



UNIVERSITÀ
DEGLI STUDI
DI PADOVA

UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Ingegneria Industriale DII

Corso di Laurea Magistrale in Ingegneria dell'Energia Elettrica

THE IMPACT OF END PLATES ON THE UNIFORMITY OF ELECTRODE
COMPRESSION IN FLOW BATTERY STACKS FOR REDUCING OHMIC
LOSSES DUE TO CONTACT RESISTANCES

Relatore: Prof. Andrea Trovò

Correlatore: Ing. Matteo Rugna

Laureando:

Leonardo Magliani

2103352

Anno Accademico 2023/2024

Vanadium vanadiorum, et omnia vanadium.

ABSTRACT

Technological development plays a crucial role in the successful realization of large Flow Battery (FB) systems and can significantly contribute to reducing capital costs. FBs represent a sophisticated technology necessitating complicated engineering. In this work a numerical modelling of the mechanical behaviour of Vanadium Flow Battery (VFB) stack with COMSOL Multiphysics® has been performed. During VFB stack assembly, it is crucial to apply adequate external pressure to properly seal the cells and prevent leaks, ensuring optimal electrical performance of the battery. Achieving homogeneous compression of the stack helps reduce ohmic losses by establishing effective contact between the cell components and minimizing contact resistance. In particular, this study examines the influence of end plate geometry and bolt configuration on the distribution of contact pressure within an industrial scale VFB stack. The pressure is applied using metal end plates secured with bolts. The thickness of the end plates, along with the number and placement of bolts, is optimised to achieve uniform contact pressure across the active area. This research offers a framework for the optimal mechanical design of VFB stacks. More generally, the same analysis can be applied to all Fuel Cell (FC) reactors utilising the stack configuration outlined below.

KEY WORDS

Vanadium Redox Flow Battery, Design Considerations of Vanadium Redox Flow Battery, End Plate, Finite Element Analysis, Numerical Optimisation Techniques

CONTENTS

List of Figures	ix
List of Tables	xii
List of symbols and acronyms	xii
Nomenclature	xiii
1. Introduction	1
a. The need of stationary energy storage	1
b. Vanadium redox flow batteries	2
i. General overview	2
ii. Structure and functioning of a VFB	4
iii. State of the art of VFBS	6
c. Alternative chemistries	7
i. All-iron flow batteries	8
ii. Organic flow batteries	9
2. Performances and losses in a flow battery	11
a. The hydraulic design and the electrolyte distribution	11
b. The electric design	14
c. Internal resistance losses	17
i. Activation overpotential	17
ii. Ohmic voltage drop	18
iii. Concentration overpotential	19
d. Shunt current losses	20
e. Cross-over losses	22
f. Hydraulic losses	23
g. Impact of the electrode compression	25
h. Importance of the mechanical design on the electrode compression and hydraulic sealing	47
3. Mechanical and fluidodynamic simulations of a single laboratory cell VBF	53
a. Structure of the simplified simulated cell	53
b. Setting up the simulation on COMSOL Multiphysics®. Description of the process	55
i. Creation of the geometry	55
ii. Assignment of materials	63
iii. Physics setting	65
1. Description of the mathematical models	65
2. Description of the boundary conditions	68
iv. Setting the mesh	69
v. Setting the solver	78
vi. Analysis of the results	80
c. The hydraulic simulation of the cell	84
d. The mechanical simulation of the laboratory cell	91
4. The optimisation	101
a. The EP thickness	101
b. The number of bolts	104
c. Aluminium and steel	106

d. The best choice	108
5. Conclusions	111
6. Bibliography	113

LIST OF FIGURES

Figure 1.1: The curve of the day-to-day energy demand.	1
Figure 1.2: The cost shares of a VFB, considering both a) the investments costs and b) the shares of total life cycle cost, in the cases of an industrial implant with either large or small cell, and of a residential implant.	2
Figure 1.3: Which materials are considered critical by the EU criticality assessment. Two thresholds are to be met: the economic importance score, and the supply risk score. Vanadium is slightly above both of them, thus is critical.	4
Figure 1.4: A cell of the stack of a VFB. Two consecutive electrodes are separated by an ion-exchange membrane and fed by the tanks hosting the electrolytes.	5
Figure 1.5: The losses in the $9kW/27kWh$ industrial scale VFB of the University of Padua for different current values. The addendum of the ancillary services is further analysed in a dedicated pie chart.	7
Figure 1.6: Schematic representation of an all-iron battery. It is conceptually very close to a VFB.	8
Figure 1.7: Possible molecules to be used as catholyte (blue) and anolyte (orange) in organic redox flow batteries.	10
Figure 2.1: The carbon felt microstructure: it is possible to see the carbon strains and the empty spaces amongst them: the pores.	11
Figure 2.2: The electrode structure and the electrolyte velocity in three different compression cases.	12
Figure 2.3: The invasion pattern of the electrolyte inside the electrode. The colours indicate which pores are progressively flooded: blue first, white last.	13
Figure 2.4: The various channel configurations for the flow-by (serpentine, parallel, interdigitated) and the flow-through case.	14
Figure 2.5: The polarisation curve during discharge.	16
Figure 2.6: The concentration overpotential VS the current.	20
Figure 2.7: The equivalent electric circuit of the stack, depicting the shunt currents paths.	21
Figure 2.8: a) The friction factor versus the Reynolds number, with magnifications focusing on b) low and c) high Re .	24
Figure 2.9: The relationship between the compression ratio and the compressive pressure: it is not linear.	26
Figure 2.10: The distribution of the pore diameters for different levels of compression: a) the frequency and b) the normalised frequency.	26
Figure 2.11: The diameters of the pores for different values of compression: a) their total volume, b) their average diameter, c) the maximum diameter.	27
Figure 2.12: The felt throughout the compression: the impact on the porous structure is visible.	28
Figure 2.13: An invasion pattern of the electrolyte in the electrode for different compression ratios.	29
Figure 2.14: Semilogarithmic plot of the relationship of the permeability κ with the porosity ε .	30
Figure 2.15: Dependence of the permeability κ of two electrodes felts (R5 and P5) on the compression ratio CR .	31
Figure 2.16: The pressure drop in the electrode VS the flow rate, for various compression ratios. The dashed lines represent a wide AA, the continuous lines a smaller AA.	33
Figure 2.17: The pressure drop VS the current density for various electrode thicknesses, i.e. compression ratios.	34
Figure 2.18: The pumping power versus the compression ratio for two electrode felts: P5 and R5.	35
Figure 2.19: The dependency of the pumping power on the Reynold's number for different compression ratios.	35
Figure 2.20: The electrical conductivity of the electrode versus its thickness, reporting experimental and expected data.	36
Figure 2.21: The electrical resistivity ρ versus the compression ratio CR .	37
Figure 2.22: The electric resistance of the electrode versus the compression.	38
Figure 2.23: The ASR and η_{act} VS the electrode thickness for various specific surface areas α_0 .	39
Figure 2.24: The anolyte and catholyte concentrations throughout charge and discharge, versus time, for	

two values of compression ratio. b) has ten times the diffusivity constant of a).	40
Figure 2.25: The voltages VS capacity during charge and discharge for different compression ratios, and for a)-d) rising current densities and flow rates values.	42
Figure 2.26: The cell potential versus time during charge and discharge for different current densities and compression ratios: a) 0%, b) 10%, c) 20%, d) 30%.	43
Figure 2.27: The intrusion of the carbon strains of the felt in the channels due to compression.	44
Figure 2.28: The cross-sectional area of the channel and the area of the intrusion.	45
Figure 2.29: The relationship between the compression ratio CR and the intrusion ratio λ .	45
Figure 2.30: The porosity of the three regions versus the compression ratio.	46
Figure 2.31: The electrode performances a) before and b) after the compression.	46
Figure 2.32: The stress-strain curve of Al 6061-T651 for different operating temperatures.	48
Figure 2.33: The stress-strain curve of EPDM, for different strain rates.	49
Figure 2.34: Stress-strain curve of the graphite felt of the electrodes.	50
Figure 2.35: The stress-strain curve for PVC.	51
Figure 2.36: The stress-strain curve of graphite sheets.	51
Figure 3.1: The explosion of the cell for its mechanical simulation, drawn in Comsol.	53
Figure 3.2: The bottom cell of the stack: only two out of the four electrolyte manifolds reach the bottom electrode.	54
Figure 3.3: The rib following the internal parameter along the AA of the FF. Another one is present on the bottom side of the FF.	55
Figure 3.4: a) The external structure of Gk1 first, b) the drawing of the AA and of the manifolds too, c) the final result after the Boolean subtraction.	56
Figure 3.5: The subtraction setting concerning the dominions.	57
Figure 3.6: Example for the use of the Boolean difference settings.	57
Figure 3.7: More subtraction operations can be necessary to describe particular cases.	58
Figure 3.8: The geometry directory so far.	59
Figure 3.9: a) Gk1 and FF1 with ribs, b) their magnification from the side, to highlight the ribs.	59
Figure 3.10: The whole cell seen a) from the top, b) from the bottom.	60
Figure 3.11: The mesh of the same object when set as a) a union, b) an assembly.	61
Figure 3.12: The effect of a) not selecting the Create imprints command, and b) of selecting it.	61
Figure 3.13: The difference between the application of a) identity pairs, rather than b) contact pairs. The top parallelepiped is pushed to the left, the bottom one to the right.	62
Figure 3.14: The impact on the mesh given by selecting a) an assembly with identity pairs, rather than b) a union.	63
Figure 3.15: The evaluation of the density of a steel as function of its temperature in Comsol.	63
Figure 3.16: The EPDM structure.	64
Figure 3.17: The molecular structure of PVC.	64
Figure 3.18: The cell cut with a plane in one of its symmetry directions.	67
Figure 3.19: Where the pressure is applied on the BP, in the case of the single cell simulation.	69
Figure 3.20: The fourth portion of the cell, cut along the symmetry planes.	70
Figure 3.21: The boundaries where the symmetry setting is applied.	70
Figure 3.22: The cell with ribs and AA projected in the gaskets and BPs.	71
Figure 3.23: The use of the work planes to cut the cell and describe the perimeters.	72
Figure 3.24: The external perimeter of the cell partitioned with a work plane.	72
Figure 3.25: The cell after all the partitions.	73
Figure 3.26: The mapping of the upper boundary of the inner corner of the inner perimeter.	74
Figure 3.27: The exponential distribution along an edge of the mapped boundary, a) when only one edge is selected, b) when both of them are.	74
Figure 3.28: The extrusion of a FF and its ribs, and of the bottom BP.	75
Figure 3.29: The extruded a) external and b) internal perimeter of the cell.	75

Figure 3.30: A normal free tetrahedral mesh around the manifold, and an extremely coarse one in the remaining domains.	76
Figure 3.31: A very coarse mesh around the manifold.	76
Figure 3.32: The mapping of the circumference of the manifold can be compatible with an extremely coarse mesh around it.	77
Figure 3.33: The complete mesh.	77
Figure 3.34: The convergence of a three-step segregated solver.	80
Figure 3.35: The view of a) a cut plane, b) eight cut lines, from the dataset of the study.	81
Figure 3.36: The view of the volume results in a 3D plot group.	81
Figure 3.37: The view of the surface results in a 3D plot group.	82
Figure 3.38: The view of the line results in a 3D plot group.	82
Figure 3.39: The view of the surface results in a 2D plot group.	83
Figure 3.40: The view of the line results in a 1D plot group.	83
Figure 3.41: The structure of the simple single cell.	84
Figure 3.42: The mesh of the single simple cell.	85
Figure 3.43: The result of the fluidodynamic simulation: the pressures in a) the catholyte and b) anolyte case.	85
Figure 3.44: The convergence of the mesh sensitivity analysis.	86
Figure 3.45: The result of the fluidodynamic simulation: the velocities in a) the catholyte and b) anolyte case.	86
Figure 3.46: The structure of the greater single cell.	87
Figure 3.47: The result of the fluidodynamic simulation of the greater the cell: the velocities in a) the catholyte and b) anolyte case.	88
Figure 3.48: The result of the fluidodynamic simulation of the greater the cell: the pressures in a) the catholyte and b) anolyte case.	88
Figure 3.49: The relationship between total pressure drop and flow rate.	89
Figure 3.50: The structure of a simple single cell in the flow-by interdigitated case.	90
Figure 3.51: The pressure in the flow-by configuration: a) the catholyte case with b) its magnification, and c) the anolyte case with d) its magnification.	90
Figure 3.52: The a) catholyte velocity in the flow-by configuration, with a b) magnification of its jumps from one channel to another.	91
Figure 3.53: The structure of the cell for the ultimate mechanical simulation.	91
Figure 3.54: The washers on the EP.	92
Figure 3.55: The mesh at the contact pair between the FF rib and the BP.	95
Figure 3.56: The a) -x, b) y and c) z displacements of the electrode for a lighter load.	96
Figure 3.57: The symmetry boundary over the top surface of the component (without the electrode).	97
Figure 3.58: The vertical displacement of the top surface of the FF.	97
Figure 3.59: The simulation results: comp1.w over the contact surface between the electrode and the BP.	98
Figure 3.60: The lines along which the mechanical pressure is to be checked to ensure the hydraulic sealing is achieved: a) and b) between the BP and the FF rib, c) and d) between the FF rib and the gasket, and between the gasket and the symmetry plane.	98
Figure 4.1: The vertical displacement over the electrode base for various EP thicknesses.	102
Figure 4.2: The hyperbolic fit of the vertical displacements.	103
Figure 4.3: The vertical displacement over the bottom electrode surface for various EP thicknesses: a) 5 [mm], b) 10 [mm], c) 15 [mm], d) 20 [mm], e) 25 [mm], f) 30 [mm].	103
Figure 4.4: The component with a) 4 bolts, b) 6 bolts and c) 8 bolts.	104
Figure 4.5: The simulation results of Table 4.2. The various lines represent the different bolts numbers considered.	105
Figure 4.6: The comp1. w evaluated for 5 [mm] with a) 4 bolts, b) 6 bolts and c) 8 bolts.	106
Figure 4.7: The studies' results and their fittings for EPs in a) steel and b) aluminium.	107

Figure 4.8: The comparison between the results for steel (in red) and aluminium (in blue) for the various bolts number: a) four, b) six, c) eight.	107
Figure 4.9: The application of Figure 4.7 for the choice of EP_z and number of bolts using steel in the flow-by design.	109

LIST OF TABLES

Table 1.1: The oxidation states of vanadium.	4
Table 1.2: The estimated price of the various components of an all-iron flow battery.	9
Table 2.1: The values of the parameters for the evaluation of the activation polarisation.	18
Table 2.2: The conductivities of the vanadium species.	22
Table 2.3: The permeability of the P5 and R5 felts and the electrolyte viscosity for different temperatures.	31
Table 2.4: Activation overpotential η and ohmic losses IR for different electrode thicknesses d .	38
Table 2.5: The effect of the compression ratio on the discharge capacity, on the efficiencies and on the cell resistance at different current densities.	41
Table 2.6: The stack components and their materials.	48
Table 3.1: The spatial dimensions of the components of the cell.	54
Table 3.2: The spatial dimensions of the ribs of the FF.	55
Table 3.3: The enhanced mechanical properties of the Al6061-T651 of the EPs, with a comparison with Al6061-O.	93
Table 3.4: The different equations that describe an elastic or a hyperelastic material in Comsol.	94
Table 3.5: The maximum Von Mises stress in some of the stack components and their yield strengths.	99
Table 4.1: The minimum, average and maximum vertical displacement for various EP height, and 8 bolts.	101
Table 4.2: The optimisation results for various EP thicknesses and bolts numbers: from top to bottom for each cell the minimum, average and maximum vertical displacement, in $[mm]$.	104
Table 4.3: The difference between the values of the 6 bolts configuration and the average between the 4 and 8 ones, in percentage.	106
Table 4.4: The different parameters between Al 6061-T651 and AISI 4340.	107
Table 4.5: An example of use of the results of Figure 4.7.	108

LIST OF SYMBOLS AND ACRONYMS

AA:	Active area
ASR:	Area specific resistance
BP:	Bipolar plate
CC:	Current collector
EP:	End plate
EPL:	Equal Path Length
ESS:	Energy storage system
FF:	Flow frame
OCV:	Open circuit voltage
SOC:	State of charge
VFB:	Vanadium flow battery

NOMENCLATURE

A	$[m^2]$	area
ASR	$[\Omega \cdot m^2]$	area specific resistance
a_0	$\left[\frac{m^2}{m^3}\right]$	specific surface area
CE	[1]	coulombic efficiency
CR	[1]	compression ratio
c	$\left[\frac{mol}{m^3}\right]$	concentration
D	$\left[\frac{m^2}{s}\right]$	diffusivity
d	[m]	thickness, diameter
E	[V]	potential of the reaction
E	[Pa]	Young's modulus of elasticity
EE	[1]	energy efficiency
F	$\left[\frac{c}{mol}\right]$	Faraday's constant
F	[N]	force
f	[Hz]	frequency
f_x	[1]	pipe friction factor
G	$\left[\frac{J}{mol}\right]$	Gibbs' reaction energy
I	[A]	current
J	$\left[\frac{mol}{m^2 \cdot s}\right]$	diffusion flux
j	$\left[\frac{A}{m^2}\right]$	current density
K	[1]	Kozeny-Carman constant
K	[1]	bolt friction factor
k	$\left[\frac{m}{s}\right]$	reaction rate
k_m	$\left[\frac{m}{s}\right]$	mass transfer coefficient
l	[m]	length
M	[mol]	moles number
N	$\left[\frac{mol}{m^2 \cdot s}\right]$	solved ion flux
n	[1]	number of transferred charges
OCV	[V]	open circuit voltage
P	[W]	power
p	[Pa]	pressure
Q	$\left[\frac{m^3}{s}\right]$	flow rate
R	[Ω]	resistance
R	$\left[\frac{J}{K \cdot mol}\right]$	universal gas constant
Re	[1]	Reynolds number
S	[m ²]	cross-section
SOC	[1]	state of charge
T	[K]	temperature

T	$[Nm]$	torque
t	$[s]$	time
U	$[V]$	voltage
u	$\left[\frac{m}{s}\right]$	velocity
v	$\left[\frac{m}{s}\right]$	interstitial velocity
V	$[V]$	voltage
V	$[m^3]$	volume
VE	$[1]$	voltage efficiency
W_s	$\left[\frac{J}{m^3}\right]$	strain energy density
z	$[1]$	external charge of the ion
α	$[1]$	pump efficiency
α	$[1]$	transfer coefficient
ε	$[1]$	porosity
ε	$[1]$	strain
ζ	$[1]$	empirical minor loss coefficient
η	$[V]$	overpotential
κ	$[m^2]$	permeability
λ	$[1]$	intrusion ratio
μ	$\left[\frac{kg}{m \cdot s}\right]$	dynamic viscosity
μ	$[1]$	stoichiometric number
ν	$[1]$	Poisson's ratio
ρ	$[\Omega \cdot m]$	electric resistivity
ρ	$\left[\frac{kg}{m^3}\right]$	fluid density
σ	$\left[\frac{1}{\Omega \cdot m}\right]$	electric conductivity
σ	$[Pa]$	mechanical stress
σ	$[Pa]$	standard deviation
ϕ	$[V]$	local potential

1. INTRODUCTION

a. The need of stationary energy storage

The current climate crisis requires the industrial society to drastically cut its greenhouse gases emissions in little time: being the production of electric energy the main cause of such gases, the need to switch to alternative sources is evident. This is mostly done by adopting renewable energies, mainly wind and photovoltaic, but their intermittence is an important disadvantage. The intra-day energy demand is known and foreseen with statistical studies and follows a certain curve of load, as depicted in Figure 1.1.

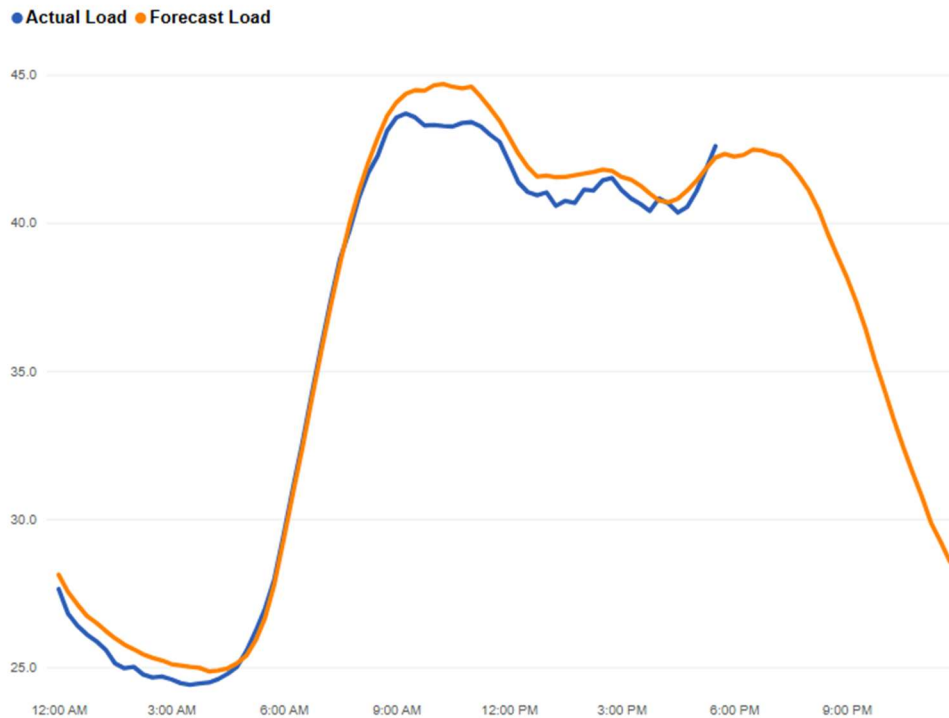


Figure 1.1: The curve of the intra-day energy demand. [1]

The production of electric energy has always followed the needs of the loads: this was not challenging since traditional energy sources are programmable, and usable at request.

Renewable energy sources instead follow other dynamics, mainly the solar irradiation and the wind speed at the current moment in the considered geographical location. They are therefore not coupled with the load, and might not be present with sufficient intensity when needed. This brings to the need of energy storage: implants where energy can be gathered whenever there is abundance of it, and then released whenever demanded. This is not only necessary to feed the loads, but also for grid stability criteria.

Being in fact the grid operated at alternate current (in Europe at the frequency of 50 [Hz] and voltage of 230 [V], in the USA of 60 [Hz] and 120 [V]), the correct values must always be kept within the corresponding tolerances. In particular, in Italy [2]:

$$f = 50 \pm 0.1 \text{ [Hz]}$$

It might be slightly different in particular cases, such as:

$$f_{Sicily} = 50 \pm 0.5 \text{ [Hz]}$$

$$f_{emergency} = 47.5 \div 51.5 \text{ [Hz]}$$

This leads to the concept of inertia of the grid: its capacity to withstand important disturbs, e.g. the sudden loss of a generation plant, faults, the connection or disconnection of heavy loads (factories, trains etc.). Historically, inertia was copiously provided by the heavy rotating motors and generators connected to the grid, which used their own great masses to keep the frequency stable. Nowadays, more and more rotating loads are being connected through inverters, which do not provide anymore the aforementioned stability, but are anyway

preferred since they allow precise control at desire. Photovoltaic and wind generation plants are connected by inverter too, so the power generation is going to contribute less to the frequency stability as well. Luckily, batteries can be set to provide a synthetic inertia, so they can give a substantial help in this issue.

Batteries are not the only energy storage system (ESS) currently available: beside them there are also many others: pumped hydro, flywheels, compressed air, buoyancy, supercapacitors, superconducting magnetic etc.

Furthermore, there are several different types amongst batteries, each with their peculiar advantages and disadvantages: lead-acid, lithium-ion, redox flow, sulphur, Ni-Cd, alkaline, carbon zinc, and much more (beside the uprising interest towards hydrogen).


b. Vanadium redox flow batteries

Amongst all the aforementioned battery technologies, this thesis focuses on vanadium redox flow batteries (VFB). They are subsequently described first in their general characteristics and then in their functioning details, to finally provide the relative state of the art.

i. General overview

Regarding VFBs, the advantages that make them a promising ESS are: [3-9]

- Scalability in power output and capacity. The energy capacity is in fact provided by the ion concentration and the electrolyte volume, which are stored in the tanks. The power output is instead dependant on the number of cells and stacks. These can even be modified in already existing implants, which is an additional advantage. The point is that energy capacity and power output can be independently set, so they can be both chosen as desired.
- Operational flexibility: VFBs are a useful system for applications like peak shaving, time shifting and frequency regulation which have beneficial effects for the grid.
- VFBs can work even for 20,000 cycles, which means that the investment costs cover a long period of time. Concerning the expenses, the vanadium electrolyte builds the greatest share, being the second one given by the power electronics. [8] A diagram of the costs is shown in Figure 1.2: it depicts three cases: industrial implant with either large or small cells, and residential implant, concerning both the investment costs and the total life cycle costs.



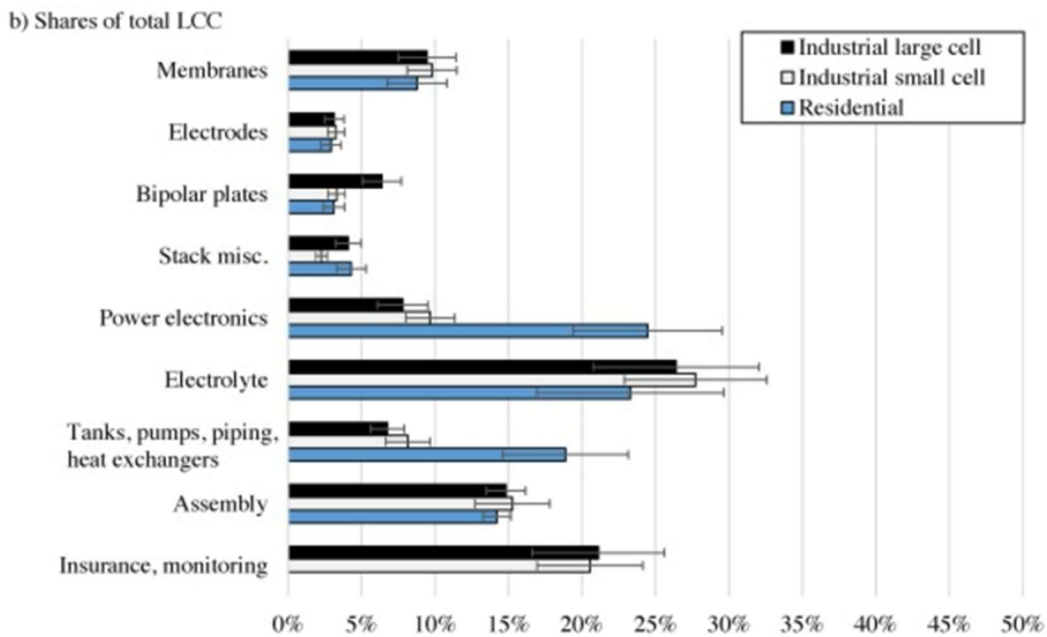


Figure 1.2: The cost shares of a VFB, considering both a) the investments costs and b) the shares of total life cycle cost, in the cases of an industrial implant with either large or small cell, and of a residential implant.

[8]

- Low self-discharge. The self-discharge is the undesired phenomenon that provokes a gradual loss of the energy stored in an ESS. The lower it is, the better. Its causes depend on the issued system, e.g. they are totally different between a VFB and a flywheel. In the case in exam, they are mainly due to the diffusion process of vanadium ions from one half-cell to the other. Such a phenomenon grows stronger with increasing flow rates, i.e. at greater powers. [6] The self-discharge is anyway very low, even negligible in large implants, since it only concerns the electrolyte present in the cells, not in the tanks.
- Response time: it indicates how fast an ESS is to activate. VFBs can do it in a matter of seconds. If they are executing the service of frequency regulation, i.e. they have to suddenly provide an important quantity of watts, the electrolyte in the cells must be fresh, not too consumed by the diffusion phenomenon. In such a case, it is necessary to change it from time to time with fresh electrolyte from the tanks.
- Environmental impact: VFBs are eco-friendly, since almost all their components can be recycled, unlike lithium-ion batteries, which are highly problematic under this point of view. In particular, the electrolyte recycling is not only possible, but even heavily advantageous with respect to the use of primary vanadium electrolyte. [5]

Concerning the vanadium supply, it is quite an abundant element on the planet, being many tonnes of it dissolved in seawater. Their extraction is not economically feasible though, so it must be found in deposits, which come in four main types: vanadiferous titanomagnetite deposits, sandstone-hosted vanadium deposits, shale-hosted deposits, and vanadate deposits, being the first kind the most important source. Vanadium can also be obtained as a by-product in other industries, such as those dealing with steel, aluminium or uranium production. [9]

Since there are no vanadium deposits in the EU area, it is considered a critical material since 2017 by the EU criticality assessment. [10] The countries that produce most of it are China, Russia and South Africa. ESSs are only a small consumer of vanadium, being most of its tonnes used by the steel industry, since it produces hard and stable carbides in nitriding steels, improving their hot resistance. [11]

The criterium the EU uses to define whether a material is critical or not consists in ranking it with two values: one represents its economic importance, the other its supply risk. If both of them are over a

threshold, the issued material is considered critical. Figure 1.3 shows which materials meet the criterium and which do not.




Figure 1.3: Materials considered critical by the EU criticality assessment. Two thresholds are to be met: the economic importance score, and the supply risk score. Vanadium is slightly above both of them, thus is critical. [10]

- Safety: they are not flammable, nor explosive, unlike lithium-ion batteries. This is due to the electrolyte being water-based.
- Low cross-contamination through the membrane, since both the electrolytes use vanadium.

The disadvantages of VFBS are dealt with in chapter 1.b.iii, since they are current object of research, to be fixed.

ii. Structure and functioning of a VFBS

The chemical reactions at the basis of the battery are the reduction and oxidation of vanadium, which is a water-based solution of sulfuric acid in which vanadium ions are dissolved. The peculiarity of this element is that it has four different oxidation states (each with a different colour, also), as reported in Table 1.1.



Table 1.1: The oxidation states of vanadium.

Oxidation state	Found as	Also named	Colour
2+	V^{2+}	$V(II)$	Purple
3+	V^{3+}	$V(III)$	Green
4+	VO^{2+}	$V(IV)$	Blue
5+	VO_2^+	$V(V)$	Yellow

Besides the aesthetics, the different colours of vanadium can help recognising the state of charge of the battery. [12]

The positive electrolyte (or catholyte) has $V(IV)$ and $V(V)$ ions, while the negative electrolyte (or anolyte) has $V(II)$ and $V(III)$ ions; their concentrations are calculated to prevent precipitation. The two electrolytes are

stored in separated tanks and pumped in an electrochemical reactor, the stack, which is made of several cells. There, redox reactions occur and turn chemical energy into electrical energy. The two half reactions are:



The left-to-right direction corresponds to discharge, while the right-to-left one to charge.

The open circuit voltage (OCV) is the sum of the two E_0 s addenda: $E_0 = 1.26 [V]$, evaluated at $T = 298.15 [K]$ and $SOC = 50\%$ for each cell. The cells are electrically connected in series, so their voltages add up, while the hydraulic connection is in parallel.

Neither in the stack do the electrolytes mix: they reach distinct porous electrodes and flood them. The ones hosting the catholyte and the others the anolyte are alternated, and the separation is given by an alternation of ion-exchange membranes and bipolar plates (BPs). Figure 1.4 shows two adjacent electrodes, separated by an ion-exchange membrane. Their hydraulic connection to the tanks features pumps, which drive the flows.




Figure 1.4: A cell of the stack of a VFB. Two consecutive electrodes are separated by an ion-exchange membrane and fed by the tanks hosting the electrolytes. [4]

Regarding the discharge reaction, $V(II)$ turns into $V(III)$ releasing an electron, while $V(V)$ turns into $V(IV)$ acquiring one. This electron travels from the first electrode to the second one crossing the BPs. This would lead to a charge imbalance, if it were not compensated by the migration of an ion through the membranes. Therefore, it will reach the first electrode from its other neighbour, and a different ion will leave the second electrode, similarly heading to the other neighbour of this one. When the migrating ion is positive, usually an H^+ as depicted in Figure 1.4, the ion-exchange membrane is a proton-exchange membrane. The stack can even work the other way around, with negative ions crossing the membranes, which in this case are called anion-exchange membrane.

For a matter of charge balance, the direction of the anions' migration is opposite with respect to the one of cations.

This was the description of the discharge process: the charge is analogous, but with the migration directions inverted.

At the two ends of the stack there are two copper-alloy current collectors (CCs), the terminals of the battery under an electrotechnics point of view. Charge and discharge determine the current direction coming from and into them. [13]

The Flow frames (FFs) are responsible for the distribution of the electrolytes inside the cells and they also provide, together with the gaskets, the hydraulic sealing. [4]

To sum up, several cells are pressed together to form a stack, inside of which liquid electrolytes flow. It is clear that an element to operate such a pressure is needed: the end plates (EPs). They are a couple of steel or aluminium-alloy plates placed at the two ends of the stack, clamped together via tie rods. They hold everything together, provide the hydraulic sealing and the electric contact, and are the focus of this thesis.

iii. State of the art of VFBs

VFBs are a promising technology, due to the previously listed advantages, but still not economically competitive with respect to other solutions, like Li-ion batteries. Regarding ESSs, the cost is usually expressed as referred to the power or energy capacity, so respectively in €/kW or €/kWh, or their multiples or equivalent with other currencies. Besides the costs, the energy and power are also usually considered with respect to the volume necessary to host them. In particular, VFBs have low energy and power densities, which means that rather big implants are needed to achieve a good number of kW and kWh. As long as the applications are stationary (i.e. not concerning moving objects, like vehicles), the energy density is not too much of a problem, while the power density is anyway something to be improved to make them more economically competitive. In fact, if the same power can be achieved with fewer materials and smaller stacks, the costs decrease.

Current research focuses on electrode modification in order to improve this power density. [14] In fact, its electrical conductivity and thickness affect the ohmic polarisation, while its electrocatalytic activity and reversibility determine the activation polarisation, and the pore structure and hydrophilicity the concentration polarisation. The ohmic polarisation is due to the internal resistance of the components (the electrode, the BPs and the membrane), the activation polarisation is due to the potential drop necessary to start the chemical reaction, with respect to open circuit conditions, and the concentration polarisation is due to the electrolyte movement, its convective motions, and the consumption of the reacting species.

