1 Overview

When we speak of the *architecture* of a system, we refer to how its components are interconnected, assembled, configured, and controlled. The outside of a system may be a flashy, industrial-designed surface with user interfaces, but the inside contains components that are interconnected in such a way as to perform the expected functions for which it was intended. Analogously, the architecture of a building is the design of the essential structure, including beams, walls, floors, and infrastructure, underneath its outer skin. This structure supports the building's functions and the myriad of human's activities as they occupy and traverse its interconnected rooms and hallways. Similarly, the architecture of energy storage affects the flow of energy and matter through a system of interconnected wires and pipes, into and out of vessels or chemical states, while supporting the customer's grid performance and efficiency.

1.1 Architecture Objectives

Ideally, the combination of optimal energy storage technology and architecture will provide the maximum benefit to the customer's grid while maintaining the highest availability and safety, and the minimum amount of lifetime cost for its operators. These are worthy goals, but in reality it is not possible to perfectly achieve every goal with a chosen technology and architecture. Some architectures will enhance a technology's reliability; some will enhance short-term performance, while others may lower the system's costs.

> *When choosing the architecture that is most appropriate for a particular project or market, system designers should decide which aspects are most important for their customers, and then select an architecture that promotes those aspects.*

The following sections describe some common architectures for the fundamental subsystems of energy storage and indicate how they achieve important application attributes, such as reliability, performance, cost-effectiveness, and safety.

1.2 Energy Storage System Subsystems

Energy storage systems (ESS) are comprised of a set of subsystems that delivers electrical power and energy services to a load or an electric grid while simultaneously ensuring proper working conditions and optimal operation of its components. The four fundamental subsystems of an ESS (depicted in Figure 1.1) are energy storage, power conversion, thermal management, and energy management.

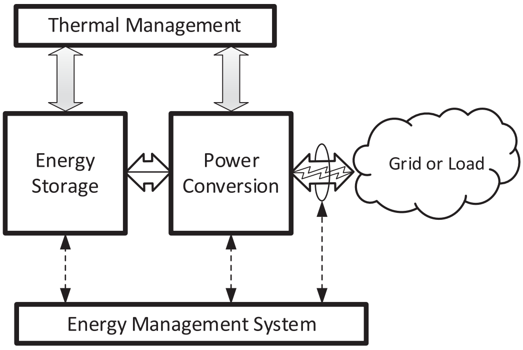

Figure 1.1 A typical arrangement and interaction of the four ESS subsystems

Energy stored in the ESS is converted to useable power by the power conversion subsystem, which also regulates the flow of stored energy to and from the grid or load. The thermal management subsystem maintains optimal operating temperatures for the ESS components by either adding or removing heat from them. Finally, the energy management subsystem monitors the other ESS subsystems and the grid conditions and controls the operation of each subsystem according to a pre-programmed plan that maximizes its benefits for its owner.

1.3 System Efficiency and Losses

Efficiency is an important aspect of energy storage since it affects the economics of the energy storage project. The more energy lost during the storage process, the more generation is required to compensate for this loss; thus, the more costly it will be to operate not only the energy storage, but the grid as well.

Losses that contribute to the reduction of efficiency can be found in every subcomponent of the system. There are conduction losses due to the resistance of every wire, connector, and switching device from the grid all the way through the last battery cell. Further losses in the batteries stem from electrochemical reluctance of ions to move through the various layers within the battery structure, as well as from other chemical reactions that take place during the storage process. The power conversion system (PCS) has its own set of losses, attributed to its switching devices, as it processes the energy going in and out of storage. The transformers (XFMR) have both conduction losses and standby losses associated with the large magnetic core's magnetization. Most of the subcomponents also have some auxiliary power losses from controls, monitoring, cooling, and heating.

