Modelling Regional Power Grids for Large-Disturbance Stability Studies

MSc thesis Electrical Power Engineering

Camiel Abraham van Altenborg

2017

Preface

Welcome, I bid you welcome... to life inside the electric circus!

– W.A.S.P., "Inside the Electric Circus", 1986

At a recruitment event, perhaps a year or two ago, I spoke to a representative of Nuon. 'Ah,' I opened, 'you are in the business of generating power!' He was quick to downplay that aspect of Nuon's business model, preferring to refer to his company as an "energy service provider" or something similarly vague. The fact that electrical energy also had to be generated somewhere, which I assumed to be the core business of a power company, seemed little more than an annoying distraction to him.

Nuon had only existed in its current form since 2009. Originally the product of a merger of provincial (i.e. publicly owned) power companies which operated transmission, distribution *and* generation infrastructure, it was forced by new legislation to break up: its transmission grid was transferred to the Dutch TSO TenneT, and its distribution grid to the newly created DSO Liander, leaving only the business of power generation for Nuon itself – and now, not a decade later, Nuon was trying to distance itself even from that.

I open with this anecdote because it clearly illustrates how, in the "electric circus" of the electricity sector – which thirty years ago had one of the most predictable programmes in the world – many performers are now looking to change their acts. Manufacturers, power companies, TSOs, DSOs, regulators, and other parties are wrestling for control over a flurry of ongoing developments, and it is far from clear who will be providing which services when the dust settles (*if* the dust ever settles). Those who stick to the same old routine for too long may be driven from the ring under a shower of jeers and rotten tomatoes, as the unpleasant fate of Delta demonstrates.

I had the privilege to enter the circus at a young age, as an intern at TenneT and then Siemens. For my MSc thesis, which you are reading right now, I returned to TenneT – fondly remembering the work and the people there, and at any rate vastly preferring a regular workplace to nine more months in academia. In the final stretches of this project, I was encouraged by the knowledge that my efforts to join TenneT as a full employee had proven successful.

But enough boasting; I am but a tiny dwarf standing on the shoulders of a veritable pyramid of giants. First and foremost, I would like to thank Frank Spaan, my daily supervisor at TenneT for the duration of this project. He tracked my progress, taught me how to use the necessary tools, and acted as a sounding board for my ideas, as a good supervisor does – but he was and is so much more than that. With his vast experience (having worked at TenneT and its predecessor, the SEP, since before I was born), Frank has also been a true mentor figure: providing insight into the inner workings of the vast organisation TenneT, gently reconciling ambition with reality when my plans for the project grew too wild, and instilling confidence when I went through phases of serious self-doubt. A soft-spoken, reflective Brabander, he perfectly complemented my own more excitable nature. Truly I could not have asked for a better supervisor.

I would also like to thank the many TenneT colleagues who lent their expertise (and sometimes crucial data!) to this project – particularly Loe von Berg, Jorrit Bos, Ruben van Dinteren, Richard de Groot, Kees Jansen and Vinay Sewdiem. I would often learn more from a twenty-minute chat with any of these gentlemen than from a month of trawling IEEE Xplore or the university library.

My thanks go out to my fellow intern Harish Krishnappa, who also finished his master's thesis within TenneT, working on a closely related project. Over these nine months we have shared laughs, shared technical tips, and shared frustrations about the sometimes byzantine bureaucracy of the TU Delft. I am cherished to know that he has also secured a job at TenneT, and hope we may be good colleagues for many years to come.

I thank my TU Delft supervisors, José Rueda Torres and Mart van der Meijden, for enabling me to graduate within the Intelligent Electrical Power Grids group and providing valuable feedback at different junctures of the project.

Many thanks, also, to my parents, without whose considerable practical and moral support I would have never pulled myself through this last leg of the journey towards a master's degree. I know that over the past twenty-four years I must sometimes have seemed a bottomless sink of time, money and hassle, but now at last before them stands a grown man and a Master of Science – I hope they will say that the investment has been worthwhile!

I would like to end on a special note of gratitude to *Doutora* Susana de Graaff, the *madrinha* of many talented young men and women within TenneT (I am honoured to be counted among them). Not only has she always set an inspiring example by her intelligence, deep theoretical interest, and force of personality; she has also helped me secure my place at TenneT, both in this project and in the trainee programme, which I am to join come January. I am eagerly looking forward to this next challenge, and to a long and fulfilling life inside the electric circus.

Contents

Chapter 1

Introduction

The reliable operation of an electrical power system is a complex task, and it is not set to get any easier as thermal power plants – with their high controllability, mechanical inertia, and large concentration of power in a single point – are gradually replaced by an array of renewable sources. A daunting literature on power systems theory exists, which can be divided into three fields of study:

- *Steady-state* or *stationary* analysis is concerned with a fixed operating point of the power system, that is, no progression over time is involved. It is performed to check that bus voltages remain within acceptable margins and branches are not loaded beyond their thermal limits. This is important when allocating transmission capacity to market parties wishing to buy or sell energy.

  - *Pseudostationary* analysis is a sub-category in which a sequence of different operating points is considered. This is useful for investigating "domino" events in which a given operating point triggers automatic control actions, thus shifting the operating point and triggering other controllers, and so on.

- *Dynamic* analysis is concerned with the time response of the power system to sudden or gradual changes in its operating conditions. Following such a change, the system may return to an acceptable operating point, but it may also fall into large oscillations or even outright collapse. If a power system is in a position to weather a disturbance and restore acceptable operating conditions, it is said to be *stable*. The variations of interest are generally below the power frequency (50 Hz) and the timeframe ranges from seconds to several minutes depending on the type of dynamics. The electromechanical behaviour of rotating machinery has long been the focus of dynamic analysis, but due to the transition mentioned above, some interest is shifting to the behaviour of power electronic converters.

- *Transient* analysis is concerned with the extremely short-term effects of a sudden disturbance, such as a lightning strike, ground fault, or switching action. These may lead to overvoltages, damaging components. Transient effects rarely spread far from the point where the disturbance occurs, so only a small part of the power system needs to be modelled – albeit in considerable detail: due to the very high frequencies under investigation, one must take into account all sorts of parasitic effects neglected in dynamic analysis, as well as the limited propagation speed of electromagnetic waves. The timeframe for transient studies is rarely longer than a power-frequency cycle or two.

This thesis falls within the field of dynamics. In the following section, we shall mark its place in this field more precisely, before outlining the practical context in section 1.2 and formulating our research question in section 1.3.

1.1 Theoretical context: Types of power system dynamics

In power system dynamics, we are interested in questions of power system *stability*. Let us take a look, then, at the various stability phenomena that occur in power systems, and define which of these fall within our scope.

Power system stability, as defined by a joint IEEE/CIGRÉ task force on the subject [1], is 'the ability of an electric power system, for a given initial operating condition, to regain a state of operating equilibrium after being subjected to a physical disturbance, with most system variables bounded so that practically the entire system remains intact.' This concept can be divided into different types of stability, based on:

- the system variables affected;

- the nature of the disturbance;

- the speed of the processes involved.

This gives us a classification tree such as in figure 1. (In much of the relevant literature, large-disturbance stability is called "transient stability"; this is an unfortunate term, as it has nothing to do with power system transients as described at the start of this chapter.)
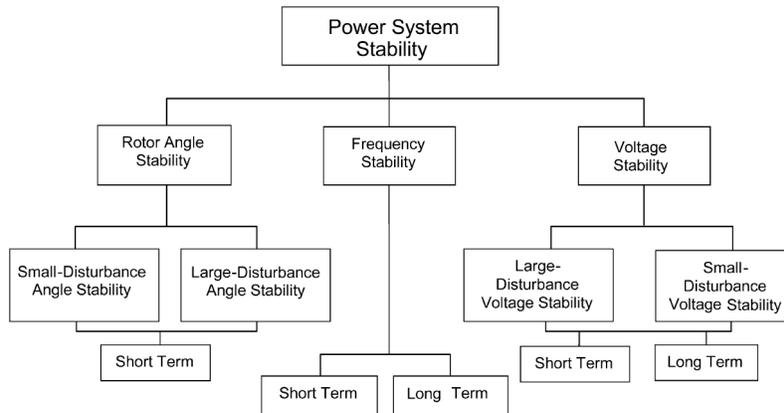

Figure 1: Types of power system stability [1, fig. 1]

As we can see, the variables of interest in power system dynamics are the system frequency, bus voltage magnitudes, and generator rotor angles. Frequency stability is fundamentally different from the other two as it is a global phenomenon; as the frequency is the same throughout an entire interconnected power system, frequency stability can only be assessed when the entire system is taken into account. Rotor-angle and voltage stability, on the other hand, can be assessed in a study with regional scope – such as ours. These two are often hard to separate in practice, especially where large-disturbance stability is concerned; a rule of thumb, however, is that rotor-angle instability is caused by the behaviour of generators, and voltage instability by the behaviour of loads [2, § 1.3; 3, § 2.3].

As for large-disturbance and small-disturbance stability, we can say that the former involves abrupt changes, e.g. the tripping of a line or generator; whereas the latter involves incremental, "straw that breaks the camel's back" changes, e.g. load variations. Obviously, then, there is a strong link between large-disturbance and short-term phenomena on the one hand, and between small-disturbance and long-term phenomena on the other. A more rigorous distinction is mentioned by Slootweg [4, § 5.2], who points out that a large disturbance is one that changes the topology of the system; small disturbances involve changes to system variables but not to topology.

The analytical tools applicable to large-disturbance and small-disturbance stability differ greatly. As shown by Kundur [5, ch. 12], the whole toolbox of state-space modelling – equilibrium points, Lyapunov functions, linearisation, modal analysis – can be brought to bear on small-disturbance stability problems. For large-disturbance stability, these methods are not relevant as they deal with stability around a certain operating point; a large disturbance shifts the operating point so far from its original position that any analysis tied to one operating point becomes meaningless.

For simple systems, large-disturbance stability can be assessed analytically by solving the generators' basic mechanical and electrical equations for the *critical clearing time* (CCT) – the fault-clearing time beyond which generators will be sent into self-reinforcing acceleration; see Kundur [5, § 13.1] or Grainger & Stevenson [6, § 16.6] for a full derivation. However, for large systems, such analyses are impractical and the only feasible way to determine the CCT is by time-domain simulation.

Given the considerable differences between small-disturbance and large-disturbance stability analysis, we have to restrict ourselves to one or the other in this thesis. Small-disturbance rotor-angle stability often involves inter-area oscillations – meaning we would have to study a large portion of the interconnected network in order to say anything meaningful about it – and depends on subtle details of generator modelling, whereas the focus of this thesis is on the behaviour of components other than synchronous generators (see section 1.3). Small-disturbance voltage stability is mostly an issue in systems where generation and load are very far apart [3, § 2.2] and thus again requires a large (part of the) network to be modelled. Moreover, any existing grid models we might have access to are intended only for short-term (read: large-disturbance) dynamic simulations, as we will discuss in section 1.2.

These are the reasons why we have opted to study the **large-disturbance stability** of regional grids in this thesis.

1.2 Practical context: Dynamics in the Dutch grid

The interested reader is referred to appendix A for an overview of the Dutch power grid, with its 220/380 kV "backbone" and seven regional grids at 110/150 kV.

Until now, the heavy meshing of this grid and the copious presence of thermal power plants have ensured a highly stable grid situation, in which dynamics were at most a minor issue. In this situation, *ad hoc* dynamic studies sufficed, and in fact this has been the approach of TenneT TSO until now – see, for example, the studies performed by Jansen et al. [7] and Spaan [8]. However, as the Netherlands slowly catch up with the sustainability transition, we can make fewer assumptions about the stability of our grid. Given TenneT's responsibility to society to ensure reliable operation of the power system, there is an increasingly pressing need to assess power system stability on a more regular basis.

What tools do we have, at present, to perform such assessments? The dynamic model of the Dutch grid currently in use within TenneT was compiled and validated by the KEMA (now DNV GL) in 2013 [9]. It is stored in a format compatible with the simulation tool PSSe. It contains dynamic models of the larger thermal generators and their controllers (governors and exciters); all other components are assumed to behave statically – that is, responding only to present conditions, without any "memory" of earlier events. Renewable generation is represented as negative load. Besides the Netherlands themselves, the grids of Germany, Belgium, Luxembourg and France are also included, albeit at a lower degree of detail. The model is intended to be valid for a study period of about 10 seconds.

One might ask if this model will continue to suffice for stability studies in a grid with less and less thermal generation, or if perhaps a higher degree of modelling detail will be necessary. That brings us to our research question.

1.3 Research question

Our research question for this thesis is:

Does a more detailed modelling of regional power grids impact their large-disturbance stability in simulations?

"More detailed modelling" concretely means that, to the KEMA grid model mentioned in section 1.2, we shall add dynamic models of the following components:

- Induction motor loads.

- Wind farms.

- Photovoltaic (PV) generation.

- Converters for HVDC links.

The specific reference to "regional grids" means that the disturbance under consideration, the monitored effects, and the added models themselves are all located within the same region of the power system (see section 4.4 for a more detailed description).

To avoid misunderstandings about the nature and aims of this project, we must stress the following points:

- We are not studying the effects of increasing wind and PV penetration itself. Although we may well use scenarios with higher and lower penetration to get a more complete picture, the comparison will always be between more and less detailed models, at the *same* level of penetration.

- Tackling four different components inevitably involves sacrificing some depth for breadth; the behaviour of each individual component and its controllers will not be investigated in great detail. This decision should be seen in the context of the project: this thesis is intended as a first step, to probe whether it is worthwhile at all to dynamically model components other than large synchronous machines. More detailed studies on the effects of any of these four components may follow; indeed, many have already been performed, as will become clear in chapter 2.

- This study is not meant to lead to hard conclusions about the stability of the regional grids under consideration; that would require far more rigorous justification of modelling assumptions. Rather, these regions (see section 4.4.1) are merely test cases in an attempt to draw conclusions about the *general* effects of increased modelling detail.

With that out of the way, the structure of this report will be as follows. In chapter 2, we shall explore the theoretical foundations of the models to be added, from which we can draw a hypothesis in chapter 3. Our method for testing this hypothesis will be described in chapter 4. In chapter 5 we shall present and analyse our simulation results, leading to an answer to our research question in chapter 6. Finally, chapter 7 will offer reflection on this study and recommendations for future research.

Chapter 2

Component theory

In order to select relevant scenarios and to properly reflect on our results, it is necessary to explore the effects on large-disturbance stability which we might theoretically expect from each of the four components selected in section 1.3. In sections 2.1 through 2.4, we shall discuss each of them in turn.

2.1   Induction motor load

The modelling of induction motors is part of a wider field of study known as *load modelling*, whose fundamental question is how the active and reactive power consumption of power system loads change as functions of bus conditions. For the purposes of this thesis we ignore frequency dependence as it is out of our scope, leaving the dependence of P and Q on the bus voltage. This can be a static dependence, only depending on the present value of voltage; the most common form of such a function is [5, eq. 7.1]:

$$P = P_0 \left(\frac{V}{V_0} \right)^a$$
$$Q = Q_0 \left(\frac{V}{V_0} \right)^b$$

(1)

where the 0 subscript indicates the base value of each quantity. It can easily be seen that $a = b = 0$ corresponds to a constant-power (P) model, $a = b = 1$ to a constant-current (I) model, and $a = b = 2$ to a constant-impedance model (Z). A commonly used choice for dynamic studies is to represent the active part of a load as entirely constant-current and the reactive part as entirely constant-impedance, i.e. $a = 1$ and $b = 2$ [5, § 12.8]. More sophisticated variations of the static load model, with non-integer exponents for different types of loads, have been proposed [2, § 4.1]; it is also possible to simply split the load into a Z, I and P component, creating the *ZIP model* [5, eq. 7.2]:

$$P = P_0 \left(p_1 \left(\frac{V}{V_0} \right)^2 + p_2 \frac{V}{V_0} + p_3 \right)$$
$$Q = Q_0 \left(q_1 \left(\frac{V}{V_0} \right)^2 + q_2 \frac{V}{V_0} + q_3 \right)$$

(2)

where $p_1 + p_2 + p_3 = q_1 + q_2 + q_3 = 1$.

For this thesis, however, we are interested in *dynamic* modelling, meaning that the load behaviour is assumed to depend on past as well as present conditions. In this regard, the induction machine, the most popular choice for motor applications [10, § 4.2.1], has similarities and differences with the generation workhorse that is the synchronous machine. The dynamic behaviour of synchronous machines is ground that has been exhaustively covered – Kundur alone devotes three chapters to it [5, ch. 3-5] – and it is not the focus of this thesis. Induction machines are fundamentally similar, in that their large-disturbance dynamic behaviour stems from the rotational mechanics of the rotor; however, there are two important differences:

- The rotor speed of a synchronous machine is locked to the system frequency, hence the name. (An exception is the brief transition period immediately following a disturbance, after which the machine either pulls out of step or settles on a new torque equilibrium at a different rotor angle; see Grainger & Stevenson [6, § 16.6].) The rotor speed of induction machines deviates from the electrical frequency by a per-unit value called the *slip* – indeed it must, in order to produce any torque at all.

- A synchronous machine can either produce or consume reactive power, depending on excitation [10, § 2.5]; an induction machine has no excitation winding and always consumes reactive power.

2.1.1   Equivalent circuit and torque-slip curve

The equivalent circuit of an induction machine, as derived in Fitzgerald [11, § 6.3-4], is shown in figure 2. Note that it is a *steady-state* equivalent circuit, meaning that it neglects the (very fast) electrical dynamics of the machine; the mechanical dynamics are still represented, however. This is a common simplifying assumption [2, § 4.3].

The most striking feature of figure 2 is that the effective rotor resistance, as seen from the machine terminals, is inversely proportional to the slip s. We thus expect the machine to draw a large, highly inductive current at high values of slip (which, in motor operation, correspond to low speeds). Some further analysis leads us to the *torque-slip curve* of figure 3.
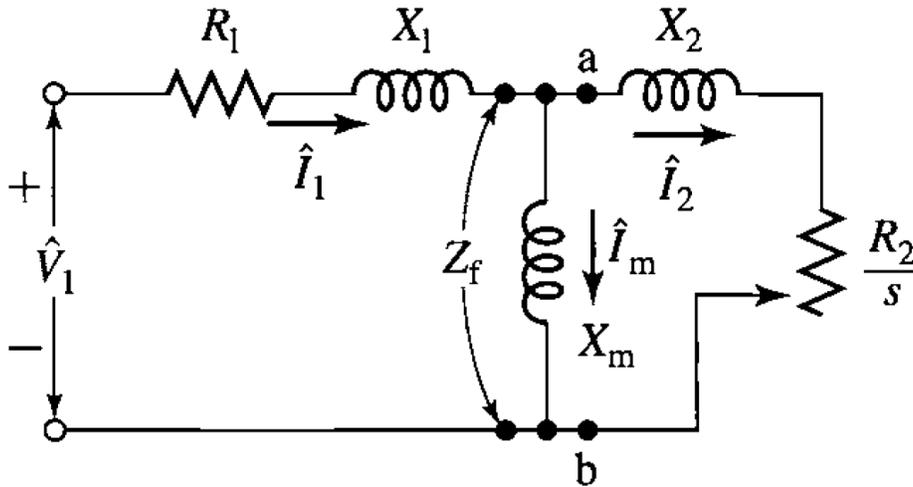

Figure 2: Equivalent circuit of an induction machine, with iron losses neglected [11, fig. 6.11]

10
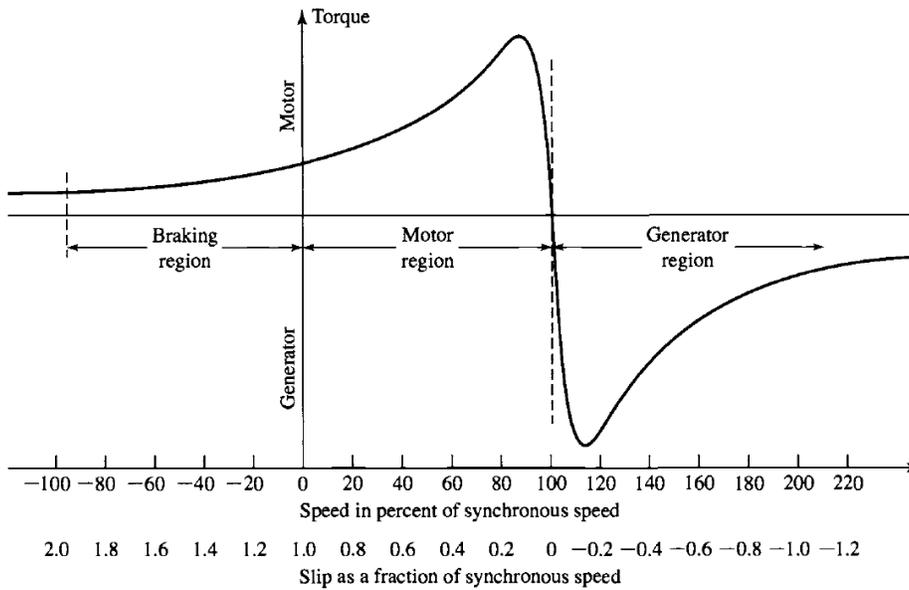
Figure 3: Electrical torque on the rotor of an induction machine as a function of slip [11, fig. 6.14]

This property of induction machines forces a compromise on the designer: a high value of R_2 is needed to keep starting currents down and avoid overheating the machine, but for efficiency in normal (low-slip) operation, a low value of R_2 is desired. Figure 4 shows the effect of R_2 on the torque-slip curve. A designer can get out from under this compromise by connecting the rotor circuit to an external, variable resistance, but this removes one of the main advantages of squirrel-cage induction machines: that they require no electrical connection between the rotor and any non-rotating elements, making them extremely robust [11, § 6.7.1].

Figure 4: The changing shape of the torque-slip curve for different values of rotor resistance [12, fig. 5.18]

11

Figure 5: The changing shape of the torque-slip curve for different values of terminal voltage [12, fig. 5.17]

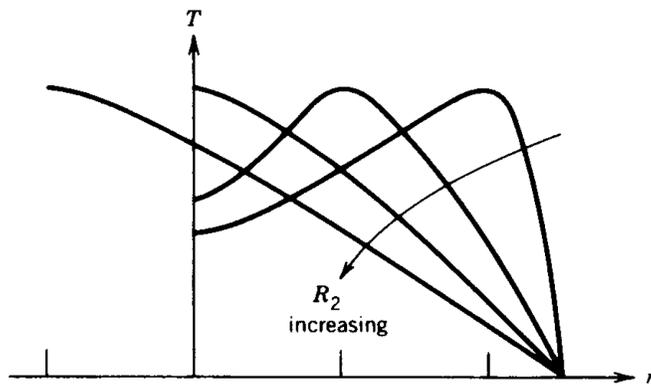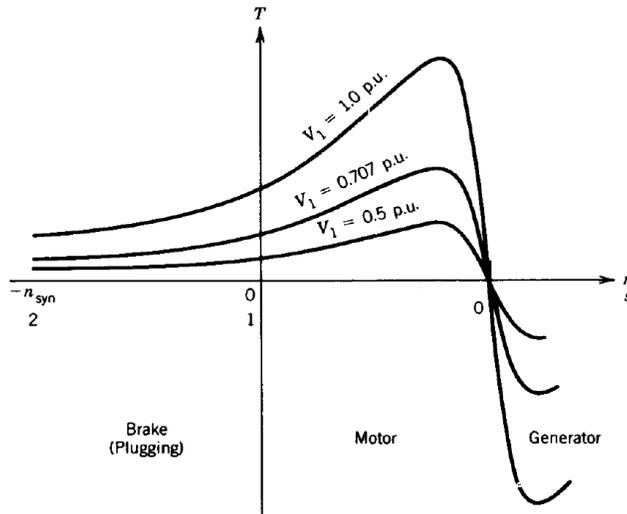In the event of a fault anywhere in the power system, the voltage on all buses drops to (nearly) zero; until the fault is cleared, little to no electrical torque acts on the induction machine rotor, as can be seen from figure 5. The induction machine is thus slowed down by the decelerating torque of its mechanical load, entering a high-slip operating region. (The lower the inertia of the machine, the faster it will decelerate.)

When the fault is cleared, accelerating torque is supplied to the machine again; what happens now depends on the mechanical load, which has a torque-slip curve of its own (see figure 6). The machine will either settle on a new torque equilibrium (intersection of the electrical and mechanical torque-slip curves) or slow to a complete stop – an event known as *stalling*. A stalled motor consumes a large amount of reactive power, significantly impairing voltage recovery; it may take very long for a stalled motor to speed up again, if it ever does. Motors that do reaccelerate can hurt voltage recovery for the same reason, in areas far from large generators; when located closer, they may improve stability by slowing down accelerated generators [13].
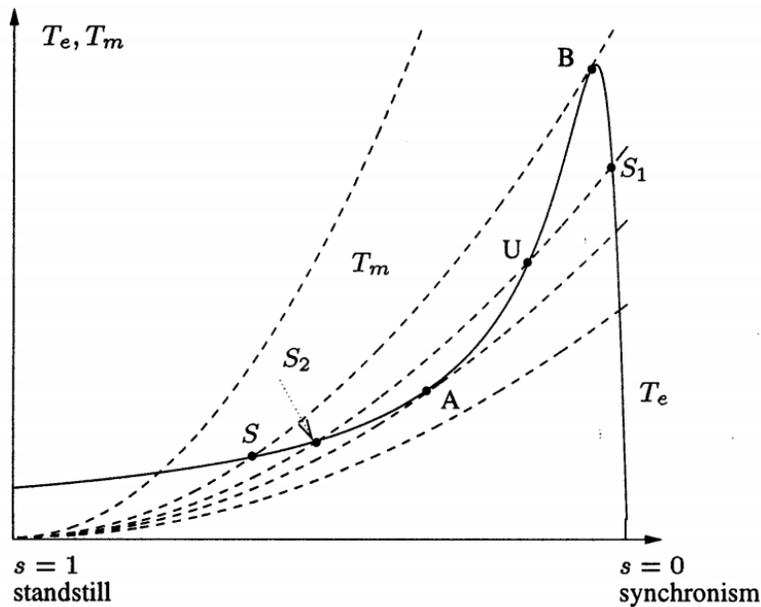

Figure 6: Induction machine torque-slip curve drawn together with several possible torque-slip curves of the mechanical load [2, fig. 4.7]

2.1.2 Protection

Besides the continuous electrical and mechanical dynamics of the induction machine, the discontinuous behaviour of protective devices must also be taken into account. The most important forms of protection for large induction motors are:

- *Undervoltage protection:* Disconnects the motor from the system if the bus voltage dips below a specified value. It can act either instantaneously or, for motors whose continuous operation is crucial, after a set delay (on the order of 2-5 seconds). Its purpose is to prevent excessive inrush currents when all motors in the system reaccelerate simultaneously after a fault, as well as preventing local hazards that might arise when a motor restarts automatically after having stopped [14, § 9.3.1; 15, § 8.11.3.2].

- *Overcurrent (stall) protection:* If a motor stalls, or settles to a torque equilibrium far below rated speed, it will draw sustained high currents which lead to overheating. For this reason, many large motors have protective devices installed that disconnect them from the system after about 10 seconds of stall [2, § 4.3; 13; 16].

2.1.3 Modelling

As with other forms of load modelling, the modelling of induction motors for power system studies almost always involves a degree of *aggregation*, in which a single large machine represents a number of smaller machines. The exact composition and structure of the load on a given high-voltage bus is rarely accurately known [13] and would introduce too much of a computational burden anyway. Nozari et al. [17] offer methods of aggregation, as well as estimates of the resulting motor parameters for different load classes.

There are three principal ways to model induction motor load for dynamic simulation [18, § 21.2.4.1]:

- Representing both its mechanical and electrical dynamics.

- Representing its mechanical dynamics and its steady-state electrical behaviour, but not its electrical dynamics.

- Representing only its static dependence on frequency and voltage.

The models available for use in this thesis – within the practical constraints given – are of the first type (representing both mechanical and electrical dynamics). They include a representation of undervoltage protection but not of stall protection [18, § 21.2.4]. Please see section 4.6.1 for a discussion of the model parameters and the values we have chosen for each of them.

2.1.4 Variable frequency drives

An increasing number of motors is connected to the power system not directly, but through an AC → DC → AC converter known as a *variable-frequency drive* (VFD). Unfortunately, established models of the dynamic behaviour of VFD-connected motors are not available [19; 20]; defining and validating such models is beyond the scope of this thesis. We shall therefore model only directly-connected induction motors and accept this as a limitation of our research.

2.2 Wind farms

2.2.1 Types of wind turbines

To convert the kinetic energy of the wind to electrical energy that can be fed into a power system, different approaches can be taken. There are four principal designs, numbered "Type 1" through "Type 4" by industry convention:

1. *Squirrel-cage induction generator:* The wind turbine rotor is connected, through a gearbox, to a squirrel-cage induction machine as discussed in section 2.1. The machine stator terminals are connected directly to the grid without any electronics; shunt capacitors compensate for the reactive power consumption of the machine. This design has the advantage of being cheap and robust, but severely limits control possibilities. The machine needs to be kept at a certain constant slip in order to be as efficient as possible; under changing wind conditions, this constant speed is (roughly) maintained by the gearbox. The shunt capacitors are typically dimensioned to supply the reactive power consumed by the generator in no-load operation, meaning that any additional reactive power consumed by the loaded generator must be drawn from the grid [21, § 2.3]. Because of the unique relationship between (absolute) slip and reactive power consumption, voltage control is not possible [4, § 5.3.1].

2. *Wound-rotor induction generator:* The wind turbine rotor is connected, through a gearbox, to a wound-rotor induction machine; the rotor terminals are connected through slip rings to an external resistance that can be varied, as mentioned in section 2.1. By varying the external resistance, the rotor currents can be controlled quite rapidly, allowing for constant power output even during gusts [22]. However, the wound-rotor design is inherently more fragile than the squirrel cage.

3. *Doubly-fed induction generator (DFIG):* The wind turbine rotor is connected, through a gearbox, to a wound-rotor induction machine (i.e. with an electrical connection from the rotor to a stationary circuit). The stator terminals are still directly grid-connected; the rotor terminals are also grid-connected − hence "doubly-fed" − through a back-to-back power electronic converter (AC → DC → AC). By manipulating the converter settings, a wide range of operating points can be realised. That is not to say that any desired (P, Q) setpoint can be achieved; the contribution to reactive power control is limited, among other things, by the current rating of the converters [4, § 3.4.4; 21, § 5.3]. The disadvantages of the Type 2 generator still apply, and the converter is an expensive and complicated addition.

4. *Direct-drive synchronous generator:* The wind turbine rotor is connected *directly* to the rotor of a synchronous machine (usually a permanent-magnet synchronous machine is used, eliminating the need for an excitation winding). Obviously the variable speed of the rotor cannot always match the 50 Hz required by the grid; hence, the machine is connected to the grid through a power electronic converter (AC → DC → AC). The main advantage of this design is that the lossy and failure-prone gearbox is eliminated; however, because of the much lower speed of the machine, a very large magnetic core is needed − especially problematic when one considers the high cost of permanent-magnet materials. Moreover, compared to the Type 3 design, the power electronic converter needs to process much more power, making it more expensive and increasing the associated losses.