Concerning the electrochemical polarisation, new materials are being studied to enhance the catalytic effect, especially when high current densities are involved. Modifications are also considered, involving the introduction of surface functional groups, the tailoring of the porous microstructure, the expansion of the active surface area. The ohmic polarisation is instead due to the internal resistances of a VFB, like those that hinder the ionic movement, and the electrical and contact resistances, respectively inside a component or between two consecutive ones. Contact resistances are reduced by ensuring a good pressure acting on the plates, but this cannot reach exaggerated values, lest the porous structure of the electrodes undergoes detrimental effects. A compromise between these two colliding necessities is to be met, therefore. Anyway, the internal resistance depends on the electrodes' thickness, which needs then to be as thin as possible. Current research focuses on the study of different materials, all carbon based anyway, to obtain the best microporous structure. To sum up, thin electrodes have short paths for the ions' and electrons' migrations, therefore low ohmic resistances, but also lower surface areas and hydraulic permeabilities, so higher electrochemical and concentration polarisations.

Regarding last concentration polarisations, they might be reduced by simply pumping the electrolytes at higher rates, but this would imply a higher power consumption by the pumps, and so lower efficiencies. The distribution can be improved by the adoption of a proper flow channel geometry, or even by tailoring the size of the pores accordingly to their position in the electrode, in order to ease the flow passage. This cannot be excessive, though, since a reduction in the active surface area would occur, leading to larger activation losses.

An important detail to consider whenever reading the bibliography is to check whether the study refers to a single cell or to an industrial scale implant: the scaling up often introduces additional challenges that might have not

been noticed in the single-cell case. Not only does this happen in the laboratory phase, but also in the computer simulations.

Regarding efficiency, the heat generated inside the stack comes from shunt currents, cell overpotentials and crossover, besides the losses due to the ancillary services. Tests conducted on the $9\text{kW}/27\text{kWh}$ industrial scale VFB of the University of Padua for different current values produced losses as show in Figure 1.5. [15]




Figure 1.5: The losses in the $9\text{kW}/27\text{kWh}$ industrial scale VFB of the University of Padua for different current values. The addendum of the ancillary services is further analysed in a dedicated pie chart. [15]

The nature of shunt currents, crossover and cell overpotential is analysed in chapter 2. Although being a smaller addendum, ancillary services constitute too a component where some gain can be obtained. This can be done by just adopting high-efficiency pumps, for example. Luckily, the higher the flow, the better the efficiencies these machines can reach, so industrial scales are advantaged with respect to laboratory implants concerning this detail. [16]

c. Alternative chemistries

As already stated, the high cost of VFBs is their main limit, so research has focused on the possibilities of other materials to substitute vanadium in its role. The basic principles stay the same (i.e. the chemical reactions of oxidation and reduction of an element dissolved in a liquid electrolyte), and so does the battery structure, though. Many possibilities have been studied: iron-chromium, bromine-polysulfide, iron-vanadium, all-vanadium, vanadium-oxygen, zinc-bromine etc. [17] The problems of the aforementioned solutions were nevertheless the expensiveness and the availability of the involved materials. This does not happen instead when dealing with an all-iron battery, being Fe abundant, besides non-toxic.

i. All-iron flow batteries

Figure 1.6 gives a schematic representation of an all-iron battery. [17]

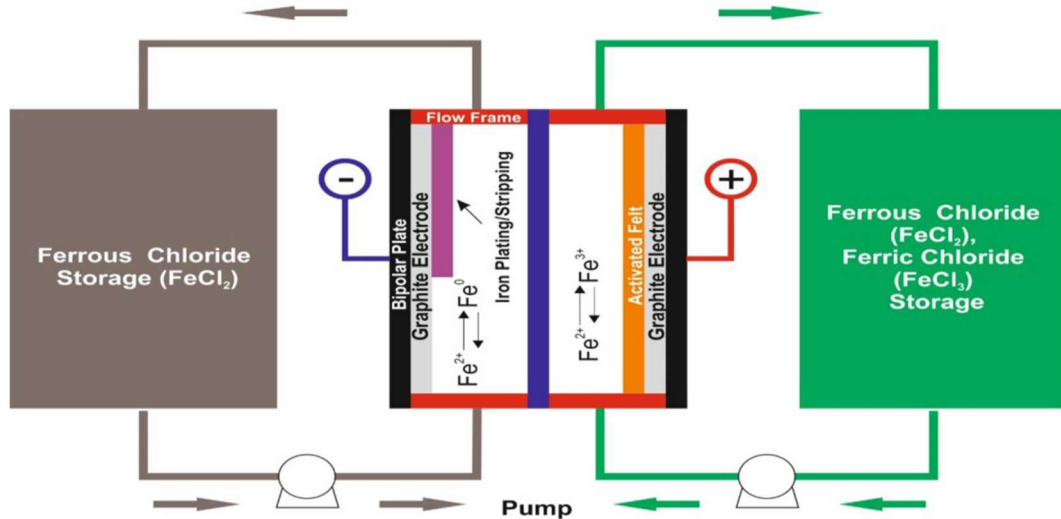



Figure 1.6: Schematic representation of an all-iron battery. It is conceptually very close to a VFB. [17]

Once again, external storage tanks host the two liquid electrolytes (both involving Cl), which are pumped through the stack, where the chemical reactions occur and convert chemical energy into electric energy:



With the discharge reaction being the left-to-right one. The sum of the two potentials gives $E_0 = 1.21 [V]$.

The reduction of Fe^{2+} to Fe^0 is also called plating, since the product is deposited on the negative electrode. Such a plate is then dissolved again in the liquid form with its oxidation.

Once again, the graphite electrodes play an important role in the redox reactions, since they host them on their surface: the dimensions of the pores matter, since they determine what can pass and what not. In particular, a microporous structure allows the migration of negative ions (chloride ions), while a nanoporous structure allows the migration of positive ions.

Slurries can improve the efficiency of the all-iron battery: they decouple the energy capacity and the power density, thus allowing the operation at high current densities. They are solid conductive particles dispersed in the electrolytes, and if present in a sufficient volume fraction a continuous electric conductive network can form: this allows redox reactions to occur on their surface. Concerning the negative electrode, the iron forms plates on their surface too: they are then pumped to the external reservoir, so that the plating does not hinder the reaction. [18] Since the electrolytes alone are the most expensive part of a flow battery, attention is paid to them, in particular to their salts. In general, it is used the ionised salt of the metal in acidic conditions. For example, NH_4Cl and H_3BO_3 are useful, since they can respectively reduce the electrolyte resistivity and inhibit hydrogen evolution. The solubility limit and the tanks capacity determine the maximum amount of stored energy, whereas the total energy output depends on the quantity of metal to be deposited on the electrode surface, since they allow the electrons' exit from the electrolyte.

Concerning therefore solubility, it is necessary to enhance the one of Fe^{2+} , but without interfering with the process of plating at the negative electrode: proper ligands are to be found. The criteria chosen to determine whether the considered ligand is valid are: [19]

- the iron-ligand solubility must be greater than $1M$;
- they need to be soluble up to $pH 3$;
- they must be electrochemically inactive in the potential window of the flow battery;

- they must not affect the $Fe(II/III)$ kinetics: its velocity and reversibility;
- the OCV must not be reduced due to their addition.

The best additive to reduce ferric precipitation is found to be glycine amongst six more others. [19] It has in fact two states: positive and negative, and can better coordinate with the iron than the other additives; its kinetics and diffusion coefficients are near those of ferric/ferrous ions without any ligand. The reaction becomes better distributed at higher current densities, which is good news considering that industrial application is the final aim of the development of these batteries. The best glycine to all iron ratio is between 0.5:1 and 1:1 in the electrolyte, with 1:1 providing stability if the pH of the electrolyte is 2.

Concerning the pumping pressures, their ratio shall follow the stoichiometric rate of reaction, in order to keep a good battery efficiency. In particular, the anolyte to catholyte ratio is function of the state of charge (SOC).

Finally, all-iron flow batteries are generally cheaper than vanadium ones, since the most expensive component (the electrolyte and its salts) are more affordable. In particular, Table 1.2 reports the estimated price of the stack components of an all-iron flow battery. [17]

Table 1.2: The estimated price of the various components of an all-iron flow battery. [17]

Stack components	Estimated price
Activated felt	\$90/m ²
Bipolar plates	\$50/plate
Flow frames	\$32/frame
Gasket	\$2/gasket
Collector plate	\$125/plate
End plate	\$175/plate
Bolt	\$12.50/bolt
PVC set	\$150/set
Ferrous chloride solution (27%)	\$0.129/L
Ammonium chloride	\$1.06/L
Boric acid	\$1.96/L
Deionized water	\$0.016/L
Electrolyte preparation	\$100.00/batch
Tanks	\$276.97/pair
Pumps	\$200.00 each
Control system	\$600.00 each

ii. Organic flow batteries

Once again in the effort of reducing the cost of redox flow batteries, other possibilities have been explored, such as those given by organic materials. Their advantage is their low cost, since they are earth-abundant and sustainable, and their production can be easily scaled up. [20]

The term “organic” specifically refers to the redox-active materials, not to the solvent nor to the electrolyte [21], but it is clear that these will differ too from the VFB case. A wide variety of candidates has been analysed, leading to a larger knowledge regarding the electrochemical properties of many new materials. Just to give an example, Figure 1.7 shows possible molecules to be used as catholyte (in blue) and anolyte (in orange) in organic redox flow batteries, either aqueous or nonaqueous.




Figure 1.7: Possible molecules to be used as catholyte (blue) and anolyte (orange) in organic redox flow batteries. [20]

Organic flow batteries can feature aqueous or nonaqueous electrolytes. Possible redox-active molecules in the aqueous case are: viologen, ferrocene, quinone, TEMPO, phenazine, alloxazine, pyrazine and their derivatives. Since the concentration of the aqueous cases is generally higher than the one of nonaqueous cases, the energy density is generally greater, [22] but is yet to overcome the one of VFBS (because of the vanadium solubility in H_2SO_4).

The advantage of nonaqueous organic batteries is that they can avoid cross-contamination, but on the other hand the redox specie is not stable enough yet. Furthermore, most of nonaqueous organic flow batteries operate at low current densities. For the research to progress, the corrosion phenomena are to be better evaluated, and so is their impact on the practical service life. In particular, it is to be considered that these tests still regard small laboratory batteries: they still are to be scaled up to industrial size ones.

2. PERFORMANCES OF AND LOSSES IN A FLOW BATTERY

a. The hydraulic design and the electrolyte distribution

Even though it does not take part in the redox reaction, the electrode is a critical component of the VFB, since it provides the reaction sites. Furthermore, the electrode has an impact on the polarisation of the stack, through the ohmic resistance and charge transfer polarisation. Concerning the choice of its material, the three-dimensional structure needs to be stable to provide a large surface area and a good conductivity, all of this at a reasonable cost. The material which has all these properties is the carbon felt. [23]

An X-ray microtomography is useful to closely inspect the inner structure of carbon felts: it reveals they are made of strains of carbon fibres, which are woven together in an irregular way, leaving amongst them many little spaces (the pores) where the electrolyte can flow through. [24]

Being the disposition of the fibres irregular, so are the pores and the whole structure too: this means that the electrolyte flow won't follow a straight pattern. The structure is shown in Figure 2.1, where the red area represents the cross-sectional area of the electrode, and the grey colour the solid material.




Figure 2.1: The carbon felt microstructure: it is possible to see the carbon fibres and the empty spaces amongst them: the pores. [24]

Since the electrodes play such an important role in VFBs, many researches focus on improving their performances. In the next chapters, the role of compression regarding this purpose will be deeply analysed, but other ways have been tried as well. For example, the partial oxidation of the carbon fibres of the structure at high temperatures or the deposition of conductive metals (*Pt*, *Au*, *Bi* etc.) which are meant to act as electro-catalysts were explored. [25]

Concerning the chemical reactions and the electrons conduction, they are hosted on the inner surface of the electrodes, on their active sites, which implies that the extension of this area plays a critical role. It is evident that this quantity is to be determined and measured in some way, so that it is also possible to understand whether a certain process can increase or decrease it. An idea would be to consider the sum of the geometric areas of all the cylindrical fibres of the felts, but this would neglect the contact surfaces between such fibres, and consequently the effects of compression, which are crucial. [24]

A possible approach might involve X-ray tomography: a huge number of projections of the felt are taken within a short time, with different angles in order to reproduce it with detail on a computer-drawn model. The 3D reconstruction is then described with a mesh of even millions of elements and analysed: the total area can be evaluated. Then, this total surface can be divided by the volume of the sample, in order to obtain the specific surface area, i.e. the available surface for chemical reactions per unit volume. [24] This approach can therefore take into account the contact surfaces amongst the carbon strains, but is not reliable yet.

In fact, the active surface area evaluated in this way mismatches the electrochemical performances of the cell. The reason behind this is that the X-ray approach does not consider that the electrolyte might not perfectly flood the whole electrode, which means that the available area might not be completely exploited. [24] Since the active surface area cannot be evaluated in a simple way, but depends on the operating conditions of the cell (in primis

the hydraulic behaviour), it will rather be found in an experimental way. In fact, its purpose is to enable the conversion of the reaction current density at the electrode surface onto a volume-dependant density for the whole-electrode. [26]

Regarding the hydraulics, the electrolyte flow is considered to be laminar, incompressible, and isothermal. [26] The tortuosity of the electrolyte path is important too, since it deeply affects the charge and mass transport dynamics. It is nevertheless difficult to evaluate, so there are many proposed methodologies to do so. [24]

It is not a surprise that the electrolyte can reach its maximum velocities in the regions where it is not blocked by carbon fibres, i.e. where the pores are large and align along clear paths; on the other way around, the velocities are the lowest where many carbon strains hinder their motion, describing a region with small pores [24] Figure 2.2 provides an example for this in three compression cases: in the top images, the structure of the electrode is shown, specifying where the inlet and the outlet are. The electrolyte flows perpendicular to them, along the x direction. The middle images show the streamlines, while the bottom ones the velocity contours on a plane perpendicular to the main flow direction. Notice that the flow will not be perfectly aligned with the x direction throughout its whole path, but will have to take many deviations to pass amongst the strains: it will generally have y and z components too. [24]

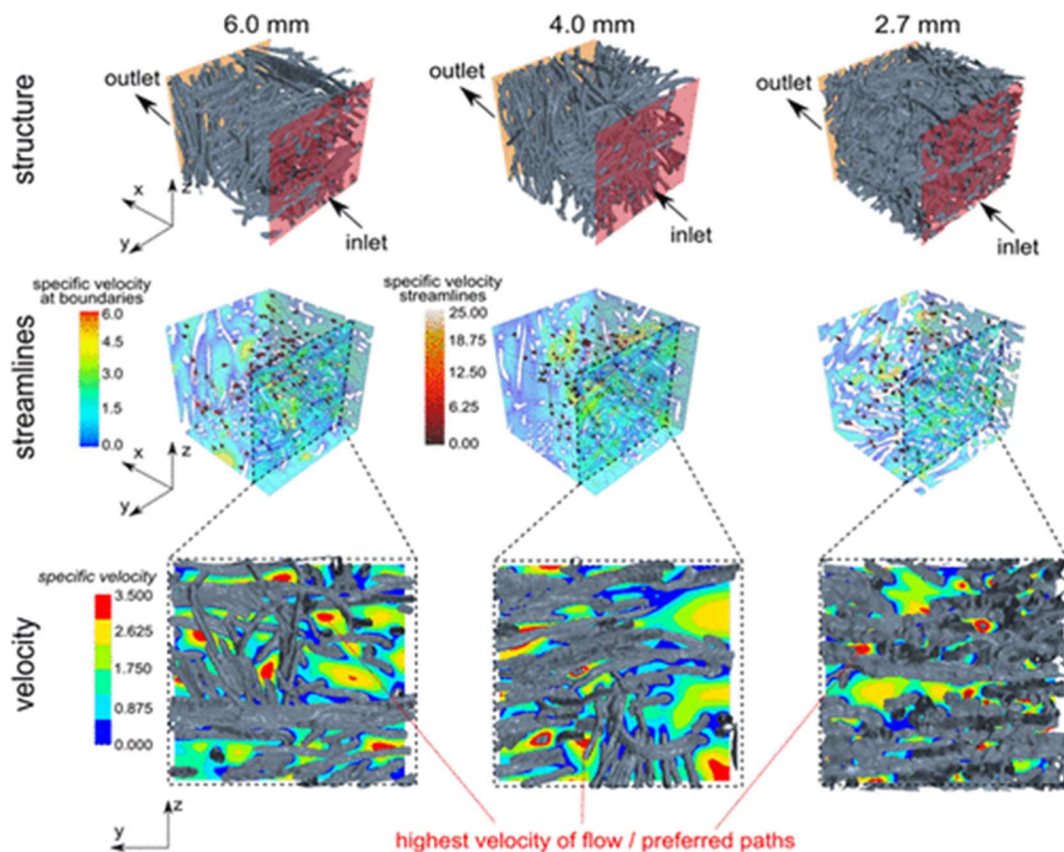


Figure 2.2: The electrode structure and the electrolyte velocity in three different compression cases. [24]

A quick glimpse to the velocity distributions makes it clear that some points are harder to reach for the electrolyte. In an electrode that is bigger than the represented one, it might even happen that some little regions are not accessible at all, thus validating the previous statement concerning the comparison between the X-ray method and the experimental evaluation.

This inhomogeneity affects the electrochemical performances of the cell, due to its hot-spots and inaccessible regions. Their dispositions depend on the electrode structure, on the compression, even on the electrolyte flow itself, so the active area (AA) available is more of a unique value for the current case. It might even change during time, due to aging phenomena. [24]

Concerning the varying velocity inside the electrode, Figure 2.3 provides an interesting example: it shows the invasion pattern, and the colours represent the order in which the pores are flooded, from blue to white.




Figure 2.3: The invasion pattern of the electrolyte inside the electrode. The colours indicate which pores are progressively flooded: blue first, white last. [27]

It is anyway interesting to point out that carbon felts are actually hydrophobic, in principle. To change this behaviour, it is necessary to heat-treat them at circa 400 [°C] for 24 hours, so that hydrophilic groups can form on the fibres' surface, enhancing the wettability. [24]

It was described how the porosity plays an important role in the cell performance, affecting the hydraulic permeability, the electric and ionic conductivity, the effective diffusion coefficient, the surface concentration of the active species in the electrode. It would be convenient to quantify its value, but it is necessary to keep anyway in mind that these pores are not actual holes in a sponge-like tissue, but rather gaps between the carbon fibres. It is therefore even harder to characterise them with a diameter, but such a formulation would anyway be convenient when dealing with hydraulics calculations. The porosity is then quantified as:

$$\varepsilon = \frac{V_{empty}}{V_{tot}} [1]$$

Given that V_{empty} is the empty volume inside the electrode, due to these pores, or better gaps. [26]

Regarding the hydraulic design, it is also necessary to illustrate how the electrolyte is led into the electrode, in the first place. Two approaches can be adopted: the flow-through and the flow-by.

In both cases, there are four manifolds: two per electrolyte: back and forth with respect to the corresponding tank. Two manifolds lead therefore the input catholyte and anolyte to each cell (where each semi-cell either receives one or the other), and then the other two collect the output ones, to bring them back to the tanks. In the flow-through configuration, the electrolyte is simply let inside the electrode, with no particular guidance, and floods it from one extremity, exiting it then from the other, where it is collected again to end in the other manifold. The flow-by configuration is instead defined with greater detail: the path of the electrolyte is accurately calculated and guided with channels, which are dug inside the BPs. These channels can have various shapes, as shown in Figure 2.4, where the flow-through situation is drawn as well.




Figure 2.4: The various channel configurations for the flow-by (serpentine, parallel, interdigitated) and the flow-through case. [28]

Each flow-by configuration features an inlet and an outlet. In the serpentine case, the electrolyte flows through a unique channel, and from that floods the electrode. In the parallel case, it has instead more options: this reduces the pressure drop and homogenises the distribution, but each path takes less time to be run through, i.e. the electrolyte has less time to flood the electrode. The interdigitated configuration distinguishes the channels connected to the inlet from those to the outlet. For the electrolyte to come from the starting point and end at the finish point it is strictly necessary to run through the electrode as well, jumping from one channel to another. The flow-through case instead, the rightmost in Figure 2.4, does not feature any channel at all. Its cell is simpler to build, since it is not necessary to engrave the channels in the graphite of the BP. This also simplifies the stacking of the cells: the electrodes are thinner (carbon papers), and the contact pressure is more uniform since there are no channels in the adjacent BPs. Given that the flow is more uniform across the electrode, and that its contact with the BP is uniform too, the flow-through performs better than the flow-by concerning the current density distribution. The advantage of the flow-by is that it performs better under the hydraulic point of view instead, thanks to the channels.

When deciding the shape of the manifolds or channels, or any part of the hydraulic path in general, it is necessary to keep in mind the laws of fluid dynamics, in order to minimise the pressure drops of the electrolytes, since they would result in losses. This comprehends avoiding sharp angles or narrow passages, for example.

The way each cell (or better, semi-cell) is fed has been explained. In particular, all the cells of the same stack are in hydraulic parallel, unlike in the electric configuration.

b. The electric design

The electric connection between consecutive cells is a series.

As it was explained at the beginning, in each semi-cell a redox reaction occurs, freeing or capturing an electron and an ion. The ion crosses the ion-exchange membrane, but cannot pass the BP, while the electron enters the BP, but cannot pass the ion-exchange membrane. The total current can exit or enter the stack from the two CCs at the extremes of the stack. Each cell has a certain voltage at its two terminals, and this quantity is not fixed.

The OCV is the voltage which the cell has in open circuit condition, i.e. of zero current.

The one with non-zero current can instead be described with: [8]

$$U_{cell} = OCV - \sum_k (IR)_k - \sum_k |\eta_k| \quad [V] \quad (1)$$

This applies to the discharging operation: in the charging case instead the signs are positive, which means the cell voltage is higher than the OCV. The IR conduction losses are due to electric and ionic conduction, while the η overpotentials to activation, diffusion, mass transport and other mechanisms. It is useful to define the ASR, the area specific resistance: [8]

$$ASR = \frac{U_{cell} - OCV}{j} \quad [\Omega m^2]$$

With j current density at the CCs. This is a convenient specific parameter to bind the operating conditions (the current density) to the effective voltage of the cell.

Going more into detail, the OCV is determined by the chemical reactions at the cathode and at the anode. The Nernst's equation helps quantifying the potential of a redox reaction: [8]

$$E_0 = E'_0 + \frac{RT}{nF} \ln \left(\frac{c_{Ox}}{c_{Red}} \right)$$

With E'_0 standard reduction potential, R universal gas constant, T [K] temperature (here the environmental one, of 293 [K]), n number of transferred electrons, F Faraday's constant, c_{Ox} and c_{Red} concentrations of the oxidating and reducing species. The cell voltage is therefore the sum of E_0 for the cathode and for the anode: [8]

$$OCV = E_{0,cat} + E_{0,an} = E'_{0,cat} + E'_{0,an} + \frac{RT}{nF} \ln \left(\frac{c_{V^{5+}} \cdot c_{H^+}^2 \cdot c_{V^{3+}}}{c_{V^{4+}} \cdot c_{V^{2+}}} \right) [V]$$

The concentration of the various ions defines the SOC: [8]

$$SOC = \frac{c_{V^{2+}}}{c_{V^{3+}} + c_{V^{2+}}} = \frac{c_{VO_2^+}}{c_{VO_2^+} + c_{VO^{2+}}} [1]$$

This means that the OCV is function of the SOC too.

Concerning instead the current density, it is due to the flux of the solved ion species, N_i , given by the Nernst-Planck equation: [8]

$$N_i = -D_i^{eff} \nabla c_i - \frac{z_i c_i}{RT} D_i^{eff} F \nabla \phi_i + \mathbf{u}_i c_i \left[\frac{mol}{m^2 \cdot s} \right]$$

The first addendum represents the diffusion, and in fact it is bound by the effective diffusion coefficient D_i^{eff} to the gradient of the concentration of the considered ion specie. The minus sign indicates that the flux takes the opposite direction with respect to the concentration gradient, i.e. goes from the region of higher concentration to the one of lower. The second addendum represents the migration, and it goes with opposite direction with respect to the potential gradient, to whom it is linearly bound by some other quantities: z_i is the external charge of the ion, the others were already previously explained. The third and last addendum represents the convection, and in fact is proportional to the ion velocity \mathbf{u}_i .

A flux N_i of charged particles implies the presence of a current: [8]

$$\mathbf{j}_i = n z_i F N_i$$

\mathbf{j}_i is the current density corresponding to a single ion specie. Summing the impact of all the species, the whole ion current density is found: [8]

$$\mathbf{j}_I = \sum_i \mathbf{j}_i \left[\frac{A}{m^2} \right]$$

The resulting ionic current can eventually be found with Ohm's law: [8]

$$\nabla \mathbf{j}_I = -\sigma_I \nabla^2 \phi_I$$

With σ_I ionic conductivity of the electrolyte.

As it was briefly illustrated in equation (1), the relation that binds the current density and the cell voltage is not simple, nor linear. The polarisation curves are useful to better comprehend this relation. As it was said, in zero-current conditions, the voltage is the OCV , and – in discharge – it decreases with increasing current. The polarisation curve graphs the applied current density to the cell potential, thus arriving to Figure 2.5.




Figure 2.5: The polarisation curve during discharge.

The higher the current density, the lower the voltage, but not according to a linear relation. Starting from the open circuit condition, i.e. for low current values, the voltage drops are mainly caused by the activation phenomenon, which means that the redox reactions themselves lower the semi-cells potentials when happening. This quantity is to be paid both at the anode and at the cathode, which means that it will bring two addenda in equation (1), which are collected in the $\sum_k |\eta_k|$ sum. After this, the rise in the current density and the additional voltage drop follow a linear relation: this region is dominated by the impact of the ohmic losses: they comprehend the contact and ionic resistances, the mass-transfer within active electrode layers. In fact, being the current nothing but the movement of charged particles (ions and electrons), this resistance represents the obstacles they have to face along their path.

Last, when a huge current is asked from the cell, limits in the mass-transfer occur: they are associated to the delivery of the bulk reagent to the electrode. The corresponding voltage is very low, which means that extracting a huge current from a cell would supply little power.

Many of the parameters above depend on the electrode porosity, which means that a mechanical compression will have an impact on the value of the currents too.

The concept of impact naturally leads to that of efficiency, whose idea is basically to compare the input and the output of any given process, i.e. the gain and the required effort. Regarding VFBs, three efficiencies can be defined: the voltage efficiency, the Coulomb efficiency and eventually the energy efficiency.

As it was said regarding equation (1), the voltage of any cell is not equal through charge and discharge: besides the effect of the *SOC*, the signs in the formula change in the two processes. This means that during discharge the cell has a lower voltage than *OCV*, while during charge it has a higher one. This is due to the activation phenomena and to the direction of the current density. The power is known to be the product of the voltage times the current, and this discrepancy implies that the charging power is higher than the discharging power (for an equal current). The voltage efficiency compares the voltages during the two phases. Since these voltages depend on the *SOC*, their values for *SOC* = 50% are conventionally taken for this evaluation: [23]

$$VE = \frac{U_{cell,di}(SOC = 50\%)}{U_{cell,ch}(SOC = 50\%)} [1]$$

Unfortunately, self-discharge phenomena occur and slowly consume the stored energy; besides this, the charge and discharge processes have efficiencies too: in the end, not all the accumulated energy will eventually be given back. This is the reason why the coulombic efficiency is introduced:

$$CE = \frac{\text{discharge capacity}}{\text{charge capacity}} [1]$$

Where the term capacity refers to the number of charged particles extracted or deposited in the system, evaluated as the integral over time of the current. [29]

Finally, the energy efficiency is the product of the two described terms: [23,29]

$$EE = CE \times VE [1]$$

The goal of the designer is to project a VFB (or in general an ESS) with a good EE , i.e. able to give back (almost) all the deposited energy, at the same voltage. The phenomena behind the inefficiencies will be described in the next paragraphs, one by one, together with possible solutions to minimise them. A detailed description will be given about compression, since it has multiple effects and is the protagonist of this thesis.

c. Internal resistance losses

In an equivalent electric circuit, according to electrotechnics laws, a real voltage generator can be represented as a series of an ideal voltage generator and a resistance. In the case of the cell of a VFB, the resistance depends on the current, as it was hinted at by equation (1) and Figure 2.5. They also introduced the addenda that comprise this perceived resistance, which is detrimental for the performances of the ESS.

i. Activation overpotential

The activation overpotential is the voltage drop tribute that the semi-cells have to pay for the electrochemical reaction to occur. Its impact is the first one to be perceived in the polarisation curves (as in Figure 2.5), i.e. is particularly visible for low current densities. It is not equal between anode and cathode, since the two of them have different reaction rates k . It can be quantified as:

$$\eta_{act,cat} = \frac{2.3RT}{\alpha^{cat}F} \log j_0 - \frac{2.3RT}{\alpha^{cat}F} \log j [V]$$

for a large negative activation polarisation at the cathode, and as:

$$\eta_{act,an} = -\frac{2.3RT}{\alpha^{an}F} \log j_0 + \frac{2.3RT}{\alpha^{an}F} \log j [V]$$

for a large positive activation polarisation at the anode. The sign of η_{act} is in fact such to contrast the OCV , i.e. negative at the cathode (which has positive potential) and positive at the anode (negative potential). R and F are always the same constants, T is always the temperature, while α^{an} and α^{cat} are the transfer coefficients of anode and cathode. j_0 is the exchange current density:

$$j_0 = Fkc_0^{1-\alpha}c_R^\alpha$$

With c_0 concentration of the reactant and c_R concentration of the product. If the two of them are equal ($c_0 = c_R = c$), then:

$$j_0 = Fkc$$

Table 2.1 provides some numerical values for the aforementioned quantities. They are dependent on the chosen electrode and electrolytes. A complete description of the symbols and abbreviations used in the table can be found in the corresponding source. [30]

Table 2.1: The values of the parameters for the evaluation of the activation polarisation. [30]

Electrodes	Electrolyte	Methods	Reaction rate k_{PE} , cm s^{-1}	Reaction rate k_{NE} , cm s^{-1}	α^a	α^c
V^{2+}/V^{3+}						
Glassy carbon (GC)	0.12 M VCl ₃ + 0.5 M H ₂ SO ₄ + 1.04 M Na ₂ SO ₄	CV		1.71×10^{-5}		0.52–0.72
Graphite reinforcement carbon (GRC)	0.05 M V(II) + 1.0 M H ₂ SO ₄	CV	4.8×10^{-3}	3.2×10^{-3}	0.52–0.59	0.33–0.35
GRC	0.05 M V(II) + 3.0 M H ₂ SO ₄	CV	9.7×10^{-3}	7.7×10^{-3}	0.52–0.59	0.33–0.35
Plastic formed carbon (PFC)	0.05 M V(III) + 1.0 M H ₂ SO ₄	CV		5.3×10^{-4}		0.5
Pyrolytic graphite(PG)	0.05 M V(III) + 1.0 M H ₂ SO ₄	CV		5.5×10^{-4}		0.5
Carbon rotating disk electrode (Carbon RDE)	1.0 M V(IV) + 2.0 M H ₂ SO ₄	LSV	5.52×10^{-3}		0.70	
Carbon felt	0.05 M V(IV) + 1.0 M H ₂ SO ₄	CV	1.54×10^{-5}		0.21	
Graphite	0.4 M V(III) + 2.0 M H ₂ SO ₄	CV	4.69×10^{-4}	1.16×10^{-3}		
VO^{2+}/VO_2^+						
GC	0.055 M V(V) + 1.8 M H ₂ SO ₄	CV	7.5×10^{-4}		0.5	
GRC	0.05 M V(IV) + 1.0 M H ₂ SO ₄	CV	9.1×10^{-4}	1.5×10^{-3}		
GRC	0.05 M V(IV) + 3.0 M H ₂ SO ₄	CV	6.1×10^{-4}	5.3×10^{-3}		
Graphite and graphite RDE	0.05 M V(IV) + 0.05 M V(V) + 3.0 M H ₂ SO ₄	CV and LSV	5.12×10^{-5}		0.71	
Graphite RDE	0.09M V(IV) + 0.11 M V(V) + 1.0 M H ₂ SO ₄	LSV	7.5×10^{-7}		0.42	
PG	0.05 M V(IV) + 1.0 M H ₂ SO ₄	CV	1.3×10^{-4}		0.5	
PFC	0.05 M V(IV) + 1.0 M H ₂ SO ₄	CV	8.5×10^{-4}		0.5	
GC	0.05 M V(IV) + 1.0 M H ₂ SO ₄	CV	6.8×10^{-5}		0.5	
Graphite	2.0 M V(IV) + 2.0 M H ₂ SO ₄	EIS	1.2×10^{-5}		0.63	
Pt RDE	2.0 M V(IV) + 2.0 M H ₂ SO ₄	LSV	1.47×10^{-3}			
Carbon RDE	2.0 M V(IV) + 2.0 M H ₂ SO ₄	LSV	7.53×10^{-14}		0.74	
Carbon felt	0.3 M V(IV) + 2.0 M H ₂ SO ₄	CV	1.83×10^{-5}	9.96×10^{-6}	0.23	0.23
Graphite	1.2 M V(IV) + 3.0 M H ₂ SO ₄	CV	2.13×10^{-3}		0.53	
Graphite	0.2 M V(IV) + 2.0 M H ₂ SO ₄	CV	2.35×10^{-4}	1.14×10^{-3}		

The fact that η_{act} s are bound to the current density j by a logarithmic relation explains the reason why these overpotentials become less visible as soon as the ohmic and concentration overpotentials reach greater values, as it can be seen in the polarisation curves.

Activation overpotentials can be reduced by easing the electrochemical reactions: this can be done with the introduction of electrocatalysts in the electrode. [31]

ii. Ohmic voltage drop

As soon as the current density rises, ohmic voltage drops become more important than the activation overpotentials.

Being the current but a flow of charged particles through a conductive medium, the resistance of the medium causes the ohmic voltage drops along the path. In the case of VFBS, two charged particles are involved: the ions and the electrons. The ions start their path from the electrode active surface, where the reaction occurs, cross the liquid electrolyte, the ion-exchange membrane and end up in the other electrolyte. The electrons instead cannot cross the membrane, but are collected by the BPs. These components establish an electric series amongst the cells of the stack, and the two CCs are at the extremities. They are the electric terminals of the stack, and from them the electric current (of electrons) departs.