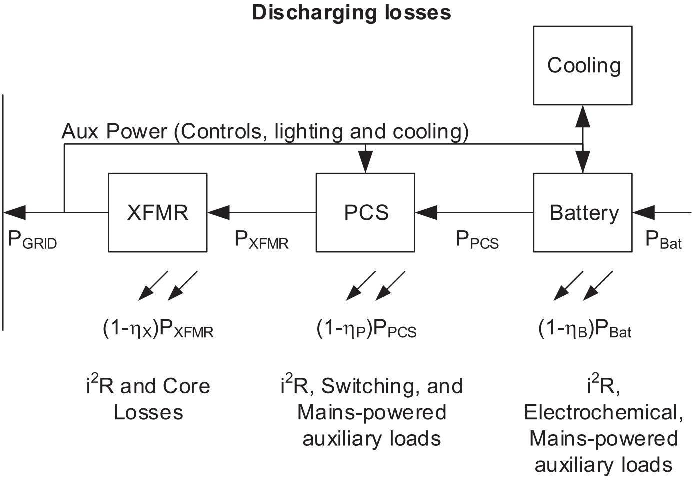

Figure 1.2 Diagram of losses during a discharge of an energy storage system

Figure 1.2 is a useful diagram for calculating the losses during discharge, in each of the elements of an energy storage system.

Looking at the diagram of losses in Figure 1.2, the total losses can be calculated using the following formula:

$$P_{loss} = P_{Bat} \times (1 - \eta_B) + P_{PCS} \times (1 - \eta_P) + P_{XFMR} \times (1 - \eta_X) + P_{Aux}, \tag{1.1}$$

where P_{Bat} is the ideal power coming from inside the battery. This simplifies to:

$$P_{loss} = P_{Grid} \times \frac{1 - \eta_{Total}}{\eta_{Total}} + \frac{P_{Aux}}{\eta_{Total}}, \tag{1.2}$$

where P_{Grid} is the power being absorbed by the grid at the point of interconnection, η_{Total} is the product of η_X, η_P, and η_B, and represents the total one-way discharge efficiency, not including the auxiliary power losses. P_{loss} includes a P_{Aux}/η_{Total} component because P_{Aux} is supplied by the battery during a discharge, and since that power goes sequentially through the battery, PCS, and transformer, the$1/\eta_{Total}$ factor is applied.

When comparing system architectures, it is important to consider the differences in the energy losses associated with each option. For example, an architecture requiring fewer controls and monitoring equipment will generally require less auxiliary power. Architectures with fewer, but larger transformers will be more efficient, because larger transformers are generally more efficient.

Batteries with higher output voltages will require less copper, for a given power level, to achieve similar efficiencies than those with lower output voltages.

2 Energy Storage Architecture

The energy storage subcomponent can be architected in several ways. Typically, the energy storage technology predisposes its architecture. For example, large, bulk energy storage dictates a unitary approach while energy storage made up of many small batteries will lend itself to a multielement parallel architecture.

2.1 Unitary Bulk Storage

Unitary bulk storage is the simplest energy storage architecture, where excess grid energy is stored in one place and is transferred to and from this place through one or more power conversion units. Several technologies employ unitary energy storage architecture. For example, pumped-hydro facilities store mechanical energy in one large water reservoir, such as behind a dam or on a mountain [1] (see Figure 2.1).

In a pumped-hydro facility, a turbine converts the potential energy of water stored at a higher elevation to electrical power when the grid needs it. When electrical power is plentiful and cheap, it powers the turbine to pump water uphill from a lower reservoir to the upper reservoir. Another example of unitary bulk storage is a compressed air energy storage (CAES) facility [2]. Although there are many variants of CAES, fundamentally, energy is stored in the form of

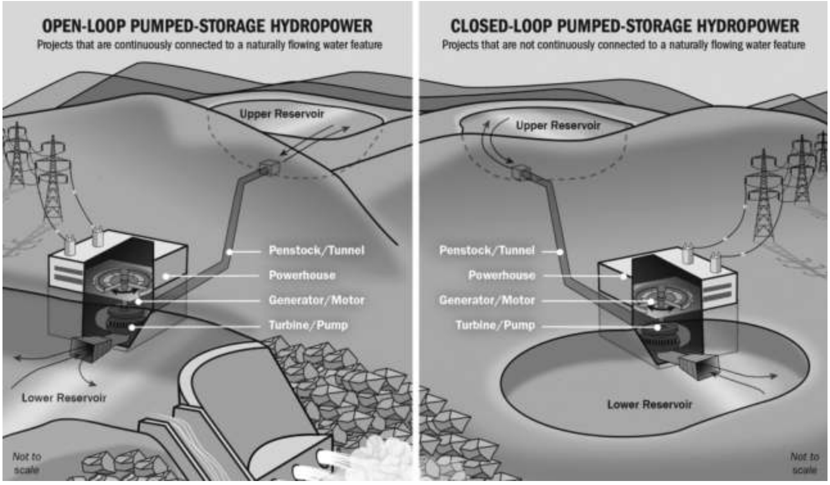

Figure 2.1 Pumped storage facilities showing upper and lower reservoirs connected by a single waterway (Source: US Department of Energy at www.Energy.gov)