Figure 7 shows each design schematically. Note that the Type 2 design − which relinquishes the simplicity and robustness of Type 1 while offering little of the flexibility of Type 3 − is rarely used; in literature it is often ignored entirely in favour of the other three [4, § 2.3.1-2; 23, § 2.1; 21]. We shall follow this practice. As the Type 1 design is rarely used anymore either [24, § 2.2], we shall focus our analysis on Type 3 and Type 4 turbines, known collectively as *variable-speed wind turbines* [4, § 2.3.1].
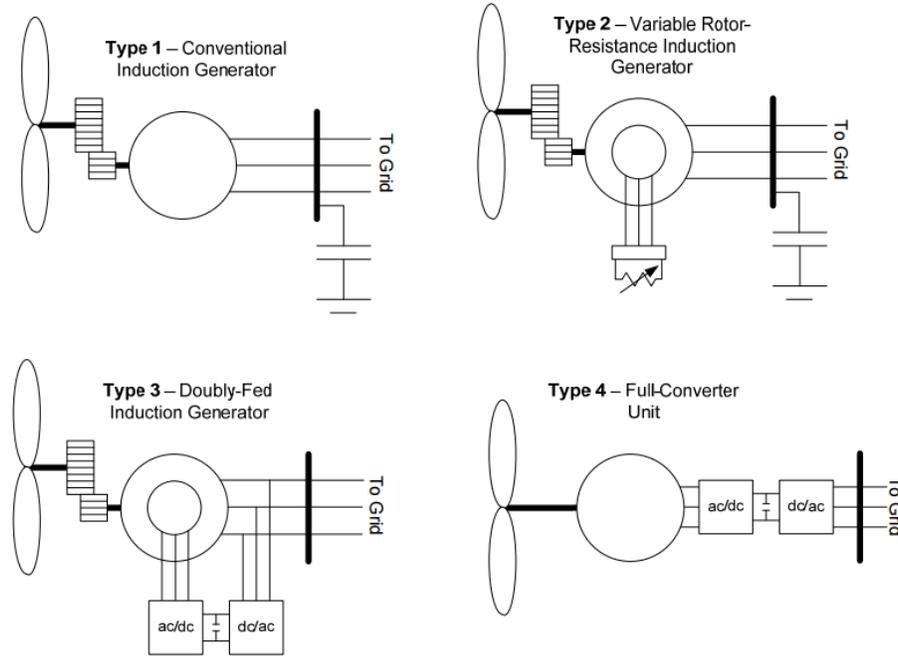
Figure 7: The four wind power generation schemes [25, fig. 1.1]

	Type 1	Type 2	Type 3	Type 4
Generator + converter	WT1G	WT2G	WT3G	WT4G
Electrical control		WT2E	WT3E	WT4E
Turbine	WT12T	WT12T	WT3T	
Pitch control			WT3P	
Pseudo-governor	WT12A	WT12A		

Table 1: Generic wind models available in PSSe [26, tab. 2]

Well-established generic models for all four wind turbine types exist, largely thanks to the efforts of the Western Electricity Coordinating Council (WECC). The models of the four wind turbine types are divided into several modules, as listed in table 1. As can be seen from the table, not all modules are included for each wind turbine type. The Type 3 model is the most complex, as this design includes both a power electronic converter and a direct grid connection to the generator. The Type 4 model is simpler, as the power electronic converter "hides" the mechanical behaviour of the turbine and generator from the grid. We shall now outline the structure of both the Type 3 and the Type 4 generic model.

First, however, it is useful to note that all model types assume constant wind speed − a reasonable assumption within the usual timeframe for short-term dynamics (10 to 30 seconds) [27]. Also, as with induction motor load (see section 2.1), wind turbine modelling usually involves aggregation: a common choice is to model all the turbines within a wind farm as a single large machine.

2.2.2 Type 3 (DFIG) generic model

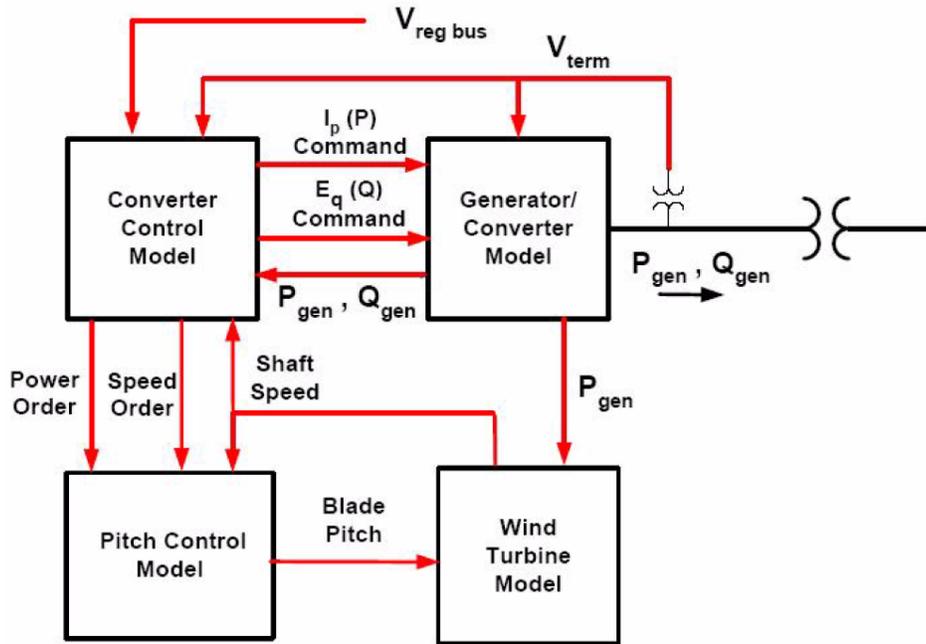Figure 8: Structure of the generic model of a Type 3 wind turbine [18, fig. 22.13]

Figure 8 shows the relationships between the four modules of the WT3 model.

- `WT3G1` (see figure 9) represents the generator and the power electronic converter. The generator's (fast) electrical dynamics are ignored; its mechanical dynamics are not represented either, as these are covered by the turbine model `WT3T1` [26, app. A; 28]. What remains is a current source: injected currents in the dq reference frame are calculated so as to track active and reactive power setpoints from the electrical controller; these current injections are then transformed to the system reference frame at an angle determined from the bus voltage using a phase-locked loop (PLL) [29, § 3.4.1; 30, § 4.1].

  Modern grid codes require *fault ride-through* (FRT) capabilities from wind turbines, meaning that turbines are required to stay online during a grid fault (within certain limits of fault severity). The sudden voltage changes associated with grid faults can induce very large currents in the rotor circuit, risking fatal damage to the power electronic converter. To be able to ride through faults without damaging the converter, a *crowbar* circuit can be used: rotor current is absorbed by a resistor rather than fed into the converter (see figure 10). When the rotor current decreases, the crowbar circuit is disconnected and normal operation resumes [29, § 3.1.2.1]. Note that this protective device is not included in the generic `WT3G1` module.

- `WT3E1` (see figure 11) represents the electrical control module, which generates the P and Q setpoints for the generator and converter to track. The P setpoint is determined by maximum-power-point tracking (MPPT) using a pre-specified power-speed curve [30, § 4.1] (see figure 12). The Q setpoint can be determined in one of three ways, depending on model settings:

  - Constant Q.
  - Q to maintain constant power factor.
  - Q control depending on voltage at a specified bus.

  Note that, although `WT3E1` is considered a single module, in fact the P and Q setpoints are determined in two separate control chains.

- **WT3T1** represents the mechanical drive train. From the blade pitch (input from **WT3P1**) and the electrical torque on the rotor (input from **WT3G1**) it calculates the rotor speed. It uses a two-mass mechanical model, with the turbine and generator rotors connected by a shaft that can be twisted, as shown in figure 14.

- **WT3P1** (see figure 15) represents the blade pitch controller. Given power and speed setpoints from **WT3E1**, it calculates the blade pitch angle and feeds it to **WT3T1** for use in mechanical calculations. Two practical limitations are taken into account [28]:

  - A time delay on pitch changes, to represent the limited speed at which large, heavy blades can be moved.
  - A minimum pitch angle, to keep the blades from "winding up" (i.e. coming full circle and beyond).
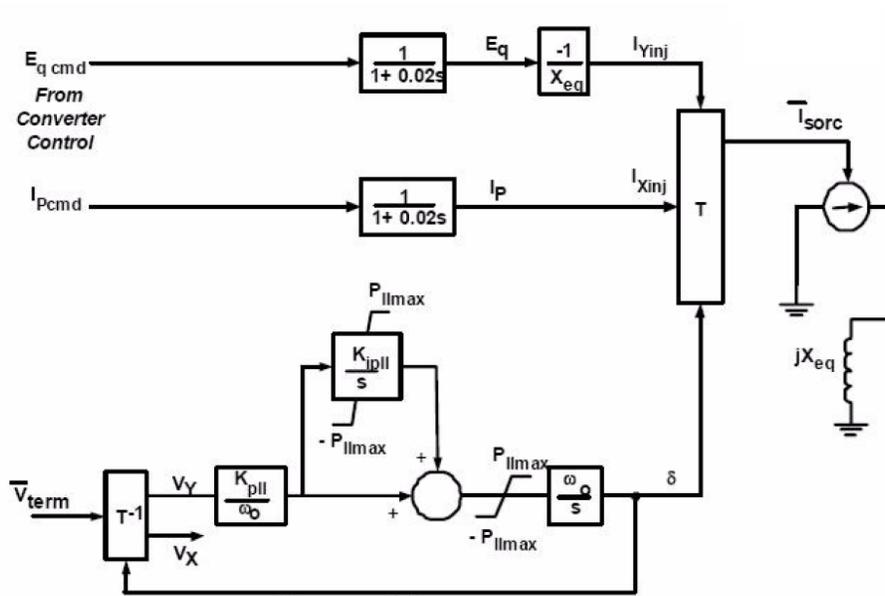
Figure 9: Block diagram of the **WT3G1** (generator/converter) module in the generic WT3 model [18, fig. 22.14]
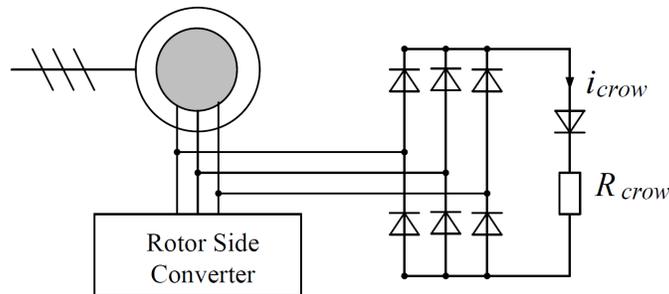
Figure 10: A crowbar circuit to protect a WT3 power electronic converter against rotor overcurrents [29, fig. 3.7]
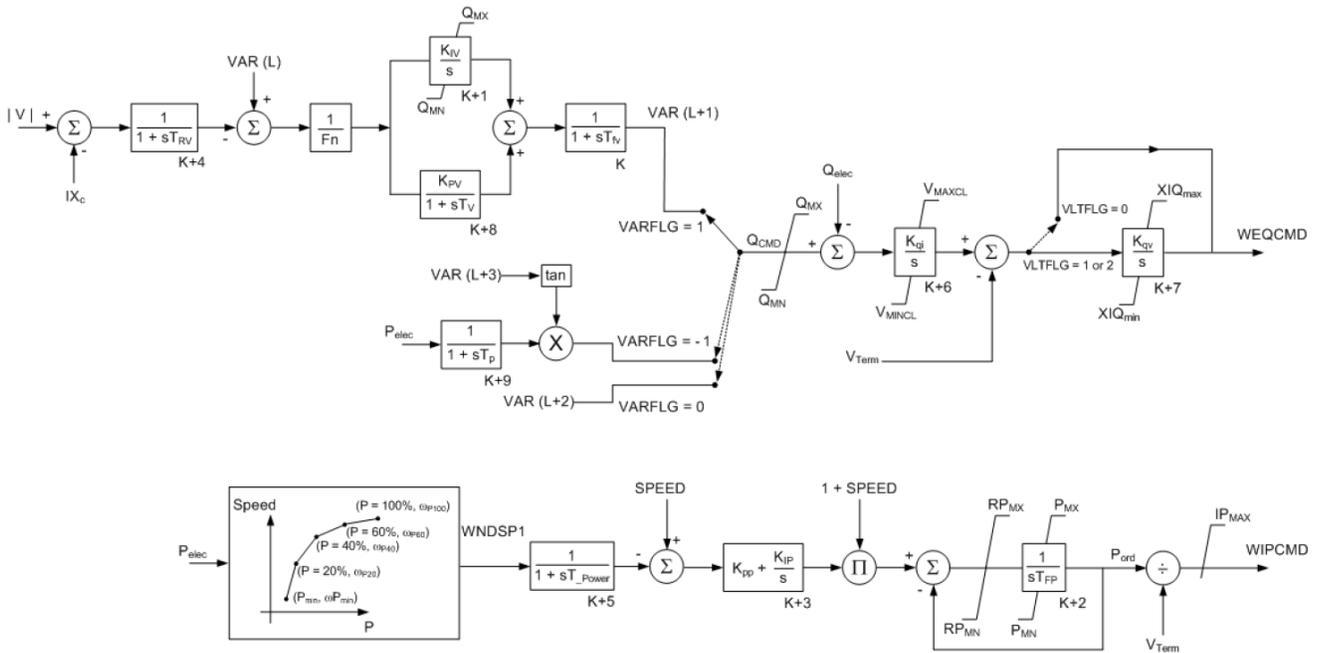
Figure 11: Block diagram of the `WT3E1` (electrical control) module in the generic WT3 model [18, fig. 22.16]
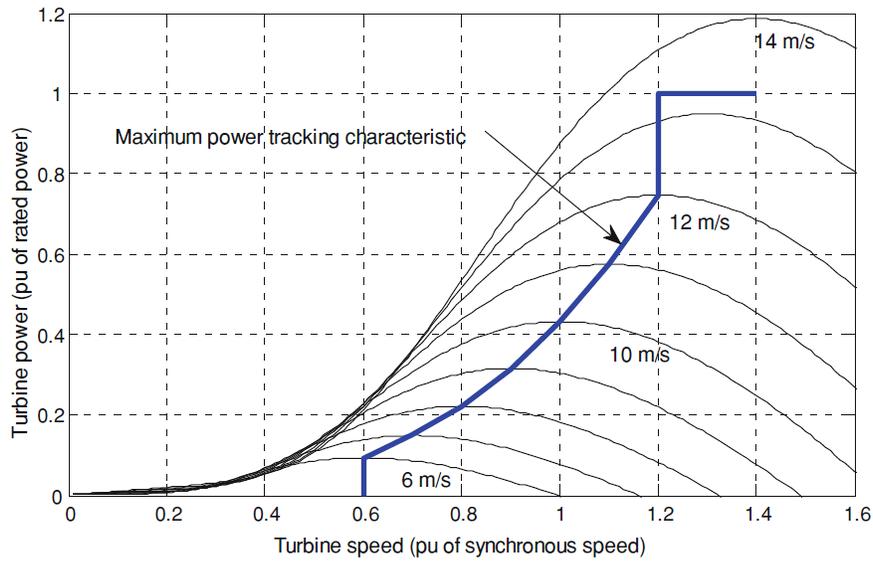
Figure 12: Illustration of MPPT for a wind turbine; the annotations on the grey curves refer to wind speeds [30, fig. 4.5]
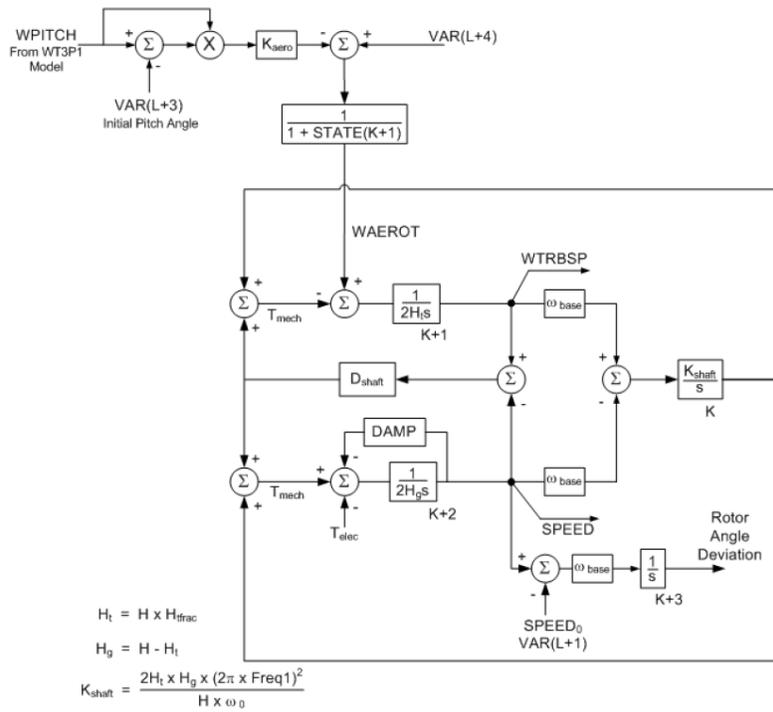
Figure 13: Block diagram of the `WT3T1` (turbine) module in the generic WT3 model [18, fig. 22.17]
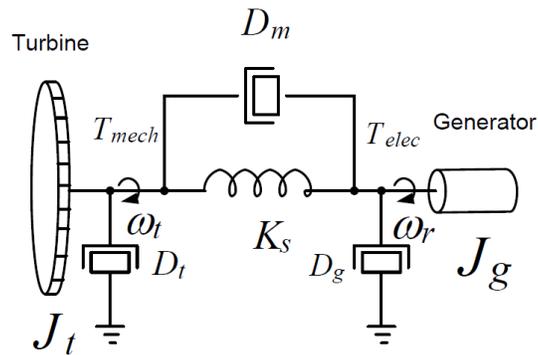
Figure 14: The two-mass model representing the mechanical drive train within the generic WT3 model [29, fig. 3.11]
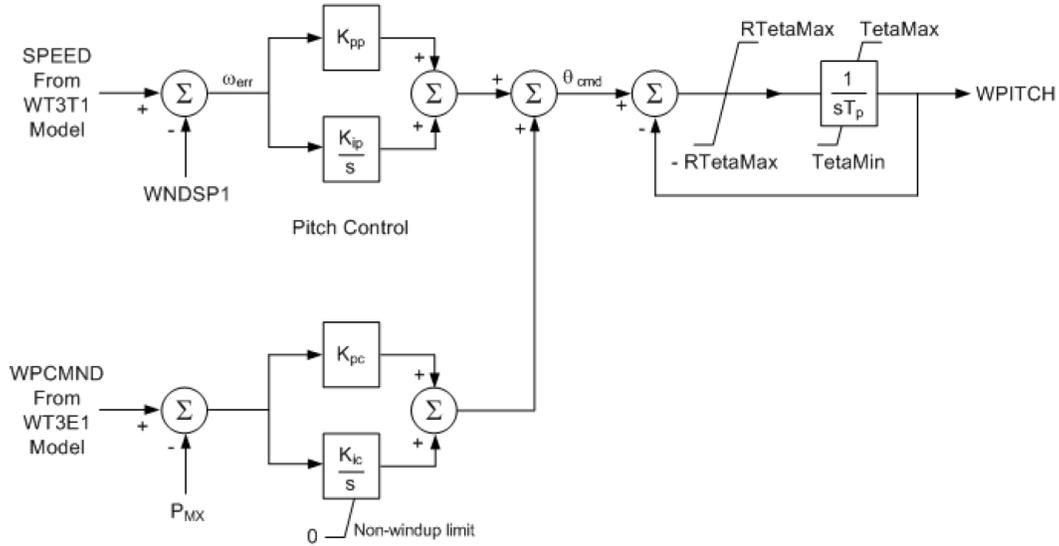
Figure 15: Block diagram of the `WT3P1` (pitch control) module in the generic WT3 model [18, fig. 22.18]

2.2.3 Type 4 (full converter) generic model

The WT4 model, as seen in figure 16, is essentially a simplified WT3 model: it lacks modules representing the mechanical drive train or the blade pitch control, as all this behaviour is isolated from the grid by the converter. The grid only "sees" the power electronic converter and its controller, represented by the `WT4G1` and `WT4E1` modules, respectively (see figures 17 and 18). These are very similar to their equivalents in the WT3 model, with two exceptions:

- The logic for calculating current injections from active and reactive power setpoints is slightly more complicated in `WT4G1` than in `WT3G1`.

- Unlike in `WT3E1`, the P and Q control chains in `WT4E1` interact. Because all current must pass through the converter in the WT4 design, active and reactive power must "compete" for the same converter capacity. WT4 model parameters include a flag specifying whether to prioritise P or Q when satisfying both setpoints would overload the converter.
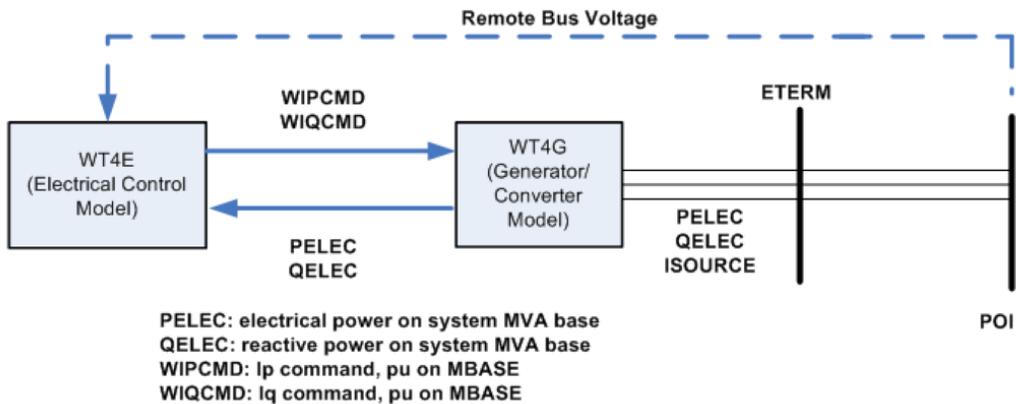
Figure 16: Structure of the generic model of a Type 4 wind turbine [18, fig. 22.24]
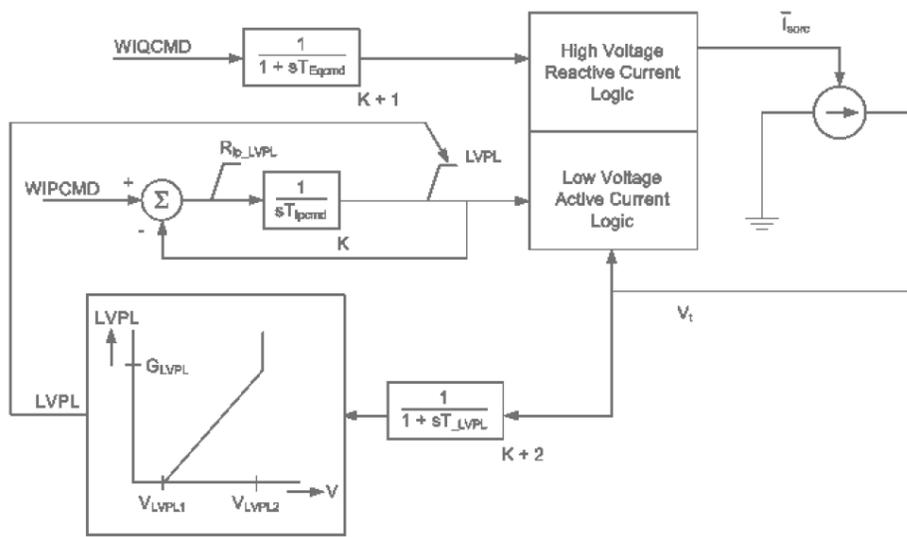
Figure 17: Block diagram of the `WT4G1` (generator/converter) module in the generic WT4 model [18, fig. 22.25]
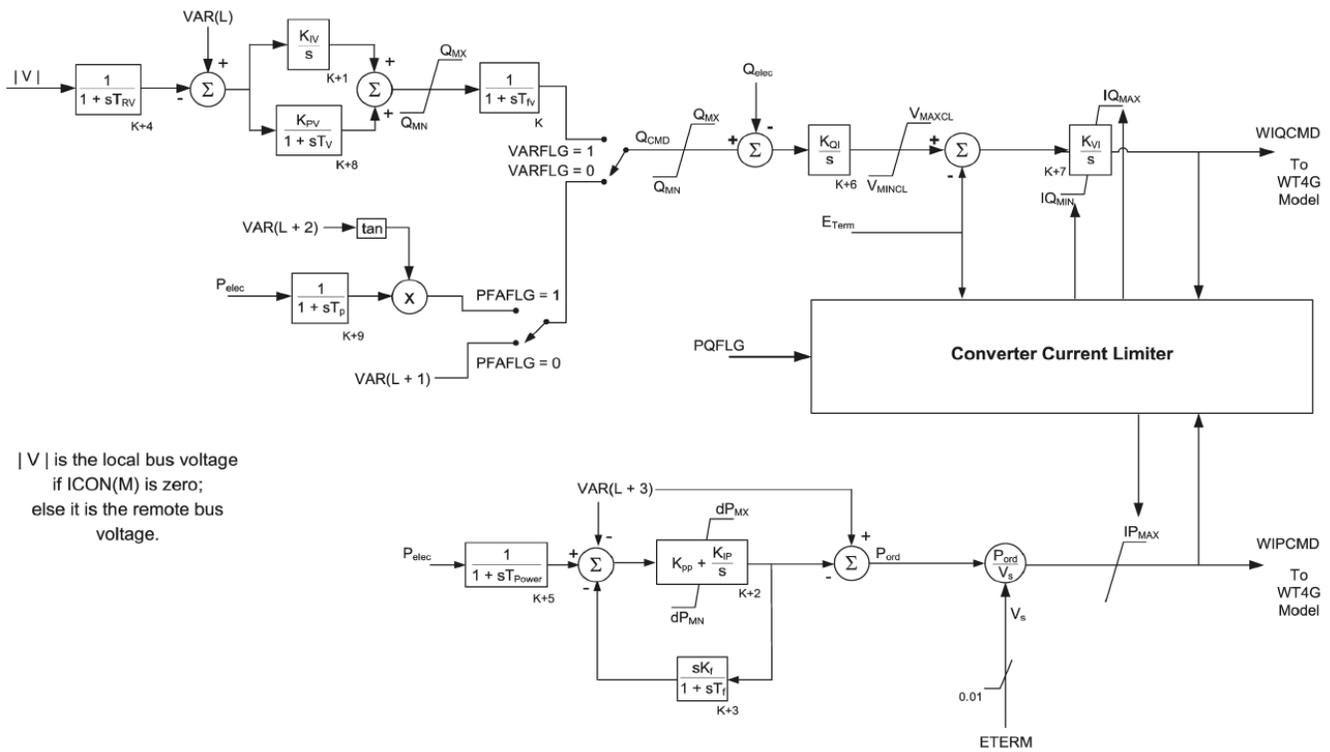
Figure 18: Block diagram of the `WT4E1` (electrical control) module in the generic WT4 model [18, fig. 22.26]

2.3 Solar PV

Although the PV panel – a semiconductor-based, DC-native source of electrical energy – is radically different from the wind turbine in terms of operating principles, from the grid perspective it is functionally identical to a Type 4 wind turbine. Both generators, after all, are connected to the grid through an inverter.

Two generic models for PV systems have been defined: one for large-scale "solar farms," and one for aggregated representation of smaller PV units [31; 32]. We choose to use the former, mostly for the prosaic reason that the distributed-PV model `PVD1` is not available in our simulation tool of choice, PSSe (see section 4.2). Focusing on large-scale plants also sidesteps a lot of the complications of aggregating small PV units, such as the representation of the distribution network [33].

The model for large-scale PV plants consists of two modules: `REGC_A` (generator/converter) and `REEC_B` (electrical control). Comparing figures 19, 20 and 21 to figures 16, 17 and 18, respectively, the similarities between `REGC_A` and `WT4G1` and between `REEC_B` and `WT4E1` are obvious. (The `REPC_A` module shown in figure 19, representing plant-level P and Q control, is optional; when it is omitted, P and Q setpoints are obtained from the power-flow solution [31, § 5.2].)

There is, however, a subtle difference in structure between `REEC_B` and `WT4E1`. Looking at figure 21, we can see that `REEC_B`'s reactive current command is a sum of two terms: one from a voltage control block (top), and one from a reactive-power control block (centre) that tracks either the Q or the φ setpoint. `WT4E1`, as can be seen in figure 18, responds *either* to the voltage control *or* to the Q/φ control signal depending on the value of the parameter `VARFLG`.
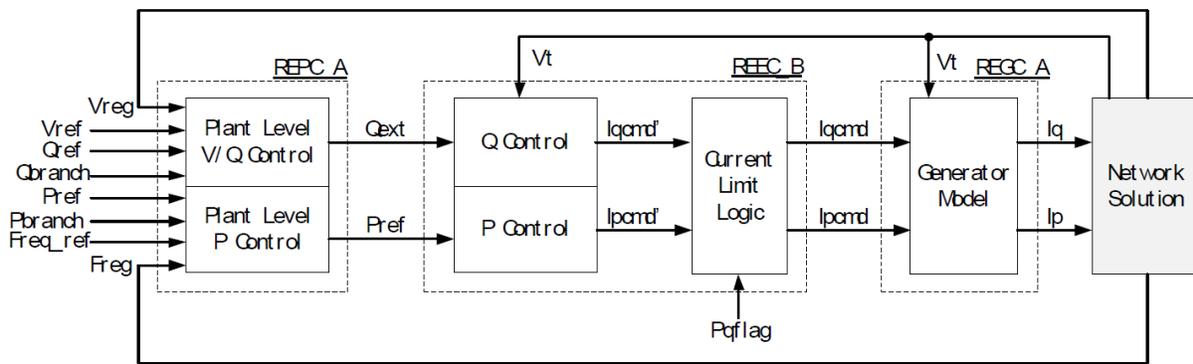
Figure 19: Block diagram of the generic model for large-scale PV plants [31, fig. 4]

Figure 20: Block diagram of the `REGC_A` (generator/converter) module in the generic large-scale PV model [31, fig. 10]

Figure 21: Block diagram of the `REEC_B` (electrical control) module in the generic large-scale PV model [31, fig. 11]

2.4 HVDC transmission

When transmitting electrical power underseas or across very long overland distances, high-voltage direct current (HVDC) transmission is a suitable choice. To connect such DC links to the wider AC power system, power electronic converters are needed, which may be called *rectifiers* (when power flows from the AC to the DC side) or *inverters* (when power flows from the DC to the AC side). Two HVDC cables are connected to the Dutch grid: BritNed to Great Britain, and NorNed to Norway. A third, the COBRA cable to Denmark, is under construction. See also appendix A.

Both the BritNed and NorNed converter stations are *line-commutated converters* (LCC), meaning that it is the sinusoidal nature of the AC-side voltages which causes commutation (i.e. the shifting of current from one thyristor pair to the next). These are also called *current-source converters* (CSC), because the DC-side current is kept constant, with the voltage fluctuating with power transfer – in contrast to the *voltage-source converter*, where the DC-side voltage is kept constant and the current fluctuates. The COBRA cable will use voltage-source converters, in order to allow for offshore wind farms to be connected along the way.

Both BritNed and NorNed use a *bipolar* DC transmission scheme, consisting of two conductors at opposite polarities. It is also possible to use a *monopolar* (one conductor) or *homopolar* (two conductors at the same polarity) scheme; the return path then is provided by a separate metallic conductor at ground potential, or by the environment (ground, sea) itself. The last option, however, may interfere harmfully with radio equipment and with oil and gas pipes [5, § 10.1]. Figure 22 shows the different HVDC transmission schemes.
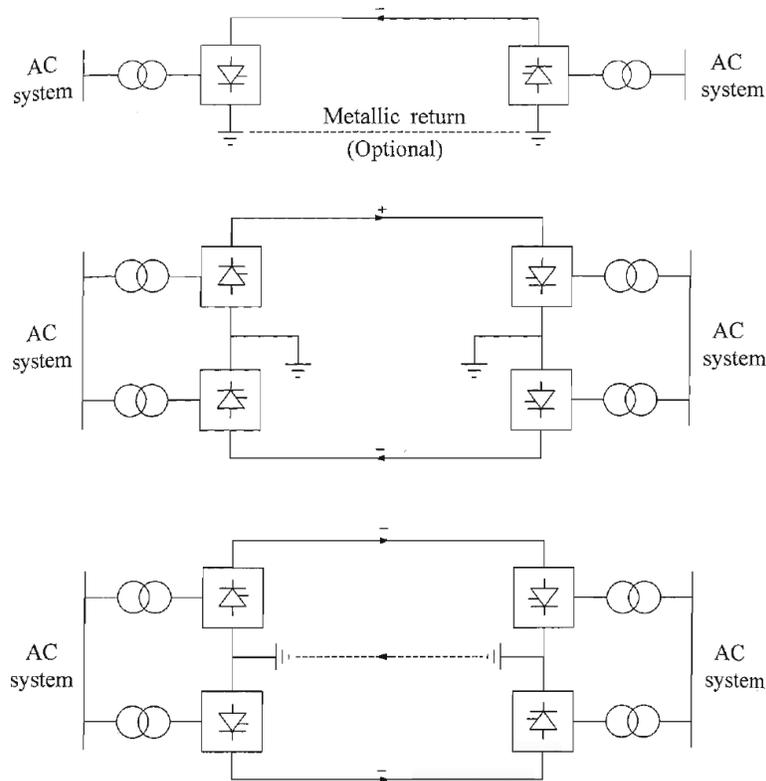
Figure 22: Monopolar (top), bipolar (centre) and homopolar (bottom) HVDC transmission schemes [5, fig. 10.1-3]

24

2.4.1 Basic converter theory

The converter "blocks" in figure 22 represent *six-pulse converters*, as shown in figure 23. The behaviour of such a converter can be controlled through the thyristor gate currents; each thyristor can conduct only when current is applied to the gate. Through the timing of these current signals, commutation can be delayed from the moment of *natural commutation* (i.e. the moment when commutation would occur if the circuit contained diodes instead of thyristors). The delay is expressed as an angle α known as the *firing angle* or *ignition delay angle*, relative to one cycle of the power frequency, which is taken as 2π.

