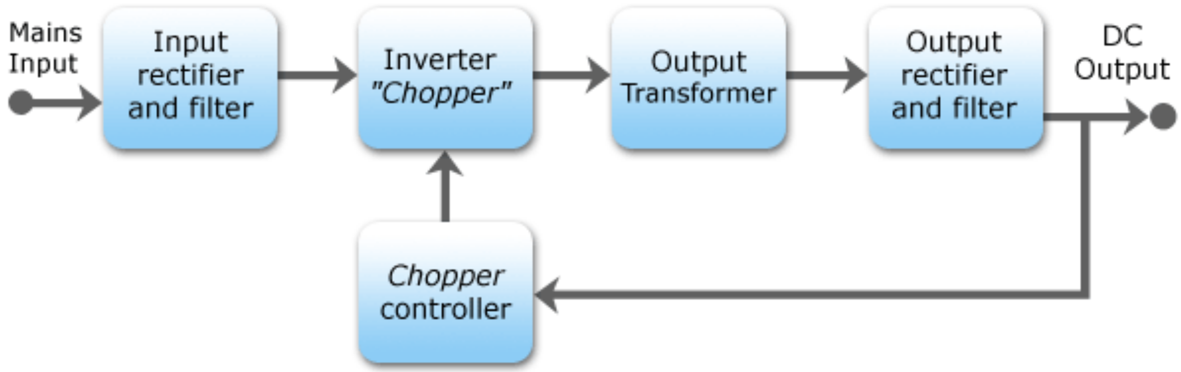

Input sampling happens at waveform peaks, shearing off the top
Of the sine wave and distorting the waveform with harmonics

sampling

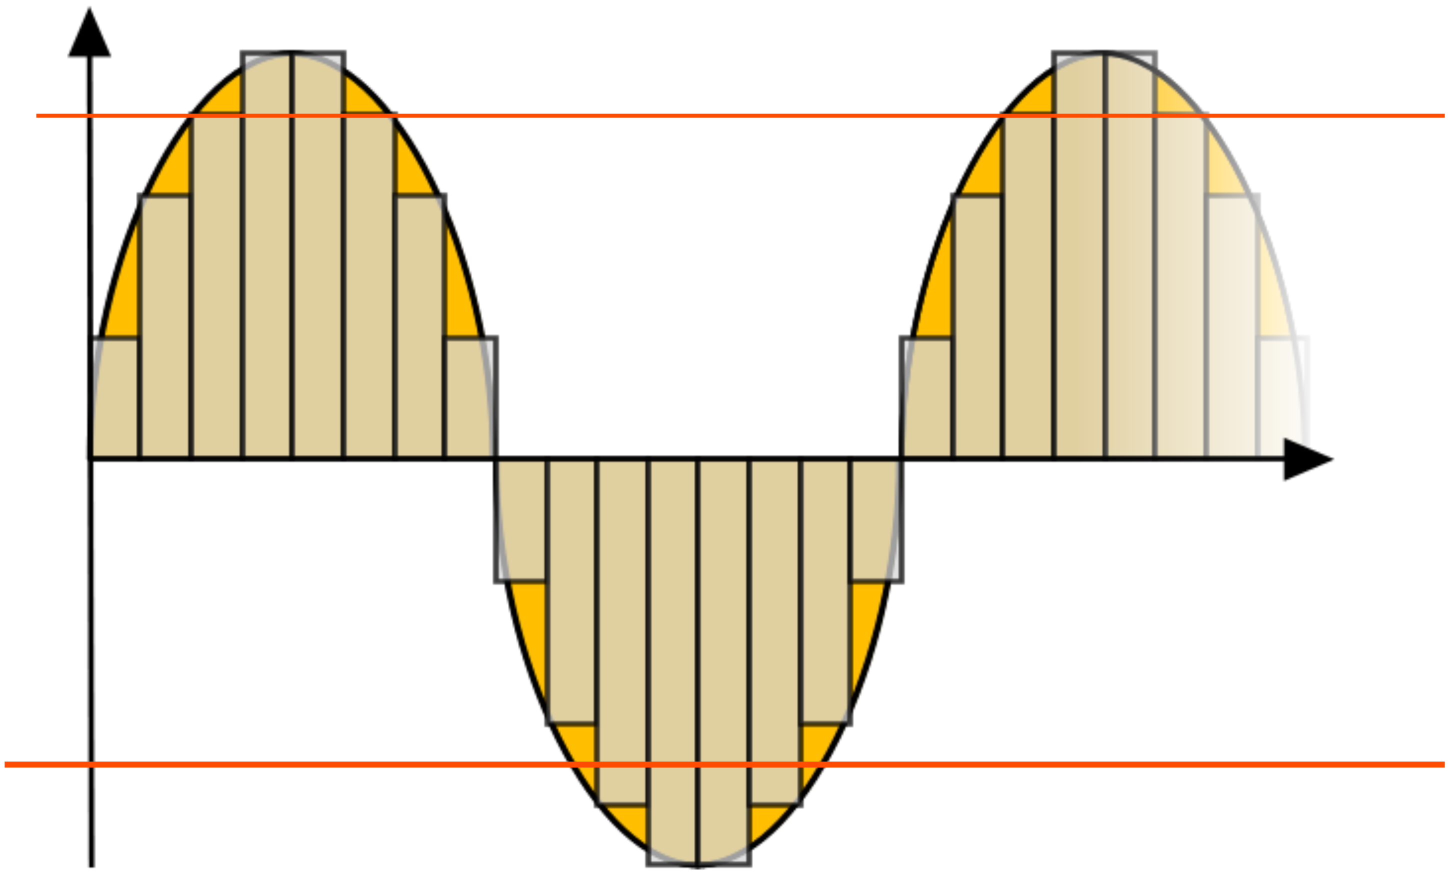

Liebert vs APC (small/medium UPS)

- Liebert takes input of 480V from building (less expensive for capital
  - But, lower operating efficiency ~ 92-94%
  - Redundant modules allow more serviceable components
- APC takes input at 208V. Higher typical cost for setup putting in 2 upstream transformers to go from 480 to 208.
  - but 98% efficiency
  - Failback tends to drain batteries

Wisdom: Daisy chaining UPSen

- Avoid connecting rack UPS to building UPS
  - If the output of the upstream UPS is not a sine-wave output (some older UPS's have square-wave output)
  - The upstream UPS should be at least 10 times larger than the downstream UPS.
  - Mostly applies to ferroresonant-type UPSen. Interaction between two such hybrid UPSen can cause voltage regulation problems.
  - Online UPSen have full rectifier/inverter systems that avoid this problem.

Power efficiency (UPS)

- 4000KW * $.13/KWH * (100%-94%) efficiency * 365.25 days/year * 24 hours/day = $273500 / year
  - Load or no load; pure overhead
  - Not including heat extraction!
- Use high efficiency transformers!

Dual source: Automatic transfer switch

- Closed-transition
  - Phase monitoring
  - Make before break
  - ¼ second overlap
  - More expensive
- Open-transition
  - Less expensive
  - Downtime
- Isolation bypass
  - Maintenance mode
  - Double the normal cost
- DC systems use large diodes (save $$)

Fuses

- Overcurrent
  - Overload – a condition produced by load where the sourced current exceeds the capacity of the circuit
  - Short circuit – insulation breakdown or wiring error, bypassing load (usually higher amperage for shorter time)
- Slow blow
  - Allows temporary and harmless inrush currents to pass without opening
  - Opens on sustained overloads and short circuits (msec).
- Fast-Acting Fuses
  - Used on circuits without transient inrush currents.
  - Short circuits very quickly (usec)
- Dual-element fuse
  - a short circuit strip, soldered joint and spring connection.
  - During overload conditions, the soldered joint gets hot enough to melt and a spring shears the junction loose. Under short circuit conditions, the short circuit element operates to open the circuit.
- Don’t forget power factor correction!

