

---

# Contents

---

<b>Introduction</b> . . . . .	ix
<b>Chapter 1. Systems in Data Centers</b> . . . . .	1
1.1. Servers . . . . .	1
1.2. Storage arrays . . . . .	3
1.3. Data center networking . . . . .	4
1.4. Components . . . . .	5
1.4.1. Central processing unit . . . . .	5
1.4.2. Graphics processing unit . . . . .	7
1.4.3. Volatile memory . . . . .	8
1.4.4. Non-volatile memory . . . . .	10
1.4.5. Non-volatile storage . . . . .	10
1.4.6. Spinning disks and tape storage . . . . .	13
1.4.7. Motherboard . . . . .	15
1.4.8. PCIe I/O cards . . . . .	16
1.4.9. Power supplies . . . . .	17
1.4.10. Fans . . . . .	18
<b>Chapter 2. Cooling Servers</b> . . . . .	19
2.1. Evolution of cooling for mainframe, midrange and distributed computers from the 1960s to 1990s . . . . .	19
2.2. Emergence of cooling for scale out computers from 1990s to 2010s . . . . .	20
2.3. Chassis and rack cooling methods . . . . .	23
2.4. Metrics considered for cooling . . . . .	27
2.4.1. Efficiency . . . . .	27
2.4.2. Reliability cost . . . . .	28
2.4.3. Thermal performance . . . . .	29

---

2.5. Material used for cooling. . . . .	31
2.6. System layout and cooling air flow optimization. . . . .	32
<b>Chapter 3. Cooling the Data Center. . . . .</b>	<b>37</b>
3.1. System cooling technologies used . . . . .	37
3.2. Air-cooled data center . . . . .	38
3.2.1. Conventional air-cooled data center . . . . .	38
3.3. ASHRAE data center cooling standards . . . . .	40
3.3.1. Operation and temperature classes. . . . .	40
3.3.2. Liquid cooling classes . . . . .	41
3.3.3. Server and rack power trend . . . . .	42
3.4. Liquid-cooled racks . . . . .	43
3.5. Liquid-cooled servers . . . . .	46
3.5.1. Water heat capacity. . . . .	46
3.5.2. Thermal conduction module . . . . .	47
3.5.3. Full node heat removal with cold plates . . . . .	48
3.5.4. Modular heat removal with cold plates . . . . .	50
3.5.5. Immersion cooling . . . . .	50
3.5.6. Recent DWC servers . . . . .	51
3.6. Free cooling . . . . .	54
3.7. Waste heat reuse . . . . .	54
3.7.1. Reusing heat as heat . . . . .	54
3.7.2. Transforming heat with adsorption chillers. . . . .	55
<b>Chapter 4. Power Consumption of Servers and Workloads . . . . .</b>	<b>65</b>
4.1. Trends in power consumption for processors . . . . .	65
4.1.1. Moore's and Dennard's laws . . . . .	68
4.1.2. Floating point instructions on Xeon processors . . . . .	72
4.1.3. CPU frequency of instructions on Intel Xeon processors . . . . .	73
4.2. Trends in power consumption for GPUs. . . . .	74
4.2.1. Moore's and Dennard's laws . . . . .	77
4.3. ACPI states . . . . .	78
4.4. The power equation . . . . .	83
<b>Chapter 5. Power and Performance of Workloads . . . . .</b>	<b>87</b>
5.1. Power and performance of workloads . . . . .	87
5.1.1. SKU power and performance variations . . . . .	87
5.1.2. System parameters . . . . .	89
5.1.3. Workloads used . . . . .	92
5.1.4. CPU-bound and memory-bound workloads . . . . .	92
5.1.5. DC node power versus components power. . . . .	93

5.2. Power, thermal and performance on air-cooled servers with Intel Xeon. . . . .	94
5.2.1. Frequency, power and performance of simple SIMD instructions . . . . .	95
5.2.2. Power, thermal and performance behavior of HPL . . . . .	98
5.2.3. Power, thermal and performance behavior of STREAM. . . . .	103
5.2.4. Power, thermal and performance behavior of real workloads . . . . .	107
5.2.5. Power, thermal and frequency differences between CPUs. . . . .	115
5.3. Power, thermal and performance on water-cooled servers with Intel Xeon. . . . .	124
5.3.1. Impact on CPU temperature. . . . .	124
5.3.2. Impact on voltage and frequency. . . . .	125
5.3.3. Impact on power consumption and performance . . . . .	127
5.4. Conclusions on the impact of cooling on power and performance . . . . .	131
<b>Chapter 6. Monitoring and Controlling Power and Performance of Servers and Data Centers . . . . .</b>	<b>133</b>
6.1. Monitoring power and performance of servers . . . . .	133
6.1.1. Sensors and APIs for power and thermal monitoring on servers . . . . .	134
6.1.2. Monitoring performance on servers . . . . .	140
6.2. Modeling power and performance of servers . . . . .	142
6.2.1. Cycle-accurate performance models. . . . .	142
6.2.2. Descriptive models . . . . .	142
6.2.3. Predictive models . . . . .	144
6.3. Software to optimize power and energy of servers. . . . .	149
6.3.1. LoadLeveler job scheduler with energy aware feature . . . . .	150
6.3.2. Energy Aware Runtime (EAR). . . . .	151
6.3.3. Other run time systems to manage power. . . . .	153
6.4. Monitoring, controlling and optimizing the data center . . . . .	154
6.4.1. Monitoring the data center . . . . .	154
6.4.2. Integration of the data center infrastructure with the IT devices . . . . .	156
<b>Chapter 7. PUE, ERE and TCO of Various Cooling Solutions . . . . .</b>	<b>159</b>
7.1. Power usage effectiveness, energy reuse effectiveness and total cost of ownership . . . . .	159
7.1.1. Power usage effectiveness and energy reuse effectiveness . . . . .	159
7.1.2. PUE and free cooling . . . . .	162
7.1.3. ERE and waste heat reuse . . . . .	163
7.2. Examples of data centers PUE and EREs . . . . .	164
7.2.1. NREL Research Support Facility, CO. . . . .	164
7.2.2. Leibnitz Supercomputing data center in Germany . . . . .	166
7.3. Impact of cooling on TCO with no waste heat reuse . . . . .	173

7.3.1. Impact of electricity price on TCO . . . . .	177
7.3.2. Impact of node power on TCO . . . . .	178
7.3.3. Impact of free cooling on TCO . . . . .	181
7.4. Emerging technologies and their impact on TCO . . . . .	183
7.4.1. Waste heat reuse . . . . .	184
7.4.2. Renewable electricity generation. . . . .	189
7.4.3. Storing excess energy for later reuse . . . . .	192
7.4.4. Toward a net-zero energy data center . . . . .	193
<b>Conclusion</b> . . . . .	195
<b>References</b> . . . . .	199
<b>Index</b> . . . . .	209