As the AC-side current is shifted by an angle α with respect to the AC-side line-to-neutral voltage, the power factor angle φ is equal to α (see figure 24). The firing angle ranges from 0 to π radians, meaning that the converter can both supply active power to the AC grid (inverter operation) and draw active power from it (rectifier operation), but that it will always draw reactive power.
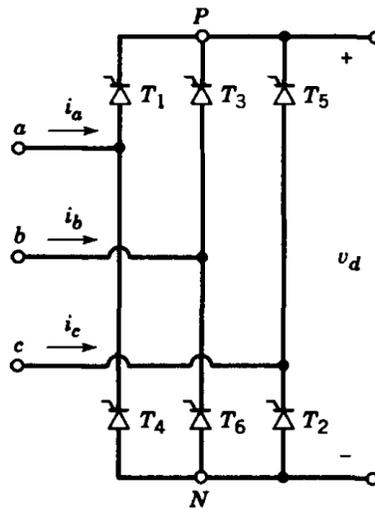
Figure 23: Six-pulse rectifier/inverter circuit [34, fig. 6.19]
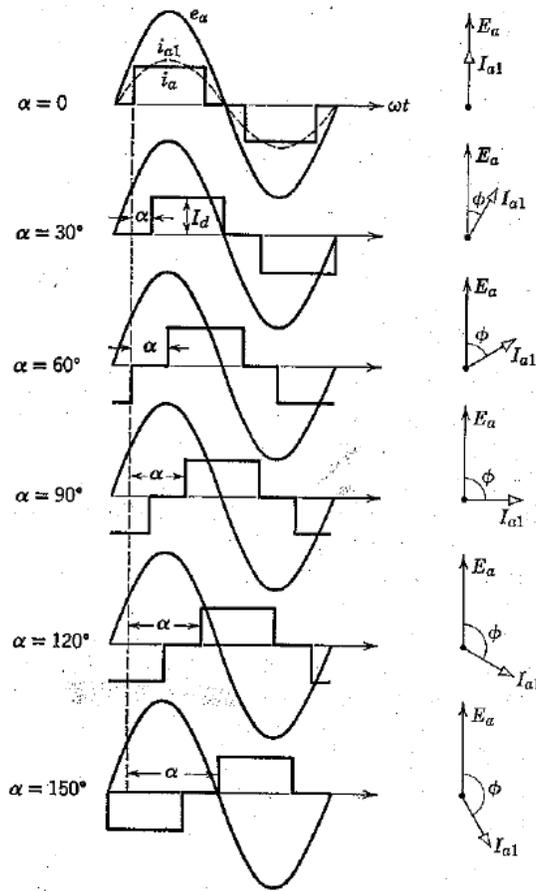
Figure 24: Firing angle and power factor of a six-pulse converter [35, fig. 3.7]

Besides the ignition delay angle α, four more angles are defined for use in converter theory – again expressed relative to one cycle of the power frequency. These are:

- μ: Commutation angle. Indicates how long commutation takes (in a practical circuit, commutation is not instantaneous due to inductances on the AC side).

- δ: Extinction delay angle. Defined as $\alpha + \mu$. The moment when commutation is completed, if the moment of natural commutation is taken to be 0.

- γ: Extinction advance angle. Defined as $\pi - \delta$.

- β: Ignition advance angle. Defined as $\pi - \alpha$.

In practice, δ is only used when discussing rectifier operation, and γ and β only when discussing inverter operation. This can be seen in figure 25, which offers a graphical illustration of these angle definitions. Note how $\alpha < \dfrac{\pi}{2}$ corresponds to rectifier operation, and $\alpha > \dfrac{\pi}{2}$ to inverter operation.
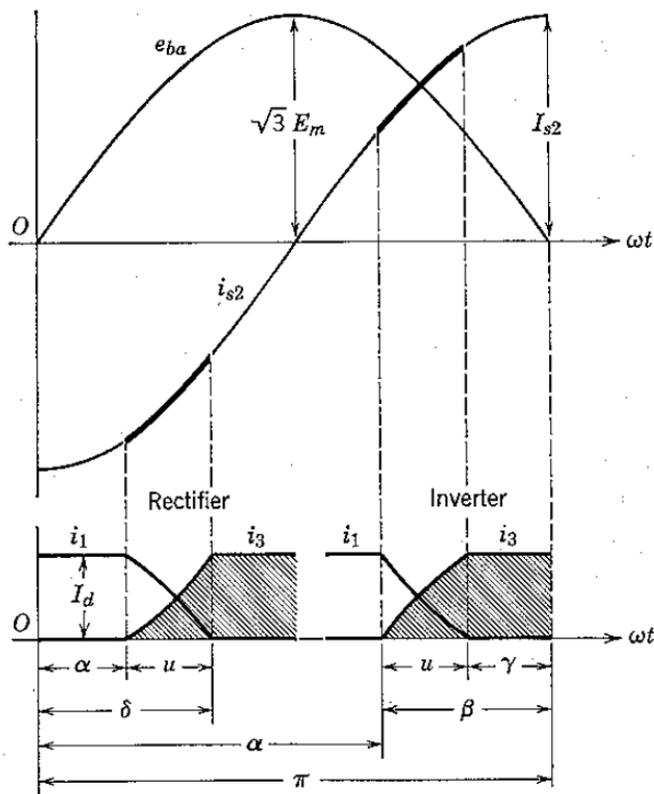
Figure 25: Angle definitions in rectifier/inverter theory [35, fig. 3.17]

Bear in mind, too, that due to commutation delay the straightforward relationship between α and φ expressed in figure 24 no longer holds; it can be modified as follows [34, eq. 6.64]:

$$\cos\varphi = \frac{1}{2}\left(\cos\alpha + \cos(\alpha + \mu)\right) \tag{3}$$

2.4.2 Control

Recall that the current-source conversion scheme tracks a power setpoint by varying the DC-side voltage while the DC-side current is kept constant. It is trivial to see that the power through the DC link is given by:

$$P = V_d I_d \tag{4}$$

The DC voltage can be written as [5, eq. 10.4]:

$$V_d = V_{d0}\cos(\alpha) \tag{5}$$

with V_{d0} locked to the (ideally constant) AC-side voltage by [5, eq. 10.3]:

$$V_{d0} = \frac{3\sqrt{2}}{\pi}V_{LL} \tag{6}$$

27

where V_{LL} is the RMS line-to-line voltage in the AC circuit. Substituting all of this into (4), we find for the power:

$$P = \frac{3\sqrt{2}}{\pi} V_{LL} \cos(\alpha) I_d \tag{7}$$

As we intend to keep I_d constant – by the definition of current-source conversion – *the only way to allow for fluctuations in DC power transmission is by adjusting the firing angle α*. This is called the *constant-current* (CC) control mode: the converter varies α with P in order to keep I_d constant. This has its limits, however:

- The minimum firing angle α_{min}, to ensure sufficient voltage across a thyristor at commutation.

- The maximum firing angle α_{max}, which follows from γ_{min}, to give each thyristor enough turn-off time before the next thyristor fires.

It is obvious from figure 25 that α_{min} will only be reached in rectifier operation, and α_{max} only in inverter operation. In the former case, the converter goes into *constant-ignition-angle* (CIA) mode, with α kept constant at α_{min}. Now, by equations (7) and (5), it is V_d which is constant, barring a small slope because of commutation resistance (which represents voltage drop due to commutation losses).

If α_{max} is reached, on the other hand, the converter goes into *constant-extinction-angle* (CEA) mode. Now γ is kept constant at γ_{min}, and the required firing angle to accomplish this is found from [35, eq. 5.12-13]:

$$\cos(\beta) = -\cos(\gamma_{min}) - \frac{2X_c}{\sqrt{3}E_m} I_d \tag{8}$$

where X_c is the commutation reactance and E_m the peak value of the line-to-neutral AC voltage.

Only one of both converters can be in CC mode, because the rectifier and inverter have different current setpoints; the difference is called the "current margin". Normally the rectifier is in CC and the inverter in CEA; if the rectifier hits its α_{min} limit, it goes into CIA and the inverter into CC. See figure 26.
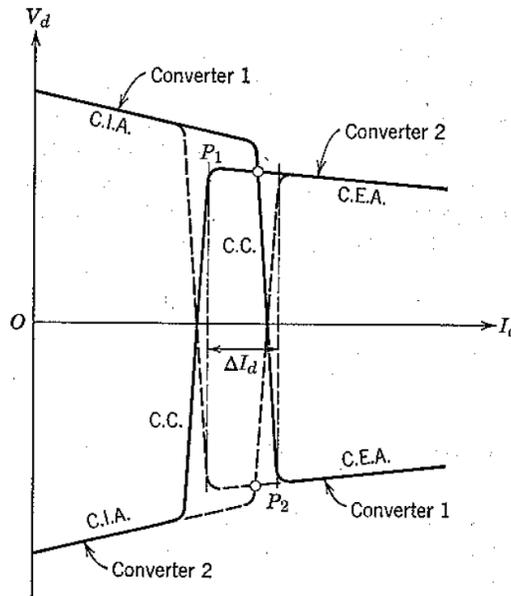
Figure 26: Voltage-current characteristics of an HVDC line with converters on both terminals [35, fig. 5.8]

A common modification of the characteristic in figure 26 is imposed by the *voltage-dependent current order limiter* (VDCOL). As illustrated in figure 27, this control mode takes a "bite" out of both converters' CC characteristics at low V_d (i.e. low power factor, by (5) and the $\alpha = \varphi$ relationship) in order to limit reactive power consumption.
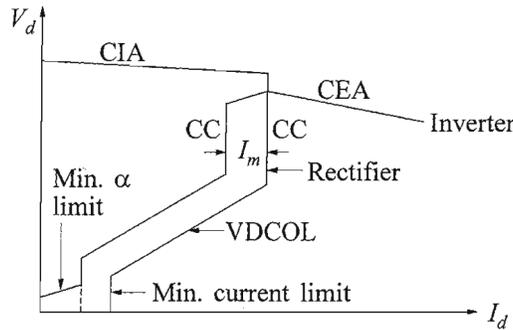
Figure 27: Effect of the VDCOL on HVDC converter characteristics [5, fig. 10.36]

Figure 28 shows how the control characteristic as expressed in figures 26 and 27 is realised in practice. Note how a single "master control" sets the current order, which is then fed to symmetrical controllers on the rectifier and inverter end of the HVDC link (they must be symmetrical, so that direction of power transmission through the HVDC link can be reversed).

As can be seen from the diagram, the master control will normally calculate the current order from the power order using the measured DC voltage on the rectifier end. However, if the DC voltage falls below a specified value, it switches to a fixed current setpoint in order to prevent overcurrent [18, § 19.2.1].
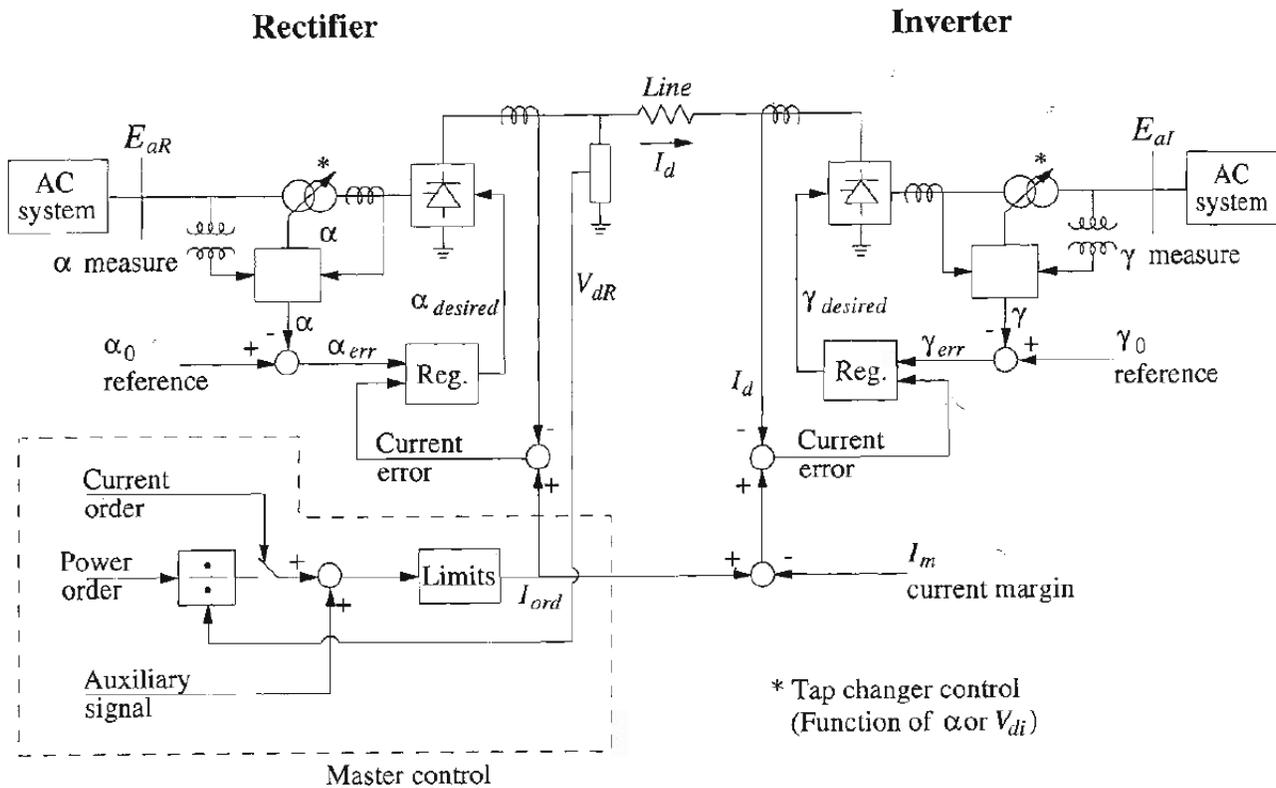
Figure 28: Diagram of the control scheme for an HVDC link [5, fig. 10.37]

2.4.3 Protection

Two protective actions can be taken in case of disturbances [18, § 19.2.2]:

- *Blocking:* The inverter and rectifier are both blocked from carrying any current. This happens in response to low AC voltage at the rectifier end.

- *Bypassing:* The inverter is bypassed, i.e. short-circuited on the DC side. This happens in response to low DC voltage at the inverter end and may be used to deal with commutation failures (see figure 29).
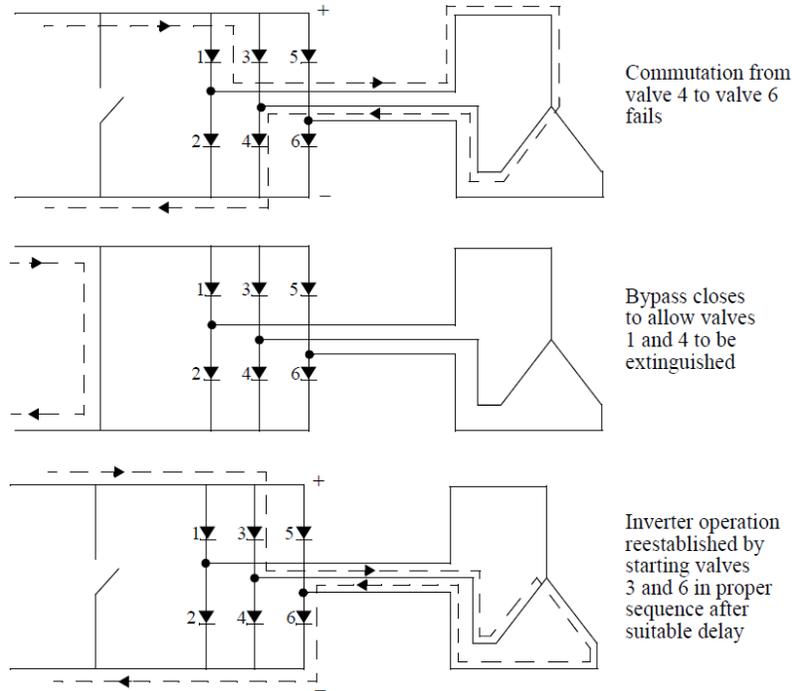
Figure 29: Bypassing the inverter after a commutation failure [18, fig. 19.5]

2.4.4 Modelling

Johnson [36] identifies three levels of modelling detail for HVDC transmission. From most to least detailed, they are:

- A fully dynamic and transient model, including representation of electromagnetic transients. This requires a very small time step, on the order of microseconds (often a "hybrid" simulation setup is used, in which the HVDC link is simulated with a smaller time step than the AC power system). As we are studying dynamics, not transients (see chapter 1), this model is too detailed for our purposes.

- A low-frequency dynamic model, including representation of the temporary situation following a disturbance when neither converter is at its α limit, meaning both are in CC mode and "tug of war" over the current ensues.

- A low-frequency dynamic, "pseudo-steady state" model; that is, the control system is 'considered to respond instantaneously to changes in AC system voltage, desired direct current, desired DC voltage or margin angle' [36]. This model cannot represent the both-converters-in-CC situation as it never occurs in steady state. The non-instantaneous recovery of DC voltage and current after a fault, while not modelled as explicitly as in the second approach, is approximated by setting maximum ramping rates for voltage and current.

The models practically available to us for use in this thesis are of the pseudo-steady state type. They include both protective actions mentioned in section 2.4.3.

The DC voltage and current ramping rates are particularly interesting parameters of these models. Taylor [3, § 8.3] warns that a slow recovery may hurt rotor-angle stability on the rectifier side: not restoring all of the "load" on the generators immediately may allow them to advance beyond their critical angles. On the other hand, a fast recovery also implies a large reactive power consumption immediately after the fault, hampering voltage recovery in the AC system [18, § 19.2.2] – with potentially similar results.

Chapter 3

Hypothesis

Based on our survey in chapter 2, we can expect the following effects on large-disturbance stability from dynamic modelling of the selected components, compared to a static representation of the same components:

- *Induction motors*, especially those with low inertia constants, are expected to slow down voltage recovery after a fault. See section 2.1.

- Variable-speed *wind turbines* (Type 3 and Type 4) are expected to support voltage recovery after a fault. See section 2.2.

- From *PV generation* we expect the same effects as from Type 4 wind turbines, as the models used are functionally identical. See section 2.3.

- The exact effect of an *HVDC terminal* depends on the current and voltage ramping rates. If programmed to recover slowly, prolonged active-power imbalance on the rectifier side may make generators more prone to pulling out of step. If programmed for a faster recovery, it may cause the same problem by slowing down voltage recovery. Either way its effect on system stability is expected to be negative. See section 2.4.

Our study cases and scenarios will be selected with a view to testing these hypotheses, as we shall describe in chapter 4.

Chapter 4

Methodology

In this chapter we shall describe the choices made in performing our research. After the process itself is outlined in sections 4.2 and 4.3, we shall discuss our choices of input data in section 4.4, 4.5 and 4.6; and our choices in selecting and analysing the output data in section 4.7.

As this is a rather complicated narrative with much potential for confusion, we shall first introduce consistent names for certain concepts in section 4.1.

4.1   Terminology

Input-side concepts:

- *Study case:* A region of the (Dutch) power system selected as a test network for our research. See section 4.4.1.

- *Parameter:* Any input property that may change between simulations but not during one simulation. For the purposes of our research, we identify three classes of parameters:

    - *Power-flow parameters,* those properties of the power system that would also be relevant for steady-state analysis, e.g. load pattern, renewable infeed, and topology. See section 4.4.
    - *Disturbance parameters,* properties of the disturbance itself, e.g. location and fault clearing time. See section 4.5.
    - *Dynamic modelling parameters,* properties of the dynamic models applied to the system. See section 4.6.

- *Scenario:* A given combination of power-flow parameters. "South Holland 2018, low load, peak RES infeed, no thermal generation" is an example of a scenario.

- *Configuration:* A given combination of dynamic modelling parameters. "Load and PV modelled statically, HVDC modelled dynamically, wind modelled dynamically as Type 3" is an example of a configuration. See section 4.6.

Output-side concepts:

- *Variable:* Any quantity whose value may change over the course of a simulation.

- *Monitor:* Any variable that is "monitored," i.e. whose values over the course of the simulation are recorded, so that they may be read and processed after the simulation is completed.

- *Output channel:* Concept from PSSe's monitoring system. Only if a variable has been assigned to an output channel will it be monitored.

- *Profile:* The total set of one variable's values over the course of a simulation, for example a "voltage profile" or a "rotor-angle profile".

- *Results:* The total set of profiles of all monitored variables resulting from one simulation.

- *Metric:* A quantitative test applied to a set of results. May be a number, e.g. "highest value of the voltage at any bus within 50 ms after fault clearing," or a boolean value, e.g. "whether or not synchronous generators stay in step."

- *Outcome:* The value of one metric for a particular set of results. For example, "in this simulation the synchronous generators pull out of step" is an outcome.

4.2 Tooling

The simulations for our research are to be carried out in PSSe. This has been the tool of choice for dynamic studies within TenneT since 2010; we can thus build on established practices and benefit from the experience of TenneT colleagues. Most crucially, the KEMA grid model [9] which serves as our base case is stored in PSSe format.

Initially, the open-source modelling language Modelica – which has drawn some recent interest from the industry [37; 38], as it allows unambiguous exchange of dynamic models among TSOs – was also considered. The ability to read and customise models on a purely mathematical level seemed a very attractive feature. However, Modelica-based simulation tools are still relatively unknown and little used within the TSO world. The effort involved in working "from scratch" (defining and validating component models; initialising and running simulations; creating grid models; etc.) was soon found to far exceed any realistic scope for an MSc thesis. Therefore, we have used Modelica purely as a means of increasing our personal insight into modelling concepts, not for the final study itself.

Another option we considered was PowerFactory; this competitor of PSSe offers neither the flexibility of Modelica nor the TenneT familiarity of PSSe, and was therefore rejected.

4.3 Workflow

For each simulation run, a number of consecutive actions have to be taken, which are outlined in this section. Thanks to PSSe's Python API, as well as other Python modules such as Matplotlib, we were able to automate the entire process so large numbers of different scenarios and configurations (see section 4.1) could be processed with a single command.
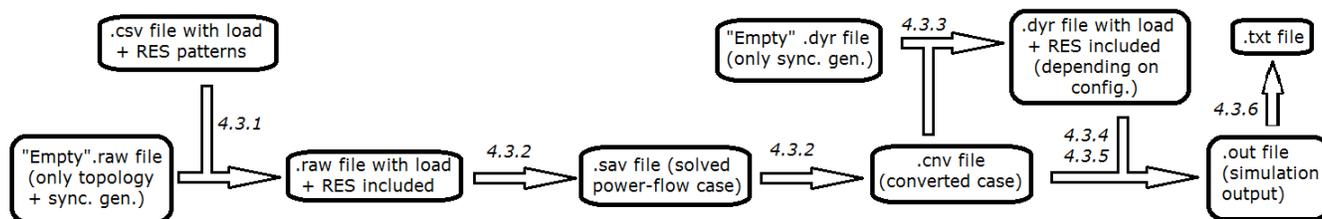
Figure 30: Flowchart showing the different files involved in a simulation run, and which of them are needed as input to create others. The numbers in italics refer to the subsection where the corresponding activity is described.

4.3.1 Creation of .raw files

The first step is to create a *.raw file* – a PSSe file format which contains the power system data necessary for running a power flow: topology, load, and generation data. Our choices for this data set are described in section 4.4.

Although the .raw file is exclusively a power-flow affair, dynamic modelling choices (see section 4.6) also have to be taken into account when creating it: depending on the configuration, renewable generation must be represented either as negative load or as generation.

4.3.2 Power-flow solution

Once a .raw file exists, the power flow must be solved so as to provide a starting point for dynamic simulations. The solved power flow is converted to a form that can be used as such a starting point, by the following steps [39, ch. 12]:

1. Converting generators to their Norton-equivalent representation (a current source behind an impedance, rather than a "power source" with a certain P and Q infeed).

2. Converting loads from constant-power to constant-current (active load) or constant-impedance (reactive load) characteristics (see section 2.1).

3. Reordering the network admittance matrix given the newly converted generators and loads.

4. Running a new power-flow solution.

Once these conversions have taken place, it is not possible to run another regular power flow in PSSe. Therefore the original solved power-flow case is stored in a *.sav file*, and the converted case in a *.cnv file*.

4.3.3 Creation of .dyr files

In PSSe, dynamic modelling data is "superimposed" on a converted power-flow case by means of a *.dyr file*, which lists the dynamic models to apply to system components, and their parameters. Such a file must be created according to the choices made in section 4.6.

4.3.4 Assigning output channels

With the .cnv and .dyr files, PSSe has all the information it needs to run a dynamic simulation. However, such a simulation will be useless to us if we cannot access its results. Hence, before running the simulation, we must decide which system variables to monitor, and assign these to output channels in PSSe. See section 4.7.1 for our selection of monitors.

4.3.5 Running the simulation

When running the simulation, it is important to note that initialising dynamic simulations in steady state (i.e. with all time derivatives equal to zero) is often difficult, due to the complicated constraints imposed by generators and control systems. The usual solution is to either tweak the power-flow outcome by hand until all constraints are satisfied, or let the simulation "settle" to steady state before applying any disturbances [40]. We choose the latter, and let each simulation run for 10 seconds before applying a disturbance.

4.3.6 Processing results

After a simulation run, the values of the output channels over time are exported to a text file. This can then be read for plotting with Matplotlib or for quantitative analysis and classification of the results (see section 4.7).

4.4 Input: power flow

4.4.1 Provinces selected as study cases

As mentioned in chapter 1, we are interested in the stability of *regional* grids; in other words, in those effects of a disturbance that occur relatively close to the location of the disturbance. We must select, then, one or more regions of the power system as study cases for our simulations. Given the structure of the Dutch power system as outlined in appendix A, we have opted to use **Zealand** and **South Holland** for this purpose.

Zealand, as mentioned in appendix A.2, is not actually operated as a standalone grid at present; it is part of the larger Zealand/Brabant grid. However, with a slight modification of the topology – disconnecting the Rilland-Woensdrecht line – we can turn Zealand into an independent 150 kV grid, with a single 380 kV connection point at Borssele.

The great advantage of this grid is that it is small and simple, containing only ten 150 kV stations. This makes it relatively easy to understand its structure and trace the effects of changes. The large concentration of thermal generation at Borssele gives us a straightforward switch to flick in order to set up different scenarios; the (mostly offshore) wind parks planned here will be useful for testing the effects of wind turbine modelling. Moreover, Zealand is quite an isolated corner of the power system, with its single 380 kV connection and two spur lines (Westdorpe-Oostburg and Kruiningen-Rilland). This is quite a distinct grid structure that can be contrasted to other, more complex grids.

South Holland (see appendix A.5) offers one such contrast. It is a large and heavily meshed grid, with a 380 kV ring of its own feeding multiple 150 kV pockets. The sheer amount of industrial load in the Rotterdam port makes South Holland uniquely suited to studying the effects of motor load modelling. Thanks to the BritNed cable, connected at the Maasvlakte 380 kV station, we can also experiment with the dynamic modelling of HVDC converters here.

4.4.2 The external grid

Given the strictly regional scope of our thesis, we will in both cases model *only* the regional grid under consideration and a small section of the surrounding 380 kV grid (the line Borssele-Geertruidenberg in the case of Zealand; the 380 kV ring Krimpen-Crayestein-Simonshaven-Maasvlakte-Westerlee-Wateringen-Bleiswijk and the line Krimpen-Geertruidenberg for South Holland). The rest of the European power system (the "external grid") will be represented as an infinite bus at Geertruidenberg. This makes it straightforward to vary the short-circuit level of the external grid in simulations, by varying the impedances of the lines Borssele-Geertruidenberg and Krimpen-Geertruidenberg respectively. Base-case values for these impedances are found from TenneT's short-circuit calculations for 2016 [41, app. C], by:

$$Z = \frac{V_b^2}{S_{sc}} \tag{9}$$

where V_b is the base voltage and S_{sc} the apparent short-circuit power. We assume an $\frac{X}{R}$ ratio of 10.

South Holland contains a second external grid: the representation of the British grid on the far side of the BritNed HVDC link. The topology and short-circuit level of this connection are copied from .raw files for BritNed provided by its manufacturer Siemens [42].

4.4.3   Load and renewable generation patterns

The installed load as well as wind and PV generation at each 150 kV substation is taken from estimates provided by TenneT's Asset Management (AM) department. We consider two reference years: 2018 and CD2035 (where CD stands for *Centraal Duurzaam*, "Centralised Sustainable" – as opposed to *Decentraal Duurzaam*, "Distributed Sustainable"). The choice for the "centralised" rather than the "distributed" scenario for 2035 was made for the practical reasons outlined in section 2.3.

The specific load and generation estimates used are confidential to TenneT. A clear pattern, however, is that the 2035 estimate – obviously – contains a much higher penetration of wind and PV generation than the 2018 estimate.

As per industry practice, we consider a "low-load" scenario to be one with 50% of peak load; the same goes for RES (wind and PV) infeed. We now have sixteen different scenarios:

- Zealand or South Holland (2);
- 2018 or 2035 (4);
- low or peak load (8);
- low or peak RES (16).

4.4.4   Thermal generators

The final parameter in our scenarios is the presence of thermal generation.

In both provinces, we would like to have the same generators in service in different scenarios (low or peak load, etc.) so that our results are not unduly distorted by different levels of inertia or different active excitation systems. With that in mind, we selected status (in or out of service), P infeed and voltage setpoints of the synchronous generators in both grids with a view to:

- producing a converging power flow in different scenarios;
- minimising voltage drops within the 150 kV grid, with the 380 kV buses at a slightly higher voltage;
- and ensuring none of the generators would be at their Q_{max} limits.

For Zealand, we succeeded with the Borssele 30 nuclear power plant and the four Elsta generators (at Terneuzen) in service and generating their respective P_{max}. For South Holland, we did not succeed entirely as the synchronous generators are concentrated in the Rotterdam port, leaving us with little ability to control bus voltages in the northern part of the grid (the area around Leiden and Zoetermeer). Eventually, with the MV-3 and CR-1 (both at Maasvlakte), two Pergen units (at Geervliet Noorddijk) and two REC units (at Rotterdam Noord) in service at their P_{max}, we had to accept a voltage drop of about 0.03 p.u. from the southern to the northern part of South Holland. The Westerlee-De Lier pocket (see appendix A.5) was at a higher voltage than either because of its close proximity to the 380 kV grid.

4.5   Input: disturbance parameters

The disturbance to be considered is a three-phase fault to ground on a heavily loaded line – a common "reference disturbance" in similar studies [43]. Fault-clearing time (i.e. the time between the start of the fault and tripping of the faulted line by circuit-breaker action) is to be varied, in order to study the effect of different modelling configurations on the critical clearing time of the system, one of the most interesting properties in power system stability (see section 1.1).

Based on branch flows from the solved power flow, one circuit of the 150 kV line Borssele-Ellewoutsdijk was selected for Zealand; for South Holland, one circuit of the 380 kV line Crayestein-Simonshaven was selected.

To make sure we do not overlook important information by considering only one disturbance type at one location, some verification scenarios are simulated with faults at other locations, as well as a different disturbance type (tripping of a large synchronous generator). See sections 5.1.2 and 5.3.4.

4.6   Input: dynamic modelling

As mentioned in section 1.3, we are interested in a comparison between static and dynamic modelling of the four selected components: motor load, wind generation, PV generation and an HVDC terminal. The static modelling is straightforward enough: we can simply use a static load model, with wind and PV generation represented as negative load. As is common practice, we assume the active part of this static load to be constant-current and the reactive part to be constant-impedance (see section 2.1).

With regards to dynamic modelling, different choices are available to us in terms of which models to use and with which parameters. These combine into many different possible configurations. To avoid confusing, excessively wordy descriptions, we use a structured label format for modelling configurations:

$$\textbf{MxWxPxDx}$$

where **M** stands for motor load, **W** for wind, **P** for PV and **D** for HVDC transmission. Each **x** represents a number – so, for example, M0W3P1D0 is a possible configuration; M4W5P0D1 is another. What the numbers for each component mean will be explained in the following subsections.

Note that the infinite bus at Geertruidenberg (see section 4.4.2) is represented in practice by a `GENCLS`-model generator with its inertia set to infinity. Although PSSe contains no true "infinite bus" model (as in, a bus that is defined to always maintain the same voltage magnitude and angle), the infinitely heavy `GENCLS` generator approximates it adequately. In fact, it is meant to be used as such [18, § 15.4.2]: 'The `GENCLS` model is intended to be used primarily as an effective short circuit current source in setting up approximate equivalents of segments of large interconnected power systems that are far removed from the area of specific interest.'

4.6.1   Motor load

For motor load we use the `CIM5BL` model, a PSSe model which represents both electrical and mechanical dynamics and undervoltage protection but not stall protection, as discussed in section 2.1.3.