Fuse vs Temperature

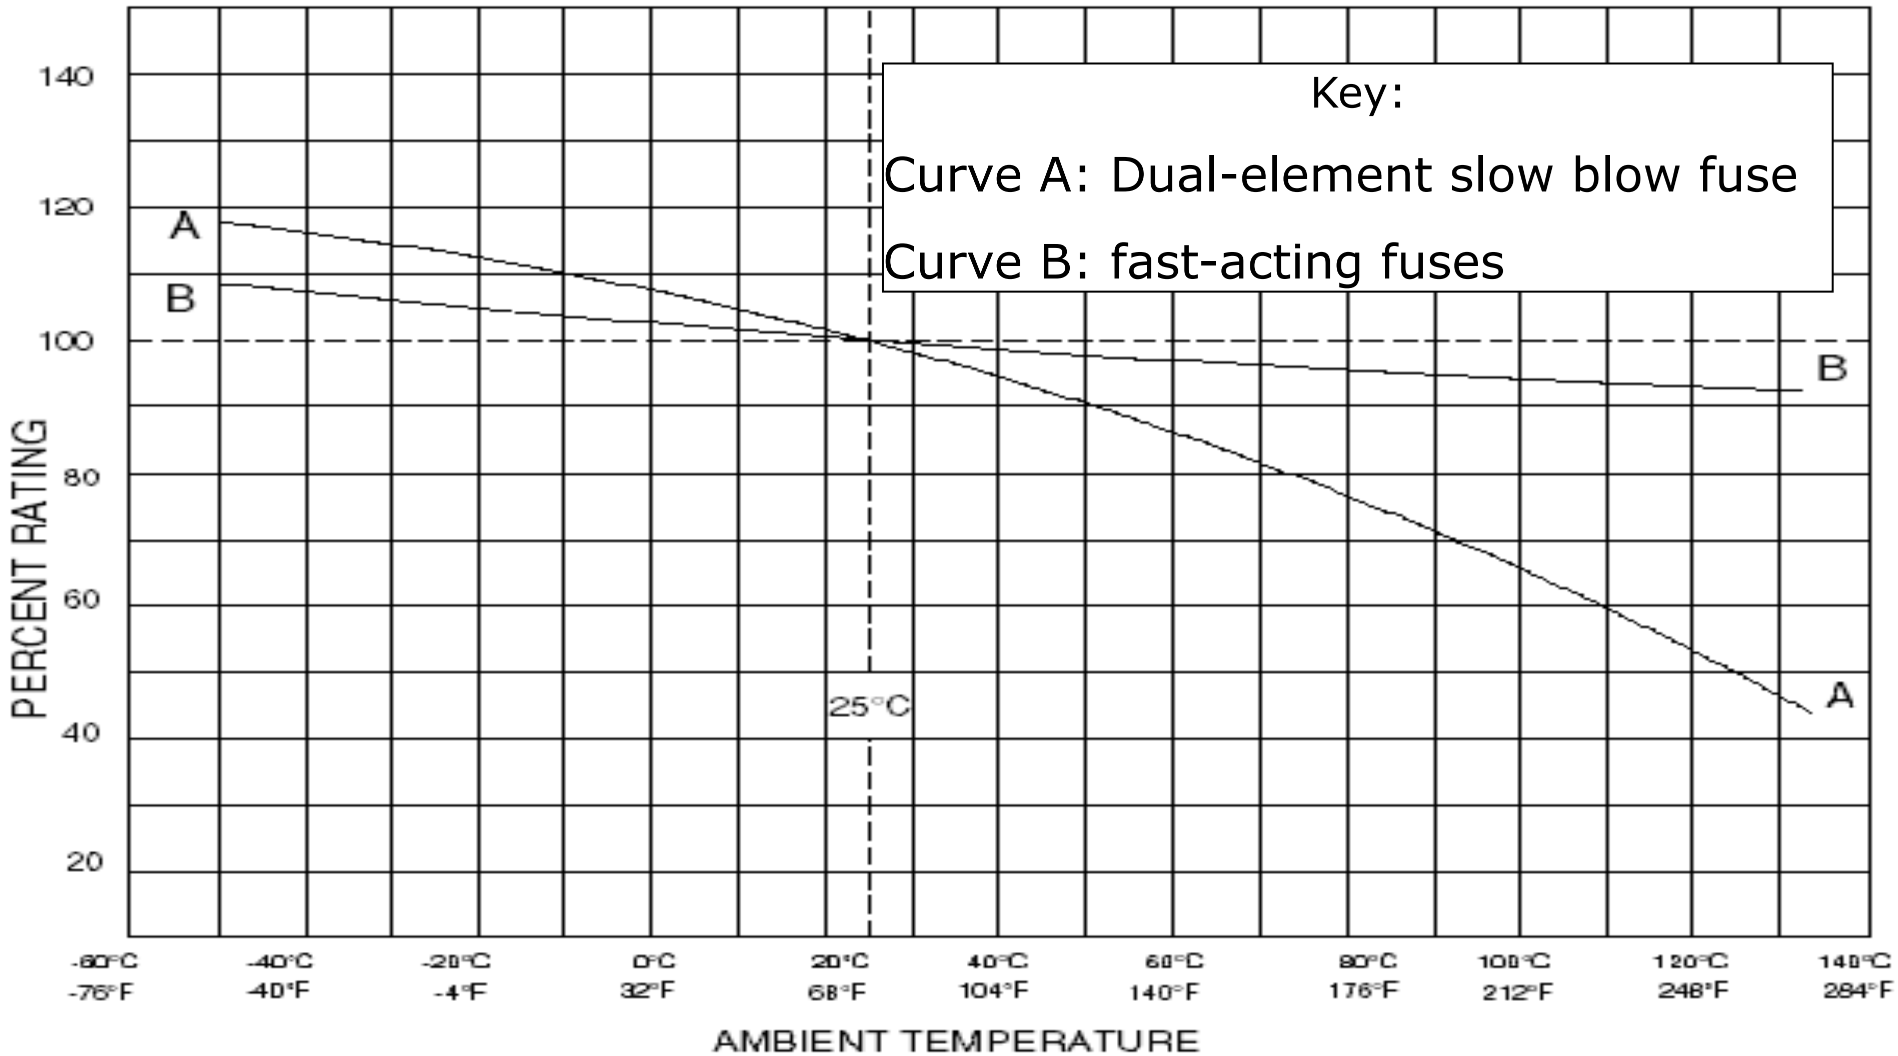

Power recommendations

- The higher the input voltage, the more efficiently the PSU will run
- Use 3 phase power distribution or high voltage DC
- Fewer transformations = greater efficiency
- Variable speed fan motors use less power
- Insist on high efficiency ‘right sized’ power supplies from your vendor (with power factor correction!)
- Use only as much redundancy as required
- Other stuff (virtualization, building automation, etc.)

Power Usage Effectiveness

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

- $6212.34 for chilled water
- 233,360 KWH * $.1049 /KWH = $24479.46
- 94.5% is cluster and UPS overhead = $23145 cluster operation electric
- $1324.34 in CRAH electric and lighting + $6212.34 in chilled water = $7536.68
- 7535.67/23145 = .33
- PUE = 1.33 (awesome!)
- (1 watt for electricity/compute work and another .33 for heat extraction)
- Sites that mix hot and cold air run in the 1.9-2.5 range (typical for colos – they bill you for their inefficiency)

Density

- How many KW per rack? (conventional colo <= 8KW)
- Hot aisle/cold aisle? Forced air return? Air mush? Inline row cooling? Centralized cooling?
  - Fluid dynamics
- Interconnect limitations? (Infiniband, Myrinet, etc)
- PDU limitations (power distribution)
- Power plant limitations
- Chilled water/cooling plant limitations

More density

- Console serving?
  - (do you need one?)
  - Cables coming out the wazoo
- Blade servers
  - Power per compute unit?
  - + Cabling advantages
- Compute units per $$
  - T2000? Sicortex? Cray 100,000 intel cores?
- Cpu clock speeds
  - Barcelona vs Harpertown (54xx) vs Clovertown (53xx) vs …
    - A note on Conroe/Wolfdale