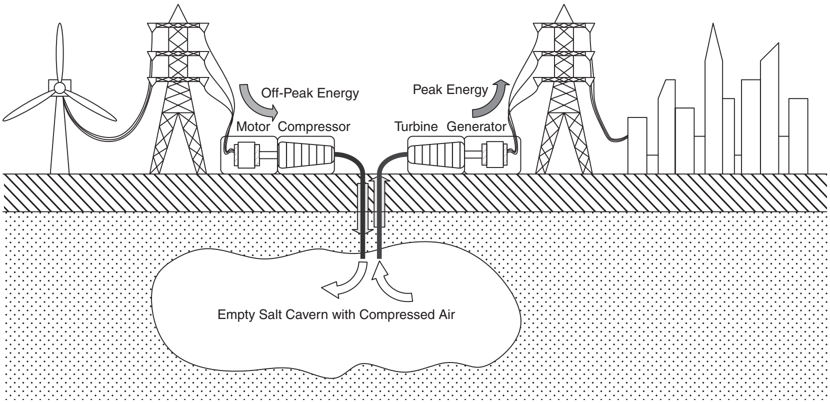

Figure 2.2 Compressed air energy storage system

compressed air in a single large, sealed, sometimes underground chamber (Figure 2.2).

Other examples of unitary energy storage include energy stored in a single tank of liquefied air [3], a tank of chemically reactive material [4], a pool of liquefied salt [5], and even a large pile of concrete blocks [6]. The need for energy storage has been rapidly ramping up and is driving innovation in the industry. Hence, there are many other innovative energy storage means in various stages of research, development, and market readiness.

2.1.1 Unitary Energy Storage with Multiple Power Conversion Paths

If a unitary energy storage facility has only one power conversion path, a failure of a critical component in that path could affect the operation of the whole site. If, for example, the pumped-hydro turbine gets jammed, or the electrical machine throws a bearing, or the electrical transformer overheats, the whole system comes to a screeching halt. Furthermore, prudent, periodic maintenance will force temporary outages throughout the system's service life. This time off detracts from a system's *availability* score.

> *"Availability" is an important characteristic of an energy storage system that defines its ability to operate when it is called upon to do so. The availability score is calculated by dividing the time that the system is fully functional by the time that it is desired to be fully operational.*

$$A = \frac{(Time_{Total} - Time_{Unavailable})}{Time_{Total}} \cdot 100\%. \qquad (2.1)$$

For example, if a single yearly repair on any one of the single-power-path energy storage components takes one month to perform, during which the entire system is inoperable, and the desired system up-time is 8,760 hours per year, the availability score, A, would be:

$$A = \frac{(8,760\ \text{hrs/yr} - 30\ \text{days/yr} \times 24\ \text{hrs/day})}{8,760\ \text{hrs/yr}} = 91.8\%. \quad (2.2)$$

One way to achieve a higher availability score is to configure a system with two or more power conversion units that operate in parallel with each other and sized such that if any one of the components fails or needs service, the other parallel system(s) can take over and maintain full operability. In systems reliability parlance, this is called N+1 redundancy, where N is the number of subsystems required for full operation, and 1 refers to the added subsystem that can take over for any one of the N units that may become inoperable in the system. The value of N is chosen to be whatever is practical and cost effective for the application.