Nozari et al. [17] provide equivalent-circuit parameters as well as inertia constants for motors representing different load classes. These can by found by estimating the share of different motor types (with known parameters) within each load class, and then applying a simple weighted-sum formula [17, eq. 1]:

$$p_{agg} = \Sigma_{j=1}^{n} \sigma_j p_j \tag{10}$$

where, for each motor type j, p_j is the value of parameter p for that motor type and σ_j is the fraction of motor load formed by that motor type.

Nozari et al. offer parameters for residential loads and for mixed residential-industrial loads, but none for entirely industrial loads. However, we can use the σ and p values provided specifically for industrial motors in the "mixed" group [17, tab. 2] and apply equation (10) to – somewhat inelegantly – find our own aggregate values for a fully industrial load class. This gives us the following parameters:

Parameter	Residential	Industrial	Mixed
RA	0.099	0.049	0.023
XA	0.112	0.096	0.085
XM	2.260	2.960	3.477
R1	0.086	0.044	0.014
X1	0.102	0.144	0.175
R2	–	–	–
X2	–	–	–
MBASE	–	–	–
PMULT	2.179	1.684	1.445
H	0.737	0.960	1.069

Table 2: `CIM5BL` parameters for different load classes (all values in per unit on the motor base).

(Note that the parameters from Nozari et al. are for single-cage motors; thus `R2` and `X2` are out of the equation. `MBASE` is an optional parameter for when `PMULT` is not specified.)

Besides these electrical and mechanical parameters, `CIM5BL` has other parameters, related to its magnetic saturation curve (see figure 31) and the settings of the protection relay. These are the same for all load classes.

Parameter	Value	Unit	Source
E1	1.0	p.u.	[18]
S(E1)	0.03	p.u.	[18]
E2	1.2	p.u.	[18]
S(E2)	0.04	p.u.	[18]
VI	0.90	p.u.	[44]
TI	100	cycles	[14]
TB	10	cycles	[45]
D	0	p.u.	[18]
TNOM	0.85	p.u.	[18]
IT	2	–	[18]

Table 3: `CIM5BL` parameters that are the same for all load classes (`VI` is the voltage limit for the undervoltage relay; `TI` is its delay time; and `TB` is the time it takes for the breaker to operate).
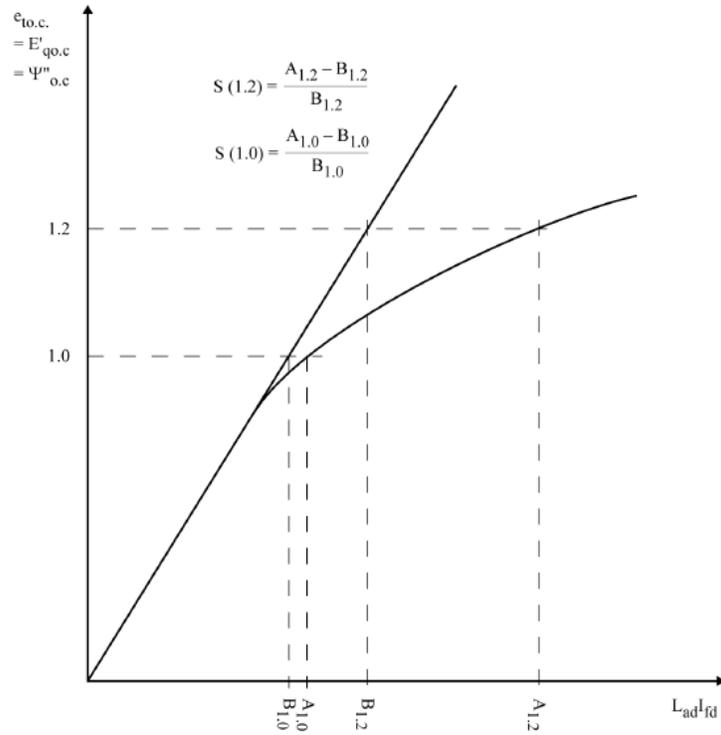
Figure 31: Saturation curve, illustrating the definition of the saturation factor S [18, fig. 15.6]

The load class of the load on each 150 kV substation is estimated based on the station's general physical location. See tables 4 and 5.

Substation	Load class
Borssele	Industrial
Goes de Poel	Residential
Kruiningen	Residential
Middelburg	Residential
Oostburg	Residential
Rilland	Residential
Terneuzen	Mixed
Vlissingen	Residential
Westdorpe	Industrial
Willem Anna Polder	Residential

Table 4: Load class estimates for Zealand

Substation	Load class
Alblasserdam	Residential
Alphen a/d Rijn	Residential
Arkel	Residential
Botlek	Industrial
De Lier	Industrial
Delft	Residential
Dordrecht Merwedehaven	Mixed
Dordrecht Noordendijk	Residential
Dordrecht Zuid	Mixed
Europoort	Industrial
Geervliet	Industrial
Geervliet Noorddijk	Industrial
Gouda	Residential
Krimpen	Residential
Leiden	Residential
Maasvlakte	Industrial
Merseyweg	Industrial
Middelharnis	Industrial
Ommoord	Residential
Oudeland	Industrial
Rijswijk	Residential
Rotterdam Centrum	Residential
Rotterdam Marconistraat	Residential
Rotterdam Waalhaven	Mixed
Rotterdam Zuidwijk	Residential
's-Gravenhage	Residential
Sassenheim	Residential
Theemsweg	Industrial
Tinte	Industrial
Vondelingenweg	Industrial
Voorburg	Residential
Wateringen	Residential
Westerlee	Industrial
Ypenburg	Residential
Zoetermeer	Residential

Table 5: Load class estimates for South Holland

The possible entries for motor load in the configuration label are now:

- **M0** − static load.

- **M1** − 10% of load MVA modelled as `CIM5BL`, with parameters as in tables 2 and 3, and load classes as in tables 4 and 5.

- **M2** − 20% of load modelled as `CIM5BL`.

- **M3** − 30% of load modelled as `CIM5BL`.

...and so on until **M9**.

4.6.2 Wind farms

The parameters for the generic models described in sections 2.2.2 and 2.2.3 vary from manufacturer to manufacturer. Having no detailed information about the wind turbines actually present in the Dutch grid, we decided to use the default parameters provided by the WECC [26]. Although they produce no accurate model of any specific wind turbine, we assume they are within the range of realistic values, and may thus lead to a good rough impression of the effects of wind turbine modelling.

The parameters for Type 3 wind turbines are given in tables 6 through 9. Those for Type 4 wind turbines are given in tables 10 and 11.

Name in PSSe	Description	Value	Unit
XEQ	Equivalent reactance for current injection	0.8	p.u.
KPLL	PLL gain	0.0	–
KIPLL	PLL integrator gain	0.0	–
PLLMAX	PLL max. limit	0.1	–
PRATED	Rated power	1.5	MW

Table 6: Default parameters for the WT3G1 model

Name in PSSe	Description	Value	Unit
TFV	Filter time constant in voltage regulator	0.15	seconds
KPV	Proportional gain in voltage regulator	18.0	–
KIV	Integrator gain in voltage regulator	5.0	–
XC	Line drop compensation reactance (pu)	0.0	p.u.
TFP	Filter time constant in torque regulator	0.05	seconds
KPP	Proportional gain in torque regulator	3.0	–
KIP	Integrator gain in torque regulator	0.6	–
PMX	Max limit in torque regulator	1.12	p.u.
PMN	Min limit in torque regulator	0.04	p.u.
QMX	Max limit in voltage regulator	0.436	p.u.
QMN	Min limit in voltage regulator	-0.436	p.u.
IPMAX	Max active current limit	1.1	p.u.
TRV	Voltage sensor time constant	0.02	seconds
RPMX	Max power order derivative	0.45	p.u./second
RPMN	Min power order derivative	-0.45	p.u./second
T_POWER	Power filter time constant	5.0	seconds
KQI	MVAR/Voltage gain	0.1	p.u.
VMINCL	Min voltage limit	0.9	p.u.
VMAXCL	Max voltage limit	1.1	p.u.
KQV	Voltage/MVAR gain	40.0	Mvar/p.u.
XIQMIN	Min reactive power order	0.5	p.u.
XIQMAX	Max reactive power order	1.45	p.u.
TV	Lag time constant in WindVar controller	0.05	seconds
TP	P_{elec} filter in fast PF controller	0.05	seconds
FN	A portion of online wind turbines	1.0	–
OMEGAPMIN	Shaft speed at Pmin	0.3	p.u.
OMEGAP20	Shaft speed at 20% rated power	0.69	p.u.
OMEGAP40	Shaft speed at 40% rated power	0.78	p.u.
OMEGAP60	Shaft speed at 60% rated power	0.98	p.u.
PMIN	Minimum power for operating at OMEGAP100 speed	0.74	p.u.
OMEGAP100	Shaft speed at 100% rated power	1.2	p.u.
VARFLG	Reactive power control mode	1	–
VLTFLG	Terminal voltage control mode	2	–

Table 7: Default parameters for the WT3E1 model

Name in PSSe	Description	Value	Unit
VW	Initial wind speed	1.25	p.u.
H	Total inertia constant	4.95	seconds
DAMP	Machine damping factor	0.0	p.u.
KAERO	Aerodynamic gain factor	0.007	–
THETA2	Blade pitch at twice rated wind speed	21.98	degrees
HTFRAC	Turbine inertia fraction (Hturb/H)	0.0	–
FREQ1	First shaft torsional resonant frequency	1.8	Hz
DSHAFT	Shaft damping factor	1.5	–

Table 8: Default parameters for the WT3T1 model

Name in PSSe	Description	Value	Unit
TP	Blade response time constant	0.3	seconds
KPP	Proportional gain of PI regulator	150	–
KIP	Integrator gain of PI regulator	25	–
KPC	Proportional gain of the compensator	3	–
KIC	Integrator gain of the compensator	30	–
TETAMIN	Lower pitch angle limit	27	degrees
TETAMAX	Upper pitch angle limit	0	degrees
RTETAMAX	Upper pitch angle rate limit	10	degrees/second
PMX	Power reference on MBASE	1	p.u.

Table 9: Default parameters for the WT3P1 model

Name in PSSe	Description	Value	Unit
TIQCMD	Converter time constant for I_{Qcmd}	0.02	seconds
TIPCMD	Converter time constant for I_{Pcmd}	0.02	seconds
VLVPL1	LVPL voltage 1	0.4	p.u.
VLVPL2	LVPL voltage 2	0.9	p.u.
GLVPL	LVPL gain	1.11	–
VHVRCR	HVRCR voltage	1.2	p.u.
CURHVRCR	Max. reactive current at VHVRCR	2.0	p.u.
RIP_LVPL	Rate of active current change	2.0	p.u./second
T_LVPL	Voltage sensor for LVPL	0.02	seconds

Table 10: Default parameters for the WT4G1 model

Name in PSSe	Description	Value	Unit
TFV	Filter time constant in voltage regulator	0.15	seconds
KPV	Proportional gain in voltage regulator	18.0	–
KIV	Integrator gain in voltage regulator	5.0	–
KPP	Proportional gain in active power regulator	0.05	–
KIP	Integrator gain in active power regulator	0.01	–
KF	Rate feedback gain	0.0	–
TF	Rate feedback time constant	0.08	seconds
QMX	Max limit in voltage regulator	0.47	p.u.
QMN	Min limit in voltage regulator	-0.47	p.u.
IPMAX	Max active current limit	1.1	p.u.
TRV	Voltage sensor time constant	0.0	seconds
DPMX	Max limit in power PI controller	0.5	p.u.
DPMN	Min limit in power PI controller	-0.5	p.u.
T_POWER	Power filter time constant	0.05	seconds
KQI	MVAR/Voltage gain	0.1	p.u.
VMINCL	Min voltage limit	0.9	p.u.
VMAXCL	Max voltage limit	1.1	p.u.
KVI	Voltage/MVAR gain	120.0	Mvar/p.u.
TV	Lag time constant in WindVar controller	0.05	seconds
TP	P_{elec} filter time constant in fast PF controller	0.05	seconds
IMAXTD	Converter current limit	1.7	p.u.
IPHL	Hard active current limit	1.11	p.u.
IQHL	Hard reactive current limit	1.11	p.u.
PFAFLG	PF control mode	0	–
VARFLG	Reactive power control mode	1	–
PQFLAG	P/Q priority flag	0	–

Table 11: Default parameters for the WT4E1 model

We now have the following possible entries for wind generation in the configuration label:

- **W0** – static: wind generation modelled as static negative load.

- **W3** – Type 3 model applied with parameters as in tables 6 through 9.

- **W4** – Type 4 model applied with parameters as in tables 10 and 11.

- **W5** – even split: half of installed wind power modelled as Type 3, the other half as Type 4.

(The numbers 1 and 2 were deliberately skipped in order to avoid the suggestion of Type 1 or Type 2 models.)

4.6.3   Solar PV

What we said about wind-turbine parameters in section 4.6.2 applies equally to PV. The parameters in tables 12 and 13 are WECC defaults [31] with some additions by the NERC [46].

An exception had to be made for certain `REECB1` parameters. Recall from section 2.3 that `REECB1` sums reactive current commands from *a)* a control block that tracks Q or φ from the power-flow solution and *b)* a control block that tracks the bus voltage. In the default parameter set, the voltage regulator is disabled entirely. Since the steady-state Q injection of wind and PV generators – and hence φ – in the power-flow solution is always zero, without the voltage regulator no reactive current would be injected at all. Hence, we have chosen off-default values for the parameters `VDIP`, `VUP`, `DBD1`, `DBD2` and `KQV`.

Name in PSSe	Description	Value	Unit
TG	Converter time constant	0.02	seconds
RRPWR	Low Voltage Power Logic (LVPL) ramp rate limit	10.0	p.u./second
BRKPT	LVPL characteristic voltage 2	0.9	p.u.
ZEROX	LVPL characteristic voltage 1	0.4	p.u.
LVPL1	LVPL gain	1.22	−
VOLIM	Voltage limit for high voltage reactive current management	1.2	p.u.
LVPNT1	High voltage point for low voltage active current management	0.8	p.u.
LVPNT0	Low voltage point for low voltage active current management	0.4	p.u.
IOLIM	Current limit for high voltage reactive current management	-1.3	p.u.
TFLTR	Voltage filter time constant for low voltage active current management	0.02	seconds
KHV	Overvoltage compensation gain used in HV reactive current management	0.7	−
IQRMAX	Upper limit on rate of change for reactive current	99.0	p.u./second
IQRMIN	Lower limit on rate of change for reactive current	99.0	p.u./second
ACCEL	Acceleration factor	0.7	−
LVPLSW	Low voltage power logic switch	1	−

Table 12: Default parameters for the `RECGA1` model

Name in PSSe	Description	Value	Unit
VDIP	Low voltage threshold to activate reactive current injection logic	0.9	p.u.
VUP	Voltage above which reactive current injection logic is activated	1.1	p.u.
TRV	Voltage filter time constant	0.02	seconds
DBD1	Voltage error dead band lower threshold	-0.01	p.u.
DBD2	Voltage error dead band upper threshold	0.01	p.u.
KQV	Reactive current injection gain during over and undervoltage conditions	10.0	−
IQH1	Upper limit on reactive current injection I_{qinj}	1.1	p.u.
IQL1	Lower limit on reactive current injection I_{qinj}	-1.1	p.u.
VREF0	User defined voltage reference	1.0	p.u.
TP	Filter time constant for electrical power	0.02	seconds
QMAX	Max. limit for reactive power regulator	0.4	p.u.
QMIN	Min. limit for reactive power regulator	-0.4	p.u.
VMAX	Max. limit for voltage control	1.1	p.u.
VMIN	Min. limit for voltage control	0.9	p.u.
KQP	Reactive power regulator proportional gain	0.0	−
KQI	Reactive power regulator integral gain	1.0	−
KVP	Voltage regulator proportional gain	0.0	−
KVI	Voltage regulator integral gain	1.0	−
TIQ	Time constant on delay s_4	0.02	seconds
DPMAX	Power reference max. ramp rate	99.0	p.u./second
DPMIN	Power reference min. ramp rate	-99.0	p.u./second
PMAX	Max. power limit	1.0	p.u.
PMIN	Min. power limit	0.0	p.u.
IMAX	Max. limit on total converter current	1.1	p.u.
TPORD	Power filter time constant	0.05	seconds
PFFLAG	PF or Q_{ext} input flag	1	−
VFLAG	Voltage or Q control flag	1	−
QFLAG	Voltage or PF control flag	0	−
PQFLAG	P/Q priority flag	1	−

Table 13: Parameters used for the `REECB1` model

For PV, we have these possible entries in the configuration label:

- **P0** − static: PV generation modelled as static negative load.

- **P1** − `REGCA` and `REECB` models applied with parameters as in tables 12 and 13, respectively.

4.6.4 HVDC transmission

A PSSe model developed specifically for the BritNed HVDC link exists: `SIEBNC` [42]. However, it is far too complicated for meaningful analysis within the scope of this thesis. Requiring a simpler, generic model, we decided on `CDC4T` − a "pseudo-steady state" model as described in section 2.4.4.

Unfortunately, BritNed parameters for `CDC4T` (or any other generic PSSe model) were not available. After considering parameter sets used in other studies, such as Stray [47] and El Chehaly [48], we eventually chose to use the parameter set from the `CDC4T` example use case in the PSSe documentation [18, § 19.2.6].

One remaining complication was that some of the `CDC4T` parameters must be entered in absolute units (ampère, kilovolt, etc.) rather than in per unit; given that the ratings of the PSSe example system differed significantly from those of BritNed, simply copying those values would not have been realistic. We decided to "scale" these absolute-unit parameters to the BritNed ratings according to table 14.

Rating	Value (PSSe HVDC example)	Value (BritNed)	Conversion factor
DC voltage [kV]	525	450	0.857
Power per pole [MW]	1500	513	0.342
Current per pole [A]	2857	1140	0.399

Table 14: Ratings of the PSSe example system for `CDC4T` [18, § 19.2.6] compared to those of the actual BritNed link [42]

This leads us to the following parameter set for use in the `CDC4T` model:

Name in PSSe	Description	Value	Unit
`ALFDY`	Minimum α for dynamics	5	degrees
`GAMDY`	Minimum γ for dynamics	15	degrees
`TVDC`	DC voltage transducer time constant	0.05	seconds
`TIDC`	Direct current transducer time constant	0.05	seconds
`VBLOCK`	Rectifier AC blocking voltage	0.6	p.u.
`VUNBL`	Rectifier AC unblocking voltage	0.65	p.u.
`TBLOCK`	Minimum blocking time	0.1	seconds
`VBYPAS`	Inverter DC bypassing voltage	0.6	p.u.
`VUNBY`	Inverter AC unbypassing voltage	0.65	p.u.
`TBYPAS`	Minimum bypassing time	0.1	seconds
`RSVOLT`	Minimum DC voltage following block	171	kV
`RSCUR`	Minimum direct current following block	380	A
`VRAMP`	Voltage recovery rate	5	p.u./second
`CRAMP`	Current recovery rate	5	p.u./second
`C0`	Minimum current demand	160	A
`V1`	Voltage limit point 1	257	kV
`C1`	Current limit point 1	399	A
`V2`	Voltage limit point 2	429	kV
`C2`	Current limit point 2	1197	A
`V3`	Voltage limit point 3	429	kV
`C3`	Current limit point 3	1197	A
`TCMODE`	Minimum time in switched mode	0.1	seconds

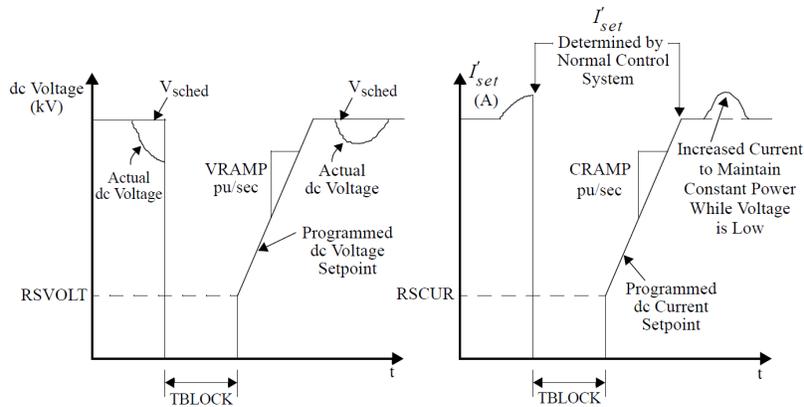Table 15: Parameters used in the `CDC4T` model in this study

Figure 32: Graph illustrating the meaning of the parameters `RSVOLT`, `RSCUR`, `VRAMP` and `CRAMP` [18, fig. 19.6a]
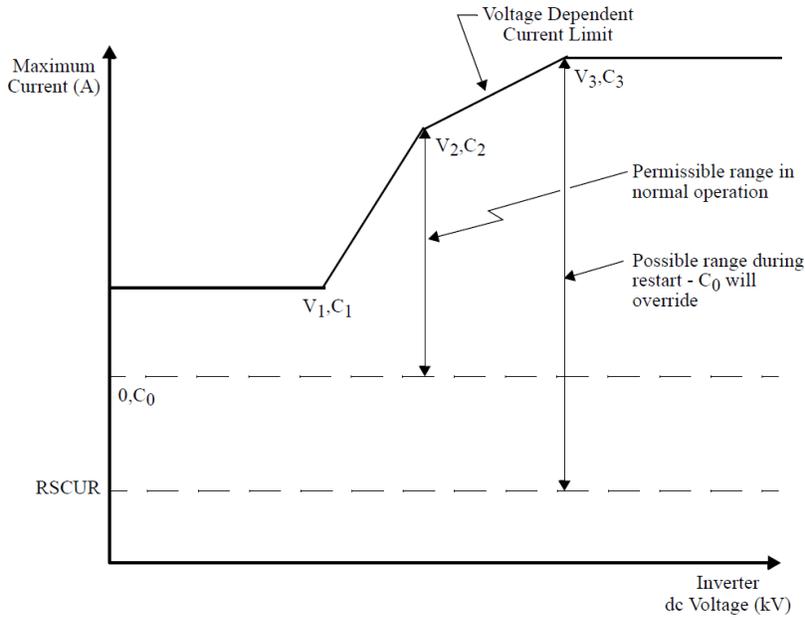
Figure 33: Graph illustrating the meaning of the parameters C0, V1, C1, V2, C2, V3 and C3 [18, fig. 19.7]

The possible entries for HVDC in the configuration label are thus:

- **D0** – static: BritNed poles modelled as static loads (active part constant-current, reactive part constant-impedance).

- **D1** – CDC4T model applied with parameters as in table 15.

As mentioned in section 2.4.4, we would also like to investigate the effect of different values of VRAMP and CRAMP. Please see section 5.4.2 for this.

4.7   Output

4.7.1   Selection of output variables for monitoring

We choose the following variables to monitor in our simulations.

- We are interested first and foremost in bus voltage magnitudes, as these give the clearest picture of the system's response at a glance. We monitor the voltage at all 150 kV buses.

- The rotor angles of synchronous generators also provide useful information: how do these machines cope with the disturbance?

- For a more detailed look at disturbance response, the slip of the motor loads we have modelled is also an interesting variable. It allows us a closer look at how motor loads affect the system's behaviour. The reactive power consumption of loads is also important: do we see loads consume more reactive power after a disturbance when motor modelling is applied, as we would expect?

- To analyse how wind and PV modelling influences stability, the active and reactive power production or consumption of wind machines and PV installations are essential monitoring variables. Do these machines provide reactive power support (as we would expect from Type 3 wind turbines and Type 4-modelled PV) or slow down voltage recovery by consuming reactive power (as we would expect from Type 1 wind turbines)?

- Active and reactive power flow into the BritNed HVDC terminal are important variables to monitor if we want to study the effects of HVDC modelling and of different ramping rates (see section 2.4.4).

4.7.2   Classification of results

Perhaps the most crucial metric in assessing the system's dynamic performance for one simulation run is whether or not the synchronous generators pull out of step (i.e. lose synchronism). PSSe models loss of synchronism as an ever-increasing rotor angle, reaching tens of thousands of degrees by the end of a simulation run. We can thus simply check the generators' rotor angles by the end of the simulation against a threshold (we choose 360°) to determine whether or not the generators have pulled out of step.

Once loss of synchronism has happened (such as in figure 34) the simulation is no longer valid, as the protection systems that disconnect out-of-step generators are not modelled. We therefore cannot further classify simulations where loss of synchronism occurs in any meaningful way – nor is that very interesting from a TSO's point of view.
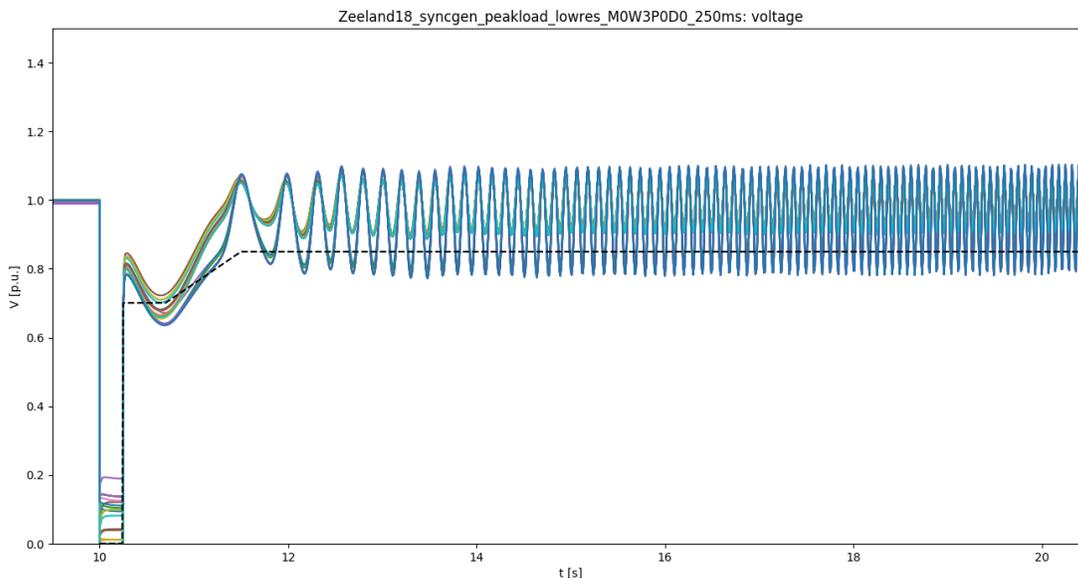
Figure 34: Bus voltage profiles in a simulation where the generators pull out of step

Simulations where the synchronous generators do remain in step can be classified in more sophisticated ways. An important metric in this regard is the voltage envelope defined in the European Union network code [49]. Synchronous generators are required by the network code to remain on-line as long as the voltage at their bus clears (remains above) this envelope.
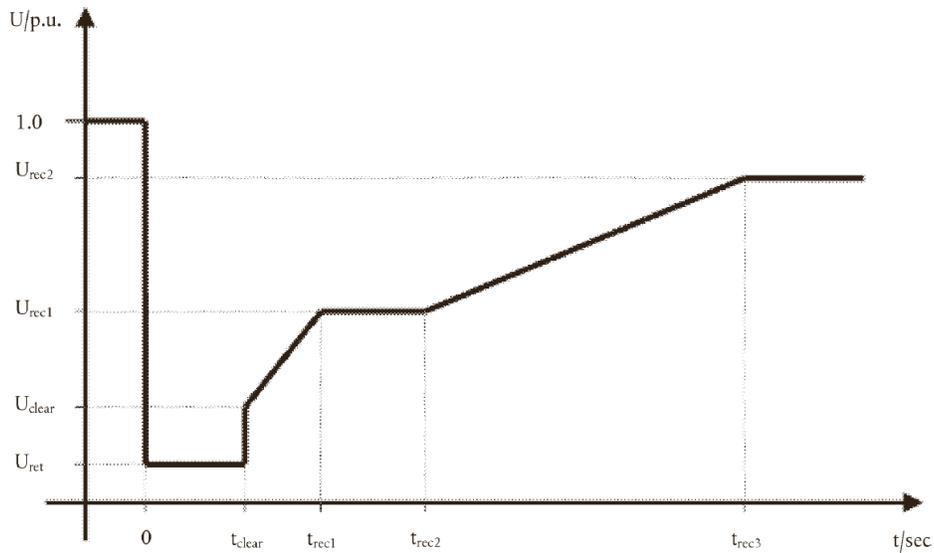
Figure 35: Voltage envelope used for fault-ride-through requirements in the EU network code [49, fig. 3]

Parameter	Value	Unit
U_{ret}	0.00	p.u.
U_{clear}	0.7	p.u.
U_{rec1}	0.7	p.u.
U_{rec2}	0.85	p.u.
t_{clear}	–	seconds
t_{rec1}	t_{clear}	seconds
t_{rec2}	0.7	seconds
t_{rec3}	1.5	seconds

Table 16: Parameters used for the envelope in figure 35 (t_{clear} varies between simulation runs) [49, tab. 3.1]

Our metric, then, is: do all 150 kV bus voltages clear the envelope, as defined in figure 35 and table 16?

To be able to detect more gradual variations than we can with such a blunt true/false test, two other metrics are used. If all bus voltages clear the envelope, we are interested in the lowest point reached by any bus voltage in the first 3 seconds after fault clearing; see figure 37. If not, the severity of the transgression can be expressed by the p.u.-second integral under the envelope as shown in figure 36 (we consider only the bus voltage with the "worst," i.e. largest, integral).
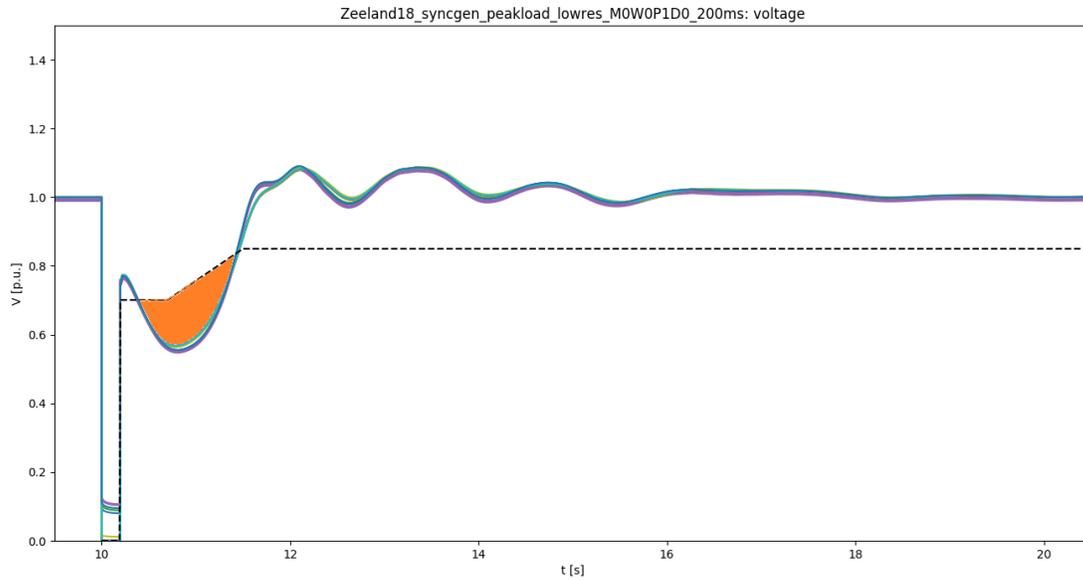
Figure 36: Illustration of the "envelope integral" metric in a bus voltage plot
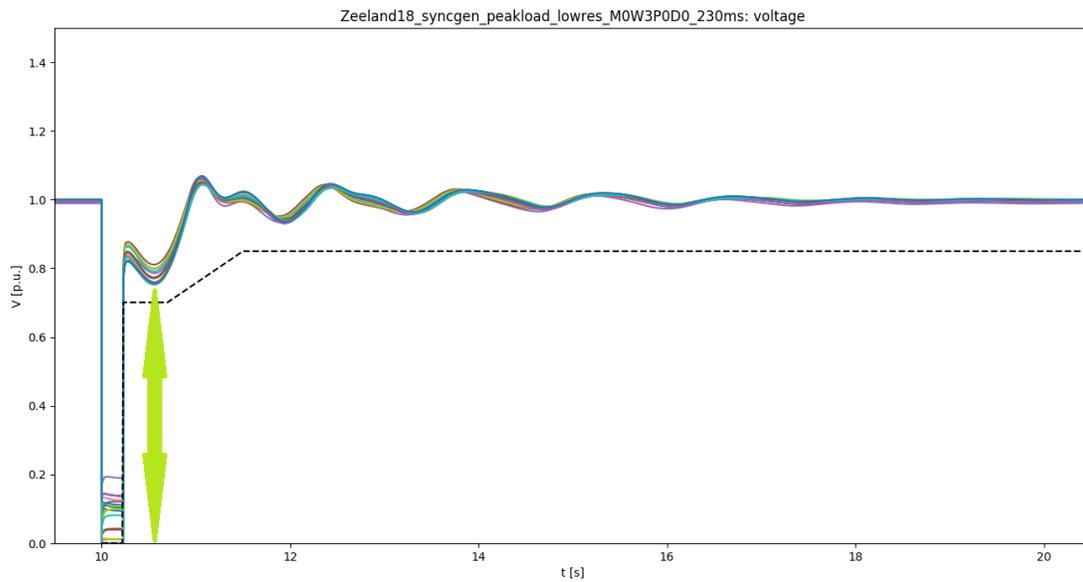
Figure 37: Illustration of the "lowest point" metric in a bus voltage plot

Note that by focusing on the most critical bus in all metrics, we lose any information about bus-to-bus variation. However, this is an acceptable loss as in all simulations we have performed, the voltage profiles at different 150 kV buses are extremely similar from the moment of fault clearing on – see figures 34, 36 and 37, or any of the voltage plots in chapter 5.