It is possible to recognise what hinders the charge motion: the electrolyte resistance, the contact resistance between distinct elements (such as between electrode and BP), the internal resistance of solid pieces. In particular, the ionic conductivity is two orders of magnitude lower than the electrical one, which makes it the most important term in the ohmic losses. [26]

The contact resistance between consecutive elements can be reduced by simply applying a good pressure. In fact, the electrode is made of a huge number of tiny carbon strains, with a lot of empty space amongst them. This leads to the fact that the contact of the electrode with the BP and with the membrane happens on the surface of the external strains, the ones facing the adjacent object. The higher the pressure, the more strains are pressed, the higher the area of contact on a microscopic point of view.

Concerning instead the other resistances, they can simply be evaluated with the famous

$$R = \rho \frac{l}{A} [\Omega] \quad (2)$$

Which means that the longer, the narrower the piece to cross, the higher the resistance faced by the moving charge. The BPs are thin sheets of graphite: wide and conductive. The electrolyte is instead slightly more complicated, since it is a liquid. Its ionic conductivity can be evaluated as: [26]

$$\sigma_I = \frac{F^2}{RT} \sum_i z_i^2 D_i^{eff} c_i$$

The term D_i^{eff} witnesses the dependence on the diffusion phenomenon. Once the conductivity is found, the resistivity is simply its inverse, and then the good old (2) can be applied. It is immediate to get how the path length l is important, to the point that the compression of the electrode can bring great results concerning the reduction of the ionic resistance.

iii. Concentration overpotential

For the redox reaction to occur, reactants are of course consumed and turned into products, which are initially found at the AAs. This means that a concentration gradient of the reactants arises: there is lower availability of them where the reactions should occur, and anyway some of the useful space is already occupied by products. This phenomenon hinders the reactions, and provokes the concentration overpotential, the third type of voltage drop. It becomes relevant as soon as very high current densities are asked from the cell, to the point that the electrolyte flow is not enough anymore to provide the necessary reactants and to wash away the products. Looking once again at Figure 2.5, it is visible how the cell voltage sharply drops to zero (i.e. the power becomes nil), thus identifying a maximum value for the current. It can be approximated as:

$$I_{limit} = AnFk_m c_r$$

With k_m local mass transfer coefficient:

$$k_m = a \cdot u^b$$

With u flow velocity and a and b fitted parameters. It is clear how a greater velocity yields to a greater local mass transfer coefficient, which yields to a greater current limit. Acting on the velocity requires a price, though, since a greater power has to be spent in the pumps to provide it, in case. [32] Another parameter that can be useful in this issue is the reactant concentration c_r , but it mainly depends on the *SOC*. A charged VFB has a greater I_{limit} than a discharged one, but this was arguably foreseeable as well.

Anyway, for a given electrolyte velocity and *SOC*, it is possible to find a limit current, to which corresponds a concentration overpotential, according to: [30]

$$\eta_{con} = \frac{RT}{F} \ln \left(1 - \frac{I}{I_{limit}} \right)$$

The behaviour of η_{con} VS I is plotted in Figure 2.6.




Figure 2.6: The concentration overpotential VS the current.

Recalling the initial (1):

$$\begin{aligned}
 U_{cell} &= OCV - \sum_k (IR)_k - \sum_k |\eta_k| \text{ [V]} = \\
 &= OCV - \sum_k (IR)_k - |\eta_{act,cat}| - |\eta_{act,an}| - |\eta_{con,cat}| - |\eta_{con,an}| \text{ [V]}
 \end{aligned}$$

All the addenda have been described. Amongst them, the ohmic losses related to the electrodes were dominant: the felts through-plane resistances and the contact resistances with the BPs. [25]

d. Shunt current losses

Shunt currents arise from the fact that the cells are in an electrical series, while in a hydraulic parallel. Since their voltage is not perfectly the same, points with different potential are connected by a conductive path, so a current can flow. This phenomenon happens both during normal operation, and during stand-by, thus provoking self-discharge.

An ideal design is able to both prevent shunt currents and to minimise the pressure drops along the electrolyte paths. The requirements of these two desires collide, though, since the first would need long narrow channels, while the latter short wide ones. An optimisation is needed, and anyway the results will favour the shunt currents reduction, since their impact is usually ten or twenty times stronger than the pressure drops' one. [7]

When dealing with the electric physics, it is generally useful to introduce an equivalent electric circuit, in order to describe the phenomena with the laws of the electrotechnics. Figure 2.7 reports an interesting one. [33]




Figure 2.7: The equivalent electric circuit of the stack, depicting the shunt currents paths. [33]

In the circuit, there are several voltage generators along a straight line: they represent the cells of the stack. Their voltages V are not necessarily equal. Each one of them is in series with a resistance R , which represents the cell internal resistance. V and R depend on the SOC of the battery.

Were the shunt currents absent, the circuit would just be a long series of voltage generators and internal resistances, whose total current would be the one found as output of the battery. There are many other elements in the circuit, though. Every couple $V - R$ (but the first and last) in fact is part of four distinct small rings, two comprehensive of the series $R_c^+ - R_m^+ - R_c^+$, and the other two of the series $R_c^- - R_m^- - R_c^-$. They depict the possible paths for shunt currents. The superscript $+$ or $-$ indicates whether the catholyte or the anolyte is involved in such a path, while the subscript c or m indicates what part of the hydraulic circuit that resistance represents: either a channel or a manifold. The internal resistance of the cell was already taken into account by R . Looking at any cell (but the first and last), there are two electric circuits available for the catholyte, and two for the anolyte. They represent the path coming from and returning to the tanks. Indeed, the leftmost cell (the one at the highest potential) only features the catholyte paths, while the rightmost (the one at the lowest potential) only the anolyte paths.

The knots of the rings of any cell are connected to those of the adjacent ones, in fact shunt currents can arise between two distinct cells and follow even longer paths, at the price of encountering more resistances.

Under an electrotechnics point of view, one might try to simplify the above circuit, but its abundance of resistances is useful to represent all the different possibilities for the shunt currents to flow. Furthermore, resistances are not of a fixed value: they both depend on the operating temperatures and on the SOC. In particular, according to the well-known:

$$R = \rho \frac{l}{A}$$

besides the role of the channel geometry, the electrolyte resistivity is important. Reminding that $\rho = \frac{1}{\sigma}$, it depends on the species concentration:

$$\begin{aligned} \sigma_+ &= SOC_{n+} \sigma_V + (1 - SOC_{n+}) \sigma_{IV} \\ \sigma_- &= SOC_{n-} \sigma_{II} + (1 - SOC_{n-}) \sigma_{III} \end{aligned}$$

This can be simply seen as a weighted average between the conductivities of the different stages of vanadium, being the SOC the weight. [15]

The values of such conductivities can be found in literature, and are for example here reported in Table 2.2: [15]

Table 2.2: The conductivities of the vanadium species.

V^{2+}	σ_{II}	27.5 $\left[\frac{S}{m}\right]$
V^{3+}	σ_{III}	17.5 $\left[\frac{S}{m}\right]$
VO^{2+}	σ_{IV}	27.5 $\left[\frac{S}{m}\right]$
VO_2^+	σ_V	41.3 $\left[\frac{S}{m}\right]$

This means that it will always be $\sigma_+ > \sigma_-$, and in particular their best values are found for $SOC = 100\%$.

e. Cross-over losses

The catholyte and the anolyte are meant to flood the two electrodes which constitute the two half-cells, and which are separated by an ion-exchange membrane. Such a surface will be crossed by ions (whether positive or negative), in order to maintain the charge neutrality, while electrons cross the BPs, eventually giving the output current.

The issue is that the membrane also allows for some vanadium ions to pass, so little quantities of electrolyte take an unexpected path, resulting in a loss. Such a loss comes both from the fact that some vanadium is just lost without releasing or absorbing useful energy, and from the fact that it will react with other vanadium ions, releasing energy in the detrimental form of heat.

This phenomenon is known as cross-over, and follows Fick's first law:

$$J = -D \frac{\partial c}{\partial x}$$

With:

- $J \left[\frac{mol}{m^2 \cdot s}\right]$: diffusion flux, the amount of substance that flows through a surface during a unit time interval;
- $D \left[\frac{m^2}{s}\right]$: diffusivity constant;
- $\frac{\partial c}{\partial x} \left[\frac{mol}{m^3 \cdot m}\right]$: concentration gradient.

The negative sign indicates that the flux goes towards the region with the lower concentration. It is possible to rewrite the equation in a more general way, comprehending all the three space dimensions:

$$\mathbf{J} = -D \cdot \nabla c$$

Now \mathbf{J} is a 3D vector.


\mathbf{J} alone is not enough to describe the cross-over phenomenon, though: not only has every vanadium specie its own distinct diffusivity constant D (the lower the valence, the faster the crossing [6]), but also different concentration at any given moment. This means that the above formula is actually function of many variables, above which there is also the SOC . In fact: [34]

$$SOC_+ = \frac{c_V}{c_V + c_{IV}} = \frac{c_V}{c_+}$$


$$SOC_- = \frac{c_{II}}{c_{II} + c_{III}} = \frac{c_{II}}{c_-}$$

Of course, the temperature plays a role too. [15]

Once the vanadium ions have crossed the ion-exchange membrane, they react with the other species of the electrode. The reactions are exothermic and eventually result in cross-over energy losses, which cause self-discharge and loss of efficiency. In particular, when V^{2+} arrives at the positive electrode: [6]




When V^{3+} arrives at the positive electrode:




Not only are $V(II)$ and $V(III)$ lost, but also $V(V)$ is wastefully turned into $V(IV)$.

When instead VO^{2+} arrives at the negative electrode:



When VO_2^+ arrives at the negative electrode:



Not only are $V(IV)$ and $V(V)$ lost, but also $V(II)$ is wastefully turned into $V(III)$. [15,35]

It might also happen that $V(II)$ reacts with $V(IV)$ in the positive electrode, or that $V(V)$ reacts with $V(III)$ in the negative electrode, but such reactions are of negligible importance, since their products further react in ways which are already described by the other reactions above. [35] Amongst the four species, $V(II)$ has the highest D . [15]

The battery energy reduction rate due to $V(II)$ cross-over can be described by:

$$P_{co} = -\Delta G \frac{dM_{II}}{dt} = zFE_0 \frac{A}{d} \sum_{n=1}^N (D_{II,n}c_{II,n-} - 2D_{V,n}c_{V,n+} - 2D_{IV,n}c_{IV,n+})$$

With:

- $\Delta G = zFE_0$: total Gibbs reaction energy (E_0 cell Nernst potential);
- A : membrane AA;
- d : membrane thickness.

Not only does the mass loss over time depend on the considered specie's migration, but also to the others': the less $V(IV)$ and $V(V)$ are available, the less the escaped $V(II)$ can actually provoke an energy loss.

The only way to reduce the cross-over losses is to choose a membrane which can hinder the vanadium migration. The problem is that this would also slow down the H^+ passage. In the end, a membrane with a high ion conductivity is more important than one with reduced vanadium crossover [4,36], so this aspect can only be optimised up to a certain point.

f. Hydraulic losses

Since the electrolytes are liquid and stored in external tanks, pumps are necessary to move them to the stack and back. Nevertheless, they introduce an energy consumption, and they have an efficiency themselves which plays its part in the overall analysis of the VFB. Luckily for industrial scale plants, the efficiency of a pump is usually the higher the bigger the pump itself, and this aspect should be kept in mind whenever dealing with a laboratory scale study.

Similarly to the voltage, the pressure drops along its path too. In general, it is possible to express the net hydraulic power as:

$$P_h = Q_+ \Delta p_+ + Q_- \Delta p_-$$

With

- Q : electrolyte flow rate;
- Δp : pressure drop;
- $+$ and $-$ subscripts to indicate whether the value is referred to the catholyte or anolyte.

Each pressure drop can then be considered as the sum of two addenda: a drop in the tank, and one in the pipes. They furthermore depend on the regime of the liquid, therefore on the flow itself, so the above expression could be rewritten as:

$$P_h = Q_+ \Delta p_+(Q_+) + Q_- \Delta p_-(Q_-)$$

The flow regime in the pipes can be identified to be either laminar or turbulent, respectively for Reynolds numbers lower or higher than 23,000. [15] Recalling the laws of fluid dynamics, this number can be evaluated as:

$$Re = \frac{\rho u_x d_x}{\mu} [1]$$

With:

- $\rho \left[\frac{kg}{m^3} \right]$: fluid density;
- $u_x \left[\frac{m}{s} \right]$: fluid velocity in the x^{th} straight segment;
- $d_x [m]$: hydraulic diameter of the segment;
- $\mu \left[\frac{kg}{ms} \right]$: dynamic viscosity.

Once the Reynolds number is known, it is possible to proceed with the pipe friction factor:

$$f_x = \begin{cases} \frac{64}{Re} & \text{if } Re < 23,000 \\ \left[-1.8 \cdot \log\left(\frac{6.9}{Re}\right) \right]^{-2} & \text{if } Re > 23,000 \end{cases}$$

In any case, the regime in a VFB is usually of the laminar type, therefore with Re below 23,000.

Figure 2.8 shows a representation of $f_x(Re)$, with a magnification of low Reynolds numbers (laminar regime) and high ones (turbulent regime).




Figure 2.8: a) The friction factor versus the Reynolds number, with magnifications focusing on b) low and c) high Re .

With the friction factor, it is finally possible to find the distributed term of the pressure drop along the straight part of the pipes:

$$\Delta p_{di\pm} = \frac{\rho}{2} \sum_x f_x \frac{l_x}{d_x} u_x^2 \text{ [Pa]}$$

The plus or minus sign still indicates whether the evaluated quantity refers to the catholyte or to the anolyte. The other pressure drop addendum to be found is the one that takes into account localised (or concentrated) losses:

$$\Delta p_{cn\pm} = \frac{\rho}{2} \sum_x \zeta_x u_x^2$$

With ζ_x empirical minor loss coefficient, to be found in literature or from data sheets.

The sum of $\Delta p_{di\pm}$ and $\Delta p_{cn\pm}$ eventually gives the total pressure drop Δp_{\pm} .

Since this value is to be minimised, it is necessary to identify which parameters can be acted upon. Velocity u_x is clearly a main character, but it is chosen accordingly to the output power desired at the given moment, so it cannot become a constraint. Anyway, the higher the velocity, the more turbulent the flow, the higher the losses. Neither can ρ be freely modified, since it depends on the chemical composition of the electrolytes, which is crucial for the chemical reactions.

It is preferable to operate with low viscosities μ , but this depends on the electrolyte composition too. Anyway, lower viscosities also ease the operation of the pumps. It is interesting to notice that μ slightly decreases at higher operating temperatures, though. [37]

A path with fewer turns or anyway spots where localised pressure drops can arise is suggestable: this is a matter of geometry and of management of the available spaces.

The last two parameters that can play a role are the path length and diameter: a short thick pipe causes less pressure drops, since it eases the fluid movement. The point is that, as previously seen, this shape would reduce the electric resistance faced by the shunt currents, thus increasing that other type of losses. A compromise is to be found between the two requests, with an optimisation.

Anyway, the pressure drops inside the electrode are more important than those in the pipes, so the reduction of its resistance to the electrolyte flow cannot be neglected. Darcy's law helps evaluating the value of the pressure drop inside a porous electrode: [38]

$$\Delta p_{felt} = \frac{\mu \cdot l \cdot Q}{\kappa \cdot A}$$

With μ viscosity, Q flow rate, l length, A cross-sectional area, κ permeability, found as:

$$\kappa = \frac{d_f^2}{16K} \frac{\varepsilon^3}{(1 - \varepsilon)^2}$$

With d_f fibre diameter, K Kozeny-Carman constant and ε porosity. A very low porosity ($\varepsilon \approx 0$, i.e. a compact electrode) yields a low permeability κ , and eventually a huge pressure drop Δp_{felt} . On the other way around, a high porosity eventually leads to a low pressure drop, and in fact it means that the fibres do not hinder the electrolyte flow. [25]

g. Impact of the electrode compression

After all the previous paragraphs, it is evident how the electrode plays a pivotal role in the performances of the VFB cell. Its porosity impacts on the electrolyte flow, thus on the hydraulic losses and on its velocity. The limit current depends on this velocity, and therefore the transport losses too. Furthermore, the compression of the electrode influences its length, regarding the ion current path, and the contact with the BP and the ion-exchange membrane, thus the ohmic resistances too.

It is clear that a diligent description of the electrode behaviour cannot therefore neglect the importance of its compression. It all starts with the porosity concept, which was already introduced:

$$\varepsilon = \frac{V_{empty}}{V_{tot}} [1]$$

The porosity is the ratio between the empty space left amongst the carbon strains in the microstructure of the felt and its total volume. The higher the porosity, the more empty space there is. This value can be expressed either as a unit or as a percent, of course.

It is useful to refer to this empty space with the term porosity, to imagine it as composed of spheric pores, but it is not accurate. In fact, the carbon felt is made of many little strains, woven together in an irregular pattern, and their cylindrical shape does not cut out spheric holes, but rather an irregular space where the electrolyte flows. It is therefore even difficult to determine the diameter of the pores, if anyway one wanted to refer to them to describe the emptiness. Such an irregular structure also leads to the fact that the same carbon felt can behave in different ways under different circumstances, so a unique description of its properties is even more complicated. [24]

The key concept of compression can be quantified by the compression ratio (CR), which is nothing but the unit minus the ratio between the final and the initial thickness of the electrode:

$$CR = 1 - \frac{d_{fin}}{d_{in}}$$

Maybe expressed as a percentage. The relationship that instead directly binds the porosity ε to the CR is not linear. In fact, the volume can always be considered the sum of two addenda: the empty one (given by the empty spaces amongst the strains: the pores) and the full one (the strains themselves). When the compression occurs, the empty space alone is pressed, the full one is not (as long as the compression does not reach very high values, where the strain-stress relation of the graphite starts to be involved as well). Calling ε_0 the porosity of the uncompressed felt, it is possible to evaluate the porosity ε of the same felt for a given compression with: [39]

$$\varepsilon = \frac{\varepsilon_0 - CR}{1 - CR} [1]$$

For example, if the starting porosity is $\varepsilon_0 = 80\%$, after a halving of the volume ($CR = 50\%$) it becomes $\varepsilon = 60\%$. Assuming a unitarian initial volume, in fact, the initial empty space is 0.8, while the full space the remaining 0.2. The compression halves the whole volume, i.e. removes 0.5, but this quantity is taken away from the empty space alone, which then becomes 0.3. So, in the end there is 0.3 empty space over a total of $0.3 + 0.2 = 0.5$: the 60%. This formula works as long as the CR percentage is smaller than the ε_0 percentage, since negative values of ε do not make sense. Such high compressions imply in fact that the carbon strains are deformed as well: a phenomenon that is ignored in this formula.

This compression requires a pressure to be applied on the electrode, usually coming from the BPs.

Due to the irregular microstructure of the felt, the relationship that binds the compression ratio to the applied pressure is not linear: it is shown in Figure 2.9.




Figure 2.9: The relationship between the compression ratio and the compressive pressure: it is not linear. [27]

An explanation of this behaviour is found in the fact that the pores are different in size, and the bigger ones will also be the first to be compressed. Once they are squeezed, the compression acts on the new bigger ones, thus resulting in the trend of Figure 2.9, where an increasing effort is needed to keep reducing the thickness of the electrode. Keeping in mind that the empty space is not actually composed of pores, it is anyway possible to estimate their (equivalent) diameter, and find the corresponding distribution, as shown in Figure 2.10.

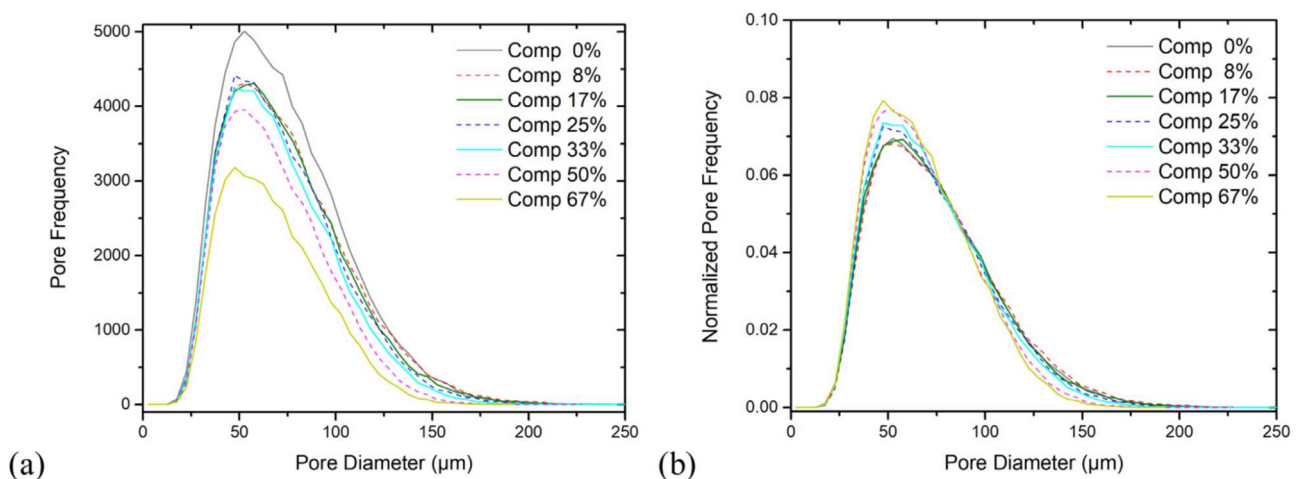




Figure 2.10: The distribution of the pore diameters for different levels of compression: a) the frequency and b) the normalised frequency. [27]

Raising the pressure and therefore the compression, it is clear that the amount of large pores immediately decreases, until there are not any more of that size. Once they are squeezed, they belong to the group of pores with smaller diameter, thus further increasing the proportion of these ones. It is interesting to notice how the

frequency of the smallest pores almost does not change, i.e. they are not affected by the compression. It is not too much of a surprise, nevertheless: after all they are the most difficult to compress. The normalisation highlights how the curve is gradually sharpened to the left, and made steeper on the right. Figure 2.11 proposes another way to analyse the same distribution.






Figure 2.11: The diameters of the pores for different values of compression: a) their total volume, b) their average diameter, c) the maximum diameter. [27]

Since the total pore volume corresponds to the empty space within the electrodes, it is not strange that its relationship with the compression in Figure 2.11 a) is linear. Since the pore volume is related to the diameter accordingly to the formula of the volume of the sphere, it is natural that the mean diameter decreases with the compression too, as in Figure 2.11 b). Last, the maximum diameter must decrease too, as it was shown back in Figure 2.10, and in fact it does, as in Figure 2.11 c).

A prettier representation is proposed in Figure 2.12, where it is possible to directly look at the sponge-like tissue of the felt throughout the compression.




Figure 2.12: The felt throughout the compression: the impact on the porous structure is visible. [27]

This porous structure that changes with the compression ratio is crossed by the electrolyte flow. The way in which this happens depends on the hydraulic configuration: in the flow-through case, the electrolyte just passes from one extremity to the other through the porous structure; in the flow-by, by jumping from one channel to the other, crossing a small portion of electrode. In any case, larger pores ease the passage of the electrolyte, while smaller ones make it difficult, but on the other side usually offer more surface area for the redox reaction. It might even happen that a portion of the electrode is not accessible at all, due to local obstacles for the flow, so the invasion pattern might be quite irregular.

The highest velocities are reached where the carbon fibres do not oppose the electrolyte flow, while the regions where the strains are very close (due to compression or not) hinder it: the maximum velocity can even be 25 times higher than the one at the inlet. [24] It would be useful to have a formula for the evaluation of the tortuosity of the electrolyte path, but such a value is very complex to determine, since it depends on a lot of variables. The easiest way to find it is by experimental data or simulations. Figure 2.13 shows a possible invasion pattern for the electrolyte in the electrode for different compression ratios.

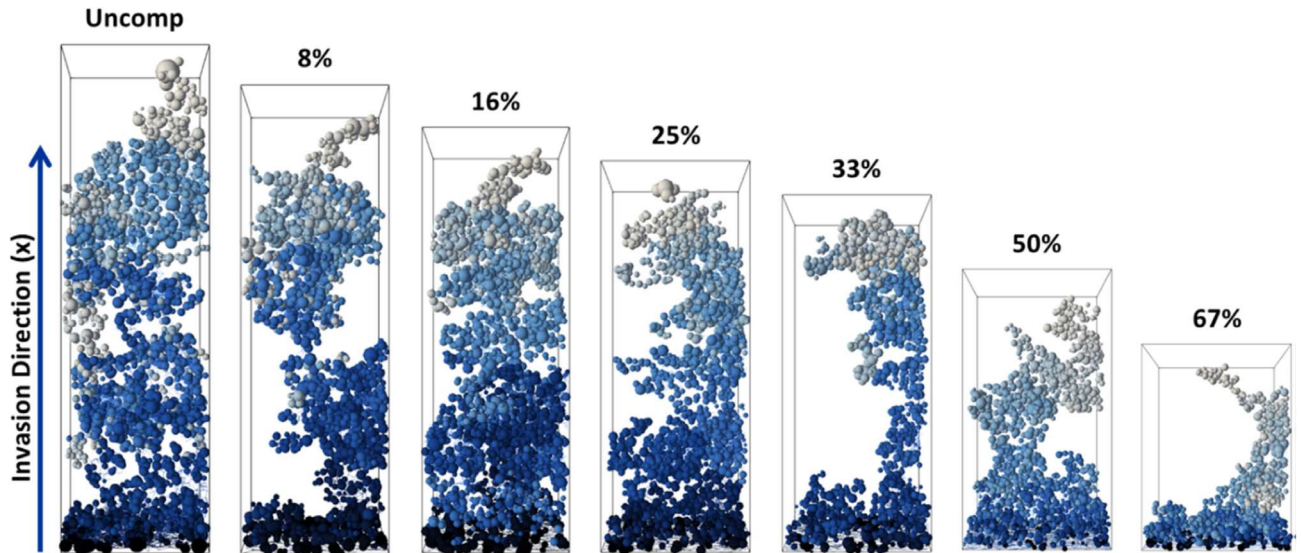


Figure 2.13: An invasion pattern of the electrolyte in the electrode for different compression ratios. [27]

The different colours represent the progressive instants: the dark blue dots the initial positions of the electrolyte, while the white ones the final positions.

The higher the number of phases in the diffusion, the more difficult it is to foresee their development.

The permeability clearly is a useful tool to better describe the behaviour of the electrode while the electrolyte is passing through it. Unsurprisingly, it depends on the porosity ε , besides other quantities: [25]

$$\kappa = \frac{d_f^2 \varepsilon^3}{16K(1 - \varepsilon)^2}$$

As it was already introduced in paragraph 2.f. The above dependence of κ on ε is shown in Figure 2.14.




Figure 2.14: Semilogarithmic plot of the relationship of the permeability κ with the porosity ϵ .

And since the porosity itself ϵ depends on the compression ratio CR , it is possible to directly link κ and CR , as shown in Figure 2.15: there, the permeability values of two different felts are shown: R5 and P5. [25]




Figure 2.15: Dependence of the permeability κ of two electrodes felts (R5 and P5) on the compression ratio CR. [25]

It is possible to observe how a gradual increase in the compression leads to a monotonous decrease in the permeability of the felt, due to the reduction of the dimensions of the pores and to a probable reorganisation of the microstructure of the felt itself. Besides the CR and its consequences, the reduced thickness of the electrode plays a role too in the determination of the final hydraulic resistance the electrolyte has to face along its path in the felt.

It is also interesting to notice that the permeability is not influenced by the operating temperature of the cell; since the electrolyte is less viscous at higher temperatures, the pressure drop in the felts decreases anyway, following a linear relationship: [25]

$$\Delta p = \frac{\mu l Q}{\kappa A}$$

Where μ is the electrolyte viscosity, l the electrode length, Q the electrolyte flow rate, κ the permeability and A the cross-sectional area. Table 2.3 provides an example for the values of permeability of felts P5 and R5 at different temperatures, together with the viscosity of the electrolyte.

Table 2.3: The permeability of the P5 and R5 felts and the electrolyte viscosity for different temperatures. [25]

Temperature °C	Electrolyte viscosity mPa s	Permeability κ for felt P5 m ²	Permeability κ for felt R5 m ²
15	6.97	1.78×10^{-10}	2.64×10^{-10}
25	5.47	1.79×10^{-10}	2.64×10^{-10}
40	4.28	1.78×10^{-10}	2.66×10^{-10}

Alike the temperature, a different choice of electrolyte does not influence the permeability of the electrode. [29]
The interstitial velocity is instead: [25]

$$v = \frac{Q}{A\varepsilon}$$

The formula of the flow $Q = vA$ is elementary in the hydraulics context, but here it is clear the impact the porosity has on it. If $\varepsilon = 1$ (i.e. $V_{empty} = V_{tot}$) all the space to cross is free of obstacles and $Q = vA$ simply, but this does not apply to the electrode case. The higher the porosity ε , the faster the electrolyte, keeping the flow Q at a constant value, since it has to pass through narrow paths amongst the carbon fibres: this reminds to Bernoulli's principle.

Not only does the felt provoke a pressure drop inside the electrode, but it also imposes a condition for the electrolyte to enter. This condition can be represented by the breakthrough pressure, which is the minimum one required for a liquid to overcome the capillary pressure in the channels with the largest pores. It linearly increases with an increasing compression ratio. [27] The electrolyte consequently needs to be pumped with a stronger pressure in order to pass the whole electrode width, in the flow-through configuration. In the flow-by one, instead, it is just sufficient for the electrolyte to reach the membrane from the channels. [27] Unlike the breakthrough pressure, though, the pressure required to fulfil the felt does not vary with compression, since the smallest pores are almost not affected by it. [27]

Concerning the electrolyte flow, it can be attributed to the phenomena of viscosity and diffusivity. The viscous transport is due to the internal shear stresses: it depends on the permeability, so it decreases with higher compression ratios. The diffusion is instead due to a concentration gradient, and the correspondent transport is less dependent on the compression ratio. This means that processes which mostly rely on diffusion are less affected by compression than those that mostly rely on viscosity. [27] This all (the change in the transport dynamics and the reduction in the pores' dimensions) can negatively affect the homogeneity of the flow, together with the reorganisation of the felt microstructure. Furthermore, the reduced permeability is also detrimental for the electrolyte convection. [23,40] It might happen that some parts of the electrode are not reached by the electrolyte at all, thus some of the AA is not used and just wasted, with a reduction of the cell performances. [24] This detail is also influenced by the type of chosen molecule: surprisingly, organic electrolytes perform better than vanadium-based ones. One might rightfully expect the smaller vanadium-based molecules to move better, due to their minor dimensions, with respect to the bigger organic ones. This is counterbalanced though by the higher rate constant of organic electrolytes: they have a faster charge transfer, which eventually results in better current values. [24] This is the reason why the AA is not the main parameter in the cell design, but rather the ASR. [24]

In the end, the pressure drop in the felt depends on a variety of parameters: the compression ratio (mostly due to the breakthrough pressure), the AA, the electrolyte flow etc.

Figure 2.16 provides an example of these dependencies. [40]




Figure 2.16: The pressure drop in the electrode VS the flow rate, for various compression ratios. The dashed lines represent a wide AA, the continuous lines a smaller AA.

There is a monotonous dependence of the pressure drop on the flow rate, which is deeply influenced by other factors. For example, the higher the compression ratio, the higher the pressure drop, due to the contribution of the breakthrough pressure. The AA of the electrode plays a crucial role too: a wider one means that the carbon straws are not in close contact one with the others, so there is more space for the liquid to flow, leading to lower pressure drops (continuous lines). A smaller AA means instead that the straws are more pressed, with narrower spaces amongst them and therefore with a more difficult passage for the liquid: a higher pressure drop. In particular, the effect of the electrode compression over the pressure drop is not linear, as shown in Figure 2.17. [26]

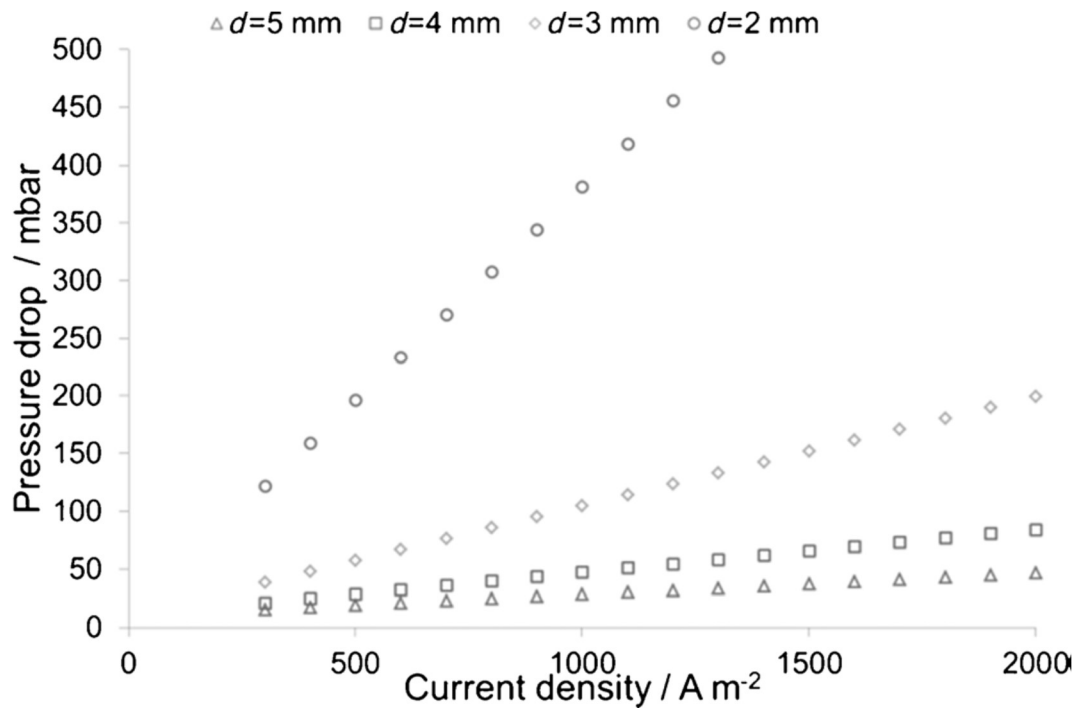


Figure 2.17: The pressure drop VS the current density for various electrode thicknesses, i.e. compression ratios. [26]

The example of Figure 2.17 shows that a compression from $d = 5$ [mm] to $d = 4$ [mm] does not particularly increase the pressure drop, while one from $d = 3$ [mm] to $d = 2$ [mm] has a way stronger impact.