automation

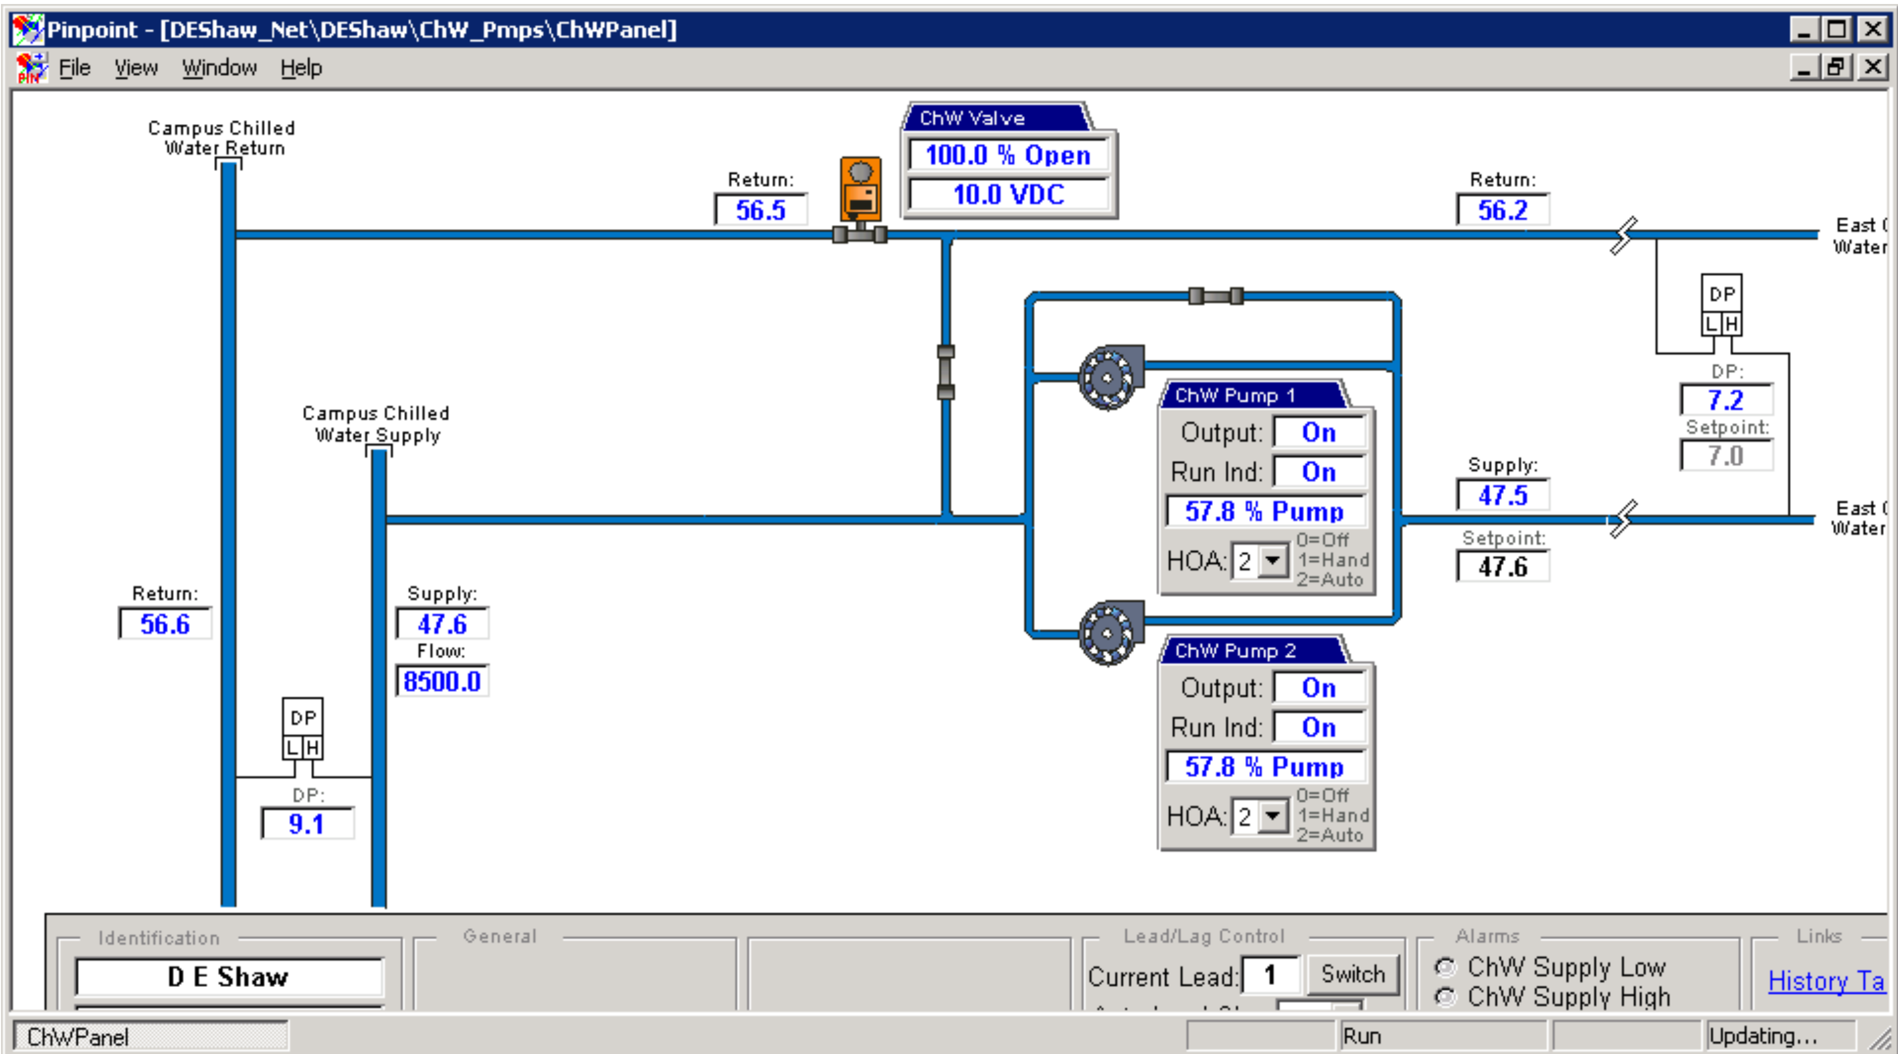

Pinpoint - [DEShaw_Net\DEShaw\ChW_Pmps\ChWPanel]
File View Window Help
Campus Chilled Water Return
Return:
56.5
ChW Valve
100.0 % Open
10.0 VDC
Return:
56.2
East (
Water
DP
L H
DP:
7.2
Setpoint:
7.0
Campus Chilled Water Supply
ChW Pump 1
Output: On
Run Ind: On
57.8 % Pump
HOA: 2
0=Off
1=Hand
2=Auto
Supply:
47.5
Setpoint:
47.6
East (
Water
Return:
56.6
Supply:
47.6
Flow:
8500.0
DP
L H
DP:
9.1
ChW Pump 2
Output: On
Run Ind: On
57.8 % Pump
HOA: 2
0=Off
1=Hand
2=Auto
Identification
D E Shaw
General
Lead/Lag Control
Current Lead: 1
Switch
Alarms
ChW Supply Low
ChW Supply High
Links
History Ta
ChWPanel
Run
Updating...

Problems happen

- Cooling problems
  - Water pressure
  - Water temperature
  - Air temperature
  - Humidity – set humidity sensitivity to 1% on Liebert CRAH
- Engineering problems
  - Drips pans when the electricity goes out – make sure drip pumps are on backup power or chilled water is shutdown upstream during power outage