To summarise our classification of results:

- If synchronous generators pull out of step: no further classification.

- If synchronous generators remain in step:

  - If any bus voltage dips below the voltage envelope: look at the envelope integral.
  - If all bus voltages clear the voltage envelope: look at the lowest voltage dip within 3 seconds after fault clearing.

In the next chapter, we shall describe the outcomes we found for different scenarios and configurations using these metrics.

Chapter 5

Results

5.1 Overview of simulations

The simulations performed for this thesis can be divided into a "main sweep," in which a large number of different scenarios and configurations are considered; and several "side sweeps," in which the effect of a single parameter is investigated more closely while keeping all other parameters constant.

5.1.1 Main sweep

The goal of this study being to investigate the effect of different modelling configurations, it is obvious we should sweep over these. All possible combinations as mentioned in section 4.6 are considered, with an unfortunate exception. Simulations with motor load modelling very often "crashed" (in the sense of becoming numerically unstable), producing no meaningful results. For discussion of this problem, what we think caused it and how we dealt with it, please see section 5.2.2; for results of a limited number of simulations *with* motor load modelling, see section 5.4.1. The consequence is that **no configurations with motor load modelling** – that is, no configurations other than **M0 – are included in the main sweep.** Looking at section 4.6, this leaves the following configurations:

M0W0P0D0	M0W3P0D0	M0W4P0D0	M0W5P0D0
M0W0P0D1	M0W3P0D1	M0W4P0D1	M0W5P0D1
M0W0P1D0	M0W3P1D0	M0W4P1D0	M0W5P1D0
M0W0P1D1	M0W3P1D1	M0W4P1D1	M0W5P1D1

Table 17: Modelling configurations used in the main sweep (for Zealand scenarios, the **D1** configurations are omitted, as Zealand has no HVDC terminal)

We consider different scenarios, to be sure that any patterns we detect hold up across different grid situations. The scenarios are the sixteen established in section 4.4, repeated here for reference:

- Zealand or South Holland (2);
- 2018 or 2035 (4);
- low or peak load (8);
- low or peak RES (16).

The Zealand 2035 scenarios with peak RES infeed were found not to converge due to too great a surplus of active power. These were omitted, bringing the number of scenarios to fourteen.

The disturbance parameters, finally, are as mentioned in section 4.5: a three-phase fault on one circuit of Borssele-Ellewoutsdijk (Zealand) or Crayestein-Simonshaven (South Holland). Fault clearing time is varied from 100 to 250 ms, in steps of 10 ms.

5.1.2 Side sweeps

Besides the main sweep, more simulations were performed. These can be categorised by their purpose:

- *Assumption verification:* To ensure we do not overlook important information because of certain assumptions about scenarios and disturbance configurations, we ran a limited number of simulations with different assumptions, to show the results were not fundamentally different. See section 5.3.
    - Scenarios with only renewable generation, i.e. all synchronous generators out of service. See section 5.3.1.
    - An N-1 scenario, i.e. with a major line out of service before the disturbance is applied. See section 5.3.3.
    - Scenarios with a weaker (i.e. higher-impedance) link to the external grid. See section 5.3.2.
    - Different disturbance locations and types. See section 5.3.4.

- *More detailed investigation:* To provide a closer insight into the effects of modelling differences, we ran a finer-grained sweeps over certain modelling parameters (including these in the main sweep, which considers fourteen different scenarios and *all* possible combinations of the modelling parameters included, would have implied an impractically large number of simulations). See section 5.4.
    - A sweep over the percentage of load modelled as motor load. See section 5.4.1.
    - A sweep over `VRAMP` and `CRAMP` values of the HVDC model, as announced in section 4.6.4. See section 5.4.2.
    - A sweep over the ratio of Type 3 to Type 4 wind generation in the network. See section 5.4.3.

5.2 Main sweep: effect of modelling on CCT

5.2.1 Overview of results

The tables in this section have been generated according to the classification in section 4.7, and should be read as such. To summarise:

- A red 'X' means the synchronous generators pulled out of step.

- An orange number means the synchronous generators remained in step, but one or more bus voltages dipped below the envelope defined in section 4.7. The number is the largest under-the-envelope integral of any bus voltage.

- A green number means the synchronous generators remained in step and all bus voltages cleared the envelope. The number is the lowest point reached by any bus voltage in the first 3 seconds after fault clearing.

- A black dash means the simulation became numerically unstable and had to be omitted.

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.86	0.84	0.82	0.79	0.75	0.71	0.01	0.07	X	X	X	X	X	X	X	X
M0W0P1D0	0.89	0.88	0.87	0.87	0.86	0.84	0.82	0.79	0.77	0.77	0.77	0.77	0.77	0.00	0.01	X
M0W3P0D0	0.98	0.98	0.97	0.97	0.95	0.94	0.93	0.92	0.90	0.89	0.86	0.83	0.79	0.75	0.00	X
M0W3P1D0	0.98	0.98	0.98	0.97	0.96	0.95	0.94	0.93	0.92	0.90	0.89	0.87	0.87	0.83	0.82	0.80
M0W4P0D0	0.95	0.95	0.94	0.93	0.92	0.91	0.90	0.88	0.86	0.83	0.80	0.76	0.71	0.01	X	X
M0W4P1D0	0.90	0.90	0.89	0.89	0.88	0.87	0.86	0.86	0.85	0.84	0.83	0.82	0.79	0.77	0.74	0.74
M0W5P0D0	0.98	0.98	0.97	0.97	0.96	0.96	0.95	0.95	0.94	0.93	0.92	0.91	0.90	0.88	0.86	0.83
M0W5P1D0	0.95	0.95	0.94	0.94	0.94	0.94	0.93	0.92	0.91	0.90	0.89	0.88	0.87	0.86	0.85	0.84

Table 18: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2018, low load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.84	0.82	0.79	0.75	0.71	0.01	0.07	X	X	X	X	X	X	X	X	X
M0W0P1D0	0.88	0.87	0.86	0.85	0.84	0.82	0.79	0.77	0.76	0.77	0.77	0.00	0.01	0.02	X	X
M0W3P0D0	0.98	0.98	0.98	0.97	0.97	0.97	0.96	0.95	0.93	0.92	0.90	0.87	0.84	0.80	0.76	0.71
M0W3P1D0	0.98	0.98	0.98	0.98	0.97	0.97	0.96	0.95	0.94	0.92	0.91	0.89	0.87	0.85	0.82	0.81
M0W4P0D0	0.96	0.95	0.94	0.93	0.93	0.91	0.90	0.89	0.86	0.84	0.81	0.77	0.73	0.00	X	X
M0W4P1D0	0.91	0.91	0.90	0.90	0.89	0.89	0.88	0.87	0.86	0.85	0.84	0.82	0.79	0.77	0.75	0.74
M0W5P0D0	0.98	0.97	0.97	0.97	0.96	0.96	0.95	0.95	0.94	0.94	0.93	0.92	0.90	0.89	0.87	0.84
M0W5P1D0	0.95	0.95	0.95	0.94	0.94	0.94	0.93	0.92	0.92	0.91	0.90	0.89	0.88	0.87	0.86	0.85

Table 19: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2018, low load, peak RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.88	0.87	0.86	0.85	0.82	0.80	0.77	0.73	0.00	0.02	0.13	X	X	X	X	X
M0W0P1D0	0.91	0.90	0.89	0.88	0.88	0.87	0.86	0.85	0.83	0.80	0.78	0.77	0.78	0.79	0.00	0.77
M0W3P0D0	0.98	0.98	0.97	0.97	0.97	0.97	0.96	0.95	0.93	0.92	0.91	0.89	0.86	0.83	0.80	0.75
M0W3P1D0	0.98	0.98	0.97	0.97	0.97	0.97	0.96	0.95	0.94	0.93	0.92	0.90	0.89	0.87	0.87	0.83
M0W4P0D0	0.96	0.96	0.96	0.95	0.94	0.93	0.92	0.91	0.90	0.89	0.86	0.84	0.81	0.78	0.73	0.00
M0W4P1D0	0.90	0.89	0.89	0.90	0.89	0.88	0.88	0.88	0.87	0.87	0.86	0.85	0.84	0.83	0.81	0.78
M0W5P0D0	0.98	0.98	0.98	0.98	0.97	0.97	0.97	0.96	0.96	0.95	0.94	0.94	0.93	0.92	0.90	0.88
M0W5P1D0	0.95	0.96	0.95	0.95	0.95	0.95	0.94	0.93	0.93	0.92	0.91	0.90	0.89	0.88	0.87	0.86

Table 20: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2018, peak load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.88	0.87	0.85	0.82	0.80	0.77	0.73	0.00	0.02	0.09	X	X	X	X	X	X
M0W0P1D0	0.89	0.88	0.88	0.87	0.87	0.86	0.85	0.82	0.80	0.78	0.78	0.77	0.77	0.00	0.00	0.01
M0W3P0D0	0.98	0.98	0.98	0.97	0.97	0.97	0.97	0.96	0.95	0.94	0.92	0.90	0.88	0.85	0.82	0.78
M0W3P1D0	0.98	0.98	0.98	0.97	0.97	0.97	0.97	0.96	0.96	0.94	0.93	0.92	0.90	0.88	0.87	0.87
M0W4P0D0	0.96	0.96	0.95	0.95	0.94	0.93	0.92	0.91	0.90	0.88	0.86	0.83	0.80	0.77	0.73	0.01
M0W4P1D0	0.91	0.91	0.90	0.90	0.90	0.90	0.89	0.89	0.88	0.87	0.86	0.85	0.84	0.82	0.79	0.77
M0W5P0D0	0.98	0.98	0.98	0.98	0.97	0.97	0.96	0.96	0.95	0.95	0.94	0.93	0.92	0.91	0.90	0.88
M0W5P1D0	0.96	0.96	0.95	0.95	0.95	0.95	0.94	0.94	0.93	0.92	0.91	0.90	0.89	0.89	0.88	0.87

Table 21: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2018, peak load, peak RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.80	0.76	0.72	0.01	0.06	X	X	X	X	X	X	X	X	X	–	–
M0W0P1D0	0.87	0.86	X	0.84	0.81	0.78	0.78	0.77	0.00	0.00	0.00	0.01	0.05	X	X	X
M0W3P0D0	0.97	0.96	0.96	0.95	0.95	0.95	0.94	0.94	0.93	0.91	0.89	0.87	0.84	0.81	0.77	0.73
M0W3P1D0	0.97	0.97	0.96	0.95	0.95	0.95	0.94	0.94	0.94	0.93	0.91	0.88	0.87	0.85	0.83	0.81
M0W4P0D0	0.94	0.94	0.93	0.92	0.91	0.89	0.87	0.84	0.81	0.78	0.74	0.00	0.01	0.05	X	X
M0W4P1D0	0.93	0.93	0.92	0.92	0.92	0.91	0.89	0.88	0.85	0.83	0.81	0.79	0.76	0.73	0.73	0.74
M0W5P0D0	0.96	0.96	0.95	0.95	0.95	0.95	0.94	0.94	0.94	0.93	0.93	0.92	0.91	0.91	0.89	0.87
M0W5P1D0	0.95	0.95	0.95	0.95	0.94	0.93	0.93	0.92	0.92	0.92	0.91	0.91	0.90	0.89	0.89	0.88

Table 22: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2035, low load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.86	0.84	0.82	0.79	0.76	0.72	0.00	0.03	0.12	X	X	X	X	X	X	X
M0W0P1D0	0.89	0.88	0.87	0.86	0.85	0.84	0.82	0.79	0.77	0.78	0.76	0.74	0.73	0.01	X	X
M0W3P0D0	0.98	0.97	0.97	0.97	0.96	0.96	0.96	0.96	0.95	0.95	0.93	0.91	0.89	0.86	0.83	0.80
M0W3P1D0	0.98	0.98	0.97	0.97	0.97	0.96	0.96	0.96	0.95	0.95	0.94	0.92	0.90	0.89	0.87	0.85
M0W4P0D0	0.96	0.95	0.95	0.94	0.94	0.93	0.92	0.90	0.88	0.86	0.83	0.80	0.77	0.74	0.00	0.02
M0W4P1D0	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.90	0.88	0.87	0.85	0.83	0.80	0.78	0.76
M0W5P0D0	0.98	0.97	0.97	0.97	0.97	0.96	0.96	0.96	0.96	0.95	0.95	0.94	0.94	0.93	0.92	0.91
M0W5P1D0	0.96	0.96	0.96	0.96	0.95	0.95	0.95	0.94	0.94	0.93	0.93	0.93	0.92	0.92	0.91	0.90

Table 23: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2035, peak load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.84	0.80	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
M0W0P1D0	0.89	0.88	0.86	0.84	0.81	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D1	0.89	0.87	0.86	0.84	0.81	X	X	X	X	X	X	X	X	X	X	X
M0W3P0D0	0.97	0.97	0.97	0.95	0.93	0.89	X	X	X	X	X	X	X	X	X	X
M0W3P0D1	0.97	0.97	0.96	0.95	0.93	0.90	X	X	X	X	X	X	X	X	X	X
M0W3P1D0	0.98	0.97	0.96	0.95	0.94	0.92	0.89	X	X	X	X	X	X	X	X	X
M0W3P1D1	0.97	0.97	0.96	0.95	0.94	0.93	0.90	X	X	X	X	X	X	X	X	X
M0W4P0D0	0.96	0.95	0.93	0.90	0.86	X	X	X	X	X	X	X	X	X	X	X
M0W4P0D1	0.96	0.95	0.93	0.90	0.87	X	X	X	X	X	X	X	X	X	X	X
M0W4P1D0	0.91	0.90	0.90	0.89	0.88	0.87	0.84	X	X	X	X	X	X	X	X	X
M0W4P1D1	0.92	0.91	0.91	0.90	0.89	0.88	0.84	X	X	X	X	X	X	X	X	X
M0W5P0D0	0.97	0.97	0.97	0.97	0.96	0.96	0.95	X	X	X	X	X	X	X	X	X
M0W5P0D1	0.97	0.97	0.97	0.96	0.96	0.96	0.95	X	X	X	X	X	X	X	X	X
M0W5P1D0	0.96	0.96	0.95	0.94	0.94	0.93	0.93	0.92	X	X	X	X	X	X	X	X
M0W5P1D1	0.96	0.96	0.96	0.95	0.95	0.94	0.93	0.93	X	X	X	X	X	X	X	X

Table 24: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2018, low load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.82	0.76	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
M0W0P1D0	0.88	0.86	0.85	0.82	0.79	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D1	0.88	0.86	0.84	0.81	0.78	X	X	X	X	X	X	X	X	X	X	X
M0W3P0D0	0.97	0.97	0.97	0.95	0.93	0.89	X	X	X	X	X	X	X	X	X	X
M0W3P0D1	0.97	0.97	0.96	0.95	0.93	0.90	X	X	X	X	X	X	X	X	X	X
M0W3P1D0	0.98	0.97	0.96	0.96	0.95	0.93	0.90	X	X	X	X	X	X	X	X	X
M0W3P1D1	0.97	0.97	0.96	0.95	0.94	0.93	0.90	X	X	X	X	X	X	X	X	X
M0W4P0D0	0.96	0.95	0.93	0.90	0.86	X	X	X	X	X	X	X	X	X	X	X
M0W4P0D1	0.96	0.95	0.93	0.90	0.87	X	X	X	X	X	X	X	X	X	X	X
M0W4P1D0	0.91	0.90	0.90	0.89	0.88	0.87	0.84	X	X	X	X	X	X	X	X	X
M0W4P1D1	0.92	0.91	0.91	0.90	0.89	0.87	0.84	X	X	X	X	X	X	X	X	X
M0W5P0D0	0.97	0.97	0.97	0.97	0.97	0.96	0.95	X	X	X	X	X	X	X	X	X
M0W5P0D1	0.97	0.97	0.97	0.96	0.96	0.96	0.95	X	X	X	X	X	X	X	X	X
M0W5P1D0	0.96	0.96	0.95	0.94	0.94	0.93	0.93	0.92	X	X	X	X	X	X	X	X
M0W5P1D1	0.96	0.96	0.96	0.95	0.95	0.94	0.93	0.93	X	X	X	X	X	X	X	X

Table 25: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2018, low load, peak RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.85	0.84	0.83	0.82	0.79	0.75	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	0.87	0.86	0.85	0.84	0.81	0.77	X	X	X	X	X	X	X	X	X	X
M0W0P1D0	0.90	0.89	0.88	0.87	0.86	0.84	0.83	0.80	0.79	X	X	X	X	X	X	X
M0W0P1D1	0.91	0.90	0.89	0.88	0.87	0.85	0.83	0.81	0.79	0.76	X	X	X	X	X	X
M0W3P0D0	0.95	0.95	0.95	0.94	0.93	0.92	0.91	0.88	0.85	X	X	X	X	X	X	X
M0W3P0D1	0.95	0.95	0.95	0.94	0.94	0.93	0.91	0.89	0.86	X	X	X	X	X	X	X
M0W3P1D0	0.95	0.95	0.95	0.95	0.94	0.93	0.92	0.91	0.89	0.87	0.84	X	X	X	X	X
M0W3P1D1	0.95	0.95	0.95	0.95	0.94	0.93	0.92	0.91	0.89	0.87	0.85	X	X	X	X	X
M0W4P0D0	0.93	0.93	0.92	0.91	0.90	0.89	0.87	0.85	X	X	X	X	X	X	X	X
M0W4P0D1	0.95	0.95	0.95	0.94	0.92	0.91	0.89	0.86	0.83	X	X	X	X	X	X	X
M0W4P1D0	0.89	0.89	0.88	0.88	0.87	0.87	0.86	0.85	0.85	0.83	X	X	X	X	X	X
M0W4P1D1	0.90	0.90	0.89	0.89	0.88	0.88	0.88	0.87	0.86	0.84	0.81	X	X	X	X	X
M0W5P0D0	0.95	0.95	0.95	0.95	0.95	0.95	0.95	0.94	0.93	0.91	0.88	X	X	X	X	X
M0W5P0D1	0.95	0.95	0.95	0.95	0.94	0.94	0.94	0.93	0.93	0.92	0.89	X	X	X	X	X
M0W5P1D0	0.93	0.93	0.92	0.92	0.92	0.91	0.91	0.90	0.89	0.89	0.88	X	X	X	X	X
M0W5P1D1	0.94	0.94	0.93	0.93	0.92	0.92	0.92	0.91	0.91	0.90	0.89	0.89	X	X	X	X

Table 26: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2018, peak load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.84	0.83	0.82	0.79	0.74	X	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	0.87	0.86	0.83	0.80	0.76	X	X	X	–	–	–	–	–	–	–	–
M0W0P1D0	0.89	0.88	0.87	0.86	0.85	0.82	0.81	0.79	X	X	X	X	X	X	X	X
M0W0P1D1	0.90	0.88	0.87	0.86	0.84	0.82	0.80	0.78	X	X	X	X	X	X	X	X
M0W3P0D0	0.96	0.95	0.94	0.93	0.93	0.92	0.89	0.86	X	X	X	X	X	X	X	X
M0W3P0D1	0.96	0.95	0.95	0.95	0.94	0.92	0.90	0.87	X	X	X	X	X	X	X	X
M0W3P1D0	0.96	0.96	0.95	0.95	0.94	0.93	0.91	0.89	0.87	X	X	X	X	X	X	X
M0W3P1D1	0.96	0.95	0.95	0.94	0.93	0.92	0.91	0.90	0.88	0.85	X	X	X	X	X	X
M0W4P0D0	0.92	0.91	0.91	0.90	0.89	0.87	0.84	X	X	X	X	X	X	X	X	X
M0W4P0D1	0.95	0.94	0.93	0.92	0.90	0.88	0.86	0.82	X	X	X	X	X	X	X	X
M0W4P1D0	0.89	0.88	0.88	0.88	0.87	0.86	0.87	0.85	0.82	X	X	X	X	X	X	X
M0W4P1D1	0.90	0.90	0.89	0.89	0.88	0.88	0.87	0.86	0.84	X	X	X	X	X	X	X
M0W5P0D0	0.96	0.96	0.96	0.96	0.95	0.95	0.94	0.94	0.92	X	X	X	X	X	X	X
M0W5P0D1	0.95	0.95	0.95	0.95	0.95	0.94	0.94	0.93	0.92	0.90	X	X	X	X	X	X
M0W5P1D0	0.93	0.93	0.92	0.92	0.91	0.91	0.90	0.90	0.89	0.89	X	X	X	X	X	X
M0W5P1D1	0.94	0.94	0.93	0.93	0.93	0.92	0.92	0.91	0.90	0.90	X	X	X	X	X	X

Table 27: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2018, peak load, peak RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D0	0.84	0.82	0.79	0.76	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D1	0.82	0.80	0.77	0.74	X	X	X	X	X	X	X	X	X	X	X	X
M0W3P0D0	0.97	0.96	0.96	0.96	0.95	0.93	0.89	X	X	X	X	X	X	X	X	X
M0W3P0D1	0.97	0.96	0.96	0.95	0.94	0.92	0.89	0.85	X	X	X	X	X	X	X	X
M0W3P1D0	0.97	0.97	0.97	0.96	0.95	0.94	0.92	0.89	0.86	X	X	X	X	X	X	X
M0W3P1D1	0.97	0.97	0.96	0.95	0.95	0.94	0.92	0.89	0.87	X	X	X	X	X	X	X
M0W4P0D0	0.98	0.97	0.97	0.96	0.93	0.90	0.86	X	X	X	X	X	X	X	X	X
M0W4P0D1	0.97	0.97	0.97	0.95	0.93	0.91	0.87	X	X	X	X	X	X	X	X	X
M0W4P1D0	0.93	0.92	0.91	0.91	0.90	0.90	0.89	0.86	X	X	X	X	X	X	X	X
M0W4P1D1	0.94	0.93	0.93	0.92	0.91	0.90	0.90	0.86	X	X	X	X	X	X	X	X
M0W5P0D0	0.97	0.97	0.97	0.97	0.96	0.96	0.96	0.95	0.94	X	X	X	X	X	X	X
M0W5P0D1	0.97	0.97	0.97	0.96	0.96	0.96	0.96	0.95	0.93	X	X	X	X	X	X	X
M0W5P1D0	0.96	0.96	0.96	0.95	0.95	0.94	0.94	0.93	0.92	0.92	X	X	X	X	X	X
M0W5P1D1	0.96	0.96	0.96	0.96	0.95	0.95	0.94	0.94	0.93	0.92	X	X	X	X	X	X

Table 28: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2035, low load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	–	–
M0W0P0D1	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
M0W0P1D0	0.75	0.75	0.71	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D1	0.72	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W3P0D0	0.97	0.96	0.95	0.95	0.93	0.91	0.88	0.85	0.81	X	X	X	X	X	X	X
M0W3P0D1	0.96	0.96	0.95	0.94	0.91	0.89	0.86	0.83	0.78	X	X	X	X	X	X	X
M0W3P1D0	0.97	0.97	0.96	0.96	0.95	0.93	0.91	0.89	0.88	0.85	0.82	X	X	X	X	X
M0W3P1D1	0.97	0.96	0.96	0.95	0.94	0.92	0.90	0.88	0.87	0.84	X	X	X	X	X	X
M0W4P0D0	1.00	0.99	0.97	0.96	0.94	0.92	0.89	0.88	0.87	X	X	X	X	X	X	X
M0W4P0D1	0.97	0.96	0.94	0.93	0.91	0.89	0.87	0.86	X	X	X	X	X	X	X	X
M0W4P1D0	0.94	0.94	0.95	0.95	0.93	0.92	0.91	0.90	0.89	X	X	X	X	X	X	X
M0W4P1D1	0.94	0.95	0.94	0.93	0.92	0.91	0.90	0.89	X	X	X	X	X	X	X	X
M0W5P0D0	0.97	0.96	0.96	0.96	0.96	0.96	0.95	0.95	0.95	0.93	0.90	X	X	X	X	X
M0W5P0D1	0.96	0.96	0.96	0.96	0.96	0.95	0.95	0.95	0.94	0.93	X	X	X	X	X	X
M0W5P1D0	0.97	0.97	0.96	0.96	0.95	0.95	0.94	0.93	0.93	0.92	0.91	X	X	X	X	X
M0W5P1D1	0.96	0.97	0.96	0.96	0.95	0.95	0.94	0.93	0.93	0.92	0.91	X	X	X	X	X

Table 29: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2035, low load, peak RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.83	0.81	0.78	0.73	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	0.83	0.80	0.77	0.72	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D0	0.87	0.86	0.85	0.84	0.82	0.81	0.79	0.76	X	X	X	X	X	X	X	X
M0W0P1D1	0.87	0.86	0.85	0.83	0.81	0.80	0.78	0.75	0.73	X	X	X	X	X	X	X
M0W3P0D0	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.90	0.87	0.83	X	X	X	X	X
M0W3P0D1	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.90	0.87	0.83	X	X	X	X	X
M0W3P1D0	0.94	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.91	0.89	0.87	0.84	X	X	X	X
M0W3P1D1	0.95	0.94	0.94	0.94	0.94	0.94	0.93	0.92	0.91	0.89	0.87	0.84	X	X	X	X
M0W4P0D0	0.93	0.92	0.91	0.91	0.90	0.89	0.88	0.86	0.84	0.80	X	X	X	X	X	X
M0W4P0D1	0.95	0.94	0.94	0.93	0.93	0.91	0.90	0.88	0.85	0.82	X	X	X	X	X	X
M0W4P1D0	0.89	0.89	0.89	0.89	0.89	0.88	0.88	0.87	0.86	0.84	0.81	X	X	X	X	X
M0W4P1D1	0.90	0.89	0.89	0.89	0.88	0.89	0.89	0.89	0.87	0.85	0.82	0.79	X	X	X	X
M0W5P0D0	0.95	0.95	0.95	0.95	0.95	0.95	0.94	0.94	0.94	0.93	0.92	0.91	0.88	X	X	X
M0W5P0D1	0.95	0.94	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.92	0.90	0.88	X	X	X
M0W5P1D0	0.93	0.93	0.93	0.92	0.92	0.92	0.91	0.91	0.90	0.90	0.89	0.89	0.88	X	X	X
M0W5P1D1	0.94	0.93	0.93	0.93	0.93	0.93	0.92	0.92	0.91	0.91	0.90	0.89	0.89	X	X	X

Table 30: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2035, peak load, low RES

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P0D1	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
M0W0P1D0	0.83	0.81	0.78	0.76	0.75	0.73	0.00	X	X	X	X	X	X	X	X	X
M0W0P1D1	0.80	0.78	0.75	0.74	0.73	0.00	X	X	X	X	X	X	X	X	X	X
M0W3P0D0	0.95	0.94	0.94	0.94	0.93	0.93	0.93	0.92	0.90	0.87	0.84	0.80	X	X	X	X
M0W3P0D1	0.94	0.94	0.94	0.94	0.93	0.93	0.93	0.91	0.89	0.86	0.83	0.79	X	X	X	X
M0W3P1D0	0.95	0.95	0.94	0.94	0.94	0.93	0.93	0.93	0.91	0.89	0.88	0.85	0.83	X	X	X
M0W3P1D1	0.95	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.91	0.89	0.87	0.85	0.82	X	X	X
M0W4P0D0	0.95	0.95	0.95	0.95	0.94	0.94	0.93	0.92	0.90	0.88	0.85	X	X	X	X	X
M0W4P0D1	0.94	0.94	0.94	0.94	0.93	0.93	0.93	0.92	0.91	0.89	0.86	0.82	X	X	X	X
M0W4P1D0	0.92	0.92	0.92	0.92	0.92	0.91	0.91	0.91	0.90	0.89	0.87	0.83	X	X	X	X
M0W4P1D1	0.93	0.93	0.93	0.92	0.92	0.92	0.92	0.91	0.91	0.89	0.88	0.84	X	X	X	X
M0W5P0D0	0.95	0.95	0.95	0.95	0.94	0.94	0.94	0.94	0.93	0.93	0.92	0.91	0.90	0.86	X	X
M0W5P0D1	0.94	0.94	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.93	0.92	0.91	0.89	0.86	X	X
M0W5P1D0	0.94	0.94	0.94	0.93	0.93	0.93	0.92	0.92	0.91	0.91	0.90	0.90	0.89	0.89	X	X
M0W5P1D1	0.94	0.94	0.94	0.94	0.94	0.93	0.93	0.93	0.92	0.92	0.91	0.90	0.90	0.89	X	X

Table 31: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: South Holland 2035, peak load, peak RES

5.2.2 Effect of motor load

As mentioned in section 5.1.1, adding motor load modelling all too often leads to numerical instability. As such failed simulations forcefully broke the flow of automated processing as described in section 4.3 – to the point of crashing the Python shell – and no class of scenarios could be found that was reliably "safe," we soon found that it was not practically feasible to include configurations with motor load modelling in the main sweep at all.

What can we say about the causes of the problem? One consistent pattern was that, all else being equal, higher percentages of motor load modelling – and thus more MVAs of motor load present in the system – were more likely to lead to simulation failures. The same applied to fault clearing times: for the same motor load percentage, longer fault clearing times often caused numerical instability where shorter ones had been simulated without trouble.

This feeds the hypothesis that numerical instability is simple a worse form of physical instability: that is, that the same factors which cause physical instability, such as long fault clearing times and large amounts of motor load (we had seen in earlier stages of the project that motor load considerably degrades dynamic performance because of its high reactive power consumption) will, when taken to extremes, cause numerical instability.

There are also hints that the issue might not be that straightforward, however. For one, not only the amount of motor load mattered, but also its spread throughout the system: a simulation with only 1% motor load modelling applied at all load buses became unstable, whereas in the same scenario, a simulation with a much higher total amount of motor load – but concentrated at a small number of buses – ran without trouble. It appears that the sheer number of equations to solve when more instances of CIM5BL are present in the system is also a risk factor for instability.

Another hypothesis is that the discontinuities introduced by CIM5BL's undervoltage relays – which may trip a large number of motors all at once – cause the problem. However, we have neither the diagnostic tools nor the mathematical expertise to properly test any of these hypotheses, so for now they remain confined to the realm of conjecture.

The side sweep in section 5.4.1 produced results showing the effects of motor load modelling in a single scenario.

5.2.3 Effect of wind modelling

In our results we can clearly see that wind modelling is a strong stabilising force. This is the most obvious in the South Holland 2035 scenarios with peak RES infeed, because of the presence of a very large offshore wind farm at Maasvlakte. However, even in the 2018 scenarios with their relatively low wind penetration, even in scenarios with low RES infeed, the difference between static and dynamic modelling of wind generation is clear. Looking at the tables, we can see that:

- Type 3 and Type 4 modelling both show much better dynamic performance than a static model;

- Type 4 performs marginally better than Type 3;

- a 50-50 mix of both types (the **W5** configurations) performs significantly better than either type on its own.