Once the pressure drop Δp is known, it is possible to finally evaluate the power spent in the pumps to have the electrolyte flow:

$$P_{pump} = \frac{\Delta p \cdot Q}{\alpha}$$

Where α is the pump efficiency.

While the use of high flow rates is generally preferable to decrease the concentration overpotential, to achieve high cell efficiencies, it also increases the pumping power, which has to be extracted from the VFB itself, eventually lowering the overall efficiency of the system. It is good news anyway that large pumps usually have good efficiencies: better than those of smaller dimensions, which are by the way utilised in laboratory scale experiments. Regarding the electrode compression, it increases the pressure drop (because of the breakthrough), and eventually reduces the overall efficiency of the battery of a value between 1 and 20%. [26]

Of course, there are also other places where pressure drops occur (manifolds, pipes), but those in the electrode felts and FF channels are the most important. [25]

Recalling that $v = \frac{Q}{A\varepsilon}$, it is necessary to keep in mind that the electrode compression will also cause a reduction of the cross-sectional area A , so the electrolyte will be forced to pass through a narrower region. If the flow rate Q is imposed, this implies higher velocities; if the velocity is instead imposed, the flow rate is reduced. [25]

Figure 2.18 reports the relationship between the pumping power and the compression ratio for two different types of electrode felt: P5 and R5. [25]




Figure 2.18: The pumping power versus the compression ratio for two electrode felts: P5 and R5. [25]

It is evident that the properties of the chosen felt are crucial: its microstructure, porosity, and all the other mechanical and textural characteristics lead to different optimal CR s. [25] Once again, it is a matter of optimisation, since the homogeneity of the electrolyte distribution and the cell performances depend on CR too. The type of flow furthermore is important as well, and it is described by the Reynold's number. The dependency is shown in Figure 2.19. [29]




Figure 2.19: The dependency of the pumping power on the Reynold's number for different compression ratios. [29]

A higher Reynold's number (i.e. a more turbulent regime) leads to higher pressure drops, and therefore pumping powers. The more compressed the electrode, the higher the power. [29]

Not only do the compression and the porosity variation influence the hydraulic properties of the cell, but also the electrical ones, as it will be now described.

The compression in fact squeezes the pores, pressing the carbon strains together, improving their electrical contact amongst them and reducing the resistance. The better the contact, the lower the resistance, but this value cannot decrease indefinitely, but rather reaches an asymptotic value, which is proper of the felt material itself. The relationship between effective electrical conductivity and porosity is the Bruggemann correction:

$$\sigma_s^{eff} = (1 - \varepsilon)^{\frac{3}{2}} \sigma_s$$

Thus leading to the graph of Figure 2.20, which also reports experimental data. [26]




Figure 2.20: The electrical conductivity of the electrode versus its thickness, reporting experimental and expected data. [26]

Being the resistivity ρ the inverse of the conductivity σ , the previous behaviour naturally leads to the one shown in Figure 2.21, which has the CR on the x axis, though, not the thickness. [27]




Figure 2.21: The electrical resistivity ρ versus the compression ratio CR. [27]

The electrode resistance is the product of the resistivity times the path length, over the cross-sectional area. This yields to a slightly different graph than the one of Figure 2.21. The resistance versus the compression is shown in Figure 2.22. [27]




Figure 2.22: The electric resistance of the electrode versus the compression. [27]

Looking at Figure.2.22, it is clear that the more the electrode is compressed, the easier it is for the electrons to cross it, resulting in a lower ohmic voltage drop. This is the result of the combined effect of the reduced resistivity (due to the enhanced contact amongst the carbon fibres) and the reduced thickness to cross.

In any case, the ionic conductivity is two orders of magnitude lower than the electrical one, resulting in its dominion concerning the ohmic losses. [26] Some numerical values for ohmic losses, in a comparison with activation overpotential, are provided in Table 2.4.

Table 2.4: Activation overpotential η and ohmic losses IR for different electrode thicknesses d . [26]

d [mm]	η [mV]	IR [mV]
5	6.9–41.3	56.5–337.5
4	6.9–41.4	48.0–286.8
3	6.9–41.6	40.5–243.9
2	6.9–42.0	33.4–205.6

The ASR is therefore dominated by the ohmic contribution.

Recalling the formulas for the activation overpotential from paragraph 2.c.i:

$$\eta_{act} = \pm \frac{2.3RT}{\alpha F} \log j_0 \mp \frac{2.3RT}{\alpha F} \log j \text{ [V]}$$

where the choice of the signs depends on whether the equation is referred to the anode or cathode. It is shown that no parameter depends on the compression of the electrode. R and F are in fact constants, the temperature T is the operative one and α is the transfer coefficient (of the anode or cathode). Regarding this last quantity, it is:

$$\alpha = \mp \nu \frac{RT}{nF} \left(\frac{\partial \ln |I_{red}|}{\partial E} \right)_{p,T,c_i^{interface}}$$

With the minus sign to refer to the cathode, plus to the anode. ν is the stoichiometric number, the derivative $\frac{\partial \ln |I_{red}|}{\partial E}$ is between the logarithm of the partial cathodic or anodic current and the electrode potential. Neither does the limit current density j_0 depend on the compression ratio. The constancy of the activation overpotential with respect to the electrode thickness was also proven in Table 2.4, whose values also make clear why the impact of the ohmic losses is heavier than that of the activation, so the firsts dominate over the seconds in the ASR. Figure 2.23 [26] shows the steepness of the functions.




Figure 2.23: The ASR and η_{act} VS the electrode thickness for various specific surface areas a_0 . [26]

The sequence of white triangles represents the values of the ASR for a certain specific surface area: it increases with the electrode thickness, i.e. it decreases with its compression. This behaviour is due, as previously explained, by the decrease in the ohmic losses with compression, even though the activation overpotential stays constant. In fact, the latter is represented by the black triangles, which stay at the same height in the graph throughout the electrode compression. The multiplication of the specific surface area by a factor of 10 leads to the lines of squares: their behaviours are the same of the triangles, but are placed lower in the graph. This means that a higher specific surface area leads to lower resistances. In fact, the redox reactions have more available space to occur, so the activation is easier. This shifts down the ASR as well, which more or less maintains the same steepness as before. A further increase in a_0 further reduces η_{act} .

To sum up: higher a_0 for lower η_{act} , while higher CR for lower RI. [26]

The compression cannot though reach exaggerated values. On one hand, this would provoke strong pressure drops for high current densities, eventually leading to a leakage of the electrolyte from the stack. This is both a security hazard, due to its acidic nature, and detrimental for the battery performances, of course. [40] A wrong compression value can even lead to the breakage of the carbon fibres of the felt, thus ruining it. [26]

As Figure 2.24 witnesses, the impact of the compression ratio on the cross-over losses is minimal. [41]

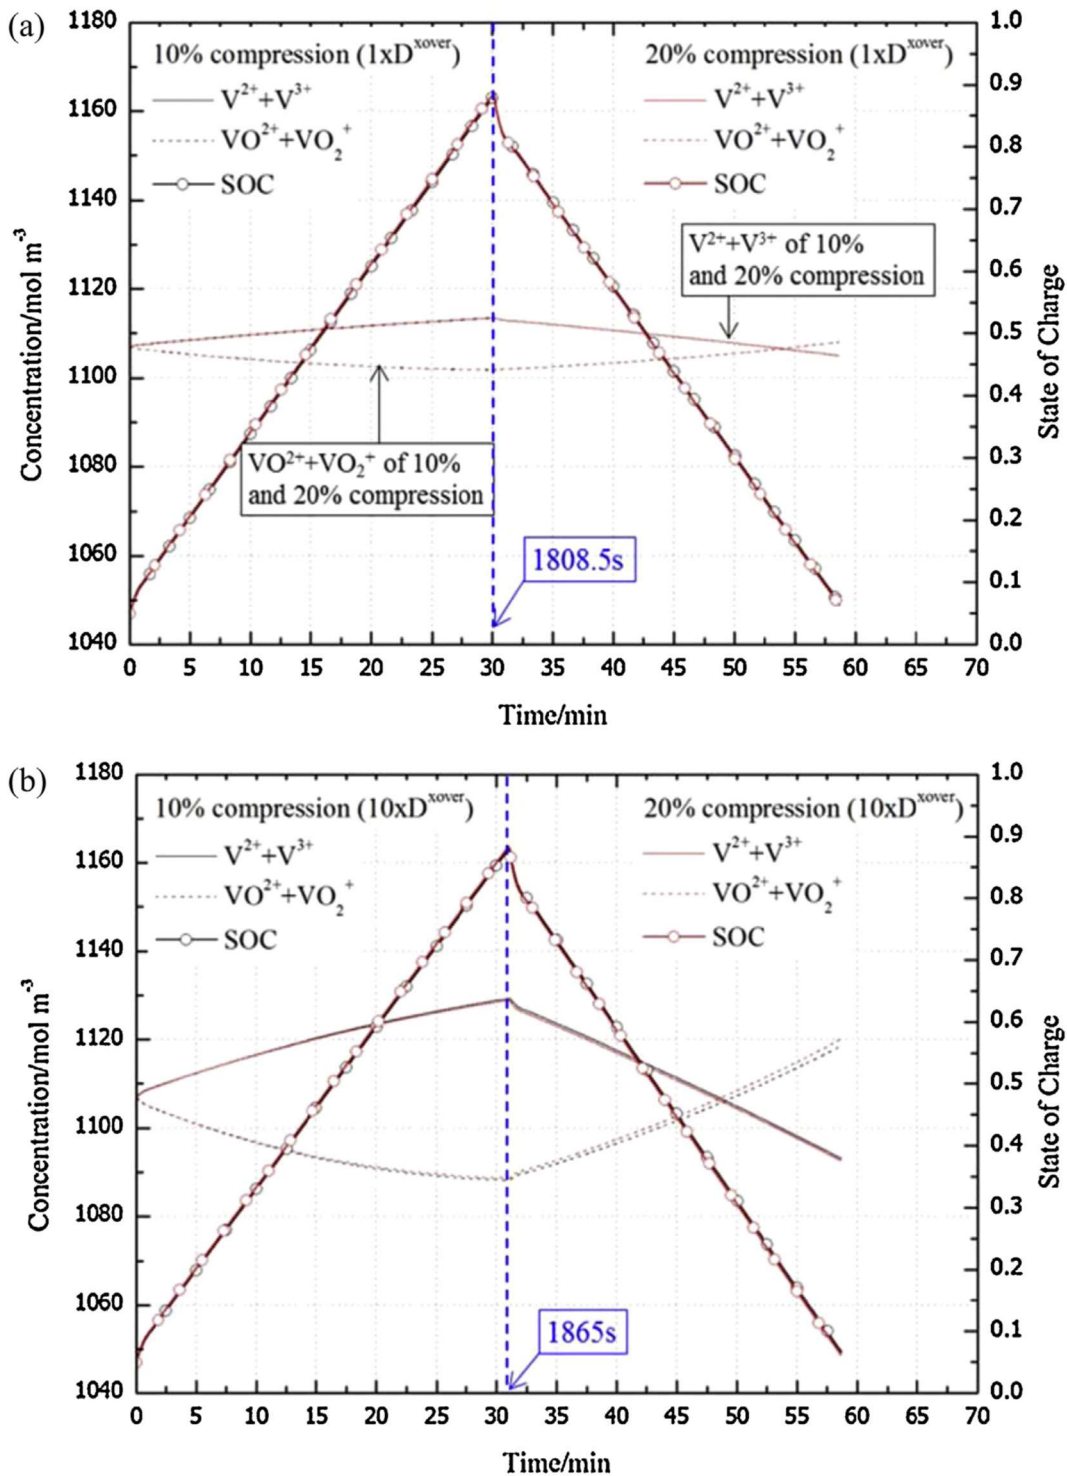



Figure 2.24: The anolyte and catholyte concentrations throughout charge and discharge, versus time, for two values of compression ratio. b) has ten times the diffusivity constant of a). [41]

The left halves of Figure 2.24 a) and b) represent the charging phase, and in fact the red lines with circles (the SOC) steadily grow. The diffusivity constant of b) is ten times that of a), so the cross-over phenomenon is stronger, thus leading to a slower charge, which anyway just takes 56.5 [s] longer. The right halves show instead the discharging phase: that of b) is slightly faster than that of a) for the same reason as before. The concentration of the anolyte (continuous line) and of the catholyte (dashed line) are not constant throughout the process, due to the cross-over phenomenon. For example, during discharge, the anolyte concentration decreases, in fact:



Which means that $V(II)$ is lost from the anolyte, migrates to the catholyte, and turns into $V(IV)$, thus changing the concentration of the two electrolytes as shown in the graph. Another similar reaction happens instead during the charge. As a result, the concentrations vary through the process, but after a full cycle come back to the original values, no matter the compression ratio. This last detail is proved by the fact that the discharge phase happens with a $CR = 20\%$, while the charge phase with $CR = 10\%$. Despite this difference, the concentration variations cancel each other, which implies that the cross-over phenomenon is (almost) not affected by the compression. To sum up, the compression has multiple effects: in the hydraulics context it increases the pressure drops, thus requiring higher pumping energy, but it also enhances the electrolyte velocity, which leads to a better penetration in the felt [40,42]; in the electric context instead, it reduces the ASR . If the best compression ratio to apply is to be found, an optimisation is necessary.

It is useful at this point to recall some quantities that might be set as variables of the optimisation: the three aforementioned efficiencies.

The coulombic efficiency CE describes how well the electrons are stored into and taken from the system, and can be used to track side-reactions that can lower the performances of the stack, like faradaic losses.

$$CE = \frac{\text{discharge capacity}}{\text{charge capacity}}$$

A full coulombic efficiency might be obtained with the adoption of an anion exchange membrane able to block the vanadium cross-over, since it has positive sign. An increase in the CR leads to a higher current density, which leads to a lower time available for the ions cross-over, which eventually leads to a higher CE . An excessive compression is detrimental, though: it reduces the porosity and eventually decreases the CE .

The voltage efficiency is instead defined as:

$$VE = \frac{\text{middle point of discharge potential}}{\text{middle point of charge potential}}$$

It considers the overpotentials during charge and discharge: their opposite signs have a negative impact on the cell potential, and eventually lead to electric losses. Since the compression lowers the ASR of the cell, the magnitude of these overpotentials is reduced, leading to a better voltage efficiency.

The energy efficiency is then but the product of CE times VE .

A higher compression also corresponds to a better discharge capacity during cycling at the same current density: the fading is attributed to the differential rates of the gaseous evolution, to the different permeability of the vanadium ions (and the consequent imbalance in the SOC between the two half-cells), to the vanadium ion precipitation, to the change in electrode polarisation and resistance. Table 2.5 reports some numerical examples. [23]

Table 2.5: The effect of the compression ratio on the discharge capacity, on the efficiencies and on the cell resistance at different current densities. [23]

Current density (mAcm ⁻²)	Percentage of compression (%)	Discharge capacity (mAh)	CE (%)	VE (%)	EE (%)	Cell resistance (mΩ)
20	0	475.0	74.9	93.3	69.9	82.1
	10	607.5	78.4	94.9	74.3	62.3
	20	634.3	82.3	95.0	78.2	56.2
	30	706.5	68.7	95.3	65.5	52.0
40	0	162.1	88.9	85.6	76.1	79.2
	10	426.7	89.1	90.0	80.1	54.8
	20	435.0	91.3	90.1	82.2	50.3
	30	636.4	82.0	91.8	75.3	46.9
60	0	110.0	92.7	78.6	72.9	74.0
	10	366.4	92.4	85.2	78.8	52.2
	20	367.6	93.8	85.7	80.3	45.8
	30	589.0	86.0	88.5	76.2	43.1
70	0	81.8	94.0	74.9	70.4	78.3
	10	336.1	93.4	83.0	77.5	50.1
	20	334.6	94.6	83.6	79.1	44.8
	30	548.7	87.5	86.9	76.1	42.1

The example of Table 2.5 further shows that for any current density, the compression ratio enhances the efficiencies (but must not become excessive, lest the CE falls down again), reduces the resistance, and improves

the discharge capacity. Considering instead a constant CR , it is also visible that higher currents correspond to lower discharge capacities, better CEs , and worse VEs .

Figure 2.25 reports the voltage values for the different capacities along the charge and discharge processes, for various current densities and compressions. [40]

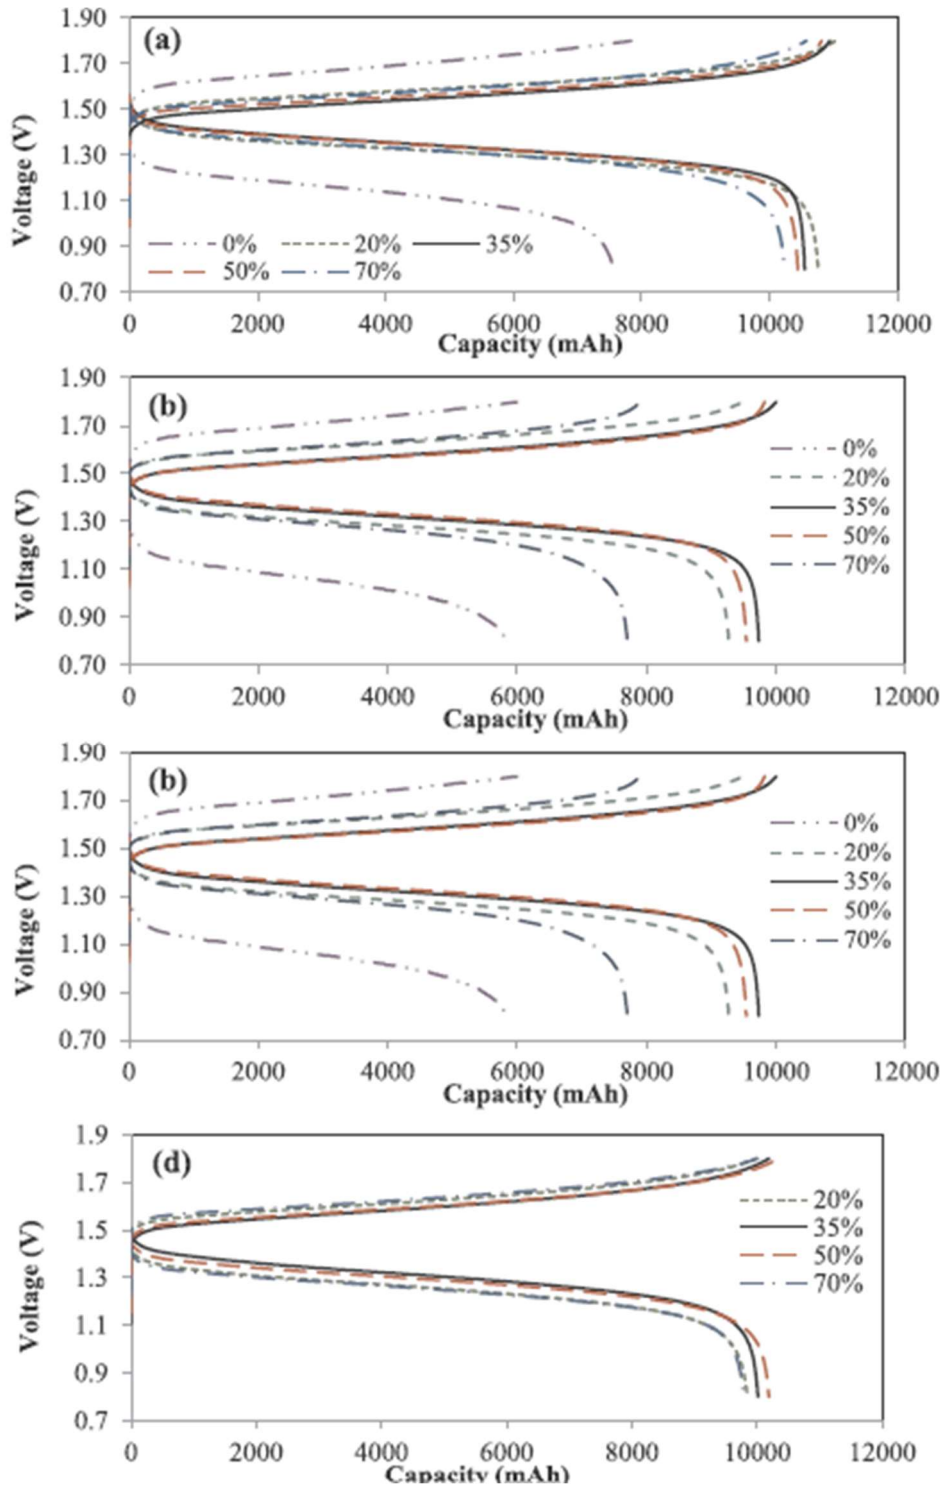


Figure 2.25: The voltages VS capacity during charge and discharge for different compression ratios, and for a)-d) rising current densities and flow rates values. [40]

As Figure 2.25 shows, the compression first raises the maximum reachable capacity, then lowers it again when excessive. In each graph, the *CR* also makes the line more horizontal, which means that the effect of the ohmic losses (due to the *ASR*) is reduced. The impact of the compression is less evident in d), where the current is the highest one.

The compression ratio which corresponds to the best efficiency is not the same that brings the highest output power though. A high power allows to design a VFB of smaller dimension, and therefore less expensive, but a higher efficiency enhances its performances, its ability to withhold the stored energy. Another trade-off, another optimisation is needed here. [23]

It is also interesting to notice how the compression increases the charging and discharging times, as understandable from Figure 2.26. [23]

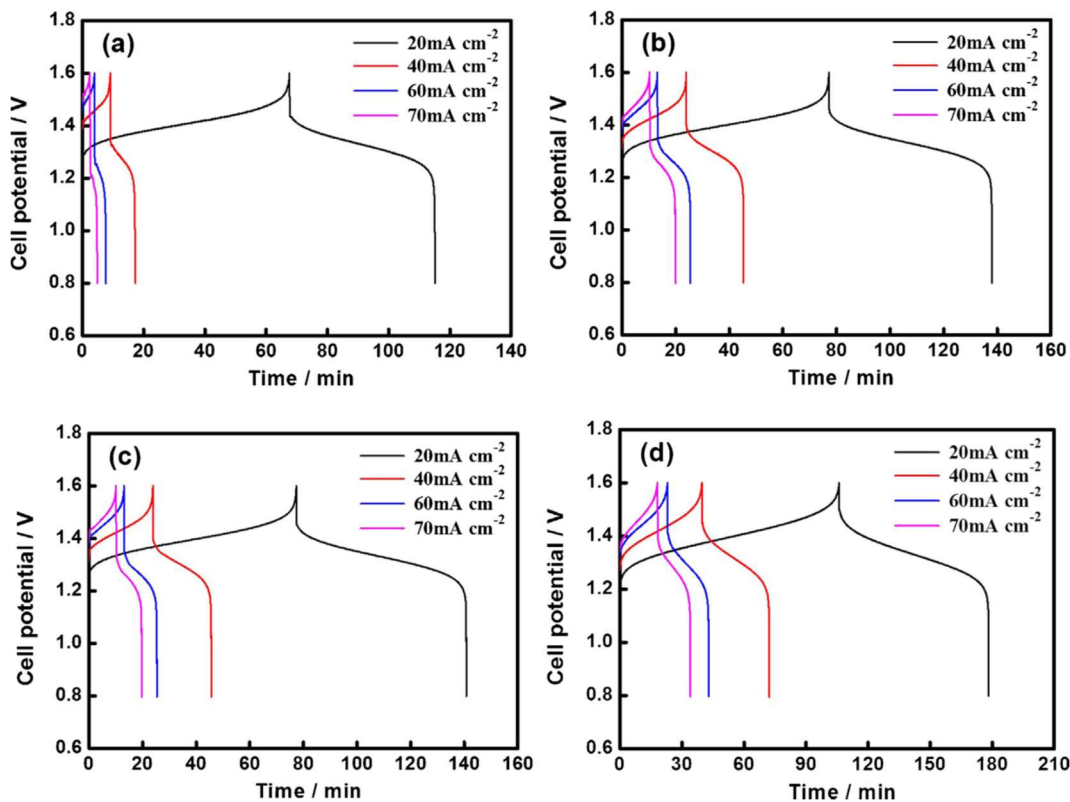


Figure 2.26: The cell potential versus time during charge and discharge for different current densities and compression ratios: a) 0%, b) 10%, c) 20%, d) 30%. [23]

It is no surprise that a lower current density makes the times longer for any compression ratio. The point is that for any current, the higher the compression, the longer the times before the complete charge or discharge is gotten. This might be due to the reduced electrical contact. [23]

Looking again at Table 2.5, it can be read that higher currents bring better coulomb efficiencies. In fact, faster times reduce the net convective cross-over, the vanadium migration, which would otherwise waste the collected charges in the VFB. [29]

To sum everything up: the compression of the electrode causes higher pumping losses and might lead to an inhomogeneous use of the felt, but also lowers the *ASR*, resulting in better efficiencies. It is obviously not possible to perfectly satisfy all the conditions above, so an optimisation is necessary.

The situation is further complicated by a last detail: the compression is not actually uniform throughout the electrode.

In the flow-by configuration in fact the electrolyte is led into the electrode from the channels, which are engraved in the BPs. This space can be occupied by some carbon strans of the felt after a compression, as shown in Figure 2.27.




Figure 2.27: The intrusion of the carbon strains of the felt in the channels due to compression. [29]

The stronger the compression, the deeper the intrusion. It is therefore useful to introduce another parameter: the intrusion ratio, λ : the ratio between the cross-sectional area of the intruded S_i and the channel cross-section S_c . [29]

$$\lambda = \frac{S_i}{S_c}$$

Figure 2.28 provides a useful visual representation to understand the meaning of these two quantities.




Figure 2.28: The cross-sectional area of the channel and the area of the intrusion. [29]

The graph of λ versus CR is shown in Figure 2.29. [29]




Figure 2.29: The relationship between the compression ratio CR and the intrusion ratio λ . [29]

Since the porosity in the intruded region is different from that found in the rest of the electrode, as visible in Figure 2.27, and since the performances of the cell depend on the porosity itself, it is clear that the intruded region will behave differently than the rest of the semi-cell. This is not too much of a concern, though, due to the small dimensions of the intruded region with respect to those of the electrode. The point is instead that the intrusion reduces the cross-sectional area for the electrolyte to flow through, thus increasing its velocity and pressure drops, eventually leading to a higher necessary pumping power.

Due to the different characteristics of the intruded area and of the electrode, it is better to separately describe these regions with distinct parameters, i.e. two porosities, two permeabilities, etc.

It is even possible to identify three regions: the intruded one in the channels, the one below the ribs that separate these channels, the one below these channels. Figure 2.30 identifies them on a photograph and shows the respective porosities versus the compression.





Figure 2.30: The porosity of the three regions versus the compression ratio. [29]

The three regions all start with almost the same porosity, but some of them are more influenced by compression than others. In particular, the intrusion region is less compressed, since its fibres have free space to occupy. The rib region is instead compressed by the ribs above, and its fibres are strongly pressed together. The channel region always takes intermediate values between the two of them: its fibres can expand upwards (pushing those of the intrusive region even more upwards), but are anyway subjected to the compression of the whole felt and of the strains from the neighbouring rib regions.

These porosity differences between channel region and rib region affect the respective permeabilities, and the whole hydraulic dynamics of the felt, as shown in Figure 2.31. [42]



(a) Concentration distribution at CR=14.2%. (b) Concentration distribution at CR=55.7%.

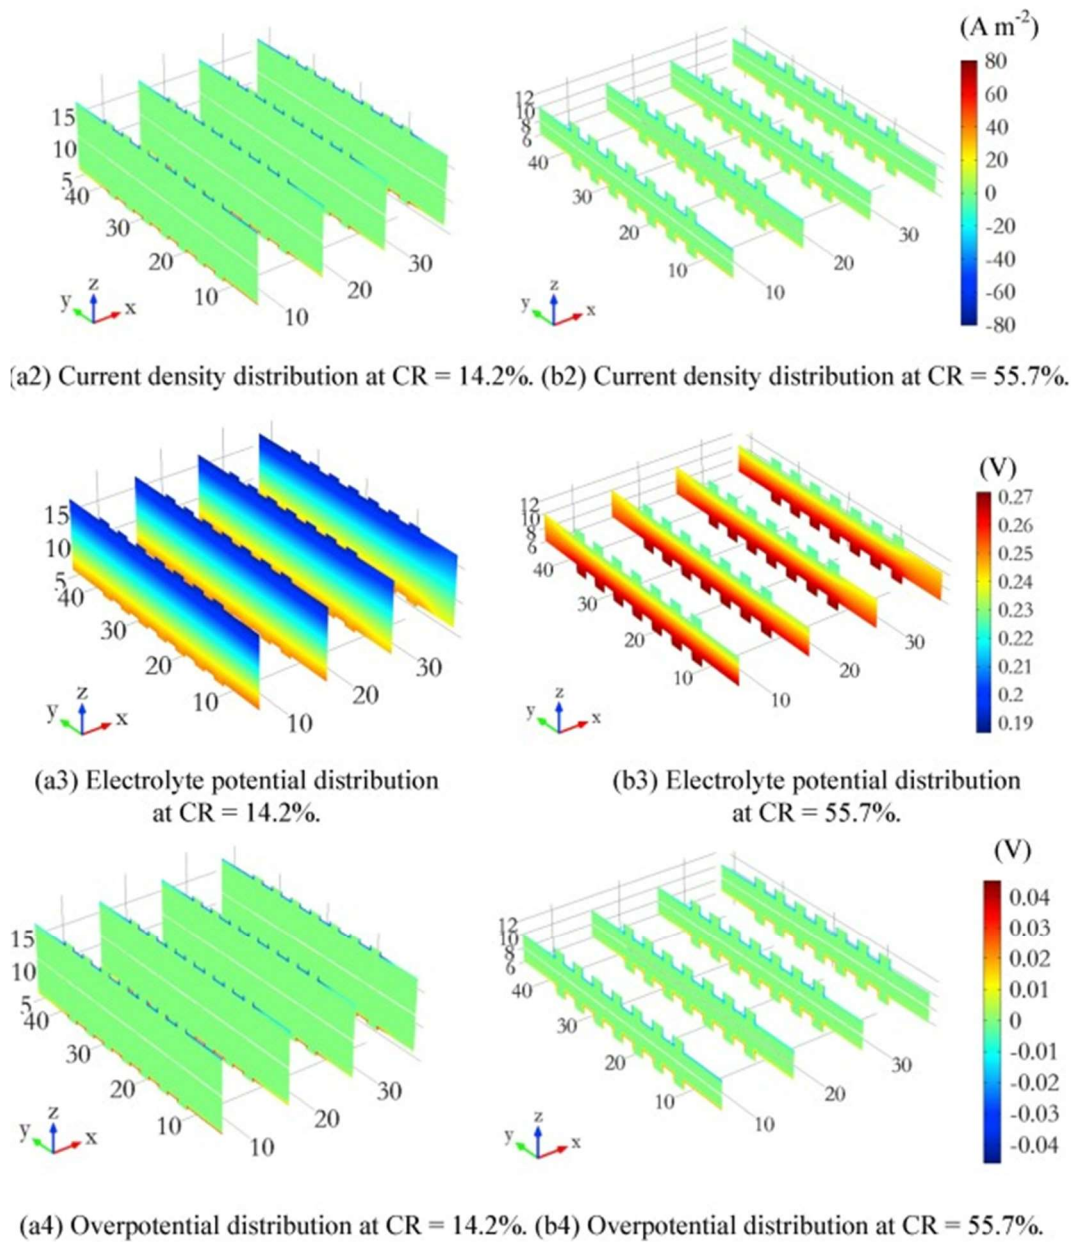


Figure 2.31: The electrode performances a) before and b) after the compression. [42]

The effect on the electrolyte potential distribution is the most evident. Its range is raised from $0.19 \div 0.25$ [V] (in a3)) to $0.23 \div 0.27$ [V] (in b3)): it is both increased and more uniform throughout the felt. It is also proved how the compression reduces the overpotential tribute the cell has to pay: it is lowered from ± 0.4 [V] (in (a4)) to ± 0.2 [V] (in b4)).

h. Importance of the mechanical design on the electrode compression and hydraulic sealing

Pressure affects the battery performances in many ways, due to several causes. Once its value to apply has been chosen, it may not be actually sustainable for the structure, though, since the mechanical properties may not be able to withstand it.

The various elements of the stack are made of different materials, which respond in different ways to strains. Table 2.6 reports them:

Table 2.6: The stack components and their materials.

End plate (EP)	Aluminium or steel
Isolator	Polyvinyl choride (PVC)
Current collector (CC)	Copper
Gasket	Ethylene-propylene Diene Monomer (EPDM)
Flow frame (FF)	PVC
Bipolar plate (BP)	Graphite
Membrane	Nafion
Bolt, washer	Steel
Electrode	Graphite felt

Regarding the behaviour of metals when subject to stress (whether expansion or compression), they first undergo the so-called elastic phase, then the plastic one. The deformation caused during the first one is reversible, since the absorbed energy is stored as elastic, and once the stress is removed, the material will return to its initial shape. There is a linear relation between the applied stress (symbol σ , in $[MPa]$), and the induced strain (symbol ε , adimensional since it is a ratio between the deformation and initial length: $\varepsilon = \frac{\Delta l}{l_0}$): the constant that joins them is the Young's modulus E , or elastic modulus, or modulus of elasticity. [11]

Graphing σ versus ε , the so-called strain-stress curve of the material is obtained: Figure 2.32 proposes the one of the aluminium alloy Al 6061-T651. [43] The curves of steel and copper are very similar to those of aluminium, since they are all metals.




Figure 2.32: The stress-strain curve of Al 6061-T651 for different operating temperatures. [43]

A high elastic modulus implies that a lot of stress is necessary to cause a deformation in the material, while a low one means that it is easier.

Not only does the deformation ε occur in the longitudinal direction, though, but also in the transversal one. In the case of the battery stack this will be important, since it will cause the plates to enlarge under the pressure action. The parameter that describes the enlargement is the Poisson's ratio:

$$\nu = -\frac{\varepsilon_{transv}}{\varepsilon_{long}} \quad [1]$$

So, during the elastic phase, a compression σ is applied: with E the longitudinal strain ε_{long} is found, then with ν the transversal one ε_{transv} is found too. These values will not be homogeneous along the cells, and neither across the same plate, but are to be evaluated little region by little region: the finite element analysis will take care of this aspect.