Returning to our pumped-hydro case, four waterways can be cut through the mountain to power four sets of turbines and electrical machines as shown in Figure 2.3. If the total power required by the site is P, then each of the power paths could be rated for a power of $P/3$. In most cases, the power in and out of the site will be distributed evenly among the four power paths, up to $P/4$ for each machine. When one of the machines fails, the power will be distributed among the remaining three power paths, up to $P/3$ on each. If by an even

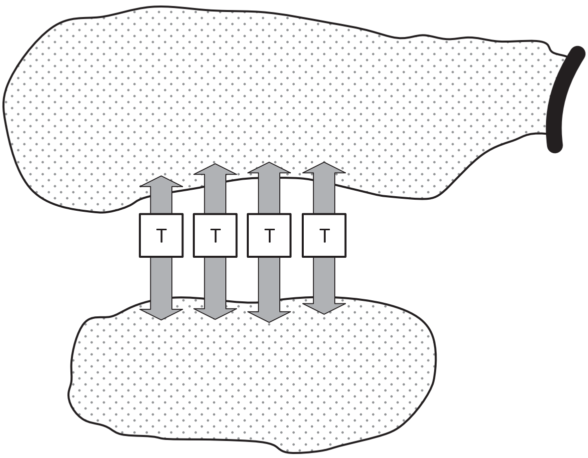

Figure 2.3 Top view of a pumped-hydro facility showing multiple turbines and waterways

smaller probability, another machine fails, the two remaining machines will be left to support the required power. Since they are not rated to handle the full power, *P*, of the site, the site will be temporarily reduced in power capacity by 33% until one of the broken machines is repaired. While a reduction in full-rated power is not ideal, partial power capability is better than no power capability, as would be the case if only one power path were responsible for the entire site's output.

Therefore, a unitary, single-energy storage, multiparallel power path architecture achieves higher operational availability by providing redundancy in its system components. If any one of the power conversion units goes down, the site will still operate at most or all of its power and energy capacity.

> *The system's availability will be a function of both the probability of individual component failure, the number of power conversion paths, and the full rated "promised" power rating of the site.*

2.2 Multielement Energy Storage Architecture

Not all energy storage consists of large unitary reserves of energy like pumped hydro or CAES do. Increasingly, energy storage is assembled from a number of smaller, discrete energy storage devices, such as battery cells [7], electrolyte tanks [8], or flywheels [9].

When energy storage is manufactured in finite and discrete sizes, the units can be combined in a way that not only scales their power and energy, but also maximizes the performance and availability of the whole ESS. We will explore several different approaches in this section, focusing on battery-based energy storage. However, many of the same approaches can be applied to other storage technologies using the same guiding principles.

2.2.1 Battery-Based Energy Storage

Battery-based energy storage uses discrete-sized components, coupled in a way that provides unified energy storage service to the grid or load. Because they are currently available in relatively small, low-voltage packages, battery cells are connected in series and/or parallel to achieve the appropriate voltage, power, and energy required to provide grid-scale storage services. One might ask:

- How are batteries configured together to achieve a system's goals? and,
- How does one determine whether to connect them in series or in parallel, or a combination of both?

Figure 2.4 Examples of energy storage cells (Source: A123 Systems and LG Chem)

***Definitions**: In this discussion, a "battery" is a collection of interconnected cells, while a "cell" is the fundamental electrochemical unit of storage. A cell is manufactured having a voltage and charge capacity, both of which vary depending on technology, manufacturer, and product.*

An electrochemical cell usually has a terminal voltage between 1 and 5 volts, depending on its chemistry, and can store any number of amp-hours, depending on its size. For example, a lead-acid battery cell's nominal voltage is 2 V, while that of nickel–metal hydride and nickel–cadmium are about 1.2 V. Lithium-ion cells such as those shown in Figure 2.4 nominally deliver between 3.2 and 3.7 V depending on which variant of lithium-ion they are.

Parallel or Series

The approximate energy content (in watt-hour, Wh) of any cell is the product of its capacity (amp-hour, Ah) and its average discharge voltage. To achieve enough voltage, power, and energy for grid services, many cells need to be connected in parallel and/or series. A battery with parallel cells will have a capacity equal to the sum of the individual cells' Ah rating. A battery with series-connected cells will have a voltage equal to the sum of the individual cells' voltage. In either case, the battery will have an energy content equal to the sum of all its connected cells, neglecting efficiency and wiring losses.