This "boost" in dynamic performance is caused by very rapidly responding voltage control functions in the wind turbine models. Compare figure 38 (static) to figure 39 (Type 3) or 40 (Type 4).
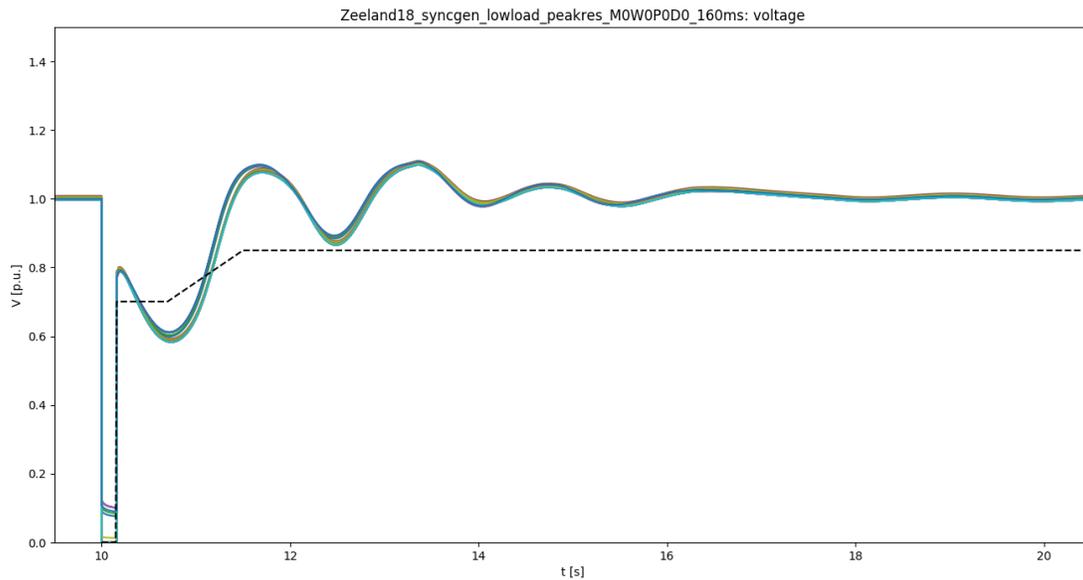
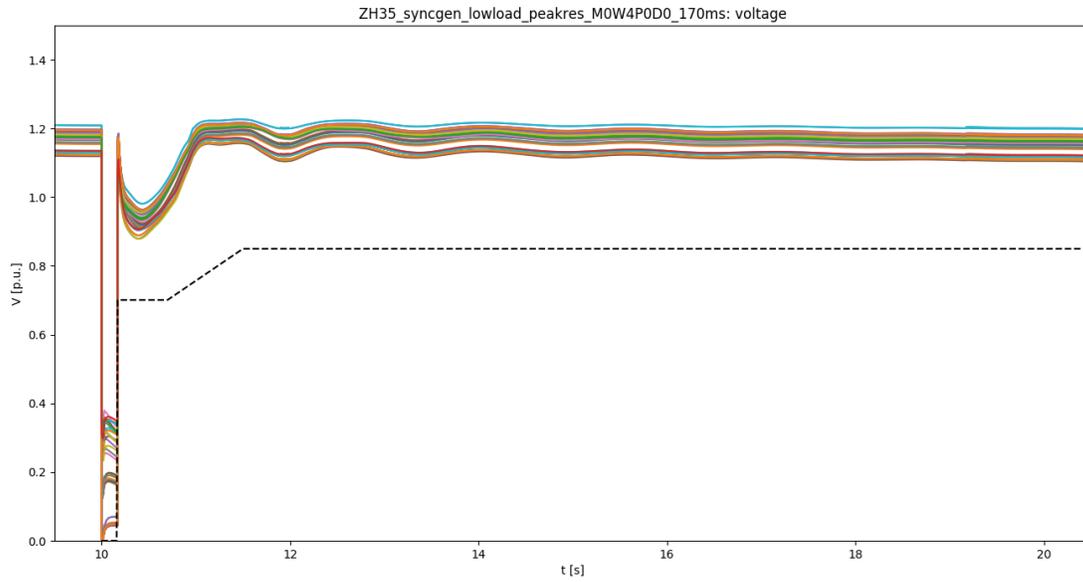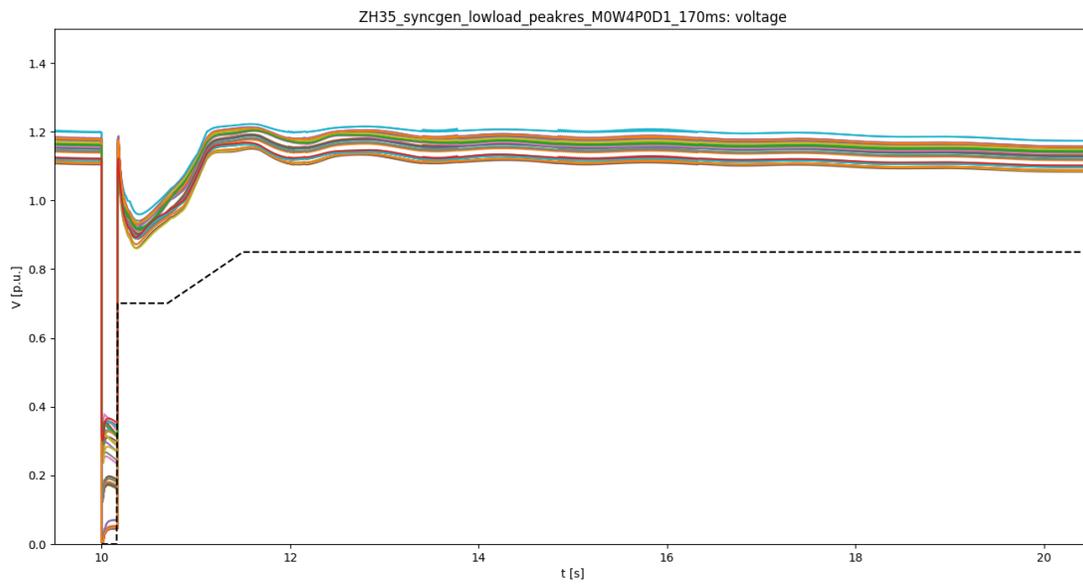Figure 38: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, low load, peak RES; configuration: M0W0P0D0
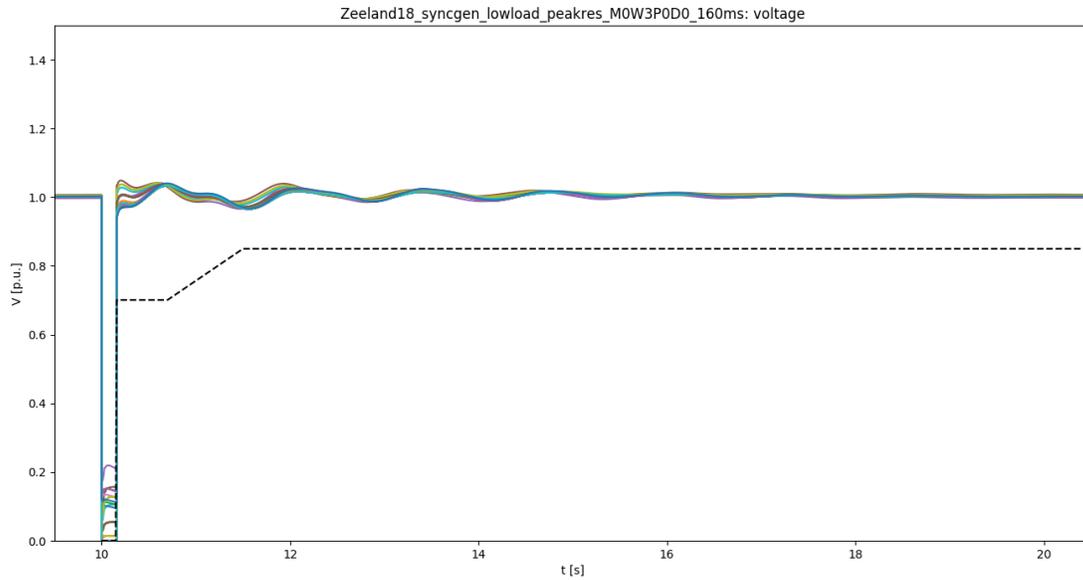
Figure 39: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, low load, peak RES; configuration: M0W3P0D0
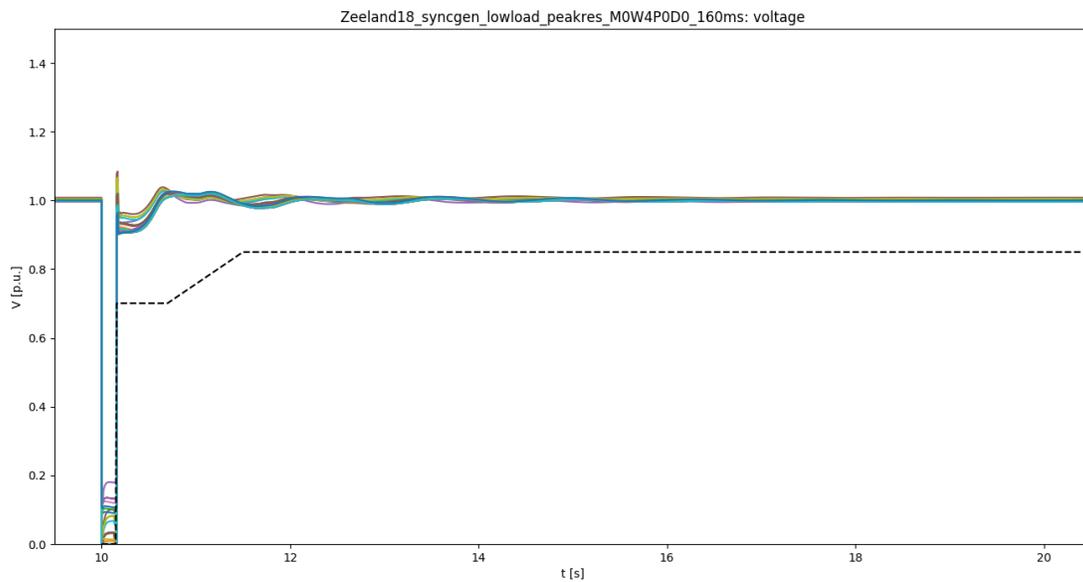
Figure 40: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, low load, peak RES; configuration: M0W4P0D0

The corresponding plots of wind turbines' reactive-power output are shown in figures 41 and 42. Notice how the voltage peak immediately after fault clearing in figure 40 corresponds to a Q peak in 42.

The heavy oscillations in reactive-power output seem to be an obvious downside of these aggressive controller settings, but they have little noticeable effect on the bus voltages.
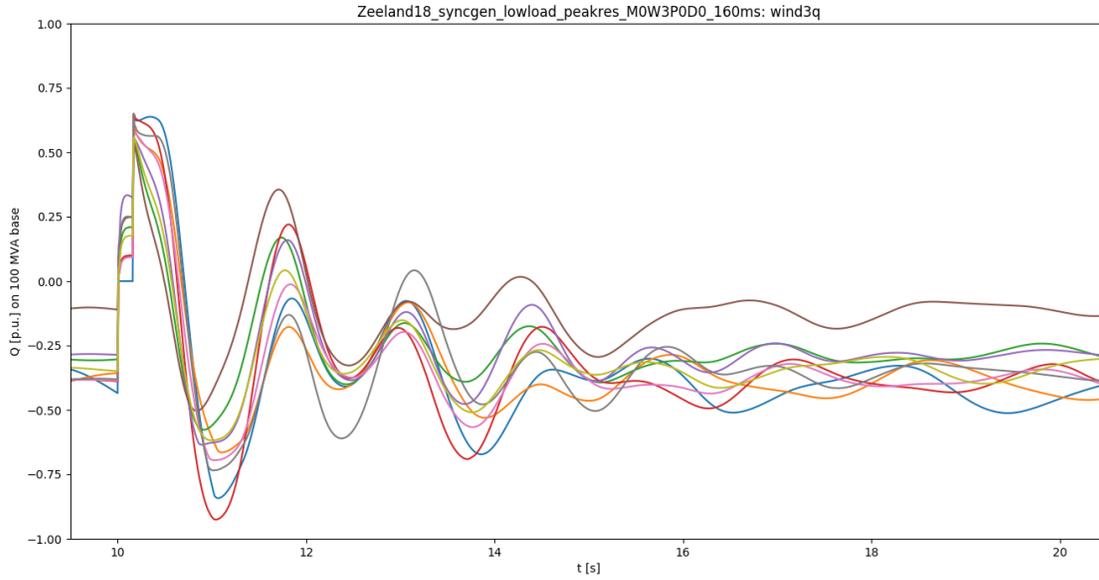
Figure 41: Wind turbine reactive power output for a 160 ms clearing time. Scenario: Zealand 2018, low load, peak RES; configuration: M0W3P0D0
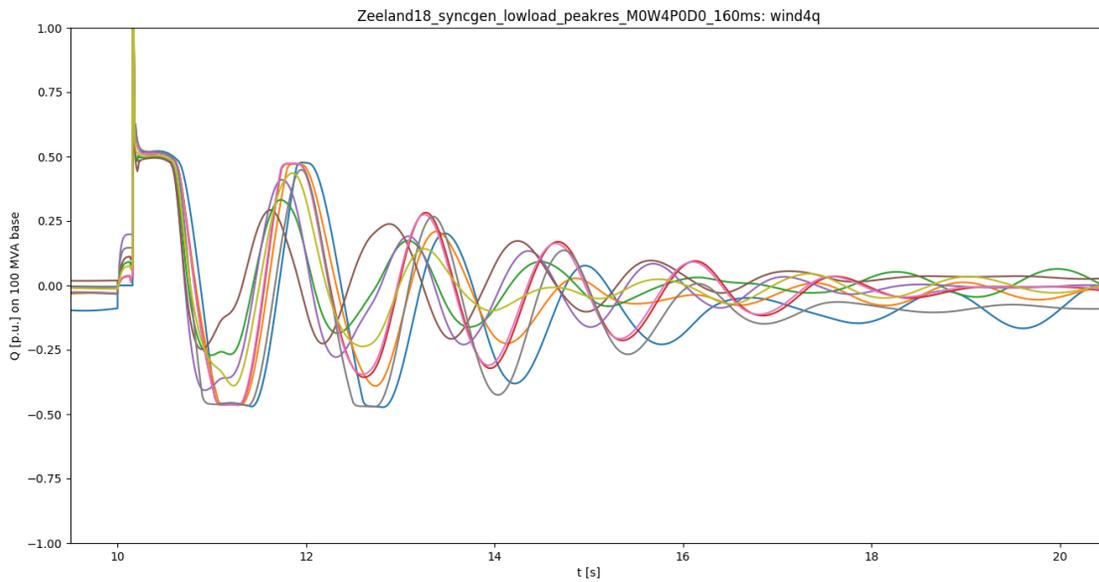
Figure 42: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, low load, peak RES; configuration: M0W4P0D0

The effects of combined Type 3 and Type 4 modelling (the **W5** configurations) are investigated further in section 5.4.3.

5.2.4 Effect of PV modelling

PV modelling, with the parameters as chosen in section 4.6.3, has a strong stabilising effect on its own; critical clearing times can be extended by up to 80 ms. When wind modelling is also applied, however, the added improvement to dynamic performance from PV modelling is rather modest – compare **M0W3P0D0** to **M0W3P1D0** configurations in the tables above. The mechanism by which PV models improve stability is the same as that of wind models: the rapidly-responding voltage control blocks command a large reactive-power injection immediately after fault clearing.

5.2.5 Effect of HVDC modelling

HVDC modelling appears to have very little discernible effect, judging by the small to nonexistent differences between **D0** and **D1** configurations in the South Holland tables. A clear "direction" cannot be discovered either: for example, in the 2035 scenario with low load and peak RES, the configuration **M0W4P0D0** gives a (marginally) longer clearing time than **M0W4P0D1**; but in the 2035 scenario with peak load and peak RES, the exact opposite applies.
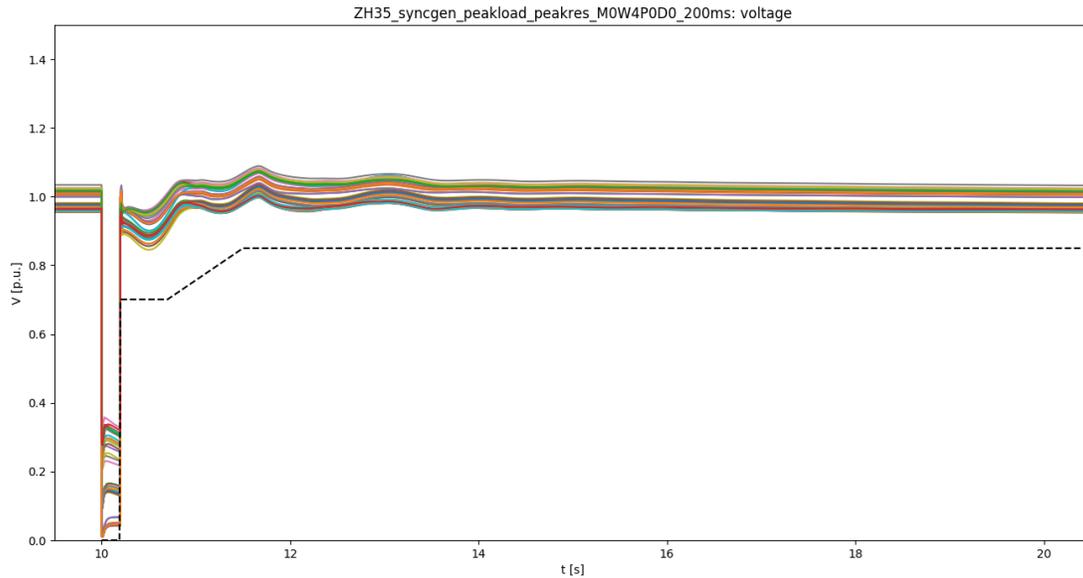
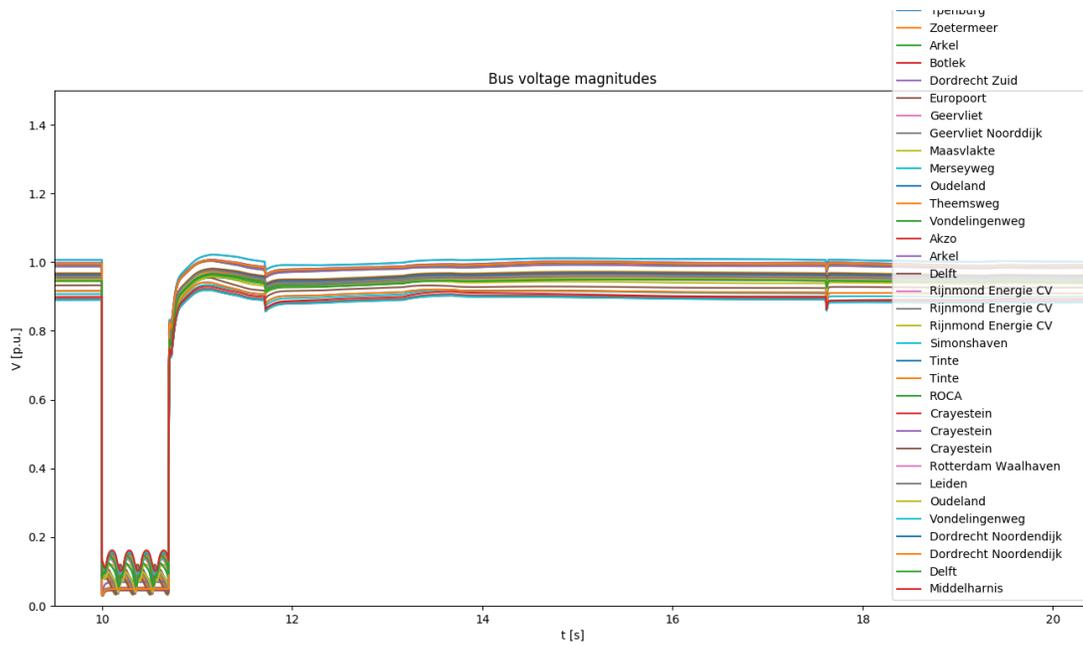Figure 43: Bus voltages for a 200 ms clearing time. Scenario: South Holland 2035, peak load, peak RES; configuration: M0W4P0D0
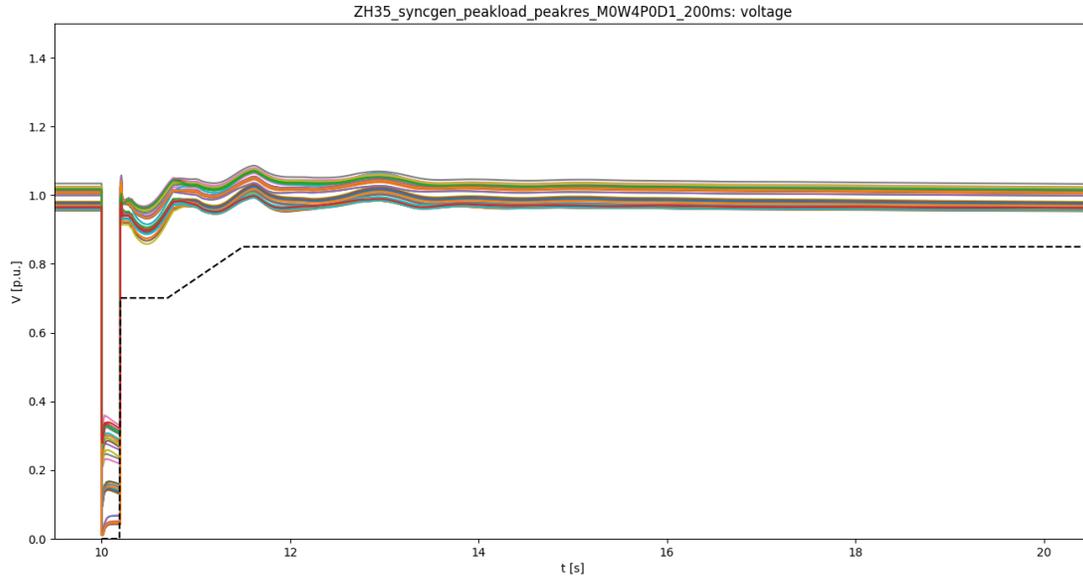
Figure 44: Bus voltages for a 200 ms clearing time. Scenario: South Holland 2035, peak load, peak RES; configuration: M0W4P0D1
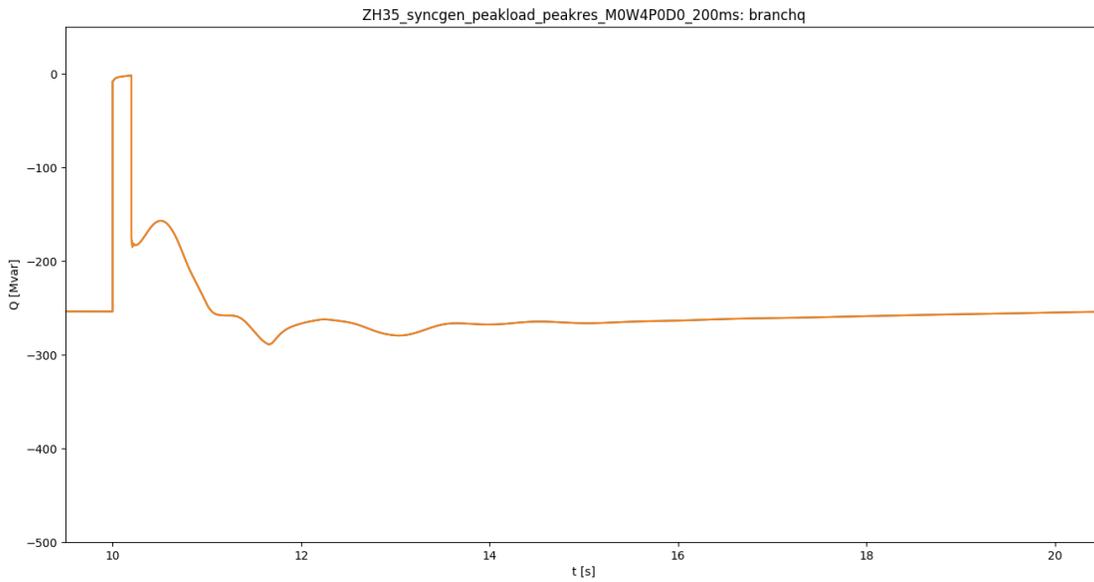
Figure 45: Reactive power flowing into the BritNed terminal (negative convention) for a 200 ms clearing time. Scenario: South Holland 2035, peak load, peak RES; configuration: M0W4P0D0
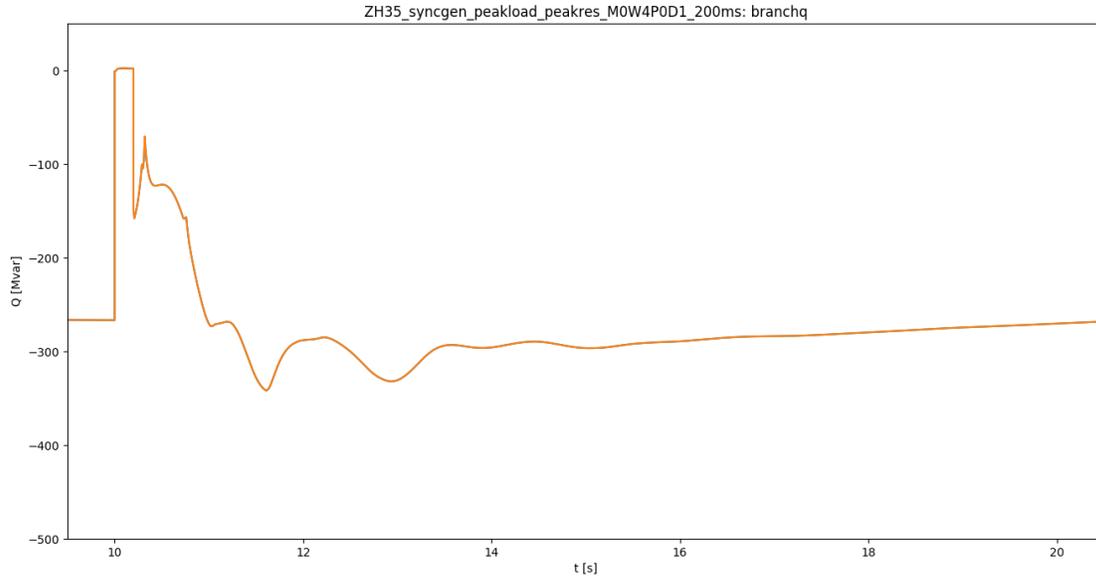
Figure 46: Reactive power flowing into the BritNed terminal (negative convention) for a 200 ms clearing time. Scenario: South Holland 2035, peak load, peak RES; configuration: M0W4P0D1

What do we see if we look at the plots of these simulations? For the peak load, peak RES scenario (figures 43 through 46), we see little difference in bus voltage profiles between the **D0** and **D1** configurations. The reactive power flows into the BritNed poles bring the difference into slightly sharper relief: in the **D1** configuration (figure 46), we see the controller suppress reactive power consumption at first, to ramp up relatively sharply later. The **D0** configuration gives us a smoother profile with less overshoot (figure 45).

The same comments apply to the low load, peak RES scenario (figures 47 through 50).

Figure 47: Bus voltages for a 170 ms clearing time. Scenario: South Holland 2035, low load, peak RES; configuration: M0W4P0D0

Figure 48: Bus voltages for a 170 ms clearing time. Scenario: South Holland 2035, low load, peak RES; configuration: M0W4P0D1
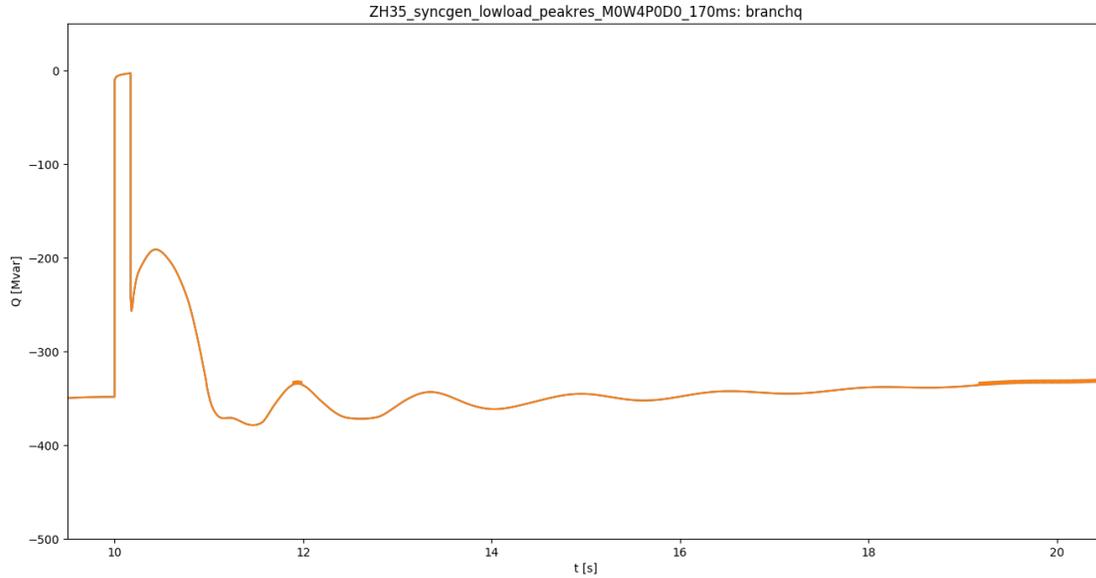
Figure 49: Reactive power flowing into the BritNed terminals (negative convention) for a 170 ms clearing time. Scenario: South Holland 2035, low load, peak RES; configuration: M0W4P0D0
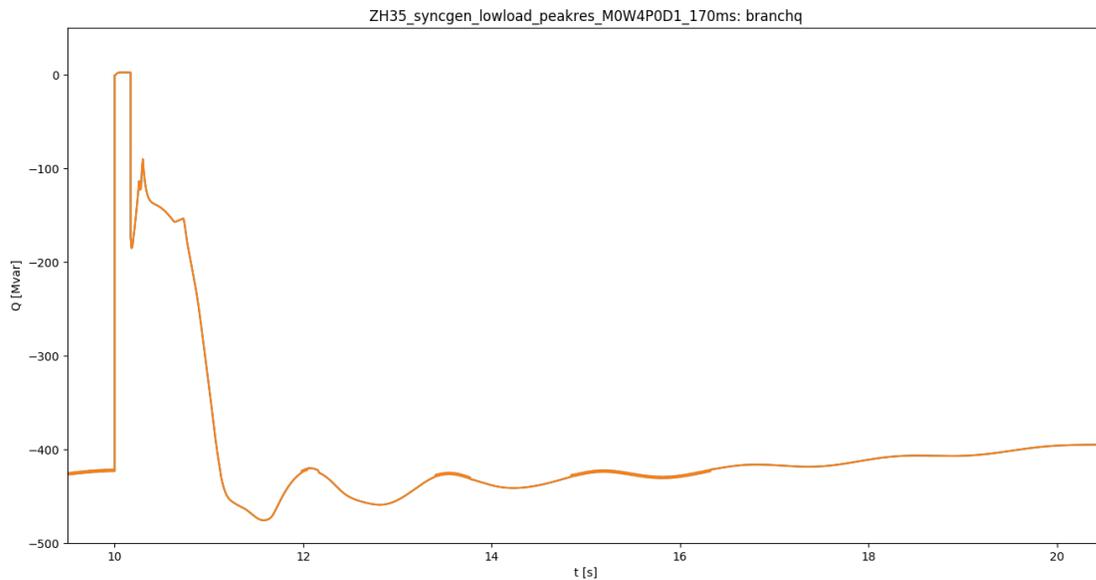
Figure 50: Reactive power flowing into the BritNed terminals (negative convention) for a 170 ms clearing time. Scenario: South Holland 2035, low load, peak RES; configuration: M0W4P0D1

One peculiarity of the **M0W0P0D1** configuration is that it often leads to numerical instability, even for very short fault-clearing times. It seems that in scenarios where the **M0W0P0D0** configuration was already very unstable, the added complexity of modelling the HVDC controls was enough to give the system the final "push" over the edge from physical to numerical instability.

5.3   Verification scenarios

5.3.1   Scenarios with only renewable generation

All scenarios in the main sweep had synchronous generators in service – the same synchronous generators, in fact, for reasons discussed in 4.4.4. Scenarios with only renewable generation were not included because, in simulations in earlier phases of the project, they showed very little variation across scenarios and modelling configurations.