This described the elastic phase, but if the applied pressure, the stress, is excessive, the reticulum of the metal is not able to go back to its initial position once the stress is removed. The caused deformation will be permanent (unless particular treatments are performed): this is the so-called plastic phase. The absorbed energy is now spent to modify the structure of the metal, and not only stored in the elastic form. This causes the $\sigma - \varepsilon$ graph to stop following a straight line, but rather to take a curvilinear direction. The plastic phase hardens the material, making its structure stiffer, but the deformations cannot be absorbed indefinitely: a maximum point is eventually reached. Once it is overcome, the piece goes towards its rupture.

Going back to the issue of finding the right pressure to apply to the VFB stack, it is necessary to pay attention not to exit the plastic phase of the materials used to build the battery. For this purpose, during the simulations the highest stress values in the pieces will be evaluated, and checked not to overcome the tensile strength, i.e. not to enter the plastic phase. Since it is a good idea to keep some margin from this limit, the safety coefficient will be evaluated: the ratio between the tensile strength and the applied stress. The higher the coefficient, the safer the situation.

When it comes to non-metals, the situation is slightly more complicated. In fact, the behaviour at compression does not necessarily work as the one at traction, and E might not be constant. The case of EPDM (i.e. of the gaskets) is shown in Figure 2.33. Its convex shape means that a linearly increasing stress is not sufficient to linearly deform the material, but rather a stronger effort is required step by step, i.e. it is harder and harder to deform it.




Figure 2.33: The stress-strain curve of EPDM, for different strain rates. [44]

The strain rate mentioned in Figure 2.33 is simply the velocity at which the deformation occurs, in fact it is defined as $\frac{d\varepsilon}{dt}$. High strain rates refer to stresses which are applied fast. It is interesting to notice how a slower deformation is therefore easier: it requires less stress to achieve the same strain.

Concerning the behaviour of the graphite felt which constitutes the electrodes, Figure 2.34 reports its stress-strain curve.




Figure 2.34: Stress-strain curve of the graphite felt of the electrodes. [29]

Like the EPDM, graphite felt has an elasticity modulus that grows with the compression ratio CR. This is due to its internal structure of carbon fibres: first the spaces amongst them are pressed, and successively the fibres themselves. [45]

The PVC, which is the material of the FFs and of the insulators, has a quite particular curve, shown in Figure 2.35. [46]




Figure 2.35: The stress-strain curve for PVC. [46]

It is interesting to notice how tension and compression behave differently from other materials after the elastic phase. During it they have instead quite the same steepness, but the limit of compression is superior.

The BPs are in graphite as the electrodes, but in the form of sheets, not porous felt. The material is therefore chemically very close, but its structure deeply different, which leads to a different stress-strain curve: Figure 2.36. [47]




Figure 2.36: The stress-strain curve of graphite sheets. [47]

A quick comparison between Figure 2.34 (graphite felt) and Figure 2.36 (graphite sheet) is sufficient to see how they behave differently, mostly because the sheets have no empty space inside them, but are rather solid. They show a linear relationship between σ and ϵ until a certain point, where the stress becomes nil. A zero stress necessary to achieve a high deformation (over 25%) means that the piece has broken.

It is therefore clear that when the stack is pressed from the EPs at the extremities, the mechanical behaviour of its cells and inner components is not simple: the pressure will not uniformly distribute across the whole cell cross-section.

The point of the optimisation of this thesis is instead to determine the best EP design in order to obtain a homogeneous pressure inside the electrodes, to enhance their performances, and to reduce the losses which were described in the previous paragraphs.

Besides the electrochemistry, the mechanical strength of the components, the losses and efficiencies, the hydraulics plays a role too when it comes to pressure. Since the electrolytes flow through the manifolds, they must provide the sealing. For any given cross-section, the pressure of the electrolyte inside is isotropic and homogeneous, which means that the manifold must withstand an equal outer force along all its surface.

3. MECHANICAL AND FLUIDODYNAMIC SIMULATIONS OF A SINGLE LABORATORY CELL VFB

VFBs have been thoroughly described so far: their working principles, the behaviour of the electrolytes and of the electrodes, the various types of losses: what causes them, and how to reduce them. It was seen how the mechanical and hydraulic designs play a crucial role in the final result.

All the taken choices should have a justification behind them: e.g. the length and width and height of any component shall be chosen considering all the aforementioned phenomena, with optimisations. Such a mathematical and engineering process is far too complicated to be solved by hand on paper, though, since the object in issue (the battery, the stack) is large, with many components, and more than one physic to take into account: electrochemistry, electrotechnics, mechanics, hydraulics...

This is the reason why computer-aided simulations are necessary. In particular, this one is done using COMSOL Multiphysics® (in the following, just named Comsol): a useful CAD program where complex objects can be described and simulated, even coordinating more physics. It also allows a parallel use of other programs (such as MATLAB®), but in this case it was not necessary.

This chapter 3 first illustrates how Comsol works: its principles and its logic, and the way the simulation is to be set, paying attention to the numerous tiny commands that actually play a crucial role. This tutorial-like part describes a simplified cell, which is then enriched with details when dealing with the real study.

a. Structure of the simplified simulated cell

Concerning the mechanical simulation, the cell is composed of:

- a gasket (Gk1);
- a flow frame (FF1);
- a gasket (Gk2);
- an ion-exchange membrane (neglected, since it does not affect the mechanics, nor is the focus of the simulation);
- a gasket (Gk3);
- a flow frame (FF2);
- a gasket (Gk4);
- a bipolar plate (BP);
- the electrodes (here not represented since their impact on the mechanics is not important yet).

They are represented in this order in Figure 3.1, from the bottom to the top.




Figure 3.1: The explosion of the cell for its mechanical simulation, drawn in Comsol.

All the components but the BP have a rectangular hole in the middle, where the electrodes are placed: the AA. The two electrodes would be separated by the membrane, here absent, between Gk2 and Gk3. One would fit in the holes of Gk1, FF1, Gk2, the other in the holes of Gk3, FF2, Gk4. All components, even the BP, have four cylindrical holes at the corners: those are the back-and-forth paths for the catholyte and anolyte: the manifolds. In the generic cell, as the one shown here, they run through the whole height, but this is not true for the last cell, the one at the bottom. In its specific case, two manifolds run down to feed the bottom electrode, while the other two only reach the top one, in diagonal couples. This case is shown in Figure 3.2.




Figure 3.2: The bottom cell of the stack: only two out of the four electrolyte manifolds reach the bottom electrode.

Another difference between the bottom cell of Figure 3.2 and the generic cell of Figure 3.1 is that the first has an additional BP. This is not totally true, but rather a way to represent the object in Comsol. In fact, two consecutive electrodes are either separated by the ion-exchange membrane (that would be at mid-height if represented in the figures) or by a BP. A more complete representation would therefore depict two BPs: one above the whole cell, the other under. Such a drawing in Comsol would imply two BPs between two cells, nevertheless, so the top one alone was put. Consequently, since there is no other cell below the bottom one, the lowest BP needs to be explicitly represented.

Unlike the bottom cell, the top one is totally similar to any other generic cell of Figure 3.1. The dimensions used are reported in Table 3.1, together with their symbols and definition.

Table 3.1: The spatial dimensions of the components of the cell.

Symbol	Value [mm]	Description
Gk_x, FF_x, BP_x	520	BP, gasket, FF length (x)
Gk_y, FF_y, BP_y	450	BP, gasket, FF width (y)
Gk_z	0.5	Gasket height (z)
FF_z	3.6	FF height (z)
BP_z	1.6	BP height (z)
mnf_rho	17.5	Manifold radius
AA_x	400	AA length (x)
AA_y	160	AA width (y)

An important detail to add are the ribs. They only belong to the FF and follow its internal perimeter, along the AA, as shown in Figure 3.3.




Figure 3.3: The rib following the internal parameter along the AA of the FF. Another one is present on the bottom side of the FF.

As soon as pressure is applied, the ribs will be pressed between the FF and the adjacent Gk, providing hydraulic sealing for the electrolyte. Concerning their dimensions, they are reported in Table 3.2.

Table 3.2: The spatial dimensions of the ribs of the FF.

Symbol	Value [mm]	Description
rib_d	6	Distance of the rib from the AA
rib_w	1	Width of the rib
rib_z	0.6	Height (z) of the rib

b. Setting up the simulation on COMSOL Multiphysics®. Description of the process

It is now explained how the cell of Figure 3.1 was drawn in Comsol, not only regarding the geometry, but also the assignment of materials, the physics, the mesh.

i. Creation of the geometry

The very first thing to clarify is the difference between a component, an object and a domain in Comsol. The component is the object of the study, in this case the cell. A simulation usually deals with only one component, but this is not mandatory (even though simpler). Every component is drawn with a geometry, described with a mesh and physics, and made of materials. Different components can obey to different physics. Simulations do not usually comprehend too many physics per time, so it is not usually necessary to set up multi-component studies. In general, it is better to break down the complex studies into more Comsol files, both for the sake of simplification and for quicker evaluation time. This is also the reason why the mechanical and the hydraulic simulations were kept apart, since they do not affect each other.

A component is made of one or more objects. In the cell case, the objects are the gaskets, FFs and BP. The ribs are part of the FF they belong to: there is the suited command “form union” to tell Comsol this. As seen later in paragraph 3.b.ii, materials are then assigned to the objects.

Every object has an outer surface, divided into boundaries, separated by edges, which end in points. It can be useful, e.g. when describing the physics or drawing the mesh, to split the objects in multiple 3D parts: the domains. For example, a sphere object can be split into two semi-sphere domains. Once more domains are formed, all their boundaries, edges and points are available for any following setting.

The word “following” is actually crucial when dealing with Comsol: any command produces effects which are only available in the following commands, according to an ordinated structure. For example: the geometry describes the object where the mesh will be drawn, but the mesh cannot retroactively impact on the geometry.

This means that the ability in using Comsol stays in knowing already how to set the previous commands, in order to find a convenient situation for the following one: this idea will occur several times through these paragraphs. Concerning the geometry, Comsol considers the component as made of several basic solids (when dealing with a 3D case): parallelepipeds, spheres, cones, cylinders etc. It necessary to assign the 3D dimensions to each one of them, and also their origin in the 3D space. Such an origin can refer to their vertex, or to their centre.

Recalling the data of Table 3.1, it could be possible then to start drawing the first gasket of the cell (Gk1) as a parallelepiped whose sides are $520 \times 450 \times 0.5$ [mm]. This is not usually suggestable in informatics, though. On one hand, it becomes necessary to manually evaluate any intermediate distance whenever necessary (for example when drawing the manifolds for the electrolytes), but mostly, if something were to be modified, all the math would need to be done again from scratch. This explains the usefulness of defining all the parameters first. The very first directory of a Comsol file is in fact “Global Definitions”, where parameters can be assigned and collected in groups. It is usually convenient to write a group of geometric parameters, one for the load (in this case such a group will almost be empty, but in the case of the electric physics it can be way more abundant, for example), and maybe even one for the mesh. The subdirectory “Geometric Parameters” basically reports the values of Table 3.1 and 3.2, even allowing for adding a brief description for each of them.

Now, if the width of the stack were to be doubled, a single command alone at the very beginning of everything would be sufficient to easily correct the drawing.

Regarding the first gasket, a first parallelepiped of dimensions Gk_x, Gk_y, Gk_z is drawn. If the reference system was kept $(0,0,0)$, the gasket centre would not be in the axes’ origin. This would not be an error per se at all, but it is better to make these two points coincide, in order to simplify the exploiting of the symmetries later. This is the reason why the vertex position is set as $(-Gk_x/2, -Gk_y/2, 0)$. The pointing of the gasket centre in the axes’ origin does not comprehend the z direction: the other components are going to be piled over it, and this simplifies their positioning.

The gasket so far is only comprehensive of the external structure (the parallelepiped Gk_ext), as shown in Figure 3.4 a). The AA is then drawn, adding a second parallelepiped, whose dimensions are AA_x, AA_y and AA_z , with origin $(-AA_x/2, -AA_y/2, 0)$.

The manifolds do not work differently: instead of a parallelepiped, they are cylinders, with radius chn_rho . Their distance from the gasket external perimeter is given: 67.5 [mm] along the x direction (chn_dx), 57.5 [mm] along the y direction (chn_dy). In the symbols, the letter “d” means distance along the specified direction. A little attention is necessary when indicating the position of the centre of these cylinders: concerning the bottom left one, $-Gk_x/2+chn_dx$ and $-Gk_y/2+chn_dy$, along x and y respectively. All the other manifolds are then drawn, paying attention to their different axis centres. Figure 3.4 b) shows the result so far.

Under “Booleans and Partitions”, Comsol has the command for the Boolean subtraction, which allows to cut an object away from another one. Among those to add there is Gk_ext alone, while those to subtract are the four manifolds and Gk_AA . The final result is in Figure 3.4 c).




Figure 3.4: a) The external structure of Gk1 first, b) the drawing of the AA and of the manifolds too, c) the final result after the Boolean subtraction.

In the Boolean subtraction operation, there are some settings that may play an important role. In fact, after a standard subtraction the starting objects will no longer be available for successive operations, but only the resulting one will be. In Comsol’s logic, this means that the starting objects are forgotten, and only the final one

kept. In some cases, the subtraction operation may not be used to build a new object (which is given by the subtraction), but rather to make available a new one, described by the difference.

In other words, if the difference between objects A and B gives C, the default settings only make C available. This setting can be modified as shown in Figure 3.5.

- Mantieni oggetti da aggiungere
- Mantieni gli oggetti da sottrarre
- Mantieni i contorni interni

Figure 3.5: The subtraction setting concerning the dominions.

If the first square were to be ticked, objects A and C would be eventually available for further operations. In the gasket case, the graphics window would still show Figure 3.4 b), and it would be possible to either select the final object (of Figure 3.4 c)), or the one to add (of Figure 3.4 a)).

The second square works the other way around: B and C are available for further use. The graphic window still shows Figure 3.4 b), and it is possible to either select the object of Figure 3.4 c), or any of the manifolds or the AA.

The two commands can be simultaneously ticked: in such a case, all the components and the result are available for selection. A situation in which this all could be useful is represented in Figure 3.6.




Figure 3.6: Example for the use of the Boolean difference settings.

Figure 3.6 shows a situation where the “Keep input objects to subtract” is useful. It is possible to only apply a pressure on the external area, while keeping the whole parallelepiped as object whose behaviour under such a condition is studied. In fact, the maintained objects are available for the selection of their domains, boundaries, edges and points.

In order to obtain a situation like the one of Figure 3.7, two subtraction operation are needed.




Figure 3.7: More subtraction operations can be necessary to describe particular cases.

It is possible to subtract one object from several to add, which means that it would be possible to define the AA only once (a much taller one), and erase it from all the gaskets and FFs at once, rather than doing it separately every single time. The same applies to the manifolds.

In particular, there are four manifolds, so it would be possible to either distinctly draw each of them, or just one and then multiply it with the array command. Both the methods are valid.

The array idea is basically of a copy-and-paste: an object to multiply is selected (here the cylinder of the manifold) and the dimensions of the array are defined, in the x, y and z directions.

For example, multiplying the object with an array of (17,1,1) dimensions will give a line of 17 of it, aligned along the x direction, starting from the initial one (comprehended amongst the 17). Were the dimensions (17,3,1), there would be three rows, spaced in the y direction, each of them following the x one. Trivially for the z component. Of course, it is also necessary to indicate how far the components of the array are positioned: this value is obviously expressed with the geometric parameters.

In the gasket case, the array has a (2,2,1) dimension, and the distances are $Gk_x-2*chn_dx$ along x and $Gk_y-2*chn_dy$ along y.

It did not matter that much to define the manifolds in this way, but it will be crucial with the bolts. In fact, the optimisation will require to find out how many of them give the best conditions. Drawing them singularly every single time is not efficient (and might lead to distraction mistakes), so an array will be useful. It will be necessary to define a parameter (something like bolt_N) to vary, which will describe the array dimensions, and with whom the distances among the bolts will be evaluated.

So far, the drawing of the gasket already required two parallelepipeds, a cylinder, an array and a subtraction. A whole cell will lead to a great number of objects and operations, so some organisation may be useful. For this purpose, the objects and commands can be collected in groups. Their only function is to behave as little directories in the geometry description.

The first gasket (Gk1) is ready, and it is time to continue with the first FF (FF1): so far, the geometry directory looks like in Figure 3.8.

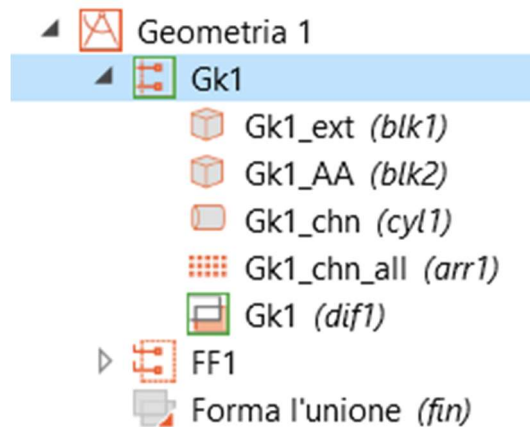


Figure 3.8: The geometry directory so far.

It is the FF's turn, now. Its basic structure is drawn in the same way as for the Gk, with the only care that its vertex position must be translated in the z by Gk_z, so that the two pieces are one on top of the other. This would neglect the ribs, though: they were shown in Figure 3.3. The vertical displacement shall therefore be of $Gk_z + rib_z$. Concerning the ribs themselves, they are not a problem to draw: each of them is just the difference between two parallelepipeds: the external one (rib_ext) is $AA_x + 2 * rib_d + 2 * rib_w$ long times $AA_y + 2 * rib_d + 2 * rib_w$ wide, the internal one (rib_int) is $AA_x + 2 * rib_d$ long times $AA_y + 2 * rib_d$ wide. Attention shall be paid when setting their position too, in order to have their centres aligned and above the origin point.

A union command is finally used to join the FF and its two ribs, eventually resulting in what is shown in Figure 3.9.




Figure 3.9: a) Gk1 and FF1 with ribs, b) their magnification from the side, to highlight the ribs.

According to Figure 3.1, it is now time to draw Gk2, which might just be obtained with a (1,1,2) array from Gk1. In a normal VFB, it would be separated from Gk3 by the ion-exchange membrane, but such an element is not put in this cell model, since it does not affect the mechanic simulation: this means that Gk2 and Gk3 are identical and just one on top of the other. They are thin elements, and this corresponds to a fine mesh, which takes a long time to evaluate. Evaluating two identical (both in geometry and material) objects with a fine mesh while they could just be merged in a single one that would correctly represent their behaviour is just a waste of computational time and RAM: this is the reason why they are joined in Gk23.

Since the height of Gk23 is two times that of Gk1, an array cannot be set, so it is to be drawn separately, as FF1 and Gk1 were.

FF2 can instead be obtained with a (1,1,2) array of FF1, comprehending the ribs: the z distance is $2 * Gk_z + FF_z + 2 * rib_z$. Similar for Gk4 from Gk1: distance $3 * Gk_z + 2 * FF_z + 4 * rib_z$.

The last missing piece is the top BP. It has no AA, so that step could be neglected, but for mesh issues it is useful to define that contour anyway. For this purpose, the parallelepiped BP_AA is drawn too (in the same way the other correspondents were), but is not subtracted from BP_ext in the Boolean operation: nor it is put amongst the elements to add. In fact, were it put amongst the objects to add, it would not be selectable anymore, and

ticking the option “Keep input objects to add” would result in not digging the holes for the electrolyte. Were it instead put amongst the objects to subtract, it would contradict what was just affirmed. Ticking “keep input objects to subtract” would anyway result in not digging the manifolds as well. In the end, simply not putting BP_AA at all in the subtraction command is best.

The cell drawing is complete and reported in Figure 3.10.





Figure 3.10: The whole cell seen a) from the top, b) from the bottom.

Even though the cell is complete, there is one more command to set, which also plays an important role, so attention shall be paid: the form union VS form assembly.

All the objects drawn so far (the gaskets, the FFs with ribs, the BP), represent different pieces of the stack: they are not a single object, nor are they glued or welded to together. This means that they are actually able to slide one with respect to the others. This is an important information that Comsol needs to know, otherwise it would consider them as parts that form a unique object (even though of distinct materials), bound together. The command needed for this purpose is the Form union, which can be otherwise set as Form assembly. It is a mandatory information to give, and is at the end of the Geometry directory by default.

It is clear that the distinction between union and assembly is crucial: the union states that all the objects involved are actually part of a single object, the assembly that they each represent a distinct object. In the case of the FFs, each one of them was already described as a single object together with their ribs by a previous union command, which now finds its explication.

Besides the kinetics and mechanics implication behind the union VS assembly choice, there are also some concerning the mesh. Since the mesh describes the object, if more pieces together form a single one, their meshes must be coordinated, otherwise this is not a necessary condition. Figure 3.11 provides a useful example to understand this statement.



a)

b)


Figure 3.11: The mesh of the same object when set as a) a union, b) an assembly.

The vertices of the mesh of the top object coincide with those of the bottom one in Figure 3.11 a): this is because the mesh is actually a single one. It does not happen instead in Figure 3.11 b): two distinct objects are described by two distinct meshes, whose vertices do not necessarily coincide. In fact, the two objects might move one with respect to the other, and after such a movement any starting vertices correspondence would be lost.

Every FF must form a union with its corresponding ribs, but this was already done with previous commands, so the final setting to be chosen is Form assembly.

In the case of Form union, Comsol does not ask anything else, while in the case of Form assembly it has more questions, concerning the imprints and the couples.

If the Create imprints command is activated, the shared boundary between two adjacent objects will draw its edges on both of them, and these edges will belong to the respective meshes. Figure 3.12 provides a useful visual example to understand this process.



a)

b)

Figure 3.12: The effect of a) not selecting the Create imprints command, and b) of selecting it.

Looking at Figure 3.12 a), it is shown that the tetrahedrons of the bottom object do not lay their vertexes along the edge of the touching surface; this instead happens in Figure 3.12 b). If it were desired that all the vertexes of all the tetrahedrons of the top object laying on the contact surface coincided with those of the bottom object, the Form union command would have been to be selected.

Going back to the cell case, the Create imprints command would draw the shade of the ribs on the adjacent gaskets' touching surfaces. It will be seen in 3.b.iv that this is not sufficient though: such a shade would need to be dug down throughout the whole gasket height in order to set the mesh as later desired. This means that the Create imprints command is indifferent to select in this context.

Concerning pairs: "The alternative method [to a union] is to form an *assembly*. The software then treats the geometry as a collection of the geometry objects instead of uniting them. This means that you must use pairs to connect boundaries where a field is continuous, but it also makes it possible to use special pair conditions for applications such as contact modeling, where you can add contact pairs to model contact between geometric parts". [48]

Since the displacement field is indeed continuous between two consecutive objects, pairs are to be created. A successive command asks whether disconnected pairs should be divided: this allows for the selection of two distinct boundaries for each pair: one belonging to the top object, the other to the bottom one. This may be necessary when drawing the mesh or setting the physics. In the VFB cell case, nevertheless, the pressure will be applied by the bolts on the EPs alone, while the mesh will be set in the same way amongst the elements, so this command is not necessary. It would not cause errors anyway, but just a complication in the boundary selection later, thus resulting in possible annoying mistakes.

The last command asks whether the created pairs are of contact or of identity. This is a crucial detail. The effects of this choice are well represented in Figure 3.13.




Figure 3.13: The difference between the application of a) identity pairs, rather than b) contact pairs. The top parallelepiped is pushed to the left, the bottom one to the right.

When contact pairs are chosen, the two objects connected by such a pair are distinct and can move one with respect to the other. When instead identity pairs are used, the two objects cannot move one with respect to the other. It might appear at this point - if only two objects are involved - that there is no difference between a union and an assembly with identity pairs: in both cases the two selected objects cannot move one with respect to the other. Figure 3.14 answers this doubt, recalling the impact on the mesh of considering or not the two objects as parts of a single one.




Figure 3.14: The impact on the mesh given by selecting a) an assembly with identity pairs, rather than b) a union.

Now that all this has been described, the reasons behind the choice of the settings of the cell case are clear. It is an assembly, without imprints and with contact pairs. The disconnected ones are not divided.

If some pairs were of contact and some others of identity, it could be set too. Oddly enough, in this particular case, the command is to be found above the geometry directory, i.e. against the principle of hierarchy seen until now. Still under the Component directory, there is the one of Definitions: it has a collection of all the contact couples or contour identities defined by the Form assembly command. They can be individually selected and if desired changed into a pair of the other kind.

ii. Assignment of materials

Once the geometry is complete, comprehensive of the definition of the pairs, it is time to assign the materials to the various objects. Comsol offers a huge library to choose from. As stated in Table 2.6, the gaskets are in EPDM, the FFs in PVC and the BP in graphite.


The main parameters involved in the mechanics simulations are the density (ρ), Young's modulus of elasticity (E) and the Poisson's number (ν). These parameters are not constants, but can vary under the influence of some other factors, like the temperature: this is the reason why Comsol expresses them with a formula. Figure 3.15 reports for example the evaluation for the density of a simple steel: 1008 [steel, polished].

Partenza	Fine	Funzione
0.0	30.0	7906.743
30.0	53.0	$7905.632 + 0.06730164 * T^1 - 0.001025121 * T^2 + 4.962534E-7 * T^3 + 1.383301E-9 * T^4$
53.0	190.0	$7924.0 - 0.4306436 * T^1 + 0.002196832 * T^2 - 6.010914E-6 * T^3$
190.0	960.0	$7910.967 - 0.100271 * T^1 - 3.511597E-4 * T^2 + 1.470685E-7 * T^3$

Figure 3.15: The evaluation of the density of a steel as function of its temperature in Comsol.

The level of precision is impressive, and the ranges start from 0 [K] and go up to 960 [K] = 687 [°C]. Concerning the operating temperature of the VFB, it is considered the fixed value of 20 [°C]. In fact, it is usually between 5 [°C] and 40 [°C]:

- The upper limit is chosen in order to avoid the precipitation phenomenon of $V(V)$:



It occurs when 2.0 M $V(V)$ is kept at 40 [°C] for a couple of days. This can be avoided by using vanadium with lower concentrations (but this would imply lower energy densities), or surprisingly with way higher concentrations. [49]

- The lower limit is instead chosen to avoid the precipitation phenomenon of the other vanadium species: $V(II)$, $V(III)$ and $V(IV)$. [50]

The mechanical parameters of the involved materials do not widely change in the selected temperature range, so the average 20 [°C] was generally chosen, also because it corresponds to the room temperature.

Each material proposed in the Comsol library comes with a detailed description: every parameter from any physics is provided (mechanics, electric, chemistry, acoustic, heat transfer etc.), often function of some other variables, here the temperature. The point is that the evaluation of the necessary parameter at every single occurrence takes times, so it makes sense to simplify the situation.

Comsol offers the possibility to insert custom materials too, bypassing those from the library: in this case, of course, the user has to manually assign all the necessary parameters. In the context of this optimisation, the aforementioned ρ , E and ν are all that is needed, so it does not take too long: their values can be found in literature.

- The gaskets are in EPDM: ethylene propylene diene methylene, a synthetic rubber of the elastomer class. Being a good electrical insulator and impermeable it is often used for sealing and gasketing (like here); it is a flexible material, so not suitable for rigid parts. Its chemical structure is shown in Figure 3.16, where the groups that compose its name are indicated too. [51]




Figure 3.16: The EPDM structure. [51]

Its mechanical properties can be found in industrial data sheets, and were implemented in the Comsol description of the material: [52] $E = 0.05$ [GPa], $\nu = 0.48$ [1] and $\rho = 1100$ $\frac{kg}{m^3}$.

Its elasticity modulus is quite low, which makes it a soft material, and is in fact flexible. Such a high Poisson's ratio means that the lateral expansion is almost half the transversal one.

- The FFs and their ribs are instead in rigid PVC: it is characterised by good chemical resistance and robustness. These characteristics are necessary since it is in contact (like the EPDM) with the electrolytes, which are chemically aggressive. Its molecular structure is shown in Figure 3.17. [53]




Figure 3.17: The molecular structure of PVC.

It is a thermoplastic polymer, where each monomer is usually arranged head-to-tail with the next one, with chloride on alternating carbon centres. [53] The values to be implemented in Comsol are: $E = 3.5 [GPa]$, $\nu = 0.38$ and $\rho = 1300 \left[\frac{kg}{m^3} \right]$. Its elastic modulus is way stronger than the one of EPDM, which means that a greater stress is necessary to provoke the same deformation. The Poisson's moduli are instead close, so a similar lateral expansion should occur after an equal longitudinal expansion. This does not mean that the FFs and the gaskets will laterally expand at the same extent though: as just said, the lateral expansions are evaluated as a consequence of the longitudinal ones, but these will be different.

- The BP is in graphite: a conductive carbon-based material which allows the passage of electrons from one semi-cell to the other, but blocks the passage of vanadium ions. Its downside is that it is brittle: it easily breaks, so polymers are inserted in its structure to improve its mechanical characteristics. The following parameters result: $E = 10 \cdot 10^9 [Pa]$, $\nu = 0.275$ and $\rho = 1500 \left[\frac{kg}{m^3} \right]$. Due to the elasticity modulus, BPs will deform the least under the applied pressure, and the low ν implies that they will also enlarge very little.

iii. Physics setting

Comsol now knows how the object to study looks like: it has been described in its shape (with the geometry) and parameters (with the materials), but does not know yet what the object is subjected to: mechanics, electrochemistry, acoustics etc.

This essential information is given in the next directory: the physics. Multiple physics can be inserted and evaluated at once: this is useful for example in the case of inductive heating, where an electric current is generated in an object (electric physics) in order to heat it with the Joule effect, and then the heat propagates inside the material (heat transfer physics). The case of this thesis is simpler, since either the mechanics context or the hydraulics one are to be evaluated alone.

1. Description of the mathematical models

When the solid mechanics physics is selected, some default commands are immediately applied, while many others can be added at will.

The first one is Linear Elastic Material, a condition that considers by default all the drawn objects as – in fact – linear and elastic. This means that they are described with a list of equations, written using as reference the global coordinate system, which is necessary when dealing with orthorhombic and anisotropic data, or with stresses and strains presented in local systems.

An anisotropic material's properties depend on the considered spatial direction, so they are given in matrixes. An orthotropic material has instead properties which are different in the orthogonal directions. [54]

Under Equations, Comsol lists the equations that are applied to the selected objects.

$$0 = \nabla \cdot (FS)^T + \mathbf{F}_v$$

This states that the sum of all the forces acting must be nil. In fact, according to Newton's laws, a non-zero sum would result in an acceleration, but this study is going to deal with a stationary situation. Otherwise, the equation would have been:

$$\rho \frac{\partial^2 \mathbf{u}}{\partial t^2} = \nabla \cdot (FS)^T + \mathbf{F}_v$$

the left term would have been zero anyway, in a stationary situation. Specific values are usually preferred, in fact the density ρ is used instead of the mass. In particular, FS is the first Piola-Kirchoff stress, while \mathbf{F}_v represents the volume forces. The Piola-Kirchoff stress tensor is used, rather than the Cauchy one, since it represents the stress as relative to the reference configuration, while the other to the present configuration. The difference between the two of them is anyway minimal when dealing with small deformations.

In particular, the deformation gradient F is: [55]

$$F = I + \nabla \mathbf{u}$$

S is instead the second Piola-Kirchoff stress tensor.

\mathbf{u} is the displacement field, a vector, and I the identity matrix. [56]

Concerning the second Piola-Kirchoff stress tensor, it is a sum of two terms:

$$\begin{aligned} S &= S_{inel} + S_{el} \\ S_{el} &= S_0 + S_{ext} + S_q \\ S_{inel} &= J_i F_{inel}^{-1} (\mathbf{C} : \varepsilon_{el}) F_{inel}^{-T} \end{aligned}$$

\mathbf{C} is the fourth-order stiffness tensor, the colon represents a tensor product, and the strain ε_{el} is evaluated as:

$$\varepsilon_{el} = \frac{1}{2} (F_{el}^T F_{el} - I)$$

With

$$F_{el} = F F_{inel}^{-1}$$

The tensor product gives the stress tensor σ , which represents the internal forces (per unit area) inside the material: [57]

$$\mathbf{C} : \varepsilon = \boldsymbol{\sigma} = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix}$$

Throughout all of this, the stiffness tensor is function of the Young's and Poisson's moduli:

$$\mathbf{C} = \mathbf{C}(E, \nu)$$

The second information in the physics is concerning the free materials: this option is automatically filled with all the objects which are not later subjected to any other condition (this is another exception to the general rule of the conditions being applied in the given order).

Comsol then asks about the initial values, i.e. the starting displacement field \mathbf{u} and velocity $\frac{\partial \mathbf{u}}{\partial t}$, both expressed as 3D vectors. No velocity is applied, since the stack is still. It could be useful to set an initial compression, so that the solver would be able to find the final condition, the solution, from a starting point which is closer to it, but this quantity is different from the displacement here asked for.

Since contact pairs were defined, it is time now to describe their behaviour with mathematical formulas. The section Contact deals with this, and would not be present if all the pairs had been of the identity type. The formula of the contact is:

$$T_n = if(g_n \leq 0, -p_n g_n, 0)$$

Which means that if the gap is nil or even negative, a traction T_n is present, proportional to the gap itself accordingly to a p_n factor, which is evaluated as:

$$p_n = f_p \frac{E_{char}}{h_{min}}$$

f_p is the penalty factor, equal to 1 by default, but that can be set otherwise if desired, while E_{char} is the characteristic stiffness: it is representative of the stiffness of the destination material. [58] This Contact command that has been briefly analysed might be of great importance in other contexts: in fact, it also allows to model frictions, adhesions, decohesions. Since the case here is stationary, there is no friction nor wear involved and nothing more is to be set, anyway.

A few more commands are to be implemented at this point, so that Comsol knows what kind of stress is acting upon the cell. The load is hence set.

The load can be of three types: a viscous force over a unit area, a pressure, or a total force. The pressure is a scalar quantity: it acts in the normal direction with respect to the boundary it is applied to. The corresponding equation is:

$$\mathbf{F}_A = -p \cdot \mathbf{n}$$

With such a \mathbf{F}_A , the second Piola-Kirchoff stress tensor is evaluated:

$$S \cdot \mathbf{n} = \mathbf{F}_A$$

It is immediate to recognise the equivalence $S = -p$. The pressure p is of course expressed in $[Pa]$. If a viscous force per unit area is instead applied, three values are to be inserted, each corresponding to a different spatial direction. Its unit is $\left[\frac{N}{m^2}\right]$, i.e. it is equivalent to the pressure: the key difference is that this load is a vector, not a scalar, nor necessarily normal to the boundary. Last, a total force can be simply and directly applied to the boundary, this time expressed in $[N]$, i.e. not dependant on the surface extension. Once again, three values are to be inserted in order to compose a vector. With them, Comsol evaluates:

$$\mathbf{F}_A = \frac{\mathbf{F}_{TOT}}{A}$$

And then proceeds with the second Piola-Kirchoff stress tensor.