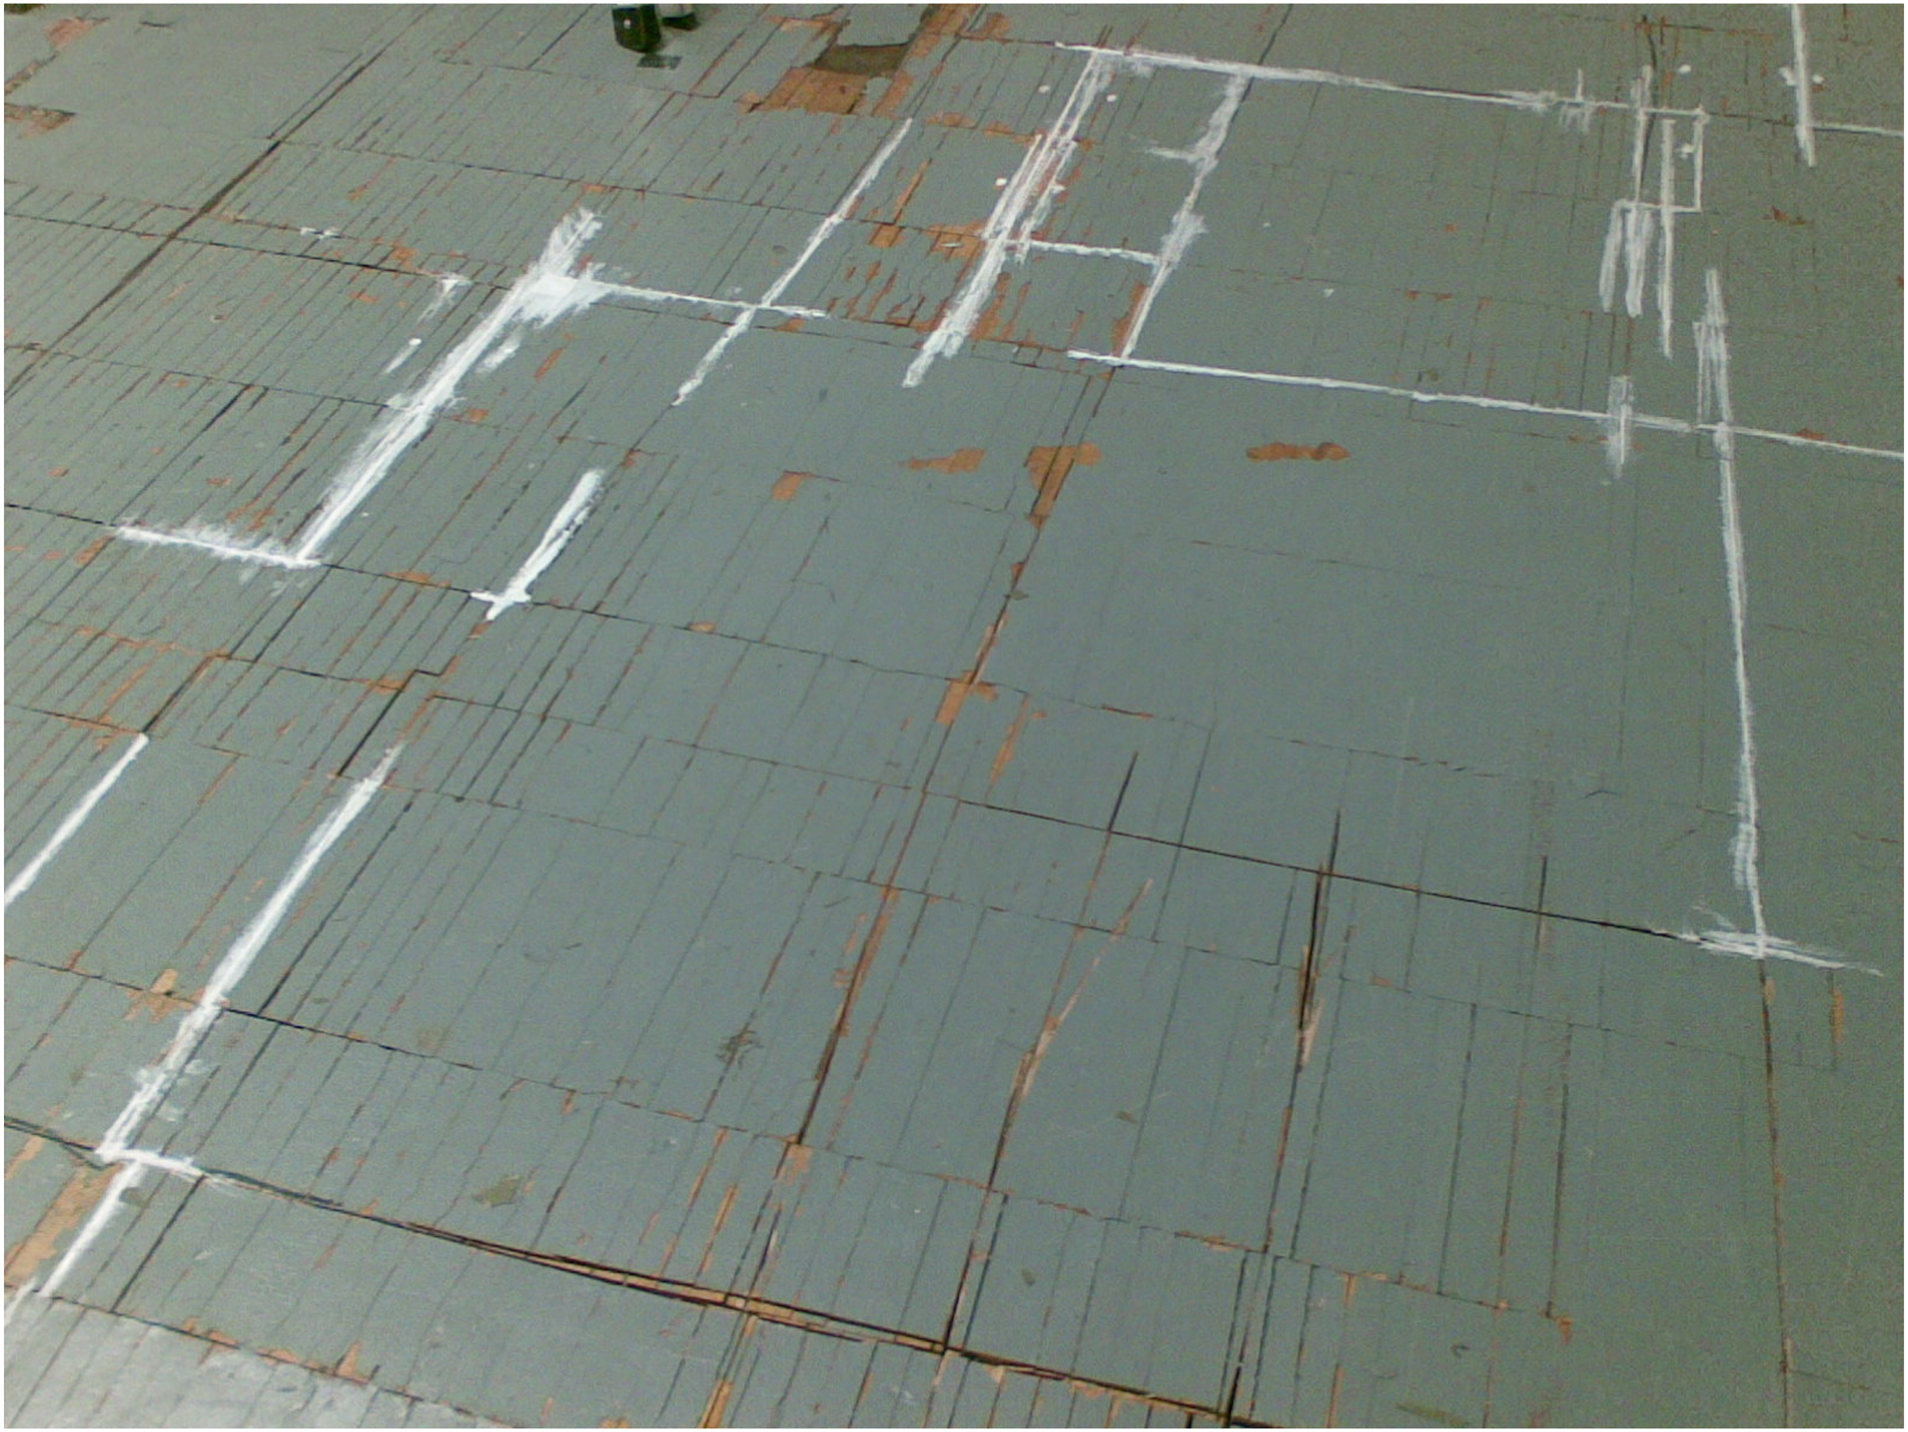

Pressure differential issues lead to heat build up in supply water

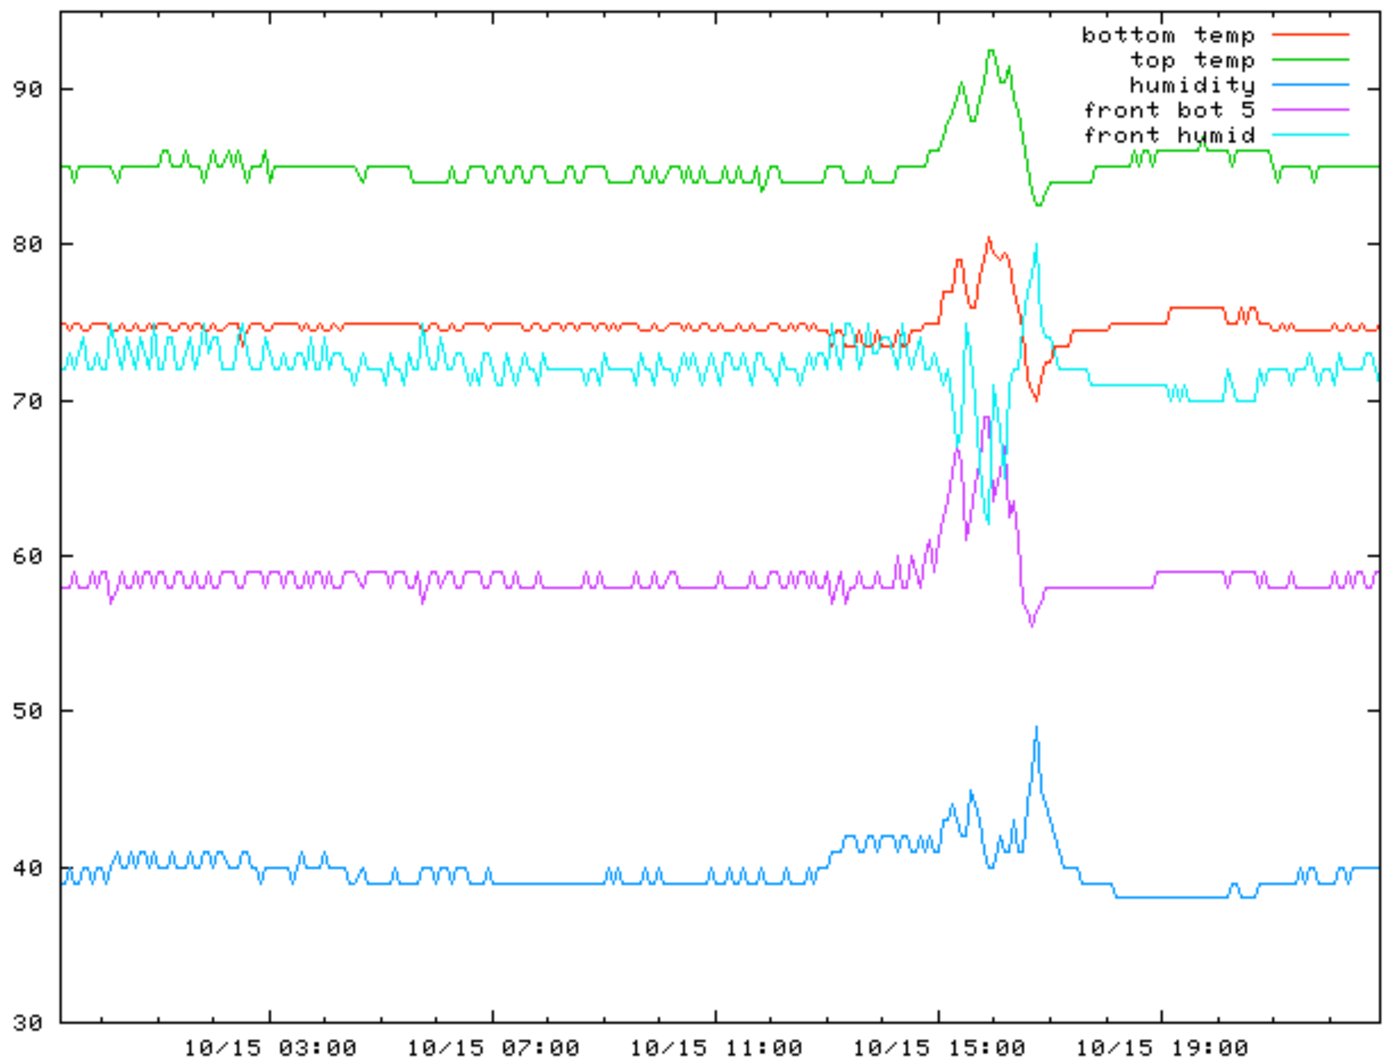

Water supply at 64F

```
# ~/bin/lieberts
liebert2:
System is  on, at 100% capacity, Cool  on Heat off hum off
dehum off
Setpoint: 78F 37%        Actual: 80F 36%

liebert3:
System is  on, at 100% capacity, Cool off Heat off hum off
dehum off
Setpoint: 81F 43%        Actual: 60F 28%

liebert14:
System is  on, at 100% capacity, Cool  on Heat off hum off
dehum off
Setpoint: 77F 45%        Actual: 83F 34%
```