One interesting feature of RES-only grids is that the concept of "pulling out of step" disappears. Wind turbines and PV panels have no equivalent to the self-reinforcing acceleration of a synchronous generator that has lost synchronism. Thus there is no such thing as a "critical clearing time" either. In the few RES-only simulations performed, we observed a smooth voltage recovery even from extremely long faults; see figures 51 and 52.
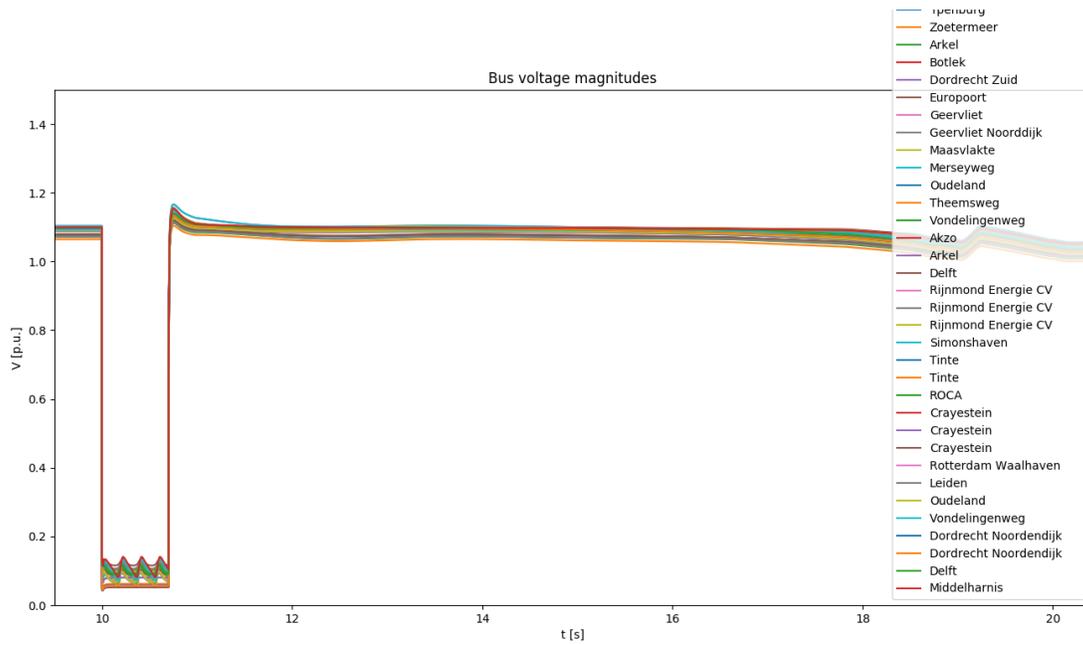
Figure 51: Bus voltages for a 700 ms clearing time. Scenario: South Holland 2035, low load, low RES, no synchronous generation; configuration: M0W3P1D0

Bus voltage magnitudes

V [p.u.]

1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0

10   12   14   16   18   20

t [s]

Legend:
Ypenburg
Zoetermeer
Arkel
Botlek
Dordrecht Zuid
Europoort
Geervliet
Geervliet Noorddijk
Maasvlakte
Merseyweg
Oudeland
Theemsweg
Vondelingenweg
Akzo
Arkel
Delft
Rijnmond Energie CV
Rijnmond Energie CV
Rijnmond Energie CV
Simonshaven
Tinte
Tinte
ROCA
Crayestein
Crayestein
Crayestein
Rotterdam Waalhaven
Leiden
Oudeland
Vondelingenweg
Dordrecht Noordendijk
Dordrecht Noordendijk
Delft
Middelharnis

Figure 52: Bus voltages for a 700 ms clearing time. Scenario: South Holland 2035, peak load, peak RES, no synchronous generation; configuration: M0W3P1D0

5.3.2 Weaker link to the external grid

Simulations were performed with a weakened (higher-impedance) link to the "infinite bus" at Geertruidenberg, which thus provides a lower level of short-circuit power (see section 4.4.2). As a reference case, we considered South Holland 2018, low load, low RES, with a 110 ms clearing time. Obviously a weaker link to the external grid makes the system less stable, as can be seen in figures 53 through 56; however, this is simply the an effect of grid topology on dynamic performance, not an effect of dynamic modelling. In our results we saw no reason to think scenarios with a weakened external-grid link would be relevant to our conclusions about the effects of modelling.
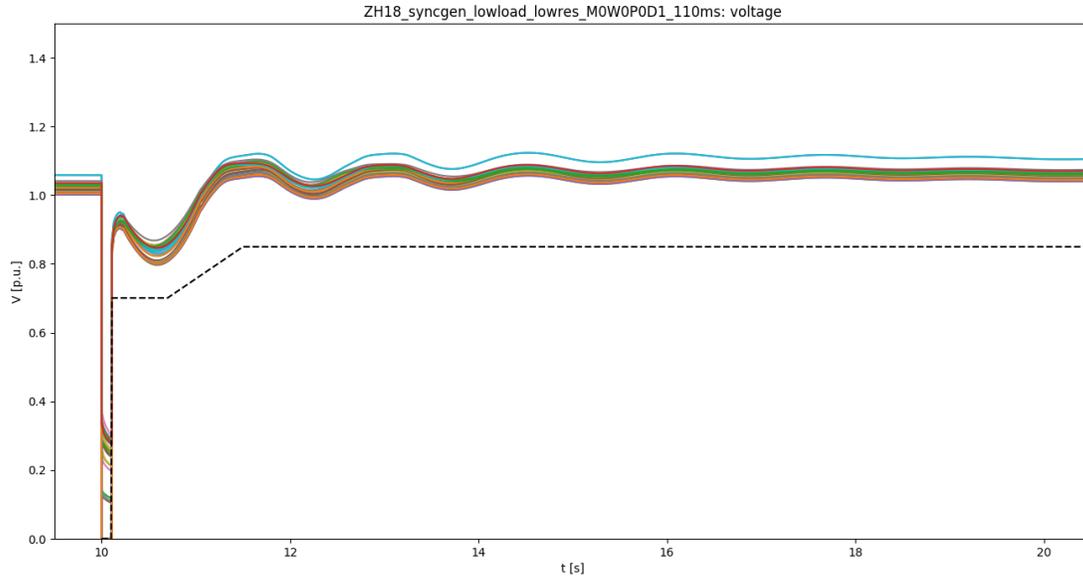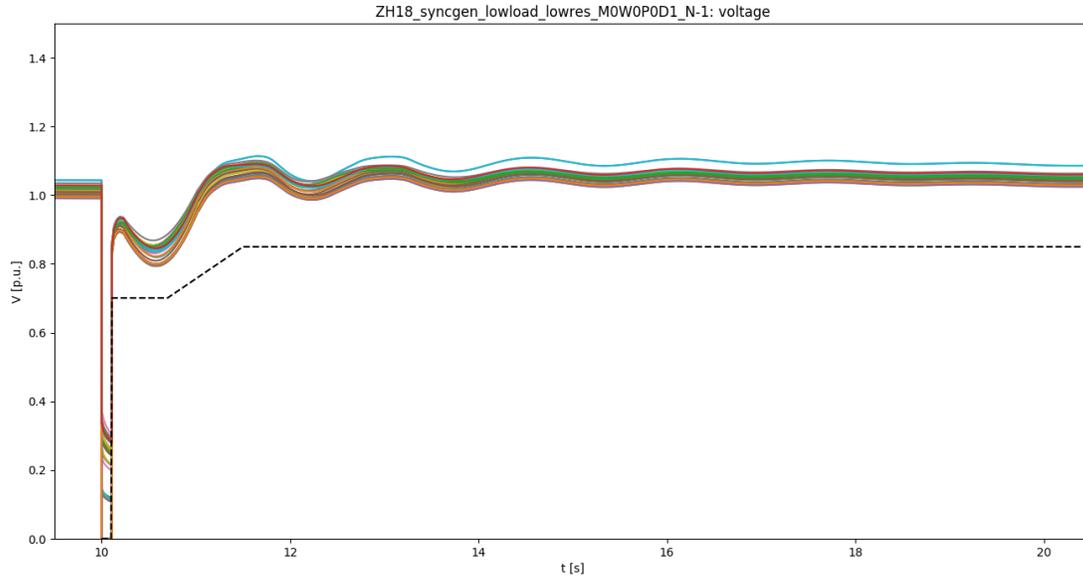
Figure 53: Bus voltages for a 110 ms clearing time. Scenario: South Holland 2018, low load, low RES, short-circuit power at reference level; configuration: M0W0P0D1
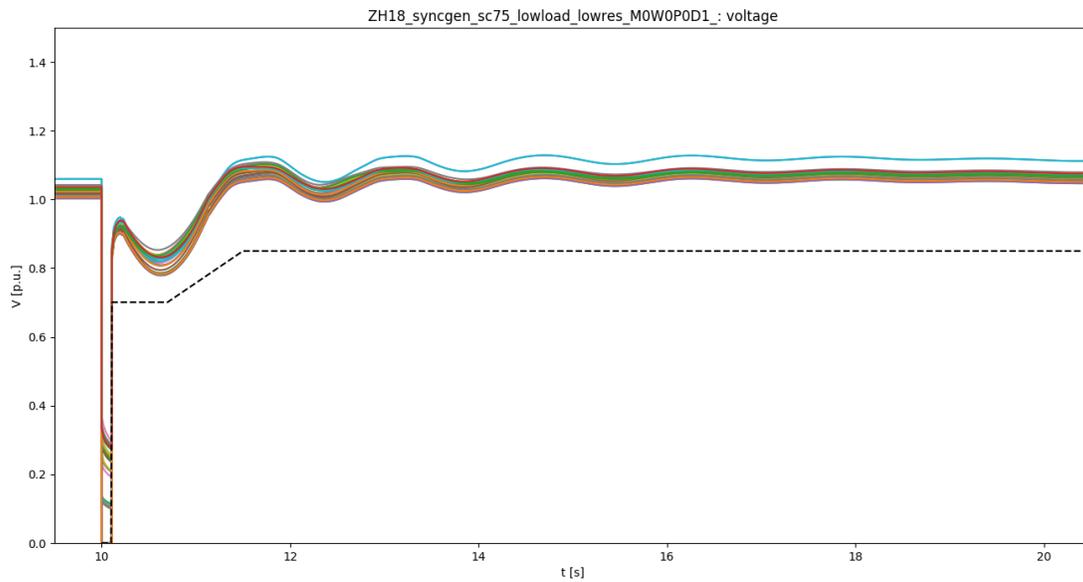
Figure 54: Bus voltages for a 110 ms clearing time. Scenario: South Holland 2018, low load, low RES, short-circuit power 75% of reference level; configuration: M0W0P0D1
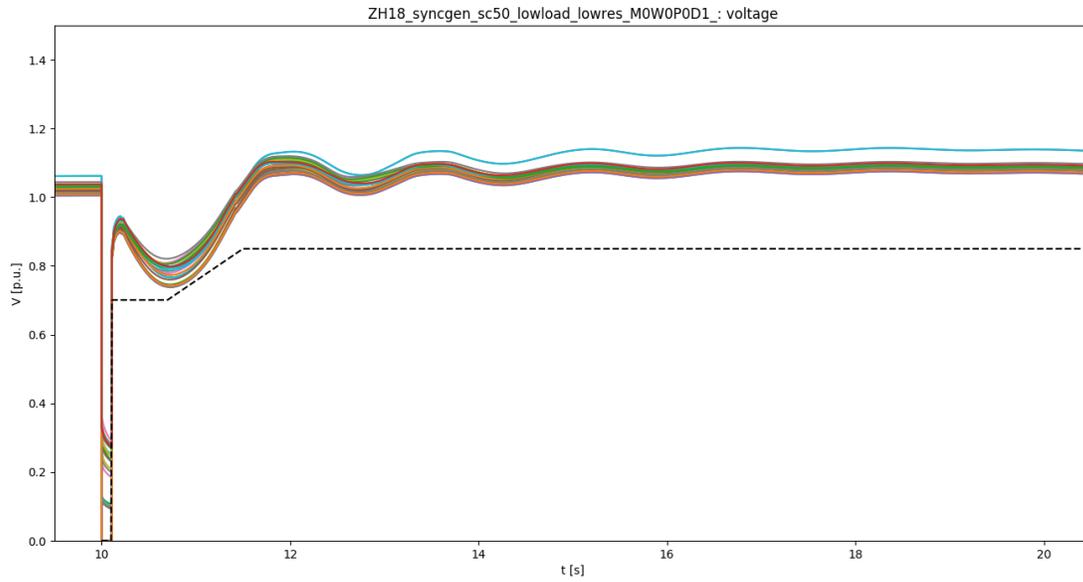
Figure 55: Bus voltages for a 110 ms clearing time. Scenario: South Holland 2018, low load, low RES, short-circuit power 50% of reference level; configuration: M0W0P0D1
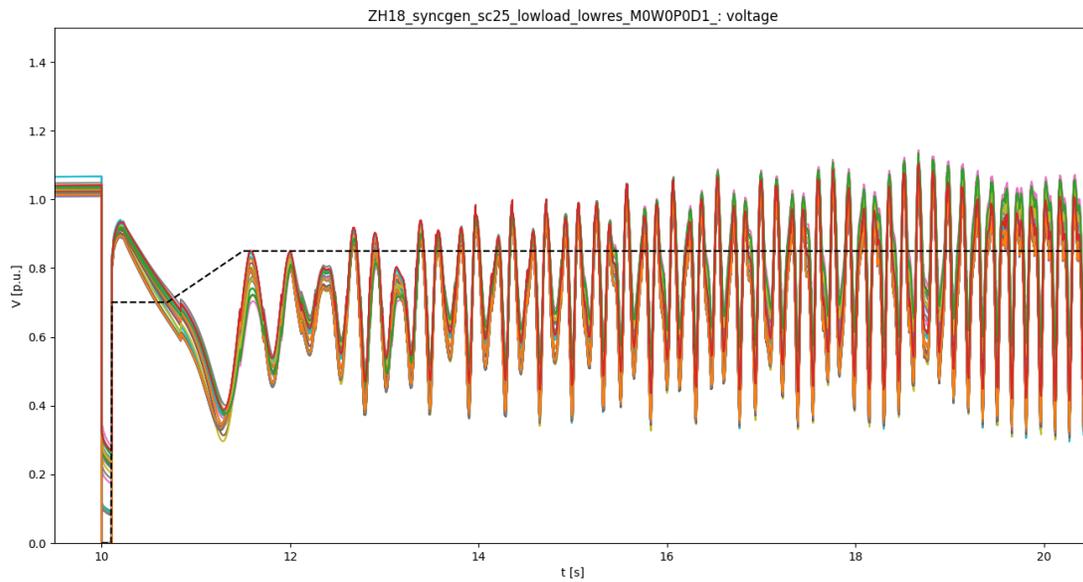
Figure 56: Bus voltages for a 110 ms clearing time. Scenario: South Holland 2018, low load, low RES, short-circuit power 25% of reference level; configuration: M0W0P0D1

5.3.3 Other line out of service for maintenance (N-1)

Another possible topology change considered was an N-1 scenario; that is, a scenario with a major line (other than the one on which the fault occurs) out of service, which often happens for maintenance reasons. We chose to put one circuit of the 380 kV line Bleiswijk-Wateringen – an important parallel link to Crayestein-Simonshaven – out of service, while still applying a fault on Crayestein-Simonshaven. The resulting voltage profiles, however, were so similar to those in the N scenario (all lines in service) that we felt justified in not considering N-1 scenarios in the main sweep. See, for example, figures 57 and 58.
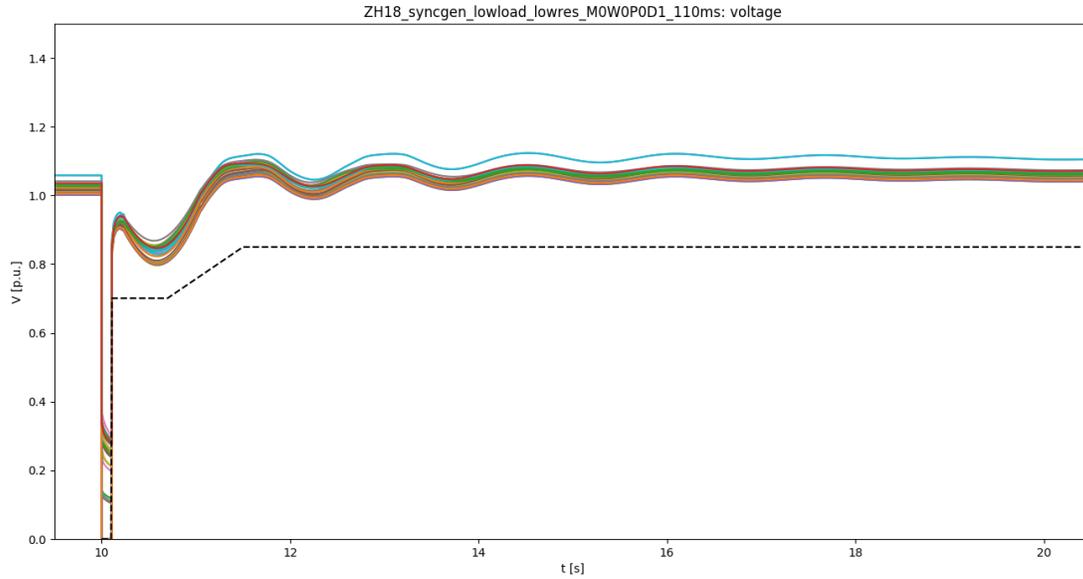
Figure 57: Bus voltages for a 110 ms clearing time. Scenario: South Holland 2018, low load, low RES, all lines in service; configuration: M0W0P0D1

Figure 58: Bus voltages for a 110 ms clearing time. Scenario: South Holland 2018, low load, low RES, one circuit of Bleiswijk-Wateringen out of service; configuration: M0W0P0D1

5.3.4 Alternative disturbance locations

Simulations with different disturbance locations than the reference locations mentioned in section 4.5 were performed. Again, we used as a reference case South Holland 2018, low load, low RES, with a 110 ms clearing time. The disturbance locations were:

- Crayestein-Simonshaven (380 kV) (reference disturbance)

- Bleiswijk-Wateringen (380 kV)

- Bleiswijk-Krimpen (380 kV)

- Alphen-Zoetermeer (150 kV)

- Alblasserdam-Dordrecht Merwedehaven (150 kV)

- Rotterdam Zuid-Rotterdam Waalhaven (150 kV)

Unsurprisingly, we found significant differences between 380 and 150 kV disturbance locations; the 150 kV lines are much less centrally located in the grid topology and carry less power, so faults on these lines have a much lower impact than faults on 380 kV lines. However, the differences *within* each voltage level (see figures 59, 60 and 61 for the 380 kV group; 62, 63 and 64 for 150 kV) were so subtle that we feel justified in considering only one disturbance location in the main sweep.
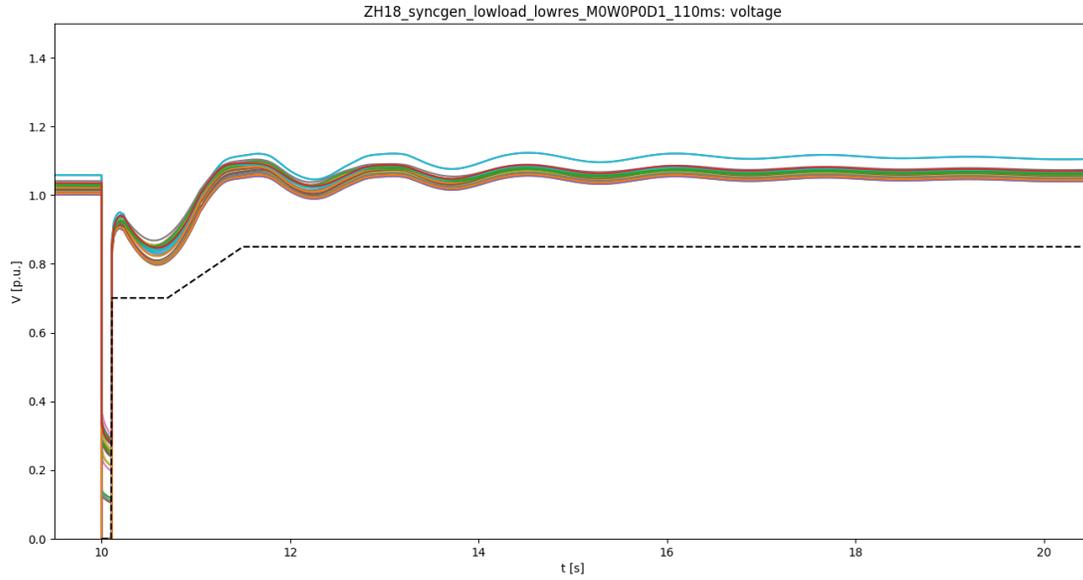
Figure 59: Bus voltages for a 110 ms fault on Crayestein-Simonshaven. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1
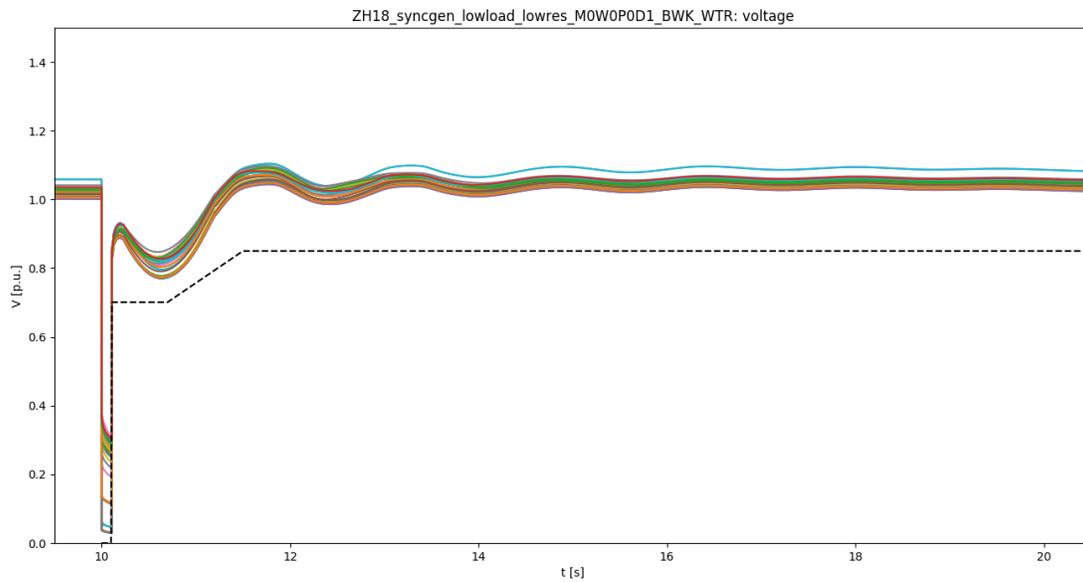
Figure 60: Bus voltages for a 110 ms fault on Bleiswijk-Wateringen. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1
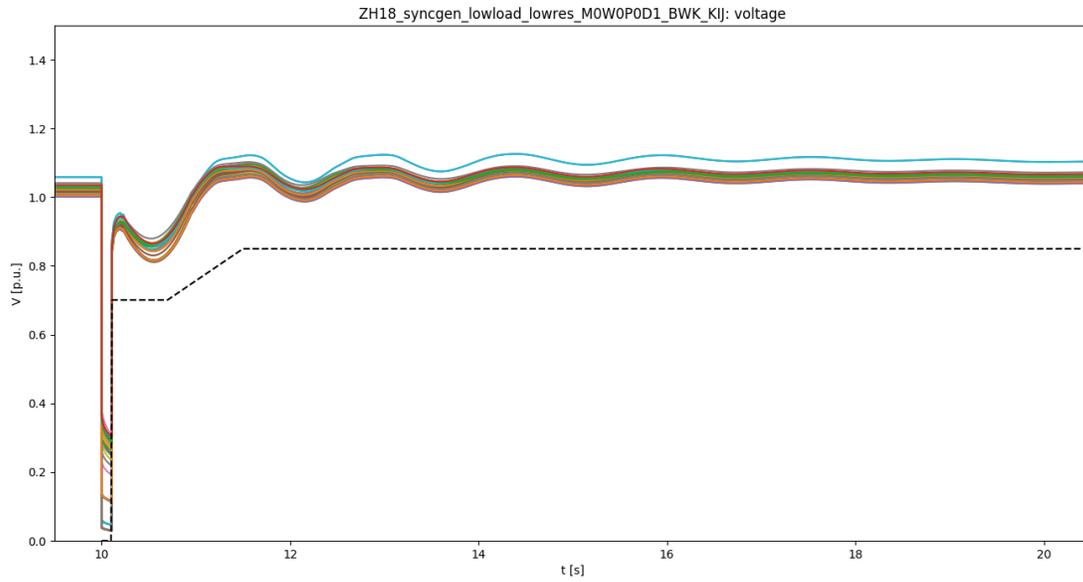
Figure 61: Bus voltages for a 110 ms fault on Bleiswijk-Krimpen. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1
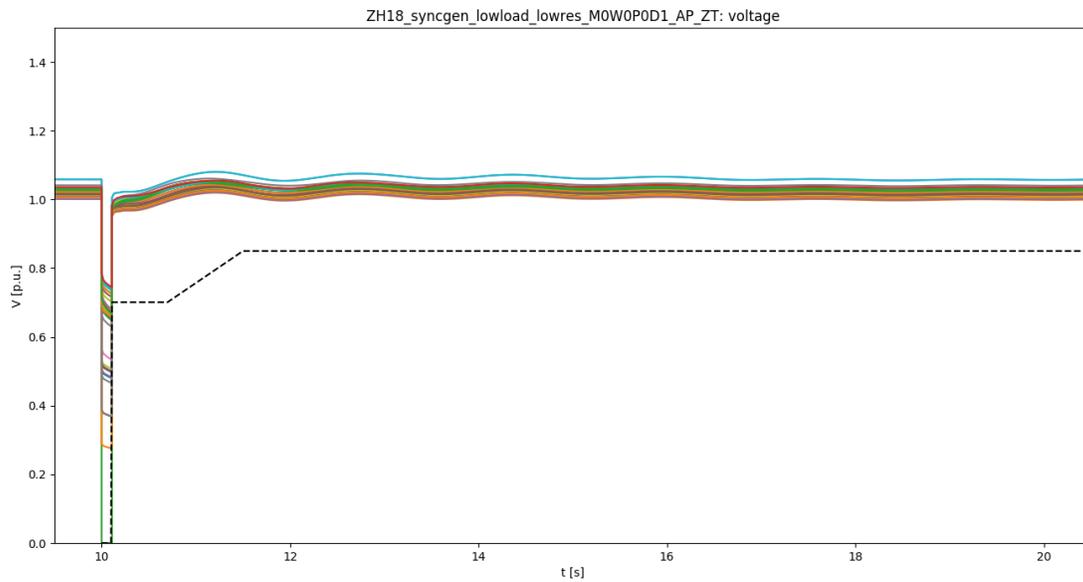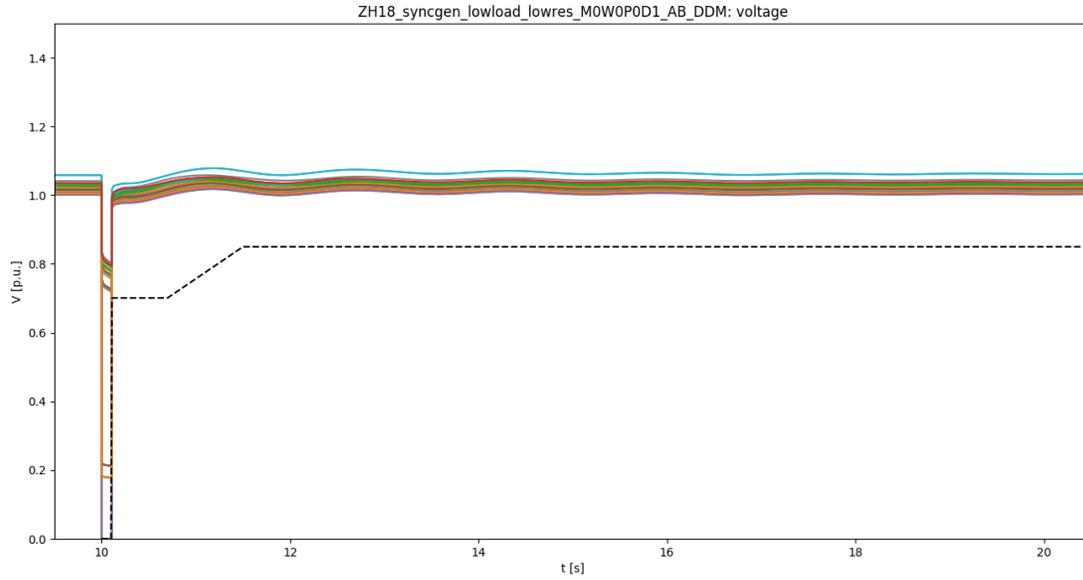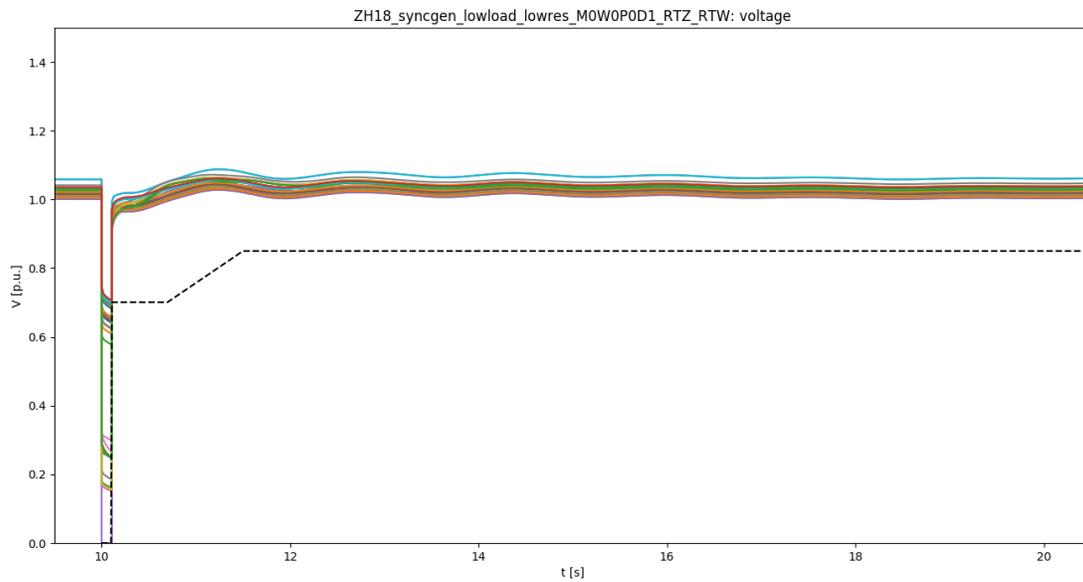

Figure 62: Bus voltages for a 110 ms fault on Alphen-Zoetermeer. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1

Figure 63: Bus voltages for a 110 ms fault on Alblasserdam-Dordrecht Merwedehaven. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1

Figure 64: Bus voltages for a 110 ms fault on Rotterdam Zuid-Rotterdam Waalhaven. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1

Finally, we considered a different type of disturbance: the sudden tripping of a synchronous generator. We tripped the Maasvlakte 3 generator – the largest in South Holland, feeding in its P_{max} of 1354 MW before the fault. As can be seen in figure 65, this disturbance barely had an impact at all. We concluded that, if the tripping of such a large generator has so little effect on the grid, generator trips are not a very interesting disturbance for us to study.
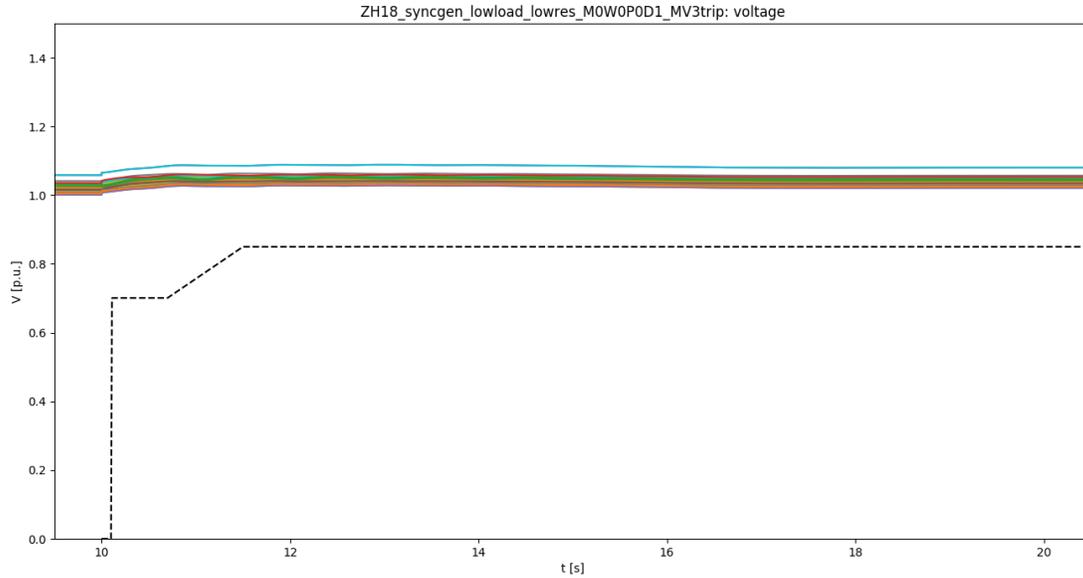
Figure 65: Bus voltages for a trip of the Maasvlakte 3 generator. Scenario: South Holland 2018, low load, low RES; configuration: M0W0P0D1

5.4 Finer sweeps over modelling parameters

5.4.1 Motor load percentage

As mentioned in section 5.2.2, too high an amount of motor load would often lead to numerical instability, and so would a too wide spread across different load buses.