Another fundamental additional description concerns the symmetries of the system. The corresponding equation is quite simple:

$$\mathbf{u} \cdot \mathbf{n} = 0$$

I.e.: no normal displacement is allowed with respect to the selected boundaries. Tangent ones instead are. This command is immensely helpful and to be exploited whenever possible. Thinking to the whole stack structure, in fact, it is – almost – symmetric with respect to every spatial axis. This means that an eighth of it alone would be enough to simulate the behaviour of the whole stack. In fact, if the stack were ideally cut with a vertical plane normal to the x direction (as in Figure 3.18), it would be clear that no cell element could displace from one half to the other.




Figure 3.18: The cell cut with a plane in one of its symmetry directions.

If the object to study can be reduced to only one eighth of its original dimensions, the computational times can be hugely reduced. As it will be seen later on, this aspect is pivotal when dealing with simulations in general, and 3D mechanics is not an exception at all, despite only – apparently - involving few values to compute.

A constraint that is similar to symmetry is the prescribed displacement: as the name suggests, it allows to explicitly define the displacement value of the selected boundary, regarding the x , y or z direction at choice. Most of the times, such a displacement is set to be zero, but it could even be otherwise.

Regarding instead the hydraulics, several physics are available: the chosen one is “free and porous media flow”. Its main equation concerning the flow (including the time variable) is:

$$\rho \frac{\partial \mathbf{u}}{\partial t} + \rho(\mathbf{u} \cdot \nabla)\mathbf{u} = \nabla \cdot [-p\mathbf{I} + \mathbf{K}] + \mathbf{F} \quad (3)$$

the left-hand-side of the equation represents the accelerations: the first addendum with respect to time, the second one to space. Their sum is equal to the total force acting upon the considered element: p is the pressure, \mathbf{K} the permeability tensor, \mathbf{F} the external forces. In stationary conditions there are no evolutions over time, so the equation (3) is simplified into:

$$\rho(\mathbf{u} \cdot \nabla)\mathbf{u} = \nabla \cdot [-p\mathbf{I} + \mathbf{K}] + \mathbf{F}$$

The fact that ρ has a unique value states that the fluid is considered to be incompressible; it is in general homogeneous. Other assumptions are that the flow is fully developed both at the inlet and the outlet, while

there is not reflux at the latter, which also perceives no external pressure. The electrolytes are at constant temperature and *SOC*, so the density and viscosity do not vary.

The second equation Comsol writes in this physics is:

$$\rho \nabla \cdot \mathbf{u} = 0$$

Which means:

$$\frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y} + \frac{\partial u_z}{\partial z} = 0$$

So any acceleration in one of the space directions must be compensated by a deceleration in at least another direction. This further simplifies (3) into:

$$\nabla \cdot [-p\mathbf{I} + \mathbf{K}] + \mathbf{F} = 0$$

Every force must be locally compensated by some other. p is the pressure, and it is multiplied times the \mathbf{I} identity matrix, so that it can be summed to the \mathbf{K} permeability tensor. The scalar product with ∇ reduces their order, making them vectors, which can be then summed to the external applied force \mathbf{F} . The permeability tensor of the fluid is evaluated as:

$$\mathbf{K} = \mu(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$$

It depends on the viscosity and on the velocity. In this way the free flow of the electrolytes in the manifolds can be described. The situation in the electrodes is more complex, since they introduce the porosity. The option "porous medium" is added to the current physics, and it applies the Brinkman equations to the domain that is described as porous. The equation that corresponds to the previous (3) has more terms:

$$\frac{1}{\varepsilon_p} \rho \frac{\partial \mathbf{u}}{\partial t} + \frac{1}{\varepsilon_p} \rho (\mathbf{u} \cdot \nabla) \mathbf{u} = \nabla \cdot [-p\mathbf{I} + \mathbf{K}] - \left(\mu \kappa^{-1} + \beta \rho |\mathbf{u}| + \frac{Q_m}{\varepsilon_p^2} \right) \mathbf{u} + \mathbf{F} \quad (4)$$

Both the addenda of the left-hand-side are now divided by the porosity ε_p , which can therefore play a role on the accelerations. The lower the porosity, the higher they are, since the electrolyte has to flow through narrow passages. Besides the pressure p , the permeability \mathbf{K} and the external forces \mathbf{F} , there is now another addendum. κ is the absolute permeability of the porous medium, β the drag coefficient, while Q_m is

$$Q_m = \rho \nabla \cdot \mathbf{u} \left[\frac{kg}{m^3 \cdot s} \right]$$

So, equation (4) in stationary conditions is simplified into:

$$\frac{Q_m}{\varepsilon_p} \mathbf{u} = \nabla \cdot [-p\mathbf{I} + \mathbf{K}] - \left(\mu \kappa^{-1} + \beta \rho |\mathbf{u}| + \frac{Q_m}{\varepsilon_p^2} \right) \mathbf{u} + \mathbf{F}$$

The sum of all the acting forces is not necessarily nil. The evaluation of \mathbf{K} is slightly more complicated too:

$$\mathbf{K} = \mu \frac{1}{\varepsilon_p} (\nabla \mathbf{u} + (\nabla \mathbf{u})^T) - \frac{2}{3} \mu \frac{1}{\varepsilon_p} (\nabla \cdot \mathbf{u}) \mathbf{I}$$

The flow can then be set as Darcian or Non-Darcian, this respectively defines a linear or non-linear relationship. Since the non-linear is due to turbulent or inertial effects, the Darcian case is selected.

The velocity at the inlet is evaluated as the ratio between the flow and the section:

$$u_{average} = -\frac{1}{A} \iint_A \mathbf{u} \cdot \mathbf{n} dS$$

2. Description of the boundary conditions

The meaning of the various possible settings has been explained, and so were the corresponding equations. Now it is time to use those tools to describe the cell of the model.

All the objects are set as linear elastic, with initial displacement and velocity equal to zero. All the pairs are of contact type, and they connect the ribs of the FFs to the adjacent gaskets, and the top gasket to the BP.

Regarding the load, in the complete stack a pressure would be applied from the bolts onto the washers, which are located on the EPs. Such a pressure would be equal amongst all these washers, which furthermore justifies the adoption of symmetry later. The point is that these components (EPs, washers, bolts) have not been described yet, since this initial phase is dealing with a generic single cell alone. This means that the pressure is to be applied somewhere else, resulting in a model that is not totally accurate. This is not a concern though: the purpose of this chapter is to familiarise with Comsol, its commands and the way it works, rather than already reaching a final

result: this will be the aim of the next chapters. According to this idea, it is justifiable to set a slightly incorrect model, as long as it works and is quite representative of the stack situation.

Considering the generic single cell, which is the issue, a pressure is applied onto it, coming from the cell which should be on top of it. A BP alone was drawn, and over it there would be a gasket, which has an AA. This means that the top gasket is pressing on the BP, but not on its whole surface: just the one around the region of the AA. This is the reason why in this context the load is applied on the external part of the BP alone, not on its core too, that would correspond to the AA. This is shown in Figure 3.19.




Figure 3.19: Where the pressure is applied on the BP, in the case of the single cell simulation.

Such a setting justifies the previous choice of drawing the parallelepiped BP_AA, which seemed unnecessary in the geometry context.

Exploring the Definition directory, it is possible to notice that no pair corresponds to the boundary between BP_AA and the rest of BP_ext: this means that Comsol considers the two of them as distinct domains of the same object. In fact, they are both inside the object *dif5* (i.e. BP), which was obtained evaluating the Boolean difference between BP_ext+BP_AA and the four BP_chns.

Since there is no contact pair, the two domains of BP would not slide one with respect to the other, so it should be possible to apply the pressure onto BP_AA too: it would not slide away. Anyway, this would not be representative of the current physics, as discussed above.

At this point, there are two possibilities: either applying the load pressure both on the top side of the BP and on the bottom side of the lowest gasket, or just on the BP, while applying the symmetry condition to the gasket. In fact, the cell is amongst many others, and pressure is coming from both the EPs of the stack. If the drawn cell anyway represents the one in the middle of the stack, it is correct to assume the symmetry condition. Once again, the purpose of this simulation is to learn to use the commands, so – under this point of view – it is advisable to implement a variety of them, as long as they will come in handy later. There is also another reason, a computational one: being the top element (the BP) different from the bottom one (the Gk1), a double load would introduce some difficulties to the evaluation: convergence would not be reached.

iv. Setting the mesh

The drawing of the mesh is a crucial part of the Finite Element Analysis. On one hand, a fine one gives a good precision when solving the equations of the applied physics, but on the other it increases the computational times, which may even become excessive. The ability of the drawer stands therefore in making one which is

tailored for the case: finer where a good precision is needed, coarser in the not important parts, or where the physics is not interesting, for one reason or another.

This is not the only main method to reduce the computational times: some tricks when dealing with the solver options can have a huge effect in this context as well. A third way to reduce the times, but important too, is the exploit of symmetries, which was already mentioned above. The idea is that if the object can be cut in two parts that behave in the same exact way, then it is superfluous to analyse the whole of it: a portion is sufficient. Looking back once again at Figure 3.10, it can be observed that the object is specular with respect to two orthogonal planes, which cut it in four slices. The one that lays in the first quadrant alone is shown in Figure 3.20.




Figure 3.20: The fourth portion of the cell, cut along the symmetry planes.

In this way only one fourth of the object is to be analysed. Of course, it is necessary to tell the program that symmetries have been applied: this can be done in the Physics directory. Figure 3.21 shows the affected boundaries by this new setting.




Figure 3.21: The boundaries where the symmetry setting is applied.

Figure 3.21 shows that symmetry is to be applied on every surface which lays on the symmetry planes that have cut the cell. The corresponding equation, as already explained in 3.b.iii.1, is:

$$\mathbf{u} \cdot \mathbf{n} = 0$$

Which means that the selected boundaries of Figure 3.21 cannot move along their normal direction, but only in the tangent ones. In fact, a normal movement would imply that their twin does it too: either they compenentrate (not admissible) or detach (not the case).

Now that the final shape of the simplified cell is obtained, it is time to draw the mesh.

When domains were introduced, it was mentioned that they may be useful for this purpose, since they could subdivide an object in many parts. The subdivision can be done in various ways, for example with the Boolean operations and with the partitions.

The starting idea is that the ribs will be described with a tailored mesh, which will be set by hand. Their extrusion will be propagated on the neighbouring objects, in order to achieve a coherent description of these regions. This means that it is necessary to draw the shade of the ribs on the nearby gaskets. The “create imprints” command that can be found in the “create assembly” is not sufficient here though: in fact, it would only draw the ribs shade on the neighbouring gasket surfaces that face them, but this shade would not be dug throughout the gaskets’ width.

A simple solution to obtain so is to draw a couple more parallelepipeds when setting the gaskets: they are concentric with respect to Gk_ext, and their length and width are equal to the external and internal ones of the ribs, in the same way that the shade of the AA is cut in the BPs too. This all is shown in Figure 3.22.




Figure 3.22: The cell with ribs and AA projected in the gaskets and BPs.

The lines which correspond to the ribs and the AA are clearly drawn in the top BP (as they are in the bottom BP and all the gaskets as well).

It will be found that the pressure concentrates on the ribs and their projection in the other objects, while the other parts will almost be not affected by it. As soon as the EPs will be introduced, it will also be seen that the external perimeter of the cell is important to evaluate too. This all implies that it is necessary to describe in detail the internal perimeter (that faces the AA, where the electrode will be inserted) and the external one.

Such a detailed description can be achieved with a hand-made tailored mesh in those regions, while the other ones (inner) can be quickly set with a free-tetrahedral. The border between the perimeters and the core can be established with auxiliary work planes, which can then be used to cut the objects.

Work planes are easy to set: it is necessary to indicate their tangent directions (x-y, y-z etc.) and the distance from the origin along the normal direction. For example, the external perimeter can be cut with two planes, one of them represented in Figure 3.23.




Figure 3.23: The use of the work planes to cut the cell and describe the perimeters.

Once the perimeters are drawn, the partition is the right command to actually cut the perimeter. The object is still one and a whole, but is now described with two domains, which can be distinctly used for any successive command: mesh drawing, physics etc. So can their corresponding boundaries, edges, and points. Figure 3.24 shows how the work plane of Figure 3.24 can partition the cell.




Figure 3.24: The external perimeter of the cell partitioned with a work plane.

The same can be done to conclude the description of the external perimeter along the other direction, and to set the inner one as well. Four more work planes are drawn to cut a square around the manifold. In fact, the manifold might show particular values of pressure, so it is better to take a closer look at it. The work planes allow to cut a square region around it, where the mesh will be set in a finer way, remaining coarse in the rest of the core. The order with whom the partitions are performed matters too, since the resulting objects from the previous can be the input of the following.

By partitioning first the external perimeter, second the square around the manifold and last the inner perimeter along the AA, it is possible to divide the stack as shown in Figure 3.25.




Figure 3.25: The cell after all the partitions.

As already said, the inner and outer perimeters will be described with a tailored mesh, the square around the manifold with a normal free-tetrahedral mesh, while the rest of the cell with an extremely coarse one.

The first step to draw an extruded mesh is to map its boundaries. The partitions that follow the inner perimeter will have the same number of elements in the horizontal directions, but different in the vertical one, since some of them are taller than others.

The first mapping refers to a FF and the corresponding ribs, and to a BP. Their vertical surfaces are mapped, and the top ones are too. Once they are selected, it is possible to tell Comsol what to do with them: if instructions were given about not selected boundaries, they would not be applied.

The edges of the selected boundaries are described with distributions, the command that tells how many mesh elements are present along the edge. This value can be set a number, or as a parameter which was already specified in the Global Definitions context. For example, the vertical edge of the ribs is divided into four equal elements, while that of the FF into six. It is not necessary to assign this division to both the vertical edges of the vertical boundary: Comsol understand that that quantity is assigned, and by default evenly spaces them. A little more attention is to be paid when drawing the corner of the inner perimeter, though: all the edges need to be selected and described, since the surface is not a simple rectangle, as shown in Figure 3.26.




Figure 3.26: The mapping of the upper boundary of the inner corner of the inner perimeter.

It is also possible not to evenly space the elements. This is adopted while getting further from the corner of Figure 3.26. In fact, the pressure concentrates on the angles of the ribs, but then assumes a more regular distribution along their lateral direction. This means that they can be described with a mesh that becomes the coarser the further from the corner. An increasing or decreasing distribution of elements can be set, and their relationship can be of either exponential type or linear. Besides the number of mesh elements along the edge, Comsol then needs to know the ratio between the larger and the shorter.

Unlike in the evenly spaced case, the edge setting needs to be applied on both the edges. If in fact only one were set, Comsol would draw what is shown in Figure 3.27 a). In order to obtain Figure 3.27 b), both the edges must be selected.




Figure 3.27: The exponential distribution along an edge of the mapped boundary, a) when only one edge is selected, b) when both of them are.

Once the boundaries that surround a domain are mapped, it is possible to extrude their mapping along the domain volume. In this way, Figure 3.28 is obtained.




Figure 3.28: The extrusion of a FF and its ribs, and of the bottom BP.

In the same way, the external perimeter is mapped and extruded too.

The horizontal surfaces of the gaskets are mapped in the same way of the FF, the only difference is the number of elements along their height. It is not necessary to set everything again for the gaskets too: it is possible to copy the mapping from the FF, then the number of elements along the vertical direction alone is still to be specified. Besides the boundary, even a whole domain mesh can be copied onto another domain, provided it has the same shaped (or is at least proportional).

After all these extrusions, the whole inner and outer perimeters are meshed, as shown in Figure 3.29.

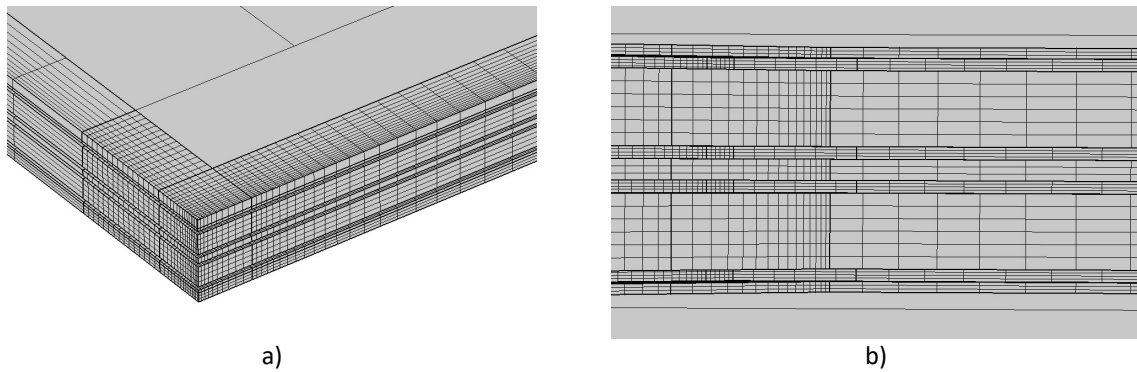


Figure 3.29: The extruded a) external and b) internal perimeter of the cell.

When choosing the number of elements along the vertical direction, a care shall be taken. In fact, the gaskets are pressed against the FFs, but since their respective Young's moduli are different, one object will be more deformed than the other. In particular, the gaskets are in EPDM, which is way more flexible than the PVC of the FFs, which means they are going to be more compressed. It is therefore better that the gasket elements in contact with the FF are described with a mesh that is finer than the one of the FFs: indicatively with the double of the mesh elements; in this way the deformations can be properly evaluated.

The rest of the cell can just be set as free tetrahedral. This free disposition should nevertheless be finer around the manifolds than in the rest of the cell, so two approaches may be taken.

A simple one would just describe the square around the manifold with a normal tetrahedral distribution, leaving the rest with an extremely coarse one, but this would result in what is shown in Figure 3.30.




Figure 3.30: A normal free tetrahedral mesh around the manifold, and an extremely coarse one in the remaining domains.

This description would not be poor per se, but quite a waste of computational power, since it would accurately evaluate some parts of the plates which do not require that much attention. Not even a very coarse mesh around the manifold would fix this issue, as shown in Figure 3.31. It would nevertheless draw a very pointy circumference around the manifold, which is detrimental.




Figure 3.31: A very coarse mesh around the manifold.

This issue can be solved instead with a mapping. In particular, the circumference of the manifold is to be mapped, specifying the number of elements that describe it. A following command that draws a free tetrahedral mesh around it would necessarily respect this condition, even if set as extremely coarse, as shown in Figure 3.32.




Figure 3.32: The mapping of the circumference of the manifold can be compatible with an extremely coarse mesh around it.

The final object is shown in Figure 3.33, with the complete mesh.




Figure 3.33: The complete mesh.

The cell is described with 75083 elements. The higher the number, i.e. the finer the mesh, the more accurate the simulation, but the more expensive it is in terms of computational times. It might even happen that the RAM of the computer is not sufficient for the calculation.

The key idea is to find a good compromise between a precise evaluation and an acceptable time. Supposing that the finer the mesh, the closer the result is to the one reachable with an extremely fine one, it is usually assumed

that this result is approached along an asymptotic path. This is what the mesh sensitivity analysis deals with, in fact. A practical example will be shown later on, when dealing with the hydraulics simulation.

v. Setting the solver

The component has been completely described with Geometry, Materials, Physics and Mesh, in this order. More components may be parallelly described and set to interact one with the other, but this is quite a complex operation, and is not necessary in this case. The current study therefore deals with one component alone.

Concerning the study, it is the next directory of the Comsol file.

The main study types deal with the frequency (autofrequency, frequency domain), the time (transitory) or stationary situations: the last option is the one to consider. Even more studies are available, depending on the selected physics, for example the solid mechanics introduces the bolt pretension, the casual vibration and much more.

As already mentioned in paragraph 3.b.iii, the stationary case eliminates the time dependency from the equations of the physics, which are therefore simplified. This means that the pressure on the cell is constant, and not gradually applied.

Under the Study directory, a lot of settings can be modified, and each of them brings some important impacts, either on the solving of the problem, or concerning the way Comsol understands the component.

A first useful tool is the auxiliary sweep, under the study extensions section. It was introduced at the very beginning how it is possible to describe the objects etc. with global parameters, which can later be easily modified. This tool further exploits this possibility, making one of these parameters at choice variate as desired. For example, it could be possible to first simulate the component with a Gk_z of 5 [mm], and then reduce it step by step until the correct value of 0.5 [mm]. This would require multiple simulations to run, but might be helpful, since the results of the previous one are set as starting point for the evaluation of the successive one. It might in fact happen that the final value (e.g. Gk_z of 0.5 [mm]) describes a situation which is very difficult to evaluate for the program, so a slower approach might be necessary.

The box "Define load cases" works in a very similar way: it particularly refers to the load which is applied to the component, as assigned in the physics. This load is slowly applied, simulation after simulation, until its full value according to load factors.

For example, if a load factor of (0, 0.5, 1) was given, three simulations would be run. The first with no load at all, useful to determine the rest position of the object. The second one would search for the solution when half of the load is applied, using as starting point of the evaluation the result found for the no-load case. The third and last simulation would start from the solution of the 0.5 case to evaluate the result for the full-load situation. Three simulations take more time than a single one, but this option can be crucial if the single one is too hard to evaluate for the program, thus not leading to a result at all (i.e. not reaching convergence).

Under the Dependent Variables subdirectory, Comsol lists the outputs it has to evaluate considering the assigned physics. In the case of solid mechanics, the Displacement Field has them all. There are three of them, along the three spatial dimensions, respectively comp1.u, comp1.v and comp1.w. The Free and Porous Media Flow would add the pressure (comp1.p) and velocity field (comp1.u2, comp1.v2, comp1.w2) as well. It can be deduced that the "comp1" part of the variable name refers to the component (which is in fact named comp1), even though there are no more besides that one, while the "u", "v", "w" are the spatial directions of the velocity or displacement. In case of possible confusion, the number 2 is added. The pressure is instead isotropic, a scalar, not a vector, so it does not need any distinction along the three directions.

The Stationary solver subdirectory contains instead all the information regarding – in fact – the solver itself. The relative tolerance is first set: its standard value is of one thousandth, but it can easily be set otherwise. It represents the relative error between two consecutive found solutions along the iterations. As soon as the found solutions are closer than the tolerance, convergence is considered to be reached and the problem is solved. Consequently, the convergence velocity represents how many iterations are needed to reach the set tolerance: a good solver can do so in few steps, where the error exponentially decreases.

There are two types of solvers: direct and iterative.

Direct solvers are based on the LU decomposition of the matrix that describes the problem. Any direct solver reaches the same solution, both in well-conditioned problem and in some ill-conditioned ones. They are robust, but usually require an important computational effort, therefore a lot of RAM. If the RAM of the computer doing the calculation is not sufficient, it may not be possible for the program to solve the problem, or it might exploit the out-of-core mode. The bigger the matrix describing the problem, the higher the computational effort.

Iterative methods instead gradually approach the solution: they are therefore slower, but also lighter to run, i.e. they do not take as much RAM as direct methods do. This is an advantage when dealing with problems of big dimensions, but they are nevertheless weaker with their convergence towards the solutions. Some of them are fitter for some physics rather than for others. [59]

Direct solvers are generally preferable for solid mechanics problems. As just mentioned, they tend to require more RAM than iterative ones, so the application of symmetries and the tailoring of the mesh were crucial steps in easing the resolution process, since they sensibly reduce the dimensions of the matrix to solve.

Three direct solvers are available in Comsol: PARDISO, MUMPS, SPOOLES and the dense matrix solver. They all work on systems of the form $Ax = b$.

- PARDISO (Parallel Direct Solver) is more suited for large sparse matrices.
- MUMPS (Multifrontal Massively Parallel Solver) is robust and efficient for large linear systems, usually employed in problems that require high precision.
- SPOOLES (Sparse Object Oriented Linear Equation Solver) is faster and more suited for smaller systems.
- The dense matrix solver is suited for very populated matrices. [60]

Comsol usually suggests PARDISO in solid mechanics problems.

If the problem dimensions are excessive for the RAM of the computer, the PARDISO solver can run anyway, but in the out-of-core mode. It means that the evaluation is not hosted in the RAM (in-core mode), but rather on the memory of the computer itself, where it can find enough space. This is advantageous, since it allows to overcome the limit of the RAM dimensions, but should anyway be avoided: the times become sensibly longer, due to the slower velocity with which the information is collected from the memory of the computer, rather than from the RAM. To avoid this operating mode, it is necessary not to operate with too big of a matrix, i.e. it is necessary to assign a smaller problem to the computer to solve.

Symmetries were already implemented for this purpose, and the mesh was tailored too. Another trick can be performed, with great results: the segregated approach.

As previously said, the problem is described by a big matrix, where all the equations are stored. The greater the dimensions of the component and the more the variables to evaluate, the bigger the matrix. This happens in the fully coupled approach, where all the variables are parallely solved together: comp1.u, comp1.v, comp1.w. The segregated approach instead unpacks this big matrix into more, each of smaller dimensions, according to the settings written by the programmer. Since each of these new matrices (which altogether describe the whole problem too) is way smaller than the full one, they are singularly extremely faster to solve. Due to their smaller dimensions, they also often guarantee to avoid the out-of-core mode using.

When the segregated approach is selected, Comsol asks what each of its segregated steps must do, i.e. what are the respective variables. They can be chosen from a list, and in this case there are the three of them mentioned above.

Recalling that the axial direction of the component is z (which corresponds to comp1.w), it is possible to set one segregated step as dealing with comp1.w, while a second step evaluating comp1.u and comp1.v together. The computation is way faster, but the simultaneous convergence of the two steps is not guaranteed. It depends whether the respective solutions collide or can coexist. The first step usually starts with an incredibly small error from the very first iteration, while the second step has a higher one. This is due to the fact that the second step has to adapt to the solution found by the first one of the other variable. Despite this initial error growth, convergence can soon be reached, if the problem is well-set.

It is possible to unpack the problem even more: three segregated steps: one for comp1.u, another for comp1.v, and the last for comp1.w. The times are even faster. Figure 3.34 provides an example for this.




Figure 3.34: The convergence of a three-step segregated solver.

The possibilities are given by the possible combinations of these three components, since two steps can involve the same variable: it is for example possible to set a first step with comp1.u and comp1.w, and a second step with comp1.v and comp1.w. Or any other combination.

A further little help can come from the scaling factor. This quantity tells Comsol the order of magnitude it can expect the solution to be: the closer to the correct one, the faster the resolution. The default value is not far from the one of this specific case, anyway, so it can be left as found.

vi. Analysis of the results

Once the simulation has run, the results can be viewed. They are stored in the Dataset section, and the right study results have first to be selected. In fact, it is possible to run different studies for the same component (stationary, transitory etc.), and is therefore necessary to specify which one's results are to be viewed.

This solution *sol1* contains all the information concerning the whole component, and it might be desired to only view a fraction of them. Cut planes and cut lines can be useful in this context. Still under the Dataset subdirectory, it is possible to select a cut plane, in fact, which is nothing but a selection of the results of *sol1* (which needs to be explicitly chosen), all laying on the indicated plane.

The plane can be of quick type, or general type. A quick plane is orthogonal to one of the three main directions, and is therefore faster to specify: it is just sufficient to indicate which type it is (xy, yz etc.) and its distance from the origin along the perpendicular direction. It is of course convenient to assign such a distance using the parametrised values. For example, if the half-height of the first FF were to be chosen, it could just be written the coordinate $Gk_z+rib_z+FF_z/2$, rather than its numerical value. This is once again convenient, were the parameters to be varied, maybe inside an auxiliar sweep. General planes do not have to be perpendicular to one of the three cartesian directions instead. They are defined with three distinct points, since Geometry teaches that – in fact – three distinct points define a unique plane in space. Since the cell's axis is parallel to the z direction, a quick plane is sufficient for the analysis of the results.

Cut lines do not work differently: they are 1D objects, specified by their two edges, whose x, y and z coordinates are to be written.

Both when dealing with planes and lines, it is also possible to select more of them at once, to view more 1D or 2D results per graph. Regarding planes, they will be of course parallel, and their distance along the normal direction will have to be specified. Were more than two parallel planes desired, the distance would need to be a vector. The additional lines work similarly, but it is also necessary to specify the 3D direction of the distance. For

example, if the first line were along x and passing through the origin, the second line might be passing through (0, 0, 1), or maybe (0, 1, 1).

If two perpendicular cut lines were desired, it would be necessary to set two distinct cut line commands.

When the drawing is shown in the Dataset subdirectory, the simulation results are not displayed yet. Comsol instead represents the considered plane, or better its intersection with the component. Figure 3.35 shows the view for a) one cut plane and b) eight cut lines.

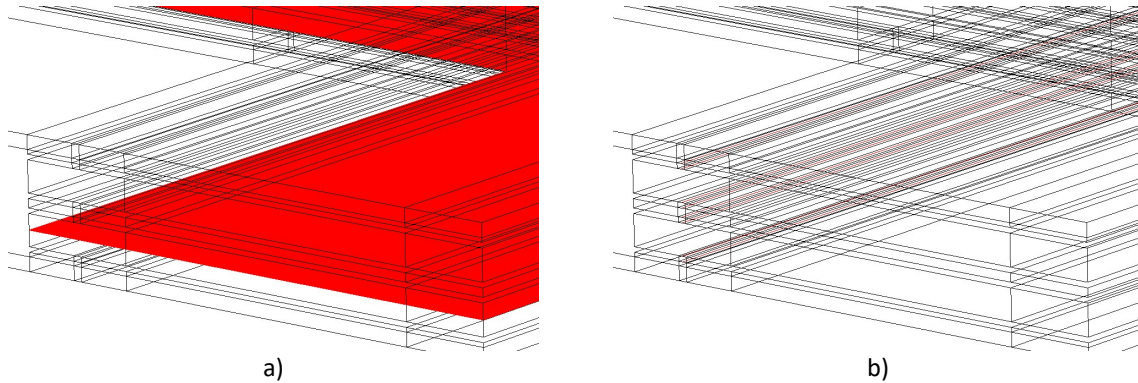


Figure 3.35: The view of a) a cut plane, b) eight cut lines, from the dataset of the study.

It is finally possible to view the results. They can be selected in a 3D, 2D or 1D context: this decision is taken with respectively selecting the 3D plot group, 2D plot group or 1D plot group.

Inside the 3D plot group, it is possible to select the results of the whole component of the study, as shown in Figure 3.36.




Figure 3.36: The view of the volume results in a 3D plot group.

It is otherwise also possible to select a single surface, which was previously defined in the dataset section: Figure 3.37.




Figure 3.37: The view of the surface results in a 3D plot group.

And the lines too: Figure 3.38.




Figure 3.38: The view of the line results in a 3D plot group.

A 2D plot group is not conceptually different. It just focuses on a specified surface. So, the same result of Figure 3.37 can be taken in another way in Figure 3.39.




Figure 3.39: The view of the surface results in a 2D plot group.

Last, the 1D plot group only considers lines. This means that the output variable can be shown in the y direction in a cartesian plane, so it might be easier to study on a numerical point of view. The eight cut lines are shown in Figure 3.40 and they represent the same result of Figure 3.38.




Figure 3.40: The view of the line results in a 1D plot group.

The way Comsol works has been thoroughly described: now it is time to face the simulation of the detailed cell and to analyse the results.

c. The hydraulic simulation of the cell

Chapter 2 dealt with the importance of the compression of the electrodes, and how this impacts on the performances of the cell under various points of view. A homogeneous compression is preferable, since it better exploits the carbon felt and its properties, so the mechanical simulation will aim to find the best EP configuration to ensure this homogeneity.

All of this might be ruined though by an inhomogeneous electrolyte flow, which would prejudice all the work: in this case the objective itself of achieving homogeneous compression would not make sense anymore. It is therefore logical to first check if the electrolyte flow is indeed homogeneous: in case it is not, the electrode compression task might need further adjustments.

The fluidodynamic simulation is first performed on a simple single cell, shown in Figure 3.41.




Figure 3.41: The structure of the simple single cell.

Its geometry is different from the one built in the previous chapter, in fact it only represents the electrode and the electrolyte channels. The first would be inserted in the AA of the gaskets and of the FFs seen before, while the channels would be engraved in the FFs. Since they only reach the extremities of the electrode, injecting into and receiving from it the electrolyte flow without particular guidance, the shown configuration is a flow-through case.

The object of Figure 3.41 represents a semi-cell: it is either flooded by the catholyte or by the anolyte. The two of them follow an analogous path, with no difference at all: the only one stays in their parameters: viscosity, density. This is the reason why two materials are inserted, one to study the catholyte, the other for the anolyte, and the two of them correspond to two distinct studies, so that the respective results can be both saved and observed, and even compared. Of course, the physic is the one introduced back in 3.b.iii.1: free and porous media flow.

Regarding its setting, the electrode is the porous medium, while the channels host the fluid (either the catholyte or the anolyte). The initial values ask what the electrolyte pressure and velocity are throughout the component: this is all set as zero. The solution will find that the outlet pressure is zero, and the inlet one will be such to compensate for all the pressure drops throughout the electrolyte path. It could have also been possible to already put the inlet pressure at such a value to ease the evaluation, nevertheless.

Another needed specification is the one regarding the behaviour of the walls: no slip is set. The other options would have been suited to represent turbulent flows, wall roughness (perhaps due to sand), but the case in analysis is of laminar flow. The corresponding equation is:

$$\mathbf{u} = \mathbf{u}_{tr}$$

i.e. the electrolyte velocity \mathbf{u} at the walls is equal to the velocity of the walls themselves \mathbf{u}_{TP} : they are still one with respect to the other. Since the object is not moving, the electrolyte is still at the walls. The flow is assigned at the inlet: since its area is known, the velocity is found by their ratio. As mentioned above, the inlet pressure is not given, but left to evaluate. The outlet pressure is set to zero as well. The flow is completely developed and laminar.

The mesh of the cell is shown in Figure 3.42.




Figure 3.42: The mesh of the single simple cell.

The mesh is drawn finer at the walls of the channels to study with precision their effects on the laminar flow. At instead the centre of the channel section the mesh is slightly coarser, since there the flow is more uniform. Regarding the electrode, the same principle was adopted. In particular, the transversal distribution was obtained with a symmetric exponential distribution, again coarser at the core, finer at the edges. Concerning the felt, it was considered as already compressed: 20%. Figure 3.43 shows the results of the simulation, in a) the catholyte and b) anolyte case, respectively.