To optimally configure the series and parallel arrangement of cells, it is necessary to understand the advantages and limitations of both approaches. Connecting all the cells of an energy storage system in parallel would provide a large amount of current, but at an impractically low voltage. Additionally, since wiring and connector losses are proportional to the square of the current, such a battery would have high operating energy losses. Alternatively, wiring cells in series would result in lower battery currents for an equivalent amount of power; however, there are limits to how many cells can be connected in series. Not only will longer strings increase the monitoring complexity, but it is crucial to keep the total string voltage within the safety limits of the intended PCS equipment and the practical limits of the factory and field handling processes.

Cell Balancing

For both parallel and series configuration, care must be taken to ensure each of the cells shares the system operating power in a manner that achieves maximum benefit with minimal operating cost. In most cases, this means that each cell should charge and discharge at the same *relative* rate with respect to each other, achieving a balanced state of charge among all cells.

When connecting cells in parallel, it is important to ensure that the current entering the parallel bus structure between cells is evenly distributed to each of the cells. This is achieved by careful design of the metal conductors and temperature control of the battery pack. Finite element analysis (FEA) tools can help simulate the direction and volume of current and heat in battery packs with complex geometries, to ensure balanced conditions under all operating modes.

A balanced state of charge will typically exist among cells in series, since the current through each cell is substantially the same. However, over time, there may be some drift caused by slightly mismatched self-discharge rates and charge efficiencies among cells. If this drift is left uncorrected, eventually it will cause enough of an imbalance in the series string that the total capacity of the string will be noticeably reduced. Since the charging process ends when the cell with the highest state of charge (SOC) in the string is fully charged, and the discharging process ends when the cell with the lowest SOC in the string is fully discharged, the maximum achievable capacity of a series string is determined by:

$$Capacity_{achievable} = Capacity_{ideal} \times [1 - (SOC_{\max} - SOC_{\min})], \qquad (2.3)$$

where $SOC_{\max}$ and $SOC_{\max}$ are numbers between 0 and 1 representing the SOC of the highest and lowest SOC cells in a series string.

To mitigate imbalanced SOC in series-connected strings, many system designs incorporate an electronically controlled cell-balancing mechanism that, over time, gradually corrects for the mismatched self-discharge among the cells.

Sizing Example – Building a Submodule

To demonstrate how a larger energy storage device can be built from smaller cells, let us consider an example of a sizing exercise for a lithium-ion battery. Suppose a cell has an operating voltage range of 3 to 4.15 V and a capacity of 27 Ah. If the nominal cell voltage is 3.7 V, there will be roughly 27 Ah x 3.7 V = 100 Wh of energy in each cell. Larger power and energy can be achieved if we connect a set of

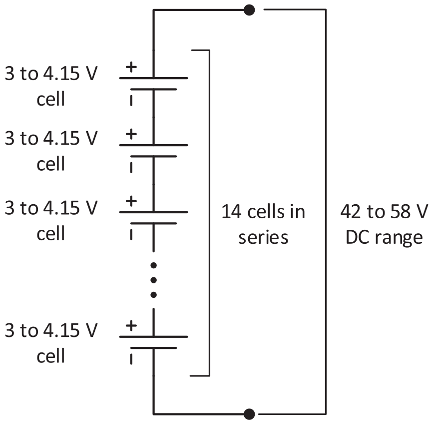

Figure 2.5 Example of a series configuration of battery cells yielding an energy rating of 1400 Wh

these cells in series. For example, as in Figure 2.5, 14 of our 3.7 V cells connected in series will yield an energy rating of 1400 Wh at a nominal voltage of 52 V.

Definition: *The arrangement of cells in a parallel or series combination is often called a submodule.*

It is important that the cell capacity of each of the series cells be matched as close as possible to each other. Without extraordinary and potentially costly energy-rebalancing means, the smallest capacity series cell will limit the useable energy of the whole string.

Selecting the Battery's Voltage

For the previous example, we chose 52 V as a nominal voltage, but the practical application of energy storage usually dictates the required working voltage of the battery. In general, the working voltage of a battery should be set to match the operating voltage range of the power equipment to which it will connect.

Furthermore, it is helpful to consider the entire operating voltage range of the cells in a real-world application. While the cells are charging and discharging, their terminal voltages vary as a function of their stored energy and rate of change of energy as well as temperature, age, and history. Therefore, any load or power supply to which they connect must be able to operate optimally under this full voltage range. If not, the system designer should adjust the number of cells connected in series to accommodate the connected system's voltage range.