To be able to show at least *some* results with motor load modelling, we chose the scenario with the most stable "base case" (i.e. the longest CCT in **M0W0P0D0**). This was Zealand 2018 with peak load and low RES. Because simulations with motor load modelling at all load buses always failed, we applied modelling only at stations classified as "industrial" or "mixed" (see section 4.6.1) and let the load at "residential" stations remain entirely static.

The outcomes are shown in table 32. As can be seen, even in this highly "stable" scenario the simulations with high percentages of motor load and/or long fault clearing times become numerically unstable.

t_c [ms]	100	110	120	130	140	150	160	170	180	190	200	210	220	230	240	250
M0W0P0D0	0.88	0.87	0.86	0.84	0.82	0.79	0.76	0.73	0.00	0.03	0.18	X	X	X	X	X
M1W0P0D0	0.85	0.84	0.83	0.82	0.81	0.78	0.75	0.71	0.01	0.07	–	–	–	–	–	–
M2W0P0D0	0.83	0.82	0.81	0.80	0.79	0.77	0.73	0.00	0.03	0.19	–	–	–	–	–	–
M3W0P0D0	0.81	0.80	0.79	0.78	0.77	0.75	0.71	0.01	0.07	–	–	–	–	–	–	–
M4W0P0D0	0.79	0.78	0.77	0.76	0.75	0.73	0.00	0.04	0.24	–	–	–	–	–	–	–
M5W0P0D0	0.77	0.76	0.75	0.74	0.73	0.71	0.02	0.09	–	–	–	–	–	–	–	–
M6W0P0D0	0.75	0.74	0.73	0.72	0.71	0.00	0.04	–	–	–	–	–	–	–	–	–
M7W0P0D0	0.74	0.72	0.71	0.70	0.00	0.02	0.11	–	–	–	–	–	–	–	–	–
M8W0P0D0	0.72	0.71	0.00	0.00	0.01	0.05	0.38	–	–	–	–	–	–	–	–	–
M9W0P0D0	0.70	0.00	0.00	0.00	0.02	0.11	–	–	–	–	–	–	–	–	–	–

Table 32: Simulation outcomes for different modelling configurations and fault clearing times. Scenario: Zealand 2018, peak load, low RES

The decrease in dynamic performance with higher percentages of motor load is apparent from table 32. The voltage profiles for a 160 ms clearing time and different motor load percentages (figures 66, 67 and 68) illustrate how motor load hinders voltage recovery with its high reactive power consumption (see figures 69, 70 and 71). This corresponds to an increase in motor slip due to the fault (figures 72 and 73).

The results of our motor load sweep include no situations where high percentages of motor load slow down voltage recovery so much that it leads to loss of synchronism – that is, there are no columns in table 32 with both green and red entries. However, simulations in earlier stages of the project have shown that this is definitely a possibility.
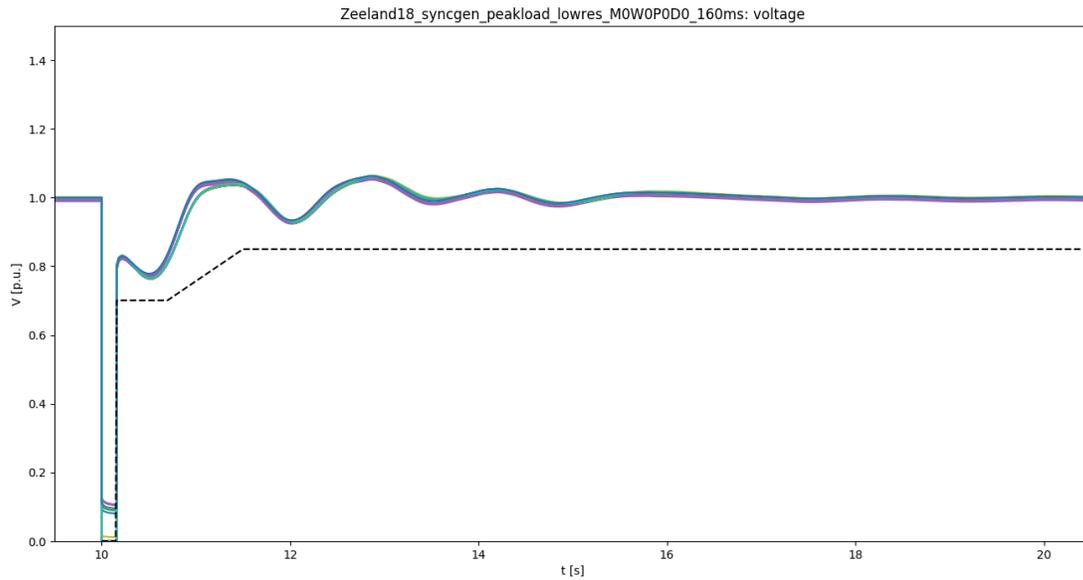

Figure 66: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M0W0P0D0
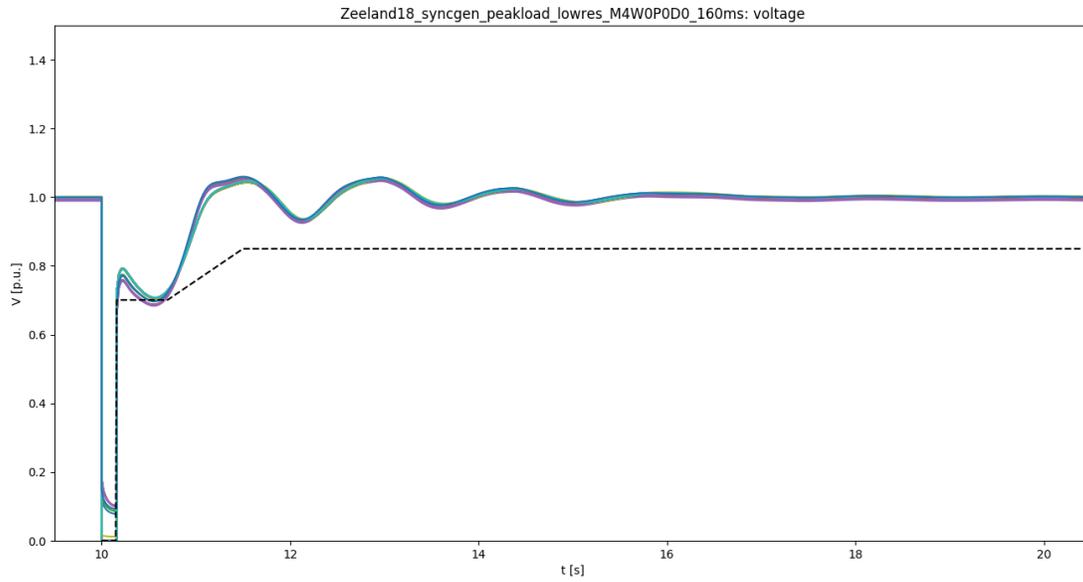
79

Figure 67: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M4W0P0D0
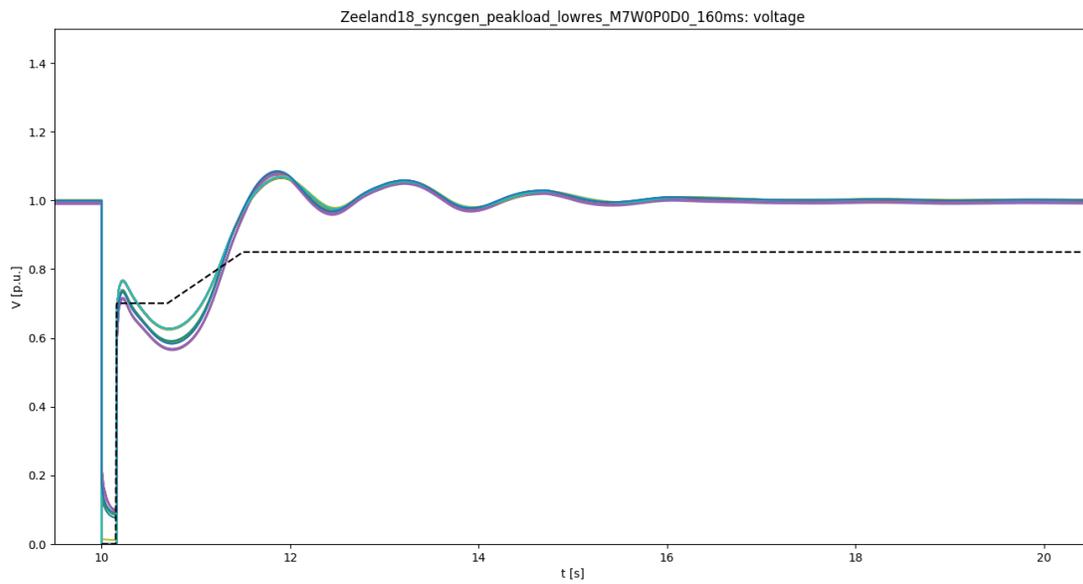

Figure 68: Bus voltages for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M7W0P0D0
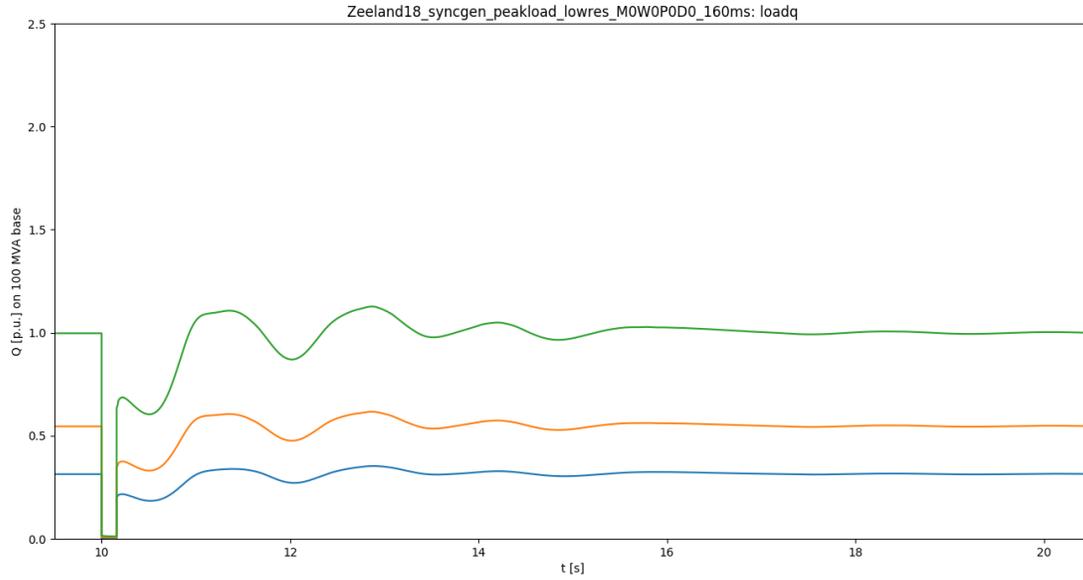
Figure 69: Reactive power consumption of mixed and industrial loads for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M0W0P0D0
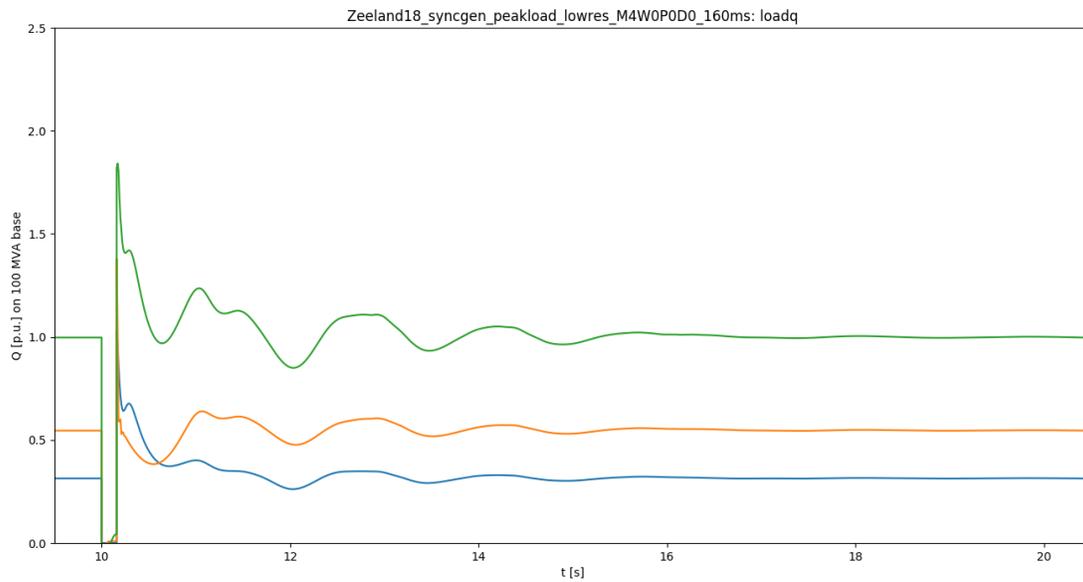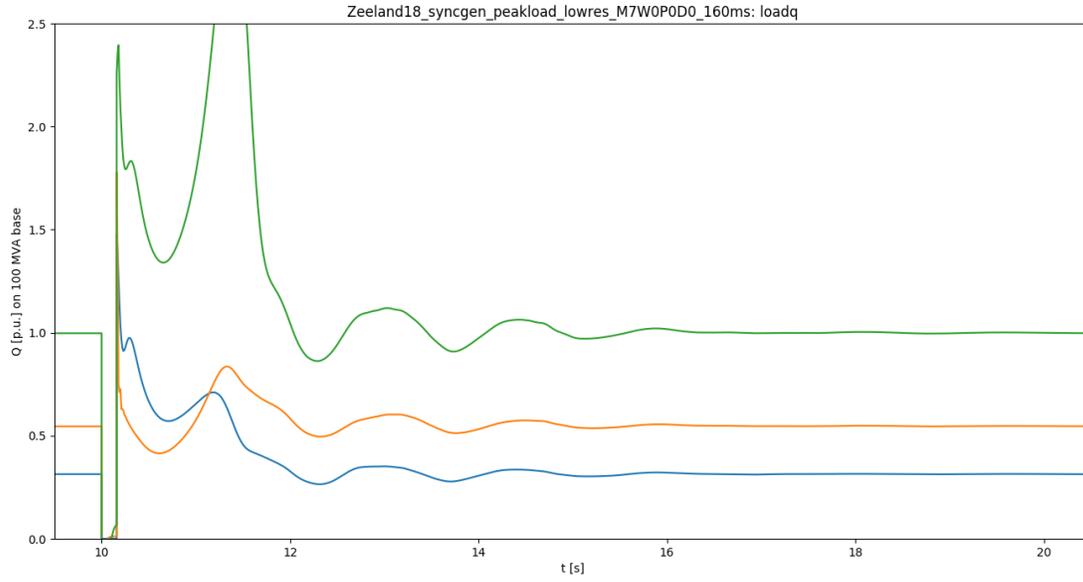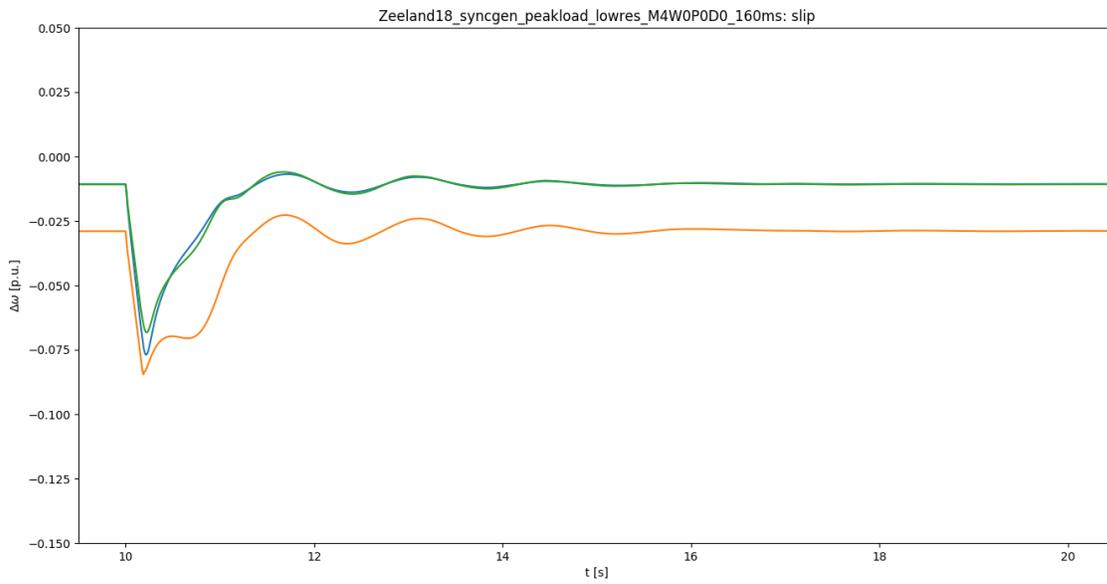

Figure 70: Reactive power consumption of mixed and industrial loads for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M4W0P0D0

Figure 71: Reactive power consumption of mixed and industrial loads for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M7W0P0D0

Figure 72: Motor speed deviations for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M4W0P0D0
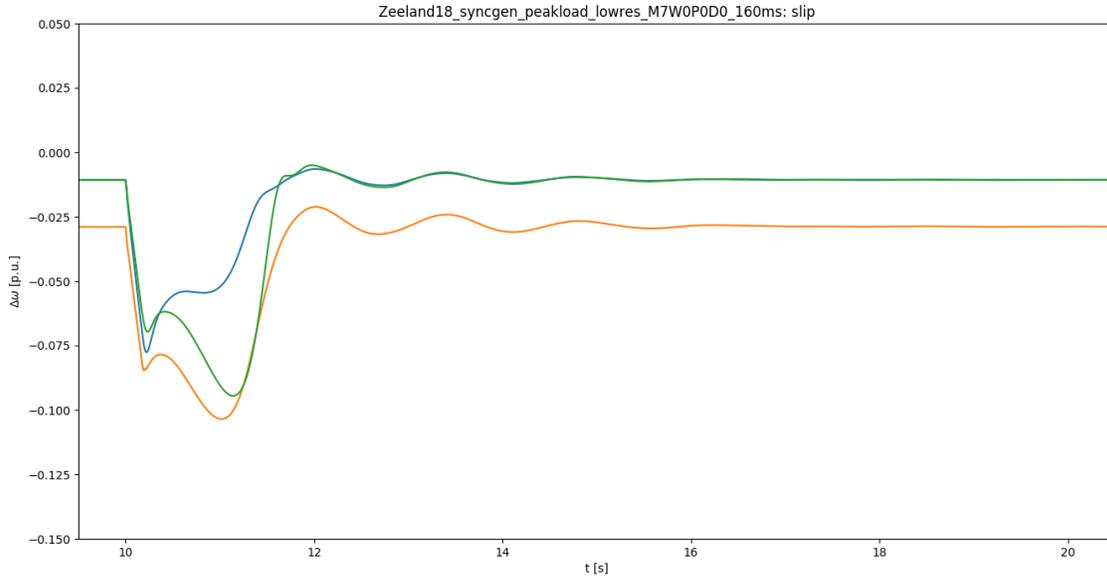
Figure 73: Motor speed deviations of mixed and industrial loads for a 160 ms clearing time. Scenario: Zealand 2018, peak load, low RES; configuration: M7W0P0D0

5.4.2 HVDC ramping rates

We tested the effect of different voltage and current ramping rates in the `CDC4T` model. For this we used the South Holland 2018, low load, low RES scenario and applied a 110 ms three-phase fault on Crayestein-Simonshaven.

As can be seen from table 33, the effects are negligible – perhaps not surprising, given the marginal difference between `CDC4T` and a static representation in the first place (see section 5.2.5). We therefore see no added value in further analysis of these results.

CRAMP ↓ VRAMP →	1.0	2.0	3.0	4.0	5.0	6.0	7.0
1.0	0.77	0.78	0.79	0.80	0.80	0.81	0.81
2.0	0.77	0.78	0.79	0.79	0.80	0.80	0.80
3.0	0.77	0.78	0.79	0.79	0.80	0.80	0.80
4.0	0.77	0.78	0.79	0.79	0.80	0.80	0.80
5.0	0.77	0.78	0.79	0.79	0.80	0.80	0.80
6.0	0.77	0.78	0.79	0.79	0.80	0.80	0.80
7.0	0.77	0.78	0.79	0.79	0.80	0.80	0.80

Table 33: Simulation outcomes for different values of `VRAMP` and `CRAMP` (both in p.u./second) in the `CDC4T` model. Scenario: South Holland 2018, low load, low RES

5.4.3   WT3/WT4 ratio

In the main sweep, we found that modelling half of all wind generation as Type 3 and half as Type 4 led to better dynamic performance than either type on its own (see section 5.2.3). To gain more insight into the interaction between Type 3 and Type 4 wind turbine models, we ran a finer-grained sweep over the ratio of Type 3 to Type 4 wind turbines present in the system (in the main sweep, the only points considered were 100% Type 3, 100% Type 4, and an even split). Using South Holland 2018, low load, low RES with a fault-clearing time of 150 ms as a reference case, we swept from 10% WT3 (and 90% WT4) to 90% WT3 (10% WT4) in steps of 10%.
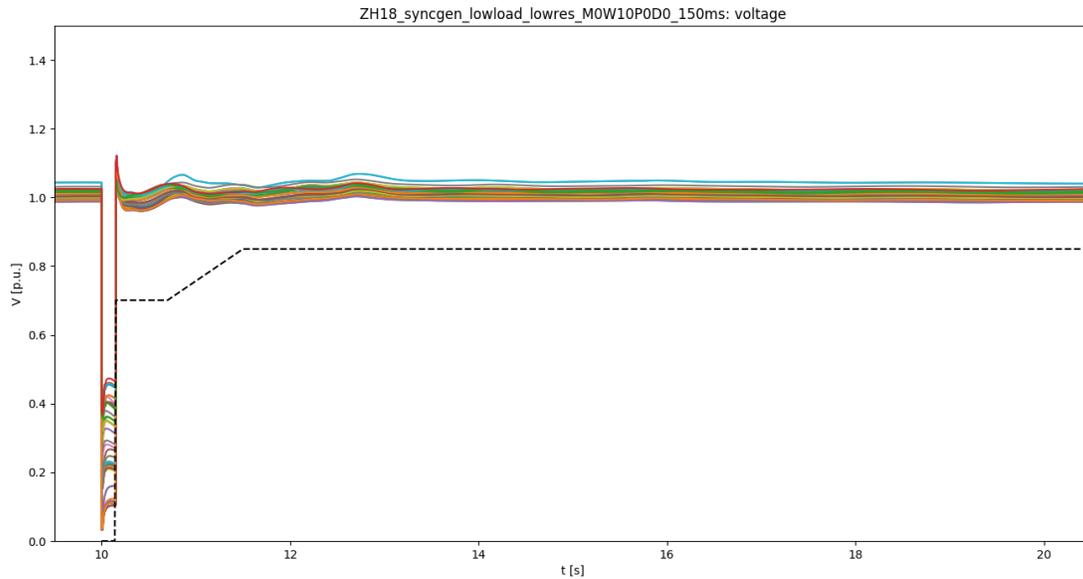

Figure 74: Bus voltage profiles for a 150 ms clearing time. Scenario: South Holland 2018, low load, low RES; 10% of wind generation modelled as Type 3, 90% as Type 4
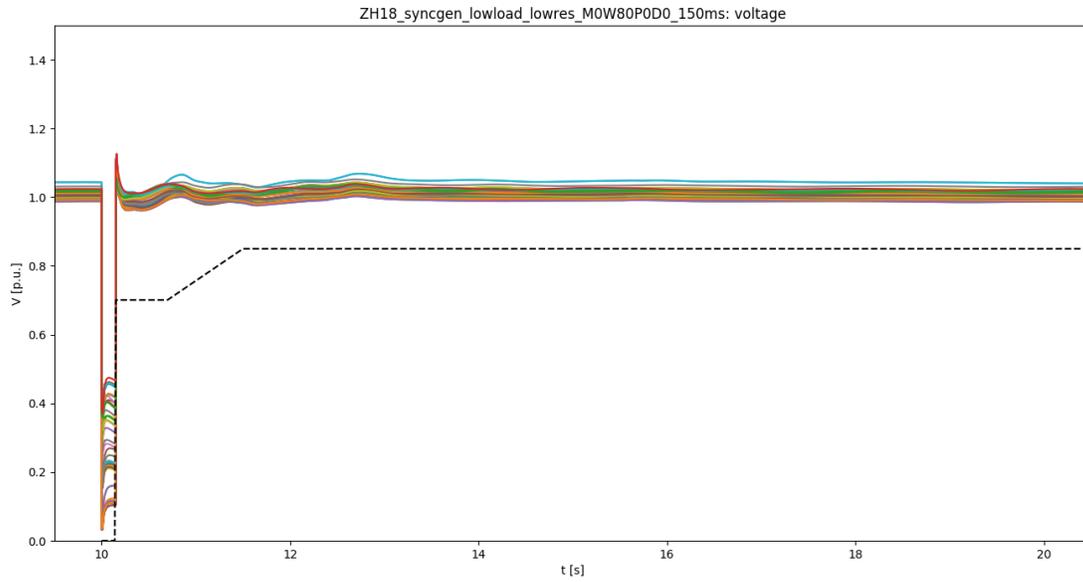
Figure 75: Bus voltage profiles for a 150 ms clearing time. Scenario: South Holland 2018, low load, low RES; 80% of wind generation modelled as Type 3, 20% as Type 4
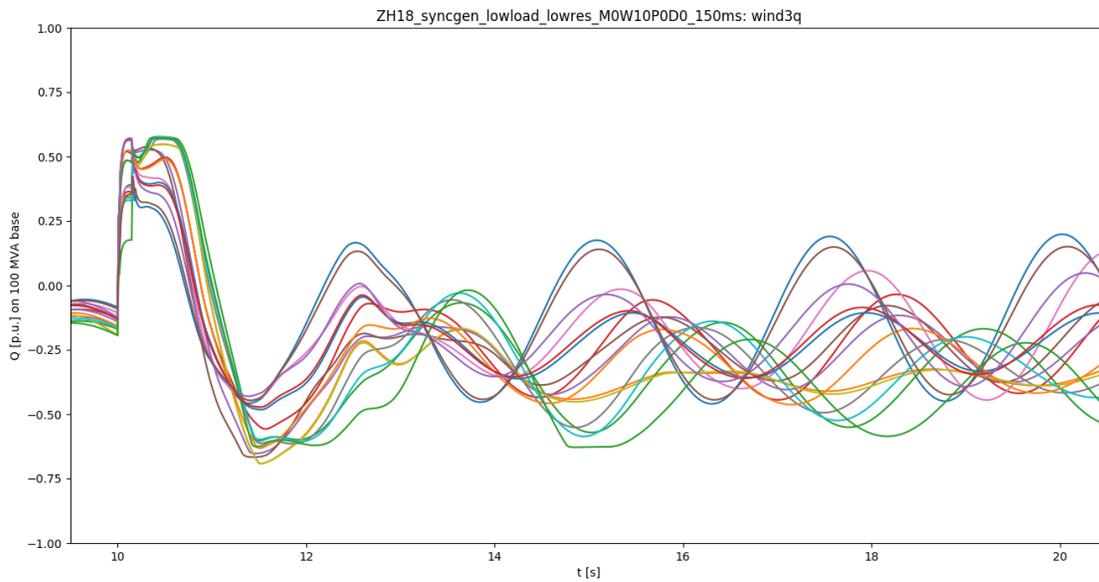

Figure 76: Reactive power output of Type 3 wind turbines for a 150 ms clearing time. Scenario: South Holland 2018, low load, low RES; 10% of wind generation modelled as Type 3, 90% as Type 4
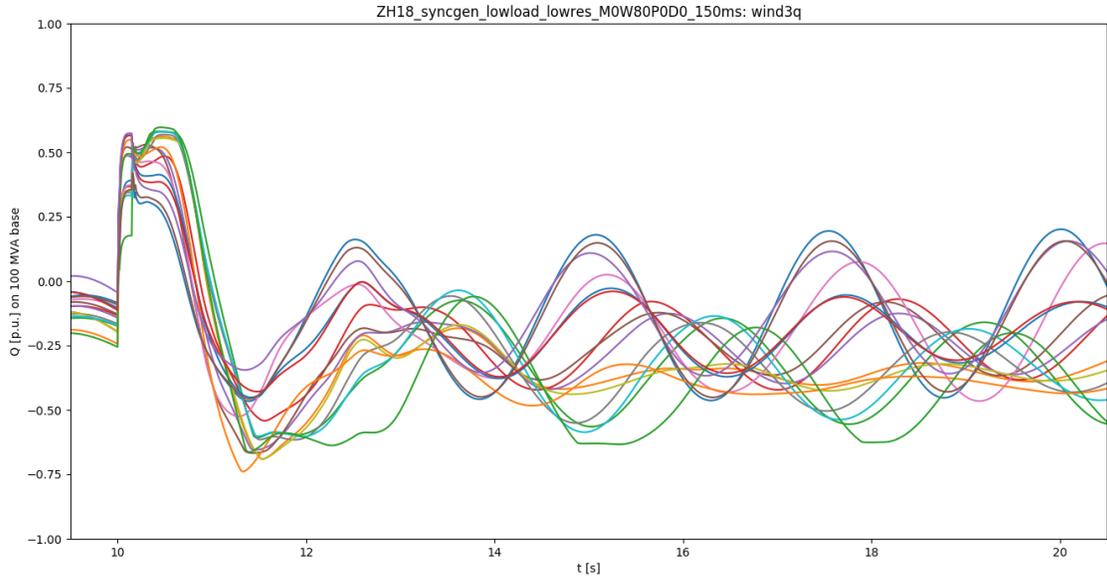
Figure 77: Reactive power output of Type 3 wind turbines for a 150 ms clearing time. Scenario: South Holland 2018, low load, low RES; 80% of wind generation modelled as Type 3, 20% as Type 4
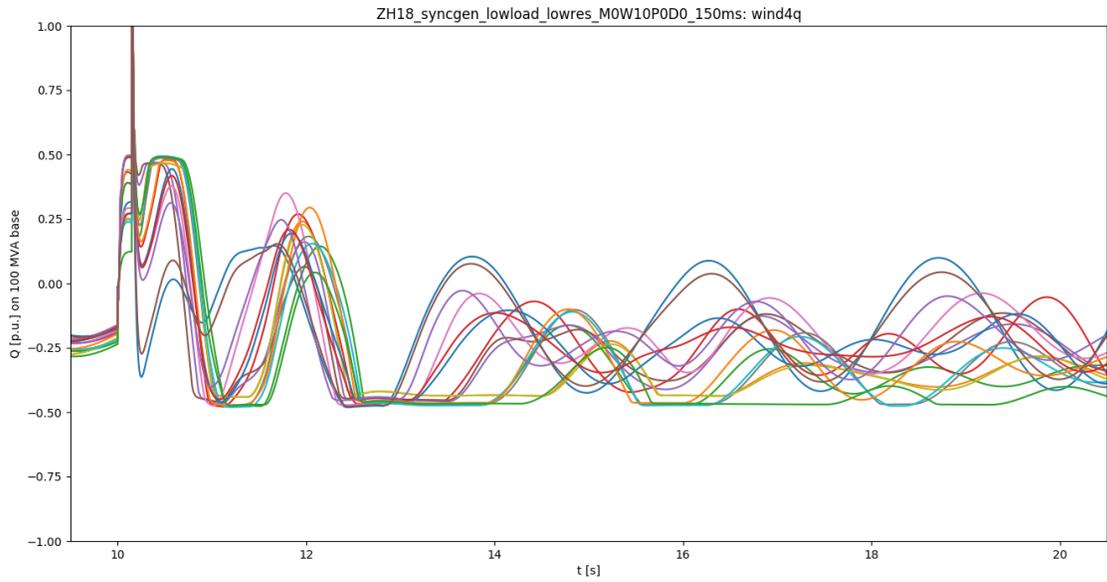

Figure 78: Reactive power output of Type 4 wind turbines for a 150 ms clearing time. Scenario: South Holland 2018, low load, low RES; 10% of wind generation modelled as Type 3, 90% as Type 4