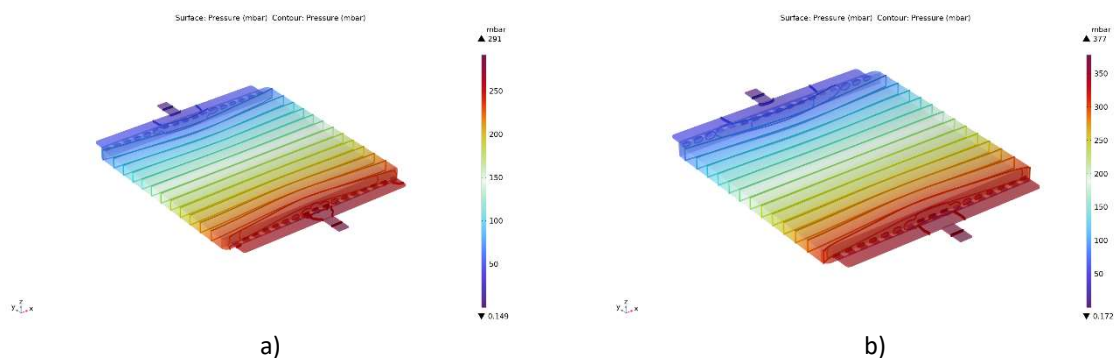


Figure 3.43: The result of the fluidodynamic simulation: the pressures in a) the catholyte and b) anolyte case.

Before proceeding with the analysis of the simulation results, it is necessary to prove that they are actually valid. It might be supposed, in fact, that the coarser mesh at the electrode core may lead to an incorrect or approximated evaluation: that it was not a legitimate assumption. The mesh sensitivity analysis deals with this issue. It demonstrates whether the coarser mesh (in some region) was feasible without loss of information, or anyway without leading to evaluation mistakes.

The idea is to perform the study with a parametrised mesh, where the number of its elements can be varied as desired. The higher this number, the finer the mesh, the more accurate the result. And the various results should converge towards the one found by the finest mesh. It is possible to represent this convergence with a graph, where the model sensitivity is shown versus the number of elements: Figure 3.44.




Figure 3.44: The convergence of the mesh sensitivity analysis.

This required a third distinct study. Since the model is indeed valid due to its convergence, looking at Figure 3.44 it is eventually possible to choose the best number of elements for drawing the mesh: enough to find an accurate result, not too many in order to limit the computational times.

The results of Figure 3.43 can also be compared with the laboratory values found, concerning the pressures inside the electrodes and channels. They are indeed coherent, even though the actual electrodes are not uniform in their porosity. Figure 2.2 showed in fact that the irregularity of the carbon strains should oblige the electrolyte to follow a tortuous path, while the one here is very regular. Comsol in fact considers the porosity to be homogeneous. In any case, the simplification of the simulation is not too misleading, when its values are compared with the ones of the laboratory: as just said, they are coherent.

Figure 3.45 shows the distribution of the electrolyte velocities through their streamlines, which are indeed quite straight inside the felt.

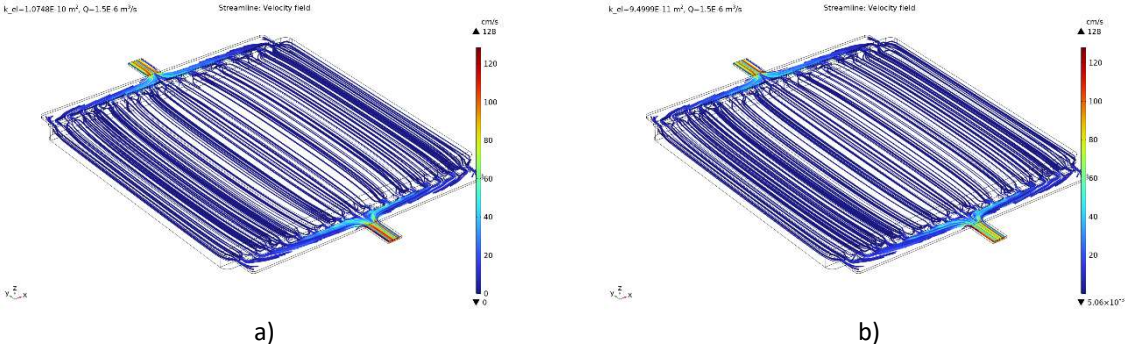


Figure 3.45: The result of the fluidodynamic simulation: the velocities in a) the catholyte and b) anolyte case.

The electrolytes flow fast through the inlet and outlet, and rapidly spread in the electrode, where they immediately decelerate and then flow uniformly. This is great news: the channels do not cause too much of a pressure drop, nor do they hinder the flow, while the felt is well exploited. This recalls the concept of active

surface area: all the wetted carbon strains can in fact host the redox reactions, and this was the principle that made the battery work. Since the product of the area times the velocity is constant for any cross-section throughout the electrolyte path, besides the fact that the inlet and outlet have fast flows (due to their smaller areas), the electrode has instead a slow flow (due to its wider area). This is convenient, because it will cause the electrolytes to spend more time inside of it, thus increasing the possibility for the vanadium molecules to react on the carbon strains. They are also invested by the flow with a constant rate, since its velocity does not decelerate much through the electrode.

The streamlines of Figure 3.45 also show that the electrolyte path is quite rectilinear, i.e. does not perform vortices, nor takes unexpected routes. This is once again positive for the electrochemical efficiency of the cell, since it further proves that the electrode is uniformly exploited. The only point where the electrolyte does not fully use the electrode is right after the incoming channels' exit, but this is a design issue, whose solution will be explored in the following.

Regarding the channels, the flow is uniform throughout them, and they do not provoke too much pressure drop. A look at the isobar lines of Figure 3.43 further shows how the electrolyte immediately decelerates after its exit from the channels into the electrode, but its spread through the volume needs some space to evolve. Then it is greatly homogenised by the felt.

It is also interesting to notice how there are no substantial differences between catholyte and anolyte case.

This simple cell, which was just shown and analysed, can be improved, thus arriving to its greater version, shown in Figure 3.46.




Figure 3.46: The structure of the greater single cell.

The greater cell has increased dimensions with respect to the simple one, which means it can host more redox reactions, thus achieving higher powers. This requires though a more complex structure concerning the channels: a single inlet may not guarantee that the electrolyte uniformly floods the electrode. It might instead happen that the flood is stronger at the centre of the felt, where the inlet channel exits, and weaker at the edges. This is not desirable, though, since it would not lead to a uniform exploit of the felt.

The solution is found by adopting three distinct channels for the inlet, and three for the outlet. They reach the electrode in three distinct spaces, and each one of them floods a third of it. They all receive (or drop) their electrolyte in a single manifold, though, so their exits are not at the same distance from this point. This would lead to three paths of different lengths: it would not bring a uniform pressure, since the longer path would provoke the higher hydraulic pressure drop. It might be useful at this point to project the longer path with a larger cross-section, to compensate for its superior length, but this might be an issue too, since they are anyway to be

engraved in the FFs. This component would in such a case present an even more irregular structure, leading to difficulties in the mechanics context. The chosen solution is instead to lengthen the shorter paths, even adding a couple more turns for each one. In this way, looking at Figure 3.46, all the three paths have the same length and the same number of turns, therefore the same pressure drop and the same electrolyte regime. This is positive for the uniformity concerning the felt use.

An artificially increased pressure drop might seem detrimental for the performances anyway: one might try to project another way to arrange these paths, for them to have the same length and number of turns, but equally reduced. It is not a total disadvantage, though: as it was introduced in chapter 2, longer and narrower paths are actually positive concerning the shunt current losses, instead: the electric resistances of their paths are increased. Furthermore, it was also seen in chapter 2 that the pressure drops are anyway more concentrated in the electrode, than in the channels [38], so the flow-through configuration of Figure 3.46 is actually acceptable. Its name is Equal Path Length (EPL).

Provided that the EQP ensures that the electrolyte enters the felt at the same velocity along all its edge, it was seen in Figure 3.45 that the flow needs a little space to spread throughout the section. In this first region, which is anyway small, the felt is not fully flooded, so slightly wasted. The project of the greater cell (Figure 3.46) solves this issue with the introduction of an intermediate region between the inlet channels' exit and the electrode bulk: a thin hollow volume. Despite its narrowness, it is sufficient for the electrolyte to spread along the transversal directions, so that it can then uniformly enter the electrode and correctly flood it in all its section from the very beginning. Homogeneity is therefore achieved, and all the AA is used.

Figure 3.47 shows the result of the simulation: the streamlines of the electrolytes' flow.




Figure 3.47: The result of the fluidodynamic simulation of the greater the cell: the velocities in a) the catholyte and b) anolyte case.

The streamlines cross the electrode following straight paths, and their uniformity is already achieved at the very beginning, thanks to the hollow region. Figure 3.48 shows instead the results concerning the pressures.




Figure 3.48: The result of the fluidodynamic simulation of the greater the cell: the pressures in a) the catholyte and b) anolyte case.

Despite the channels' length and their turns, the pressure drops they provoke are not excessive. As it was expected, the drops are instead localised along the electrode felt; being the isobars of Figure 3.48 plain, aligned and parallel, the pressure decreases in a linear and constant way along the felt.

It is also interesting to take a look at the pressure drops versus the flow rate: Figure 3.49.




Figure 3.49: The relationship between total pressure drop and flow rate.

Recalling chapter 2, since the flow regime is laminar, there is a linear relationship between pressure drop and flow rate, which can be seen in Figure 3.49. The steepness of each curve depends on the viscosity of the involved flow: the more it hinders the flow, the higher the pressure drop. The anolyte has a slightly higher angle than the catholyte, since it has a slightly higher viscosity.

This difference concerning the pressure drop means that there is a pressure difference between the two sides of the ion-exchange membrane. A pressure difference times an area gives a force, which the membrane must be able to sustain, otherwise it breaks. Figure 3.49 should therefore be used in order to determine the maximum electrolyte flow allowed: the one that corresponds to the maximum possible pressure difference between the two semi-cells for the membrane not to break.

This limit might be overcome, though, with an overload of the catholyte flow. This would in fact increase on purpose its pressure drop, reducing the difference between the membrane layers and saving it. This would of course also involve higher pumping powers (losses), but might be necessary to achieve higher output powers from the VFB. The surplus catholyte would not be wasted, anyway: its ions would just not react, and in the end come back to the tank.

It is interesting to make a comparison with the flow-by design as well. The interdigitated option is considered, and shown in Figure 3.50.




Figure 3.50: The structure of a simple single cell in the flow-by interdigitated case.

It was said in chapter 2 that the advantage of the flow-by design is that it performs better under the hydraulic point of view. In fact, the electrolyte can flood the electrode from a much larger surface, and can reach every part of its AA with an equal ion concentration. In the flow-through case instead, the concentration was gradually decreasing from the inlet to the outlet. The pressure drops are also lower, since the path through the felt is reduced in length, so the sufficient inlet pressure to feed the system is lower. This can in fact be read in Figure 3.51, which shows the analysis of the pressure drop of this flow-by case.

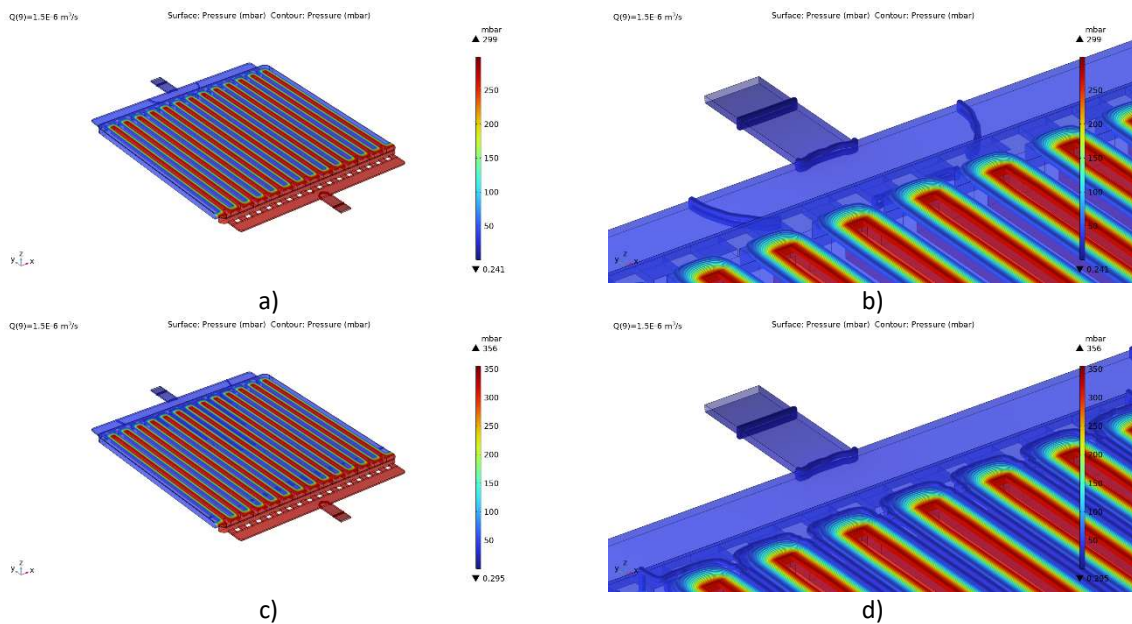


Figure 3.51: The pressure in the flow-by configuration: a) the catholyte case with b) its magnification, and c) the anolyte case with d) its magnification.

The enhanced hydraulics is nevertheless counterbalanced by a more difficult mechanics. In fact, the ribs are to be engraved in the BP, which is an additional request in the building phase. It also leads to a less uniform compression of the electrode, and causes the intrusion of the carbon strains in the engraved channels too, as it was explained in chapter 2 as well.

Finally, Figure 3.52 shows the electrolyte velocity.

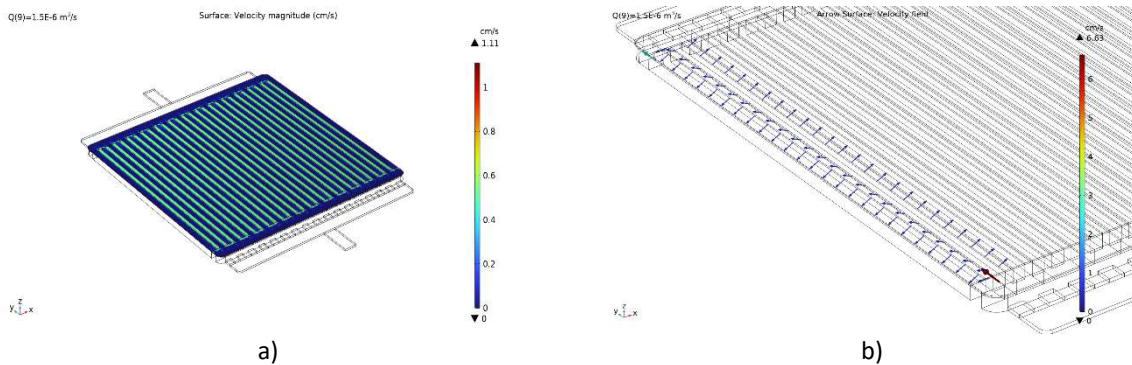


Figure 3.52: The a) catholyte velocity in the flow-by configuration, with a b) magnification of its jumps from one channel to another.

The electrolyte jumps from the inlet channels towards the outlet ones, passing through the electrode. This means anyway that the felt region below them is not uniformly crossed by the flow.

So, the flow-through configuration uniformly floods the whole electrode, but the concentration lowers along the path, and so does the pressure. The flow-by configuration instead achieves a more uniform pressure and concentration distribution, but does not equally use the whole AA. Despite these local differences, the aim to achieve a uniform compression of the whole carbon felt persists.

d. The mechanical simulation of the laboratory cell

Chapter 3.b thoroughly described how Comsol works: how the objects have to be drawn and set to launch a proper study. It used a simplified cell as example for the process, whose main purpose was not to reach a particular result, but to describe the way in which such a result is found.

Chapter 3.c proved instead that the mechanical simulation has to find an optimised structure for the cell that is able to uniform the electrode compression. The meaning behind this purpose was thoroughly shown throughout chapter 2, and further validated in 3.c.: were the electrolyte flows through the carbon felt inhomogeneous, a homogeneous compression would not have been justified.

Now that the aim of the study is justified and the process has been described, it is time to deal with the exact cell configuration: the one of the laboratory cell. It is shown in Figure 3.53.




Figure 3.53: The structure of the laboratory cell for the ultimate mechanical simulation.

A quick look at Figure 3.53 is enough to notice that there are some differences with the simplified structure of chapter 2.b. In fact, the previous one only represented a cell of the stack: two semi-cells (each made of a FF pressed between two gaskets), sided by a BP. Here the drawing shows a whole stack made with a single cell. This adds a copper-alloy CC, an insulator and an EP. Symmetry has been exploited, so only half of the cell is shown. The ion-exchange membrane has been neglected once again, since it does not affect the mechanical simulation, due to its reduced dimensions. Its mechanical requirements for integrity were anyway already presented in the previous paragraphs, when dealing with the pressures of the electrolytes.

Besides the components that are now introduced, it is possible to see how the old ones are shaped in a different way. In the simplified cell, in fact, the external and internal perimeters of the gaskets coincided with those of the FFs and of the BP. The only difference between those parallelepipeds was the respective height. In the laboratory cell it is possible instead to see how the gaskets are laterally hosted inside the FF, since they are narrower.

The FF now has three perimeters of ribs on each side: one following the AA of the electrode, another the external perimeter, and a third one in between. The gasket is laid in the space between the inner and the outer rib, and its edges are also rounded. For the sake of evaluation simplification, only the parts of the gaskets adjacent to the FF ribs are reported. The gaskets are slightly thicker than the ribs, which means that in the starting assembly the inner and outer do not press yet on the BP, but some empty space is between them. The compression will reduce and ultimately shut this gap.

Besides the improved structure of the semi-cell, the design is not complex as far as the CC, insulator and EP are concerned. In fact, the CC is nothing but a parallelepiped where the holes for the manifolds were cut, and so is the insulator. This last component is made of the same material of the FF, but its structure is simpler than its: it does not feature any ribs, nor channels. Recalling chapter 3.c, it would have been more correct to draw the channels engraved in the FF as well, but they were not reported for the sake of simplification. The model of Figure 3.53 is in fact detailed enough, and they would not have brought any substantial difference in the results, but an enhanced complexity in the setting of the mesh and of the physics. Regarding last the EP, its structure is evidently larger than that of any other components. It needs in fact to host the bolts and the tie bars, which connect it to the EP on the other extremity of the stack and are placed next to it, on the outside. Figure 3.53 reports two holes for these tie bars: considering that symmetries were applied, it means that their total number is eight. The position of the bolts is not random, but chosen in order to enhance the compression uniformity. Were they in fact positioned at the corners of the stack, the pressure would have mostly strained them. Putting them along the sides is instead preferable, since their effect better spreads through the stack section.

The tie bars are not shown in Figure 3.53, but their impact can be correctly represented anyway: they do not directly press against the EP, but upon washers, as shown in Figure 3.54.




Figure 3.54: The washers on the EP.

The fact that the ribs of the FF do not initially touch the BP, but this only happens once the compression is sufficient to push them low enough to shut the gap requires a particular setting in Comsol. The bottom surface of each rib has to be coupled with the top surface of the BP with a contact pair. The corresponding equation of

such a setting was explained in chapter 2.b.iii.1, and stated that as long as the gap between them is positive, no force is transmitted, otherwise it is. The initial contact between the gasket and the FF is instead described with an identity pair, since the two objects do not move away one from the other, but remain attached.

The EPDM of the gasket is softer than the PVC of the FF, so the latter is set as source, while the first as destination in the pair command. Concerning the materials in general, no change was applied with respect to what was previously described. The CC is in copper, the insulator in PVC (as the FF), and the EP in aluminium.

The properties of this alloy can widely change according to the ligands and the thermal treatments that were performed during its production.

Al6061-T651 was adopted. The first number determines the series, in this case the 6xxx. The alloys that belong to it present *Mg* and *Si* ligands, which make them easy to machine, weldable, and able to be hardened by precipitation. [11] The weldability can be a useful property as soon as the EP is to be strengthened with the addition of other components to its starting structure. The 6061 was chosen amongst the others since it is the most commonly used: as a consequence, it is also easy to find from sellers and to acquire at reasonable prices. The second part of the alloy name describes the treatment performed during its production. "T6" means that the metal was first heat treated, then artificially aged. In particular, "T651" adds that the stress was relieved by stretching. This can widely improve its mechanical characteristics with respect to the same alloy but with another treatment. For example, Table 3.3 compares those of Al6061-T651 with those of Al6061-O (i.e. annealed). [61,62]

Table 3.3: The enhanced mechanical properties of the Al6061-T651 of the EPs, with a comparison with Al6061-O.

	Al6061-T651	Al6061-O
Hardness, Brinell	95	30
Ultimate tensile strength [MPa]	310	124
Yield tensile strength [MPa]	276	55.2
Modulus of elasticity [GPa]	68.9	68.9
Poisson's ratio [1]	0.33	0.33
Density $\left[\frac{kg}{m^3}\right]$	2700	2700

It is not a surprise that the thermal treatment does not affect the density of the alloy, since it depends on the chemical components that form it. It is anyway interesting to notice how E and ν stay constant as well. Since the solution of the equations in Comsol does not require any other parameter besides the three of them, it would not actually be necessary to specify whether the alloy is a T651 or an annealed, anyway.

Regarding instead the physics, the applied load can finally be set in the correct way: as imposed by the tie bars. In the cell of chapter 3.b, the load was given by a pressure acting on the BPs' surface, now on the washers, where the tie bars would actually push. Their force can be evaluated as: [63]

$$F = \frac{T}{Kd} [N]$$

Where T is the tightening torque (measured from the laboratory case: 24 [Nm]), d the bolt diameter (10 [mm]), and K the bolt friction factor (0.43): it accounts for the roughness of the surface, for the deformation and the friction that occur during the tightening. This force is inserted in the Comsol file.

Still regarding the physics, symmetries are exploited once again, since they can hugely speed up the evaluation, reducing the object to only a fraction of it. On the mathematical point of view, a symmetry imposes that the considered boundary cannot move along its perpendicular direction, but only along the tangential ones.

The prescribed movement boundary is not too far from this idea: it allows to specify the movement of a component along a given direction. This instruction is to be given with respect to the three spatial directions, and any numerical value can be put, being a common one zero. A prescribed movement of zero along a direction implies that the selected surface cannot move along it: this is applied to the washers. In fact, the tie bars can only go down along the tightening direction and thus stay aligned with the holes that host them, not moving in the tangential directions. Due to symmetry reasons, this would in fact mean that the tie bar is flexing, which is not true. The tie bars were not inserted anyway, but the washers resent of this behaviour of theirs. In fact, neither can the washers move along the horizontal directions, but only following the one of the axis of the stack, i.e. of

the tightening. This also implies that the washers stay aligned with the corresponding holes and in fact cannot slide away from them.

Not only is this command correct considering the movements of the pieces of the stack, but also necessary for a correct evaluation. Were in fact not put, the washers would be able to slide along the surface of the EP. Since they are placed along its perimeter, not at the centre, they cause it to slightly flex, thus resulting in a light curvature. This curvature could cause the washers to slide away from their position, which is not correct and would lead to a divergence of the solution research.

A new setting to be added which was absent in chapter 3.b is the hyperelastic modelling. The deformation the gaskets undergo is in fact of such an extent that they cannot be approximated as normal elastic objects: hyperelastic materials are indeed a special case of the Cauchy elastic material, where the $\sigma - \varepsilon$ relationship derives from the strain energy density function. Some equations change, as reported in Table 3.4

Table 3.4: The different equations that describe an elastic or a hyperelastic material in Comsol.

Elastic material	Hyperelastic material
$0 = \nabla \cdot (FS)^T + \mathbf{F}_v$ $\mathbf{F} = \mathbf{I} + \nabla \mathbf{u}$	$0 = \nabla \cdot (FS)^T + \mathbf{F}_v$ $\mathbf{F} = \mathbf{I} + \nabla \mathbf{u}$
$S = S_{inel} + S_{el}$ $\varepsilon_{el} = \frac{1}{2} (F_{el}^T F_{el} - I)$ $\varepsilon = \frac{1}{2} [(\nabla \mathbf{u})^T + \nabla \mathbf{u} + (\nabla \mathbf{u})^T \nabla \mathbf{u}]$	$S = S_{inel} + \frac{\partial W_s}{\partial \varepsilon}$ $\varepsilon = \frac{1}{2} (F^T F - I)$ $W_s = \frac{1}{2} \mu (I_1 - 3) - \mu \ln(J_{el}) + \frac{1}{2} \lambda (\ln(J_{el}))^2$

The dependency on the strain energy density W_s is shown. The formula that describes it is not universal, but rather changes according to the chosen model: the one of Table 3.4 is given by the Neo-Hookean one. This model is suitable to describe the gasket since it is tailored for materials made of polymers chains, such as plastics and rubbers. The other available models are instead adopted for different situations.

Concerning the mesh, attention is to be paid when describing the EP. The addition of other bolts through the optimisation may in fact change the nomenclature of the domains of the component, and therefore confuse the program about which boundaries are mapped by which commands, and which domains are extruded. The EP height itself changes along each optimisation, so it would be imprecise to divide it into a fixed number of vertical elements. The free tetrahedral description is the most suitable therefore to describe the EP: the number and dimension of its elements will automatically vary according to the size and shape of the piece.

Attention is to be paid when dealing with the mesh of the surfaces that compose the contact pairs too. The ribs of the FF do not in fact initially touch the BP, but only after the compression has been applied, squeezing enough the gasket between the two of them. The lower boundaries of the FF ribs and the upper one of the BP are thus joined in a contact pair, and Comsol suggests that the number of mesh elements describing the destination (the softer material, i.e. the PVC of the FF) is at least twice that of the source (the harder material, i.e. the graphite of the BP). The mesh drawing is shown in Figure 3.55.




Figure 3.55: The mesh at the contact pair between the FF rib and the BP.

The solver is then set. Since the involved equations are linear, a linear solver can be adopted, and the choice selects PARDISO, being it the recommended one for the solid mechanics. The aggregated approach is not necessary, and can be substituted by the segregated approach, which allows for an important acceleration of the evaluation time. There is a new step, with respect to the simulation of chapter 3.b: a lumped one which evaluates the contact pressure between the FF ribs and the BP. It was seen in the physic setting that this force is nil as long as there is a gap between the two selected surfaces, otherwise present.

Load cases are implemented too, to strengthen the evaluation convergence. This means that the study does not directly evaluate the displacements and stresses that the cell would undergo with the full force, but rather for lighter intermediate cases. Many of them were considered: the load is gradually increased with steps of a fraction of its magnitude. A lighter case is easier to evaluate than a heavier one, and this strengthens the convergence. In particular, the solution of the first evaluation is used as starting point for the computation of the second one; while the third evaluation uses the solution of the second one, and so on and so forth. Smaller steps are individually easier to manage than a wider one: on one hand, this means that more configurations are to be evaluated, but on the other hand the single wide step might even not be feasible at all. Beside the use of the solution of a case with a slightly lighter load, the program can anyway be told the scale of the result that it should expect. This is provided by the scale factor, which is by default of the order of 10^8 [Pa] for pressures and 10^{-3} [m] for displacements. It is instead corrected with 10^6 [Pa] and 10^{-4} [m] respectively: closer to the expected values for the solution.

In these conditions, the simulation would not reach convergence, nevertheless: the stress acting on the electrode would be excessive and compress it to such an extent that the hypothesis of small deformations would not be valid anymore. The model therefore cannot be successfully run, but its design shows that the electrode is only pressed by the BP, while its lateral boundaries are free to enlarge as much as they need. Its upper boundary is instead blocked by the symmetry. The crucial detail to remember is nevertheless that the graphite felt has a very low Poisson's number ν , which means that small lateral expansions occur after a transversal one: their magnitude is only a tenth of it at most, and only present along the perimeter. This is shown in Figure 3.56, which represents the case of a slightly lighter load (such that the model can reach convergence).




Figure 3.56: The a) $-x$, b) y and c) z displacements of the electrode for a lighter load.

Figure 3.56 shows that the lateral expansion only occurs along the external boundaries of the electrode, and its components $comp1.u$ and $comp1.v$ (respectively along x and y) can at most reach the hundredth of millimetre circa, while $comp1.w$ (along z , which represents the vertical compression) is of the tenth of millimetre. This means that the vertical displacement $comp1.w$ is dominant everywhere, while the other two can even be neglected in the evaluation of the CR : the vertical compression is way greater than the lateral expansion, which anyway only occurs in a limited portion of the felt.

This all means that $comp1.w$ is representative of the compression, while $comp1.u$ and $comp1.v$ can be ignored. In particular, it is known that $comp1.w$ must be nil over the whole top electrode surface (as also depicted in Figure 3.56 c)), since the symmetry boundary imposed that no normal displacement (i.e. along z) was allowed, while coincides with the $comp1.w$ of the BP over the whole bottom surface, since the two objects were bounded with an identity pair.

Not only is $comp1.w$ known over the whole bottom and top surfaces of the electrode, but even in any point in between. Being the object described as linear, homogeneous and isotropic (which is a forgivable simplification with respect to the inhomogeneity a carbon strain structure would present in reality), the relations that govern it are linear too, as it is in fact the whole solid mechanics' physic. This implies that $comp1.w$ linearly varies along any arc that orthogonally connects the bottom surface (shared with the BP) with the top surface (of symmetry). This means that the knowledge of the $comp1.w$ values along the BP surface is enough to fully describe the electrode compression too. Not only is the electrode facultative in the analysis of the results, but even (surprisingly) during their computation too. In fact, the elasticity modulus E of the electrode felt is so weak with respect to the others that it does not influence the other components at all, and neither does its position in the stack, unlike the gaskets, which were soft too, but anyway important since placed between other components. This all justifies the simplification of the model to simulate: the electrode can be neglected, and the top BP surface studied instead. The hypothesis of small deformations is now valid, so the simulation can reach convergence. Once it has run, the results are finally available.

Many variables can be observed: the displacements along the spatial directions, the pressures etc. A detail that might be necessary to check concerns the displacement of the top surface of the FF. Figure 3.57 recalls in fact the way the symmetry boundary was applied to it.




Figure 3.57: The symmetry boundary over the top surface of the component (without the electrode).

The external rib of the FF does not in fact reach the same height as the gasket, so it can displace upwards. The matter is that this displacement cannot overcome the height of the gasket, since it would not make sense considering the whole stack object: a compenetration would occur with the mirrored part. This is the reason why it is first necessary to analyse the vertical displacement of the FF, which was reported in Figure 3.58.




Figure 3.58: The vertical displacement of the top surface of the FF.

As it can be seen, the FF body displaces of at most of 0.32 [mm], which is an acceptable value, since the available space before the gasket is 0.4 [mm] tall. The external rib of the FF uses too a good part of the available gap, since it displaces of ~ 0.25 [mm], having the highest value at the corner, and a limit of 0.3 [mm] before the symmetry plane.

The simulation results are shown in Figure 3.59, which reports comp1.w over the surface shared by the electrode and the BP.




Figure 3.59: The simulation results: comp1.w over the contact surface between the electrode and the BP.

A detailed description of the behaviour of comp1.w over the surface is given later, in chapter 4.a, in comparison with the results corresponding to other EP thicknesses.

Besides the main objective, it was also mentioned that the hydraulic sealing has to be provided too. In order to achieve this necessary feature, the mechanical pressure between the components that bind the electrolyte shall be superior to the hydraulic pressure of the electrolyte itself. This means that this contact is to be provided between the BP and the bottom rib of the FF, between the top rib of the FF and the upper gasket, and between this gasket and the symmetry plane. All these couples are joined by a surface, but it is not necessary that all the points of such surfaces respect this boundary concerning the pressure, but rather only those facing the electrode, as in Figure 3.60. This means that the pressure is to be evaluated along the inner lines alone of the contact surfaces, in particular if the minima of each line are higher than the hydraulic pressure, than the sealing is achieved.




Figure 3.60: The lines along which the mechanical pressure is to be checked to ensure the hydraulic sealing is achieved: a) and b) between the BP and the FF rib, c) and d) between the FF rib and the gasket, and between the gasket and the symmetry plane.

The hydraulic simulation provided the electrolyte pressures: in the case of the simple cell, it was at most around $35 \div 40$ [kPa], while in the complex cell case ~ 350 [kPa] (as shown back in Figure 3.43, 3.48 and 3.51). It is possible to check the minimum pressure values along the lines of Figure 3.60 c) and d): it is found that between the FF rib and the gasket it is 58.667 [MPa], while between the gasket and the symmetry plane 32.609 [MPa]. Since the bottom rib of the FF was coupled with the BP with a contact pair, the variable contact force is available too. Its minimum in the lines of Figure 3.60 a) and b) is 11.190 [MPa]. Since all these pressures and contact forces are greater than the highest hydraulic pressure, the sealing is achieved.

The other requirement that had to be met was about the mechanical resistance, i.e. the yield strength of any material of the stack should not be reached. A first idea might involve the check of the maximum pressure in any component of the stack, but this would not be totally accurate. In fact, ductile materials present a greater resistance when subjected to a non-uniaxial stress, and the Von Mises criterium specifically deals with this situation. Being the stress tensor σ symmetric, its components can be named as

$$\sigma = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{12} & \sigma_{22} & \sigma_{23} \\ \sigma_{13} & \sigma_{23} & \sigma_{33} \end{bmatrix}$$

And through them it is possible to arrive to the Von Mises stress σ_v :

$$\sigma_v = \sqrt{\frac{(\sigma_{11} - \sigma_{22})^2 + (\sigma_{22} - \sigma_{33})^2 + (\sigma_{33} - \sigma_{11})^2 + 6(\sigma_{13}^2 + \sigma_{23}^2 + \sigma_{12}^2)}{2}}$$

The Von Mises criterium states that the material resists as long as σ_v is lower than the yield strength. [61] Comsol can evaluate σ_v , so its minimum value in any component can be checked, as reported in Table 3.5.

Table 3.5: The maximum Von Mises stress in some of the stack components and their yield strengths.