Dehumidification consequences with CRAH

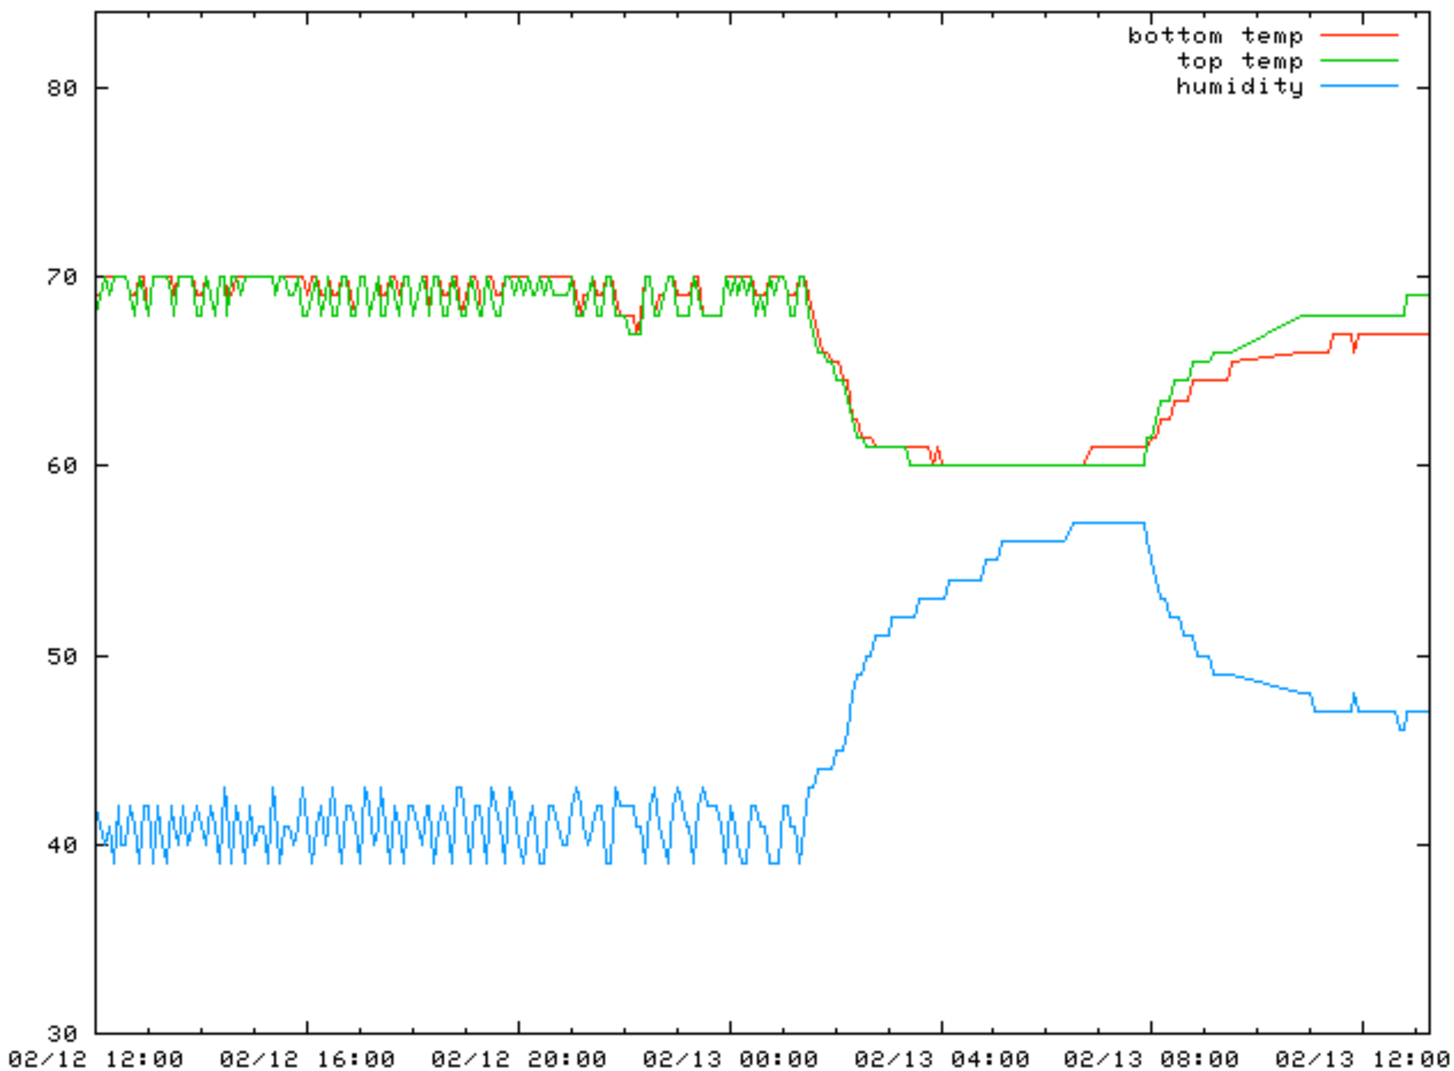

Dehumification CRAH yaw

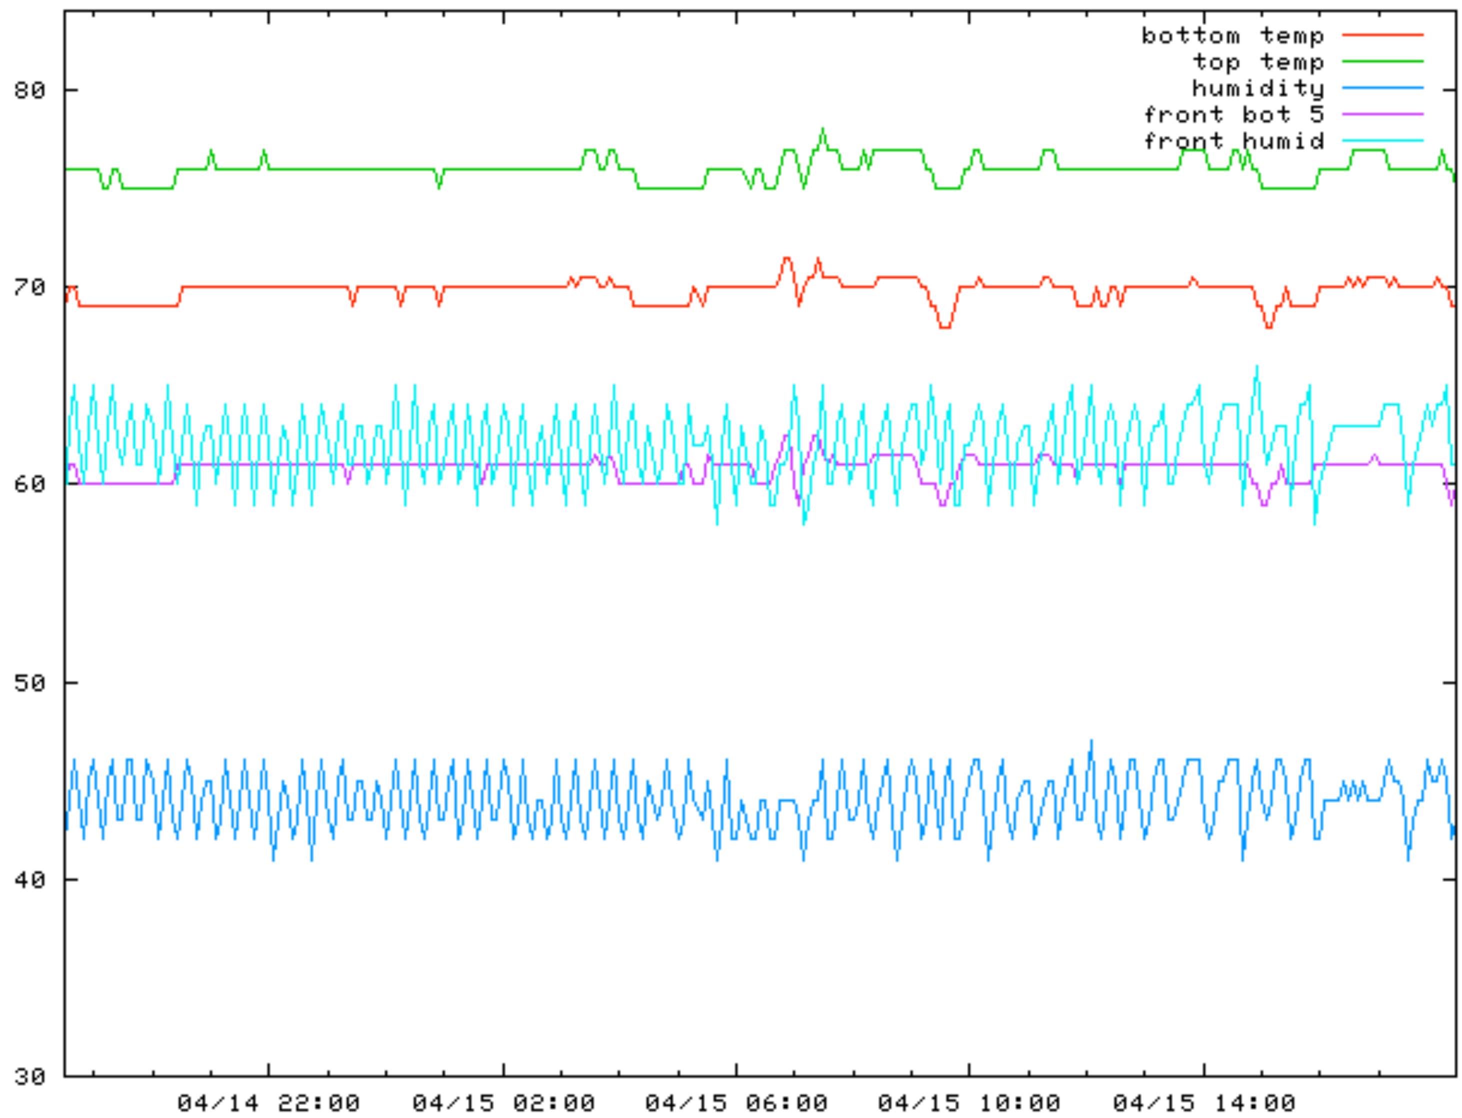

Big yaw (humidity range at +/- 5%)

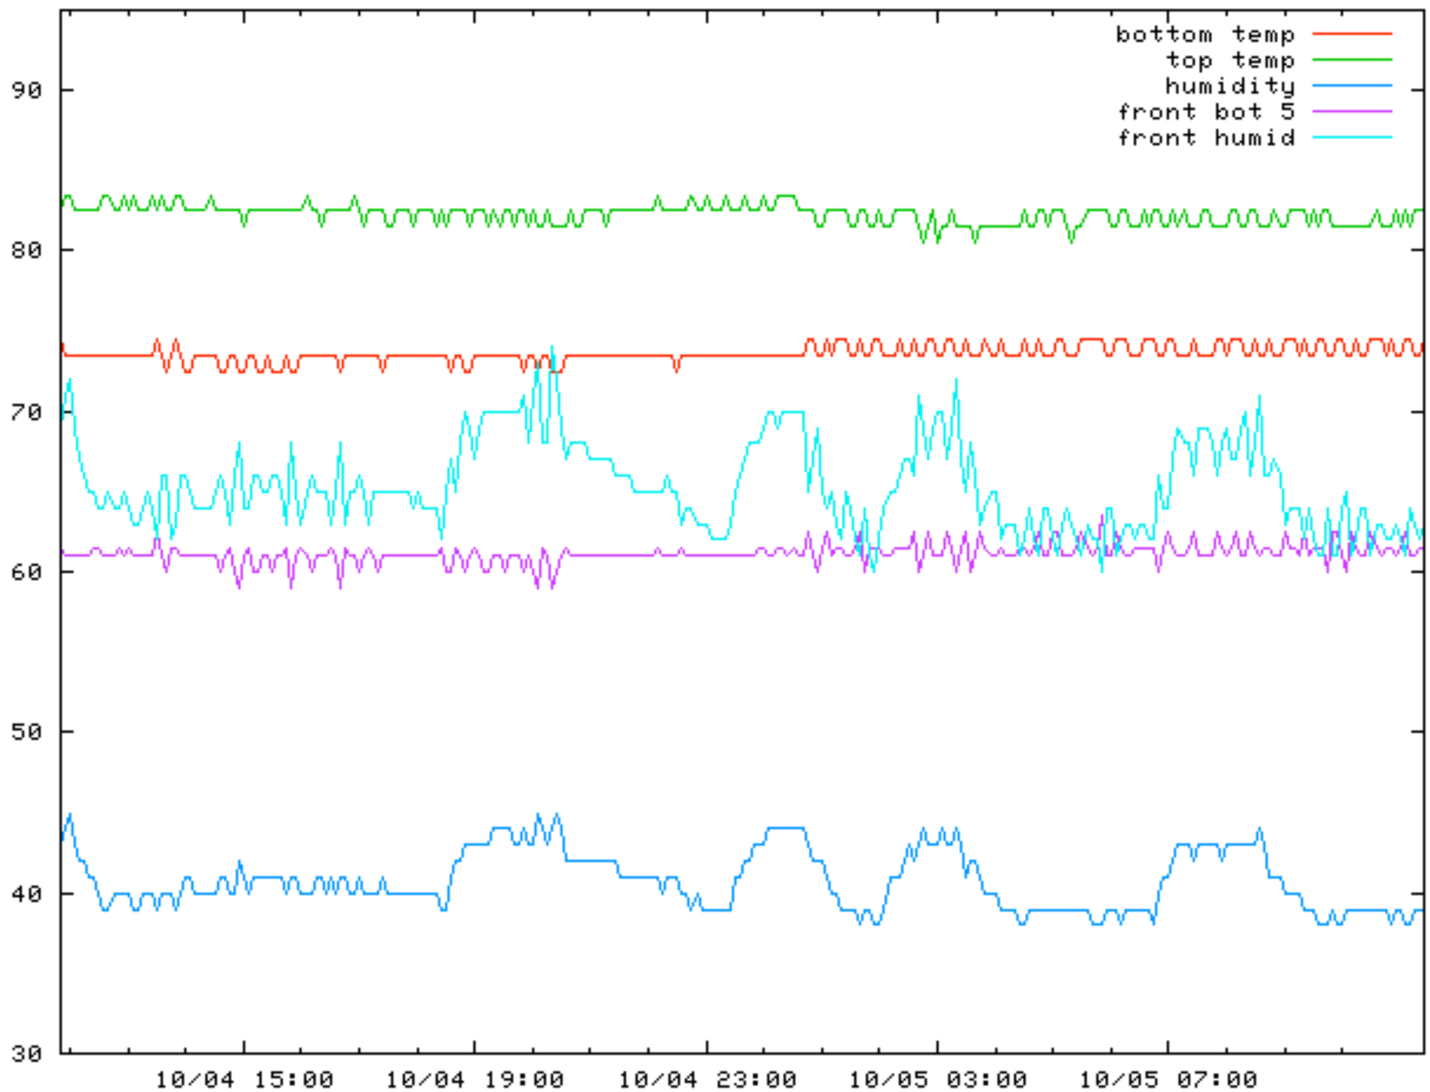

Backdraft and setpoint relativity

```
> ~/bin/lieberts
liebert2:
System is  on, at 10% capacity, Cool  on Heat off hum off dehum off
Setpoint: 78F 45%        Actual: 77F 34%

liebert3:
System is off, at  0% capacity, Cool off Heat off hum off dehum off
Setpoint: 81F 43%        Actual: 53F 24%

liebert14:
System is  on, at 100% capacity, Cool  on Heat off hum off dehum
off
Setpoint: 77F 45%        Actual: 81F 29%
```

Testing additional units

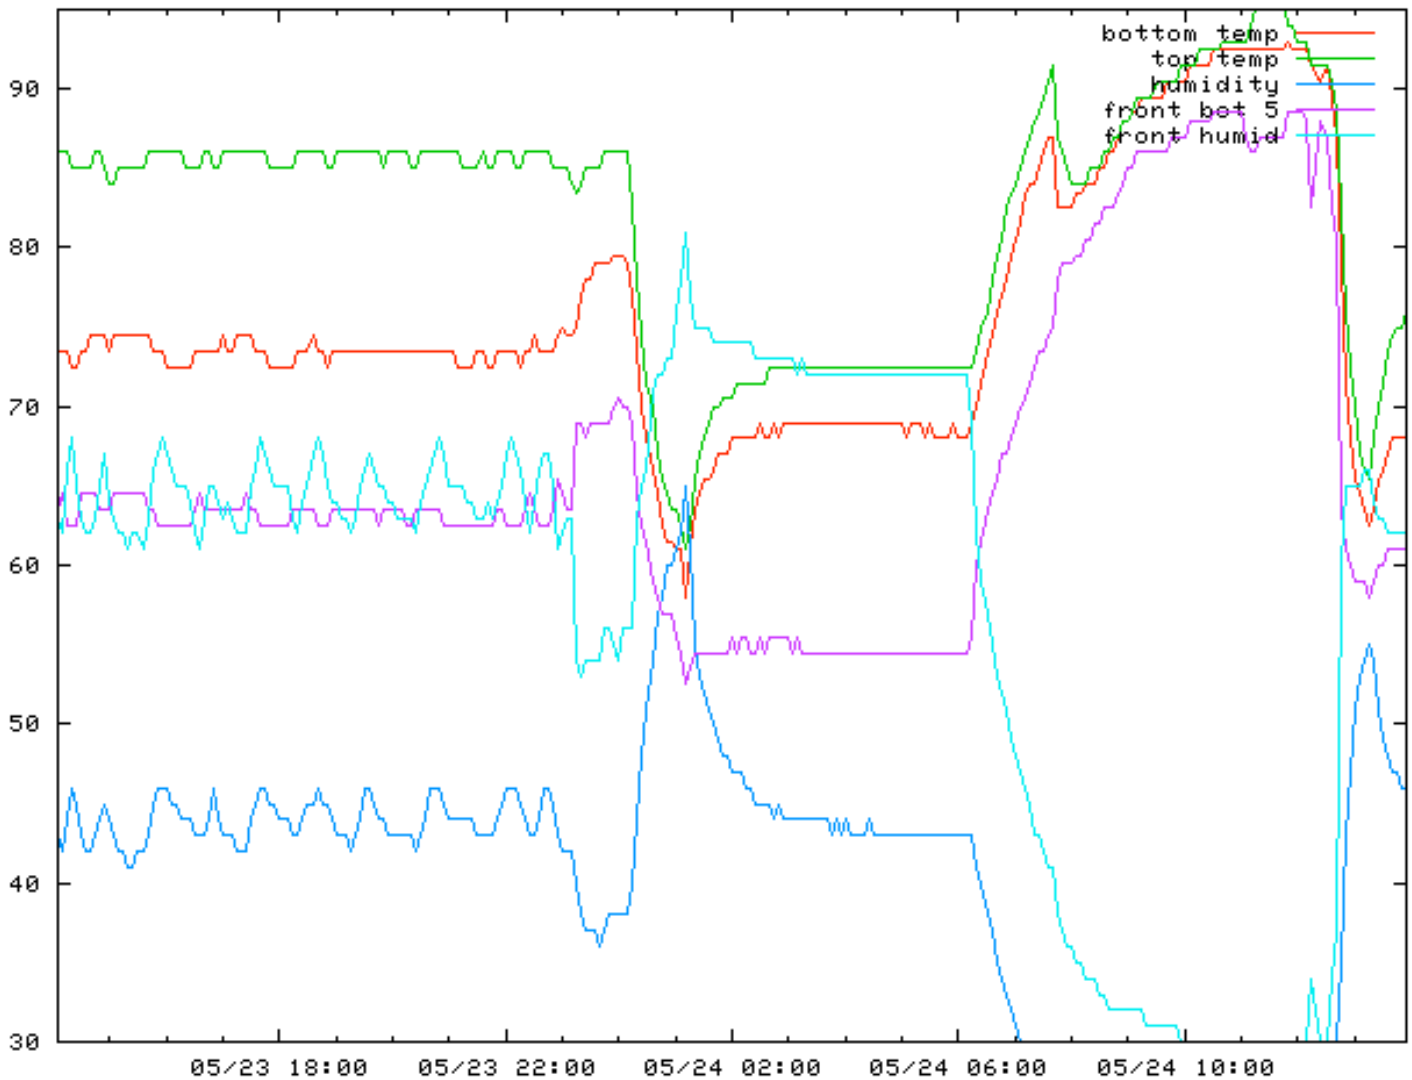

And then the water stopped…

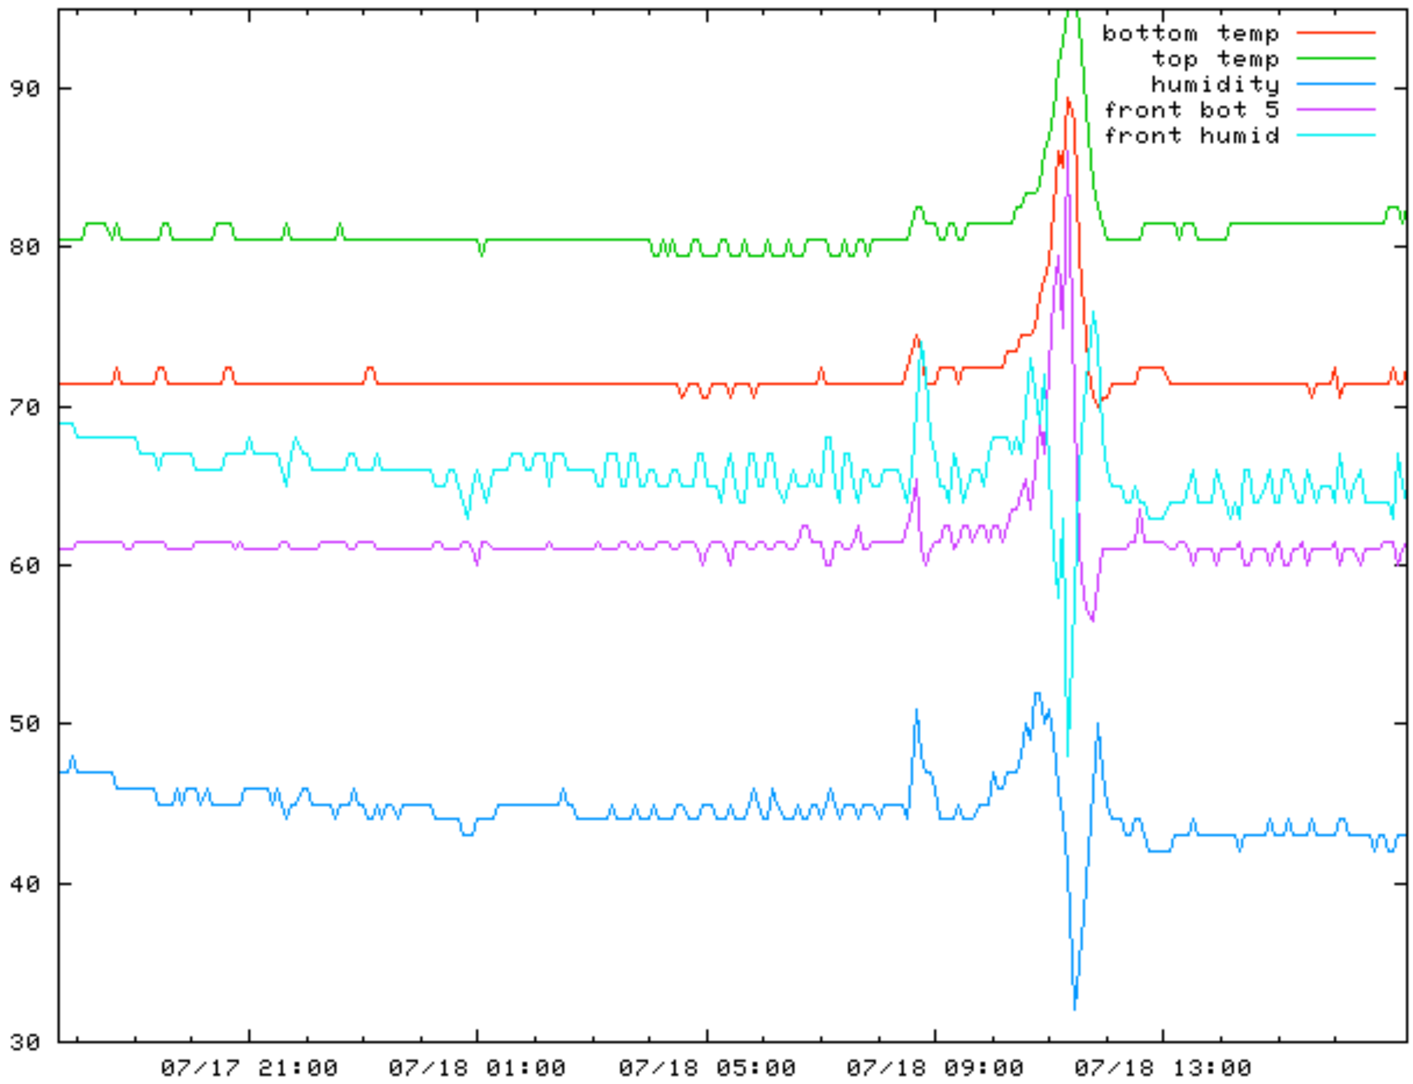

Benchmark benchmark benchmark

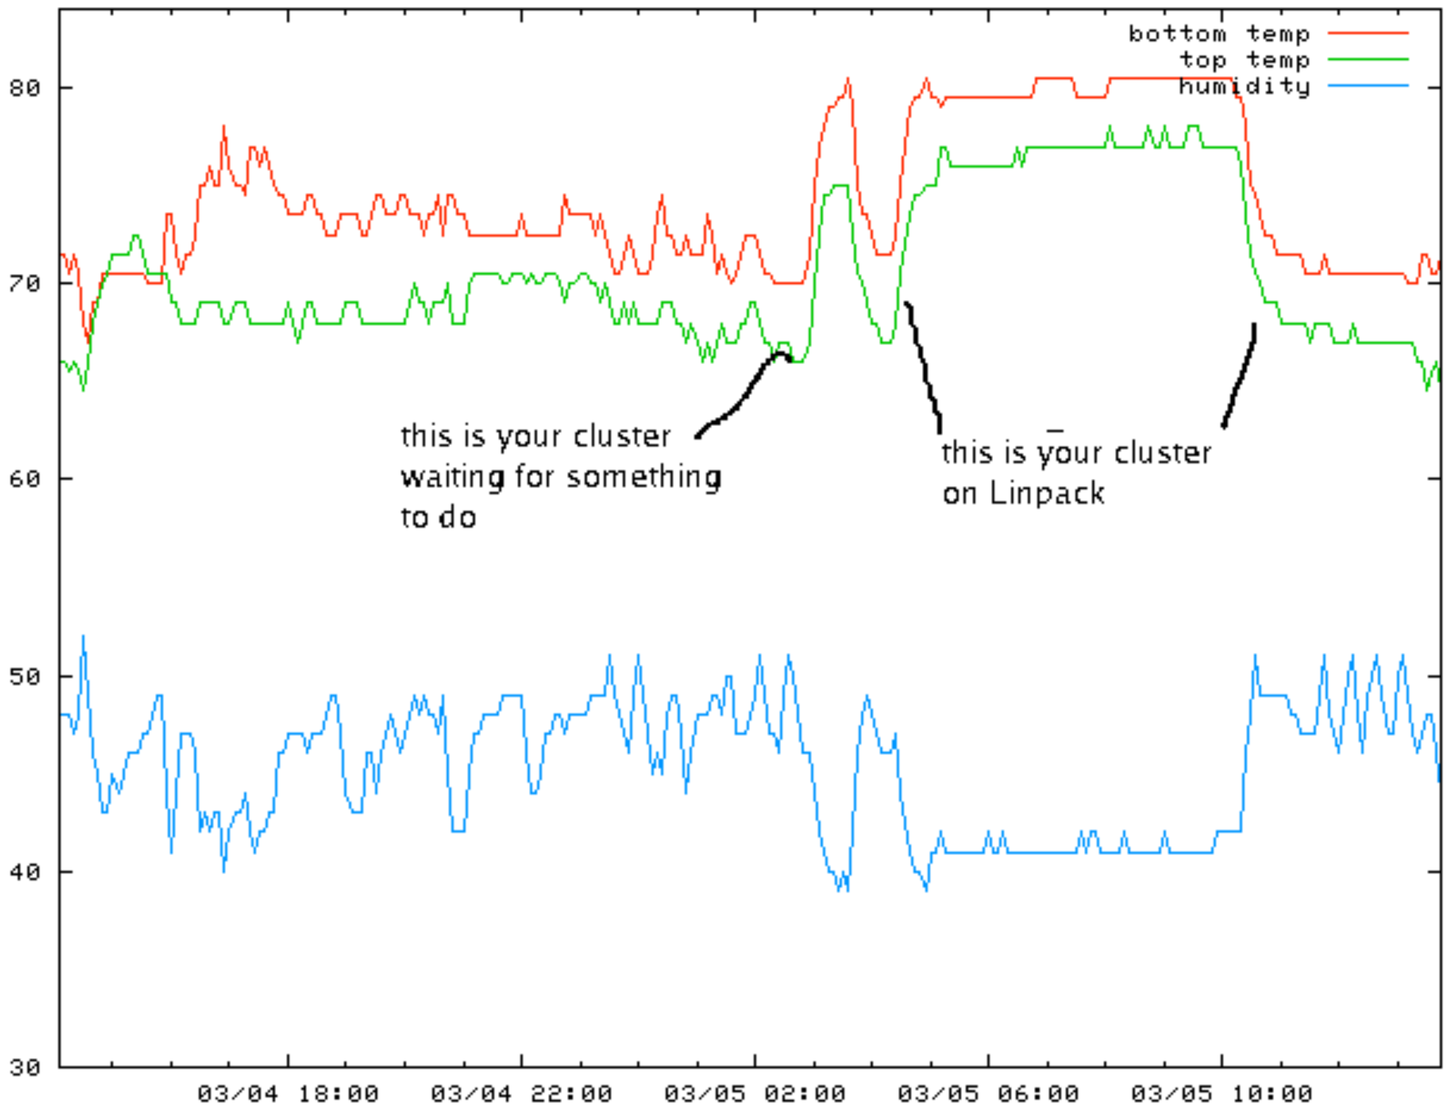

The effect of jobs on electrical load