COMPONENT	MATERIAL	YIELD STRENGTH [MPa]	MAXIMUM σ_v [MPa]	SECURITY COEFFICIENT [1]	SOURCE
EP	Al6061-T651	276	187.90	1.469	[43,62]
CC	Copper	543	125.08	4.341	[63]
Washers	Steel AISI 4340	862	501.97	1.717	[64]

A quick glance at Table 3.5 is sufficient to notice that it only contains some of the components of the stack: the EP, the CC and the washers. This happens because the three of them are in metal (aluminium, copper and steel respectively), so they distinguish an elastic phase from a plastic one, thus featuring a meaningful value for their yield strength. The ratio between this number and σ_v provides the security coefficient: the higher than the unit, the better. Were it equal to one, instead, it would mean that the piece is withstanding the maximum possible stress: a serious hazard. Were it instead lower than one, the limit would have already been overcome. The three coefficients are all anyway superior to one, so the EP, the CC and the washers can all withstand the stress.

The BP is instead a graphite plate, a rigid material that does not properly feature a plastic phase, as it was already seen in Figure 2.36. Its limit is therefore the ultimate stress under compression, that is anyway of the magnitude of 100 [GPa]: greatly sufficient for its $\sigma_v = 29.6$ [MPa]. [47]

The isolator, the gaskets and the FF are in PVC and EPDM, whose ultimate resistances are usually found in bibliography spanning large intervals. [65,66] Alike the graphite plate, in fact, they do not present a yield strength: not because they are rigid, but rather due to their polymeric nature, which allows large deformations and does not distinguish a plastic phase from an elastic one. [44,46] The precise ultimate strengths to be considered should be provided from the builder of the piece, anyway. Nevertheless, this maximum value of σ_v is not the most important parameter to consider in this context, due to the presence of the ribs: these narrow components usually concentrate in themselves high stresses, which therefore might appear superior to the material resistance, but the object might be able to withstand them anyway. It happens in fact that the corners of the ribs – which are the point where the highest stresses always gather – can spontaneously improve their own mechanical resistance, due to the behaviour of the micro cracks. This all means that the precise limits of PVC and EPDM are not trivial to find, neither anyway so crucial in determining whether the pieces can or not sustain the stress. The best and easiest way to therefore check the mechanic resistance is to perform a laboratory test. The

VFB of the University of Padua adopts the structure, materials and loads reported in this simulation: since the FFs, the gaskets and the isolators all resist, the condition is considered to be satisfied.

The graphite felt is for sure far from its limit. In fact, the compression first squeezes the pores inside of it, and no sooner are they all shut than the carbon strains deform too. This means that as long as the porosity ε is higher than zero (with some safety margin), the carbon strains are not compressed yet. The permeability κ depends on ε , and the flow of the electrolyte through the felt witnesses that $\kappa \neq 0$. This means that neither is ε equal to 0, so the electrode felt is far from its mechanic limit.

4. THE OPTIMISATION

Chapter 3 showed the results of the mechanical simulation: the vertical displacement of the bottom surface of the electrode is not perfectly homogeneous. The hydraulic simulations proved that the whole electrode felt is flooded by the electrolyte, so the displacement inhomogeneity shall be limited in order to better exploit the AA of the carbon strains.

This requirement can find in the gap between $comp1.w_{max}$ and $comp1.w_{min}$ the right tool for its optimisation, together with $comp1.w_{ave}$. For this purpose, several simulations are run, changing the EP thickness and bolt configuration and then recording the corresponding displacement gap.

a. The EP thickness

As it was first said when introducing the numerical simulations, it is good practise to parametrise the used quantities, such as the widths, lengths and heights of the components. This allows for an easy modification of them: the new value can be simply inserted at the very beginning of the program, and without further operation the whole object is correctly scaled.

The optimisation is a particular kind of study, where a parameter at choice is varied: it assumes the values inserted in a vector at desire. In particular, the quantity here is the EP height: EP_z, which will be equal to 5 [mm], 10 [mm], 15 [mm], 20 [mm], 25 [mm] and 30 [mm]. Table 4.1 reports the corresponding results.

Table 4.1: The minimum, average and maximum vertical displacement for various EP height, and 8 bolts.

EP_z [mm]	$comp1.w_{min}$ [mm]	$comp1.w_{ave}$ [mm]	$comp1.w_{max}$ [mm]
5	0.3480	0.3665	0.3950
10	0.3850	0.3975	0.4150
15	0.4010	0.4091	0.4180
20	0.4090	0.4148	0.4190
25	0.4130	0.4176	0.4200
30	0.4150	0.4190	0.4210

The values of Table 4.1 are immediately graphed in Figure 4.1.




Figure 4.1: The vertical displacement over the electrode base for various EP thicknesses.

The load is constant throughout the optimisation, but the vertical displacement is slowly rising towards an asymptotic value around 0.42 [mm]. This is due to the fact that thicker EPs bend less under the applied pressure, thus leading to a more uniform vertical displacement of the whole stack structure above, BP included. Due to its central position, a minor bending leads to a greater vertical movement of the BP surface below the electrode. This explains the reason why the average $comp1.w_{ave}$ rises. The fact that the interval $[comp1.w_{min}; comp1.w_{max}]$ gradually narrows around these averages is also due to the reduction of the bending of the EPs and of the other components subsequently. It is interesting to apply a hyperbolic fit, as done in Figure 4.2.




Figure 4.2: The hyperbolic fit of the vertical displacements.

The hyperbolic fit describes the behaviour of the data with a satisfying precision, and can be fairly useful in estimating the impact of any intermediate EP thickness. It also reports that the average vertical displacement $comp1.w_{ave}$ will tend to 0.42254 [mm]. Figure 4.3 shows the $comp1.w$ distributions through the surface for all the EP thicknesses evaluated.

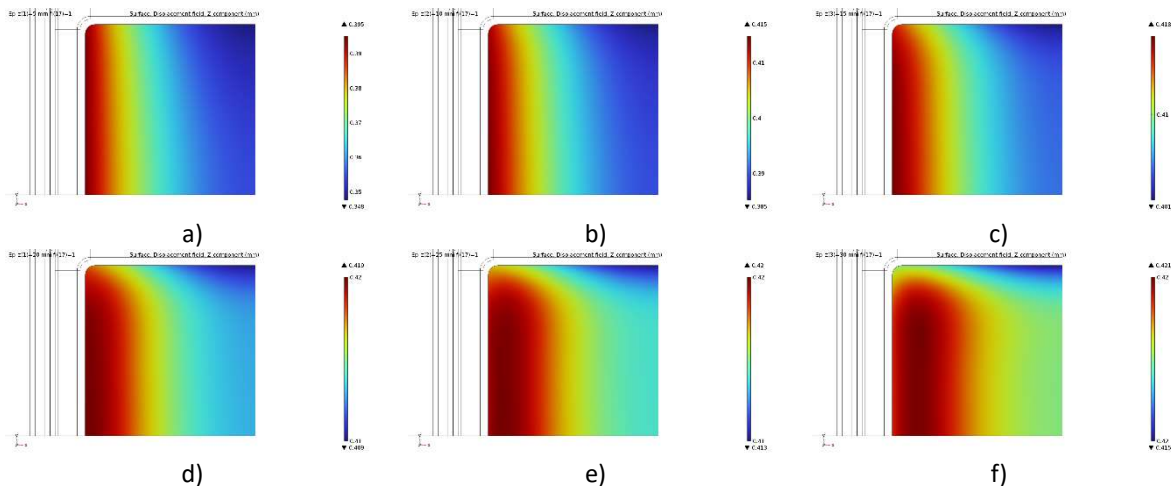


Figure 4.3: The vertical displacement over the bottom electrode surface for various EP thicknesses: a) 5 [mm], b) 10 [mm], c) 15 [mm], d) 20 [mm], e) 25 [mm], f) 30 [mm].

When looking at Figure 4.3 a-f), it is crucial to remember that the values represented by the colours are different from one image to another, since the gap between the maximum and the minimum is gradually closing. A comparison between the a) first and f) last case is particularly interesting. On one hand, the red colour invades a

way greater portion of the surface, while the blue is pushed to the right upper corner. This corner represents the middle of the edge which is not facing the bolts, and in fact that portion is the least touched by their effect. The right edge turned to green, since it eventually perceives the effect of the compression from the face which is specular to the left one. The top corner does not turn red like the left edge since the EP tends to bend. This effect is more marked at the top of the active area because there we have more material that opposes to the compression.

b. The number of bolts

Since the load propagates from the bolts, their position and number are crucial to obtain a homogeneous electrode compression. It is generally preferable to avoid putting them at the corners of the EPs, since this would mostly stress the vertical edges of the cells, but rather along the sides. Symmetry is usually preferred so, if two or more bolts are present, they should be equally spaced from the plane of symmetry, in our case the plane x-z. Figure 4.4 shows the stack for various numbers of bolts.




Figure 4.4: Image of the cell with a) 4 bolts, b) 6 bolts and c) 8 bolts.

At this point, the optimisation explained in paragraph 4.a just has to be run several times, one for every bolt number adopted. The values to record are always the same, and now shown in Table 4.2.

Table 4.2: The optimisation results for various EP thicknesses and bolts numbers: from top to bottom for each cell the minimum, average and maximum vertical displacement, in [mm].

	4 bolts	6 bolts	8 bolts
5 [mm]	0.331	0.340	0.348
	0.361	0.365	0.367
	0.400	0.399	0.395
10 [mm]	0.379	0.382	0.385
	0.398	0.398	0.398
	0.422	0.419	0.415
15 [mm]	0.398	0.400	0.401
	0.410	0.410	0.409
	0.422	0.421	0.418
20 [mm]	0.408	0.408	0.409
	0.415	0.415	0.415
	0.422	0.421	0.419
25 [mm]	0.412	0.413	0.413
	0.418	0.418	0.418
	0.422	0.421	0.420
30 [mm]	0.415	0.415	0.415
	0.419	0.419	0.419
	0.423	0.421	0.421

Since every bolt number configuration was studied for the same EP_z interval of paragraph 4.a, Table 4.2 can be considered as an extension of Table 4.1. Figure 4.5 reports all the corresponding data, and is therefore an extension of Figure 4.2.




Figure 4.5: The simulation results of Table 4.2. The lines represent different bolts number configurations.

A lot of information can be collected from Figure 4.5, since it both reports the values of Table 4.2 and also provides their hyperbolic fitting. The asterisks and the continuous vertical lines represent the values found with the optimisation: the blue colour corresponds to the configuration with 8 bolts, green with 6 and red with 4. The asterisks in particular represent the $comp1.w_{ave}$ for any EP_z and bolts number, and tend to overlap as soon as EP_z grows from its minimum of 5 [mm]: ~ 8 [mm] are already indicatively enough to make them undistinguishable. This means that, whatever the EP thickness, the bolts number does not make a difference concerning the average displacement. The error bars then connect the dots of $comp1.w_{max}$ to $comp1.w_{min}$: this gap tends to reduce when increasing the thickness of the EP. It is not a surprise that the higher the bolts number, the narrower the gap is, for any EP_z: in fact, the tie bars help achieving a more homogeneous load distribution in the stack structure. This also corresponds to the fact that a higher number of bolts reduces the maximum displacements and increases the minimum ones for any EP thickness.

Figure 4.6 reports the $comp1.w$ distribution along the contact surface between the electrode and the EP for the three analysed configurations.




Figure 4.6: The $comp1.w$ evaluated for 5 [mm] with a) 4 bolts, b) 6 bolts and c) 8 bolts.

It is immediate to notice the varying colour gradient, in particular red uniformly grows along the vertical direction when increasing the number of bolts. In fact, more tie bars help compressing the electrode along all its surface, not focusing their load in a single spot as in Figure 4.6 a). Figure 4.5 anyway witnessed that this gap is the widest for the thinnest EP: the thickest value of 30 [mm] saw instead almost no difference whatever the number of bolts. All the fittings provided the coefficients for the hyperbolic function:

$$f(x) = a + \frac{b}{c + x}$$

It is fascinating how such a simple formula can be used to interpolate the behaviour of such a complex model, and it can be expected to correctly predict any intermediate value, thanks to the linearity of the problem. For example, it can be considered that an EP 12 [mm] thick with 6 bolts will achieve $comp1.w_{min} = 0.3909$ [mm], $comp1.w_{ave} = 0.4041$ [mm] and $comp1.w_{max} = 0.4199$ [mm], which is reasonable looking at Figure 4.5.

That graph might even suggest that the values of the 6 bolts configuration are the average between those of 4 and 8 bolts, as if some linearity existed concerning their number too. It is quite easy to check this: it is sufficient to evaluate the average between the 4 and 8 bolts configuration, then its difference with the 6 bolts one. The results are reported in Table 4.3, expressed as percentages.

Table 4.3: The difference between the values of the 6 bolts configuration and the average between the 4 and 8 ones, in percentage.

	$comp1.w_{min}$ [%]	$comp1.w_{ave}$ [%]	$comp1.w_{max}$ [%]
5 [mm]	-0.050	-0.124	0.050
10 [mm]	0	-0.045	-0.100
15 [mm]	-0.050	-0.029	-0.100
20 [mm]	0.050	-0.002	-0.050
25 [mm]	-0.050	0.002	0
30 [mm]	0	0.001	0.100

Some values are approximated to be nil, and all the others are anyway very small. This is not sufficient though to prove the hypothesis that the displacements are linear with the bolts number. In fact, their effect on the structure of the stack is quite complicated, and also depends on their position. For example, the disposition of four tie bars close to the x-z symmetry plane, gives results that are likely to be very different from those obtained positioning them at the corners of the stack. Also, many more simulations would have to be run before defining a formula that can bind displacements to bolts number, also considering that (as just said) their position in the x-y plane surely plays a crucial role too.

It is reasonable to assume that a high number of tie bars helps homogenising the pressure load along the electrode surface, but this quantity cannot anyway indefinitely grow: they would, in fact, of course end up occupying the whole available space on the EP. Anyway, even if it was possible to place many of them, it would not necessarily improve the homogeneity in a substantial way. Figure 4.4 showed in fact that the thickest EPs (i.e. above 25 [mm] indicatively) are not particularly affected by the bolts number: only the fitted lines of the maxima do not overlap yet, and anyway their distance is very small.

c. Aluminium and steel

It is quite easy to perform the study for other EP materials: it is just necessary to change the material input in the Comsol settings. The previous studies concerned Aluminium 6061-T651, while now it is the turn of the steel AISI

4340. Table 4.4 reports their different parameters (similarly to Table 3.5), which have to be updated in the program.

Table 4.4: The different parameters between Al 6061-T651 and AISI 4340.

	Al 6061-T651 [43,62]	AISI 4340 [64]
Young's modulus E [GPa]	68.9	200
Poisson's modulus ν [1]	0.33	0.29
Density ρ $\left[\frac{kg}{m^3}\right]$	2700	7850
Yield strength σ [MPa]	276	862

Due to superior steel rigidity, it is expected that steel will bend less than aluminium for the same EP geometry and configuration. It is also expected that its gaps between minimum and maximum displacements will always be narrower than those of aluminium, as a consequence, thus performing better.

Running once again all the studies, it is possible to draw the correspondent of Figure 4.5 for steel, which is reported as Figure 4.7 a). For the sake of convenience, that graph is sided by Figure 4.5 again, here Figure 4.7 b).




Figure 4.7: The studies' results and their fittings for EPs in a) steel and b) aluminium.

The figures are very similar, and in both the cases the average displacements overlap for any bolts number, while the increase of this parameter narrows the gap between the minima and the maxima, for any EP thickness. A closer look at the numerical values on the vertical axes proves that minimum displacements obtained with steel always perform better than their aluminium counterparts. No substantial difference can be appreciated concerning the maxima instead, as long as the thickness is at least 10 [mm].

The results of Figure 4.7 can be reorganised as in Figure 4.8, where aluminium and steel are compared concerning one bolts number per time.

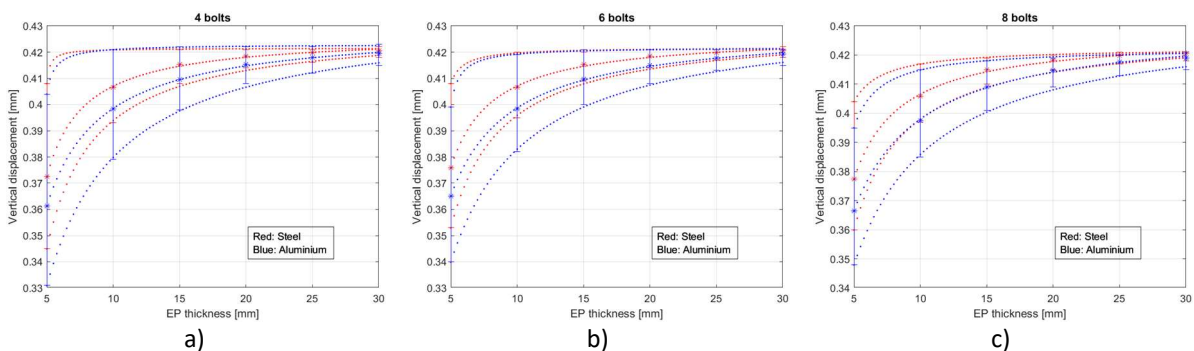


Figure 4.8: The comparison between the results for steel (in red) and aluminium (in blue) for the various bolts number: a) four, b) six, c) eight.

Figure 4.8 further proves how steel performs better than aluminium. The minimum lines of steel in fact are always very close the average lines of aluminium, while the maximum lines are not too distanced. This means that the gaps of aluminium are way larger than those of steel. A hypothesis may arise concerning a relationship between the $[comp1. w_{min} \div comp1. w_{max}]$ gap widths found for steel and their counterparts for aluminium for the same EP thickness and bolts number, as if they might be bound by some value which depends on their respective Young's moduli. A quick evaluation of these gaps shows that their ratios are not stable. Even if some kind of dependence existed, anyway, two materials alone would be way too few to discover it.

To conclude, steel appears to be superior with respect to aluminium, since its greater rigidity reduces the bending of the EPs and consequently of the BP. Stricter tolerances can be achieved with equal thicknesses.

d. The best choice

The graphs of Figure 4.7 can be used as support when it is necessary to choose the EP geometry and configuration, if all the other inputs of the problem are fixed. It is obvious in fact that the load alone would play a crucial role if it could be varied as desired, as well as the gasket shore hardness.

Supposing anyway that three degrees of freedom in total are available (i.e. the EP thickness, bolts number and material), it is possible to follow these steps to design the EPs.

- As starting point, a hydraulic configuration is given. Case A is a flow-through, therefore a carbon felt with an uncompressed thickness of 4.6 [mm] is involved, while Case B is a flow-by carbon paper electrode, with a 0.5 [mm] uncompressed thickness. Both of them have to be compressed of 0.4 [mm], which corresponds to $comp1. w_{ave}$. in this way their respective thicknesses become 4.2 [mm] and 0.3 [mm] (thanks to the BP design).
- With Figure 4.7, it is found that $comp1. w_{ave} = 0.4$ [mm] corresponds to an EP thickness of 10.5 [mm] if aluminium is used, 8 [mm] if steel instead.
- A quality is also required: the maximum error from $comp1. w_{ave}$ that can be accepted is 210 [μm] for the carbon felt and 15 [μm] for the carbon paper, corresponding to 5 % of the desired compressed thickness. This means that $comp1. w_{max} < 610$ [μm] and $comp1. w_{min} > 290$ [μm] for the carbon felt, while $comp1. w_{max} < 415$ [μm] and $comp1. w_{min} > 385$ [μm] for the carbon paper.
- Proceeding along a vertical line from the EP thickness of 10.5 [mm] or 8 [mm] in Figure 4.7 it is found that the conditions of the previous point can be satisfied by any bolts number for the carbon felt, while the carbon paper strictly needs 8 bolts if steel is used, but no solution exists if aluminium is chosen.

This example can be collected for the sake of clarity as in Table 4.5.

Table 4.5: An example of use of the results of Figure 4.7.

	A: carbon felt (4.6 [mm])		B: carbon paper (0.5 [mm])	
EP material	Al 6061-T651	AISI 4340	Al 6061-T651	AISI 4340
$comp1. w_{ave}$	400 [μm]	400 [μm]	400 [μm]	400 [μm]
EP_z	10.5 [mm]	8 [mm]	10.5 [mm]	8 [mm]
$comp1. w_{max}$	< 610 [μm]	< 610 [μm]	< 415 [μm]	< 415 [μm]
$comp1. w_{min}$	> 290 [μm]	> 290 [μm]	> 385 [μm]	> 385 [μm]
Bolts number	any	any	none	8

Figure 4.9 reports the use of Figure 4.7, concerning the EP choice for the flow-by design, featuring steel as the material.




Figure 4.9: The application of Figure 4.7 for the choice of EP_z and number of bolts using steel in the flow-by design.

The requested average displacement defines the EP_z , thus identifying the horizontal position of the vertical black line. The quality decides instead the height of the two horizontal marks of that line, which represent the required displacements: the configuration results must lay within their interval. Figure 4.9 shows how 6 or 8 bolts would satisfy this condition concerning the minimum displacement, but 8 bolts are strictly necessary regarding the maximum one, instead.

5. CONCLUSIONS

This thesis examines the issue of electrode compression in Vanadium Redox Flow Batteries (VFBs), an important phenomenon that influences ohmic losses due to contact resistances between the carbon strands and the adjacent bipolar plates (BPs), along with other parameters that define the energy storage system. A brief digression also addresses the hydraulic sealing of the stack and the mechanical integrity of its components, as these are similarly affected by compression.

This comprehensive and critical topic was investigated using solid mechanics simulations, conducted in COMSOL Multiphysics®. A simplified model of a laboratory cell, inspired by that of the Electrochemical Energy Storage and Conversion Laboratory at the University of Padua, was introduced for this purpose. The numerical results from these simulations highlighted the significant role of the end plates' geometry and configuration, not only in terms of electrode compression but also in determining its quality. In this context, finite element analysis once again proved to be a valuable tool for engineering applications.

Specifically, the study demonstrated that the thickness and material of the end plates are crucial parameters for achieving a desired compression ratio, while the number of bolts affects the uniformity of this compression.

Future work will first aim to refine the model presented here by incorporating experimentally obtained parameters. Subsequently, the corresponding numerical results will be compared with experimental simulations. A second phase will involve scaling up the model by increasing the number of cells in the stack and adjusting its overall dimensions, allowing the simulation to reflect the behaviour of an industrial-scale system. Eventually, it would be interesting to explore alternative end plate designs, including new bolt configurations, materials, and solutions such as welded frames or multi-layer structures.

6. BIBLIOGRAPHY

- [1] Link: <https://dati.terna.it/>, last view: 30/10/2024.
- [2] Benato, R. (2023). *La trasmissione dell'energia nelle reti elettriche di potenza*.
- [3] Guarnieri, Massimo, et al. "Developing vanadium redox flow technology on a 9-kW 26-kWh industrial scale test facility: Design review and early experiments." *Applied Energy* 230 (2018): 1425-1434.
- [4] Trovo, Andrea, et al. "Prospects for industrial vanadium flow batteries." *Ceramics International* 49.14 (2023): 24487-24498.
- [5] Blume, Nick, et al. "Prospective Life Cycle Assessment of Chemical Electrolyte Recycling for Vanadium Flow Batteries: A Comprehensive Study." *Energy Technology* 12.1 (2024): 2300750.
- [6] Sun, Jiawei, et al. "Investigations on the self-discharge process in vanadium flow battery." *Journal of Power Sources* 294 (2015): 562-568.
- [7] Lourenssen, Kyle, et al. "Vanadium redox flow batteries: A comprehensive review." *Journal of Energy Storage* 25 (2019): 100844.
- [8] Minke, Christine, and Miguel A. Dorantes Ledesma. "Impact of cell design and maintenance strategy on life cycle costs of vanadium redox flow batteries." *Journal of Energy Storage* 21 (2019): 571-580.
- [9] Petranikova, M., et al. "Vanadium sustainability in the context of innovative recycling and sourcing development." *Waste Management* 113 (2020): 521-544.
- [10] Commission, E. et al., 2017. *Study on the review of the list of critical raw materials : final report.*, Publications Office of the European Union. Belgium. Retrieved from <https://coilink.org/20.500.12592/wdfx0s> on 30 Oct 2024. [COI: 20.500.12592/wdfx0s](https://coilink.org/20.500.12592/wdfx0s).
- [11] Barella, S., Gruttadauria, A. (2017). *Metallurgia e materiali non metallici teoria ed esercizi svolti*. Bologna, IT: Società editrice Esculapio.
- [12] Link: [https://chem.libretexts.org/Courses/Saint_Marys_College_Notre_Dame_IN/Chem_122L%3A_Principles_of_Chemistry_II_Laboratory_\(Under_Construction_\)/08%3A_Vanadium_Rainbow#:~:text=Vanadium%20is%20one%20such%20metal,%20and%20yellow%20\(V\)](https://chem.libretexts.org/Courses/Saint_Marys_College_Notre_Dame_IN/Chem_122L%3A_Principles_of_Chemistry_II_Laboratory_(Under_Construction_)/08%3A_Vanadium_Rainbow#:~:text=Vanadium%20is%20one%20such%20metal,%20and%20yellow%20(V)), last view: 30/09/2024.
- [13] Choe, Jaeheon, and Jun Woo Lim. "Carbon-composite bipolar plate-integrated current collector for vanadium redox flow battery." *Journal of Power Sources* 589 (2024): 233751.
- [14] Liu, Tao, et al. "Progress on the electrode materials towards vanadium flow batteries (VFBs) with improved power density." *Journal of energy chemistry* 27.5 (2018): 1292-1303.
- [15] Trovò, Andrea, Francesco Picano, and Massimo Guarnieri. "Comparison of energy losses in a 9 kW vanadium redox flow battery." *Journal of Power Sources* 440 (2019): 227144.
- [16] Li, Tianyu, et al. "Cost, performance prediction and optimization of a vanadium flow battery by machine-learning." *Energy & Environmental Science* 13.11 (2020): 4353-4361.
- [17] Dinesh, Anarghya, et al. "Iron-based flow batteries to store renewable energies." *Environmental Chemistry Letters* 16 (2018): 683-694.
- [18] Petek, Tyler J., et al. "Slurry electrodes for iron plating in an all-iron flow battery." *Journal of Power Sources* 294 (2015): 620-626.
- [19] Hawthorne, Krista L., Jesse S. Wainright, and Robert F. Savinell. "Studies of iron-ligand complexes for an all-iron flow battery application." *Journal of The Electrochemical Society* 161.10 (2014): A1662.
- [20] Luo, Jian, et al. "Status and prospects of organic redox flow batteries toward sustainable energy storage." *ACS Energy Letters* 4.9 (2019): 2220-2240.
- [21] Winsberg, Jan, et al. "Redox-flow batteries: from metals to organic redox-active materials." *Angewandte Chemie International Edition* 56.3 (2017): 686-711.
- [22] Li, Zening, et al. "Recent progress in organic species for redox flow batteries." *Energy Storage Materials* 50 (2022): 105-138.
- [23] Park, Se-Kook, et al. "The influence of compressed carbon felt electrodes on the performance of a vanadium redox flow battery." *Electrochimica Acta* 116 (2014): 447-452.
- [24] Emmel, Dominik, et al. "Understanding the impact of compression on the active area of carbon felt electrodes for redox flow batteries." *ACS Applied Energy Materials* 3.5 (2020): 4384-4393.

- [25] Charvát, Jiří, et al. "Performance enhancement of vanadium redox flow battery by optimized electrode compression and operational conditions." *Journal of Energy Storage* 30 (2020): 101468.
- [26] Bromberger, Kolja, Johannes Kaunert, and Tom Smolinka. "A model for all-vanadium redox flow batteries: introducing electrode-compression effects on voltage losses and hydraulics." *Energy Technology* 2.1 (2014): 64-76.
- [27] Banerjee, Rupak, et al. "Carbon felt electrodes for redox flow battery: Impact of compression on transport properties." *Journal of Energy Storage* 26 (2019): 100997.
- [28] Zheng, Qiong, et al. "Flow field design and optimization based on the mass transport polarization regulation in a flow-through type vanadium flow battery." *Journal of Power Sources* 324 (2016): 402-411.
- [29] Wang, Q., et al. "Experimental study on the performance of a vanadium redox flow battery with non-uniformly compressed carbon felt electrode." *Applied Energy* 213 (2018): 293-305.
- [30] Zhang, Kaiyue, et al. "In-situ measurement of electrode kinetics in porous electrode for vanadium flow batteries using symmetrical cell design." *Applied Energy* 272 (2020): 115093.
- [31] Liu, Tao, et al. "Investigation on the effect of catalyst on the electrochemical performance of carbon felt and graphite felt for vanadium flow batteries." *Journal of Power Sources* 286 (2015): 73-81.
- [32] Lu, Wenjing, Xianfeng Li, and Huamin Zhang. "The next generation vanadium flow batteries with high power density—a perspective." *Physical Chemistry Chemical Physics* 20.1 (2018): 23-35.
- [33] Tang, Ao, et al. "Investigation of the effect of shunt current on battery efficiency and stack temperature in vanadium redox flow battery." *Journal of Power Sources* 242 (2013): 349-356.
- [34] Trovò, Andrea, et al. "Standby thermal model of a vanadium redox flow battery stack with crossover and shunt-current effects." *Applied energy* 240 (2019): 893-906.
- [35] Tang, Ao, Jie Bao, and Maria Sklyllas-Kazacos. "Thermal modelling of battery configuration and self-discharge reactions in vanadium redox flow battery." *Journal of Power Sources* 216 (2012): 489-501.
- [36] Oh, Kyeongmin, Seongyeon Won, and Hyunchul Ju. "A comparative study of species migration and diffusion mechanisms in all-vanadium redox flow batteries." *Electrochimica Acta* 181 (2015): 238-247.
- [37] Zhang, Cheng, et al. "Effects of operating temperature on the performance of vanadium redox flow batteries." *Applied Energy* 155 (2015): 349-353.
- [38] Tang, Ao, Jie Bao, and Maria Sklyllas-Kazacos. "Studies on pressure losses and flow rate optimization in vanadium redox flow battery." *Journal of power sources* 248 (2014): 154-162.
- [39] Guarnieri, Massimo, Andrea Trovo, and Francesco Picano. "Enhancing the efficiency of kW-class vanadium redox flow batteries by flow factor modulation: An experimental method." *Applied Energy* 262 (2020): 114532.
- [40] Gundlapalli, Raveendra, and Sreenivas Jayanti. "Effect of electrode compression and operating parameters on the performance of large vanadium redox flow battery cells." *Journal of Power Sources* 427 (2019): 231-242.
- [41] Oh, Kyeongmin, Seongyeon Won, and Hyunchul Ju. "Numerical study of the effects of carbon felt electrode compression in all-vanadium redox flow batteries." *Electrochimica acta* 181 (2015): 13-23.
- [42] Wang, Q., et al. "Numerical study on vanadium redox flow battery performance with non-uniformly compressed electrode and serpentine flow field." *Applied Energy* 220 (2018): 106-116.
- [43] Link: https://www.researchgate.net/figure/Stress-Strain-curves-of-uniaxial-tensile-testing-of-Al-6061-T6-at-different-temperatures_fig1_311508444, last view: 31/10/2024.
- [44] Link: https://www.researchgate.net/figure/Engineering-stress-strain-curves-of-the-EPDM-rubber-at-various-dynamic-strain-rates_fig4_228995770, last view: 31/10/2024
- [45] Tenny, Kevin M., et al. "A comparative study of compressive effects on the morphology and performance of carbon paper and cloth electrodes in redox flow batteries." *Energy Technology* 10.8 (2022): 2101162.
- [46] Link: https://www.researchgate.net/figure/Stress-strain-curves-of-PVC-obtained-from-tensile-and-compression-tests_fig2_233070610, last view: 22/11/2024.
- [47] Link: https://www.researchgate.net/figure/Stress-strain-curve-of-graphene-sheet_fig1_328555916, last view: 31/10/2024.

- [48] Link: https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/comsol_ref_geometry.14.016.html, last view: 31/10/2024.
- [49] Skyllas-Kazacos, Maria, C. Menictas, and M. Kazacos. "Thermal stability of concentrated V (V) electrolytes in the vanadium redox cell." *Journal of the Electrochemical Society* 143.4 (1996): L86.
- [50] Zhang, Jianlu, et al. "Effects of additives on the stability of electrolytes for all-vanadium redox flow batteries." *Journal of Applied Electrochemistry* 41 (2011): 1215-1221.
- [51] Link: https://www.researchgate.net/figure/Chemical-structure-of-EPDM-monomer_fig1_291523288, last view: 31/10/2024.
- [52] Data sheet delle proprietà Meccaniche dell'EPDM
- [53] Link: <https://greenpvc.com.vn/blog/differences-between-flexible-and-rigid-pvc-compounds/>, last view: 22/11/2024.
- [54] Link: https://doc.comsol.com/5.5/doc/com.comsol.help.sme/sme Ug_solid.07.06.html, last view: 31/10/2024.
- [55] Link: <https://www.comsol.com/multiphysics/analysis-of-deformation>, last view: 31/10/2024.
- [56] Link: <https://www.comsol.com/multiphysics/stress-and-equations-of-motion>, last view: 31/10/2024.
- [57] Link: https://doc.comsol.com/5.5/doc/com.comsol.help.sme/sme Ug_theory.06.23.html, last view: 31/10/2024.
- [58] Link: https://doc.comsol.com/5.5/doc/com.comsol.help.sme/sme Ug_solid.07.85.html, last view: 31/10/2024.
- [59] Link: <https://www.comsol.com/blogs/solutions-linear-systems-equations-direct-iterative-solvers>, last view: 01/11/2024.
- [60] Link: https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/comsol_ref_solver.27.118.html, last view: 01/11/2024.
- [61] Link: <https://www.simscale.com/docs/simwiki/fea-finite-element-analysis/what-is-von-mises-stress/#:~:text=Von%20Mises%20stress%20is%20a,ductile%20materials%2C%20such%20as%20metals,>, last view: 19/11/2024.
- [62] Link: <https://www.matweb.com/search/DataSheet.aspx?MatGUID=b8d536e0b9b54bd7b69e4124d8f1d20a&ckck=1>, last view: 11/11/2024.
- [63] Link: <https://www.matweb.com/search/datasheet.aspx?matguid=e5272c3abef6485a87a391370eab18ca&ckck=1>, last view: 19/11/2024.
- [64] Link: <https://www.matweb.com/search/DataSheet.aspx?MatGUID=fd2df45bffa54018b54989bc14092d9f>, last view: 19/11/2024.
- [65] Link: https://www.matweb.com/search/datasheet_print.aspx?matguid=69642362cb864d25b8f6eb9d02092ecf, last view: 22/11/2024.
- [66] Link: <https://matweb.com/search/DataSheet.aspx?MatGUID=64dd6091f7f6469baac44fee19140a74&ckck>, last view: 19/11/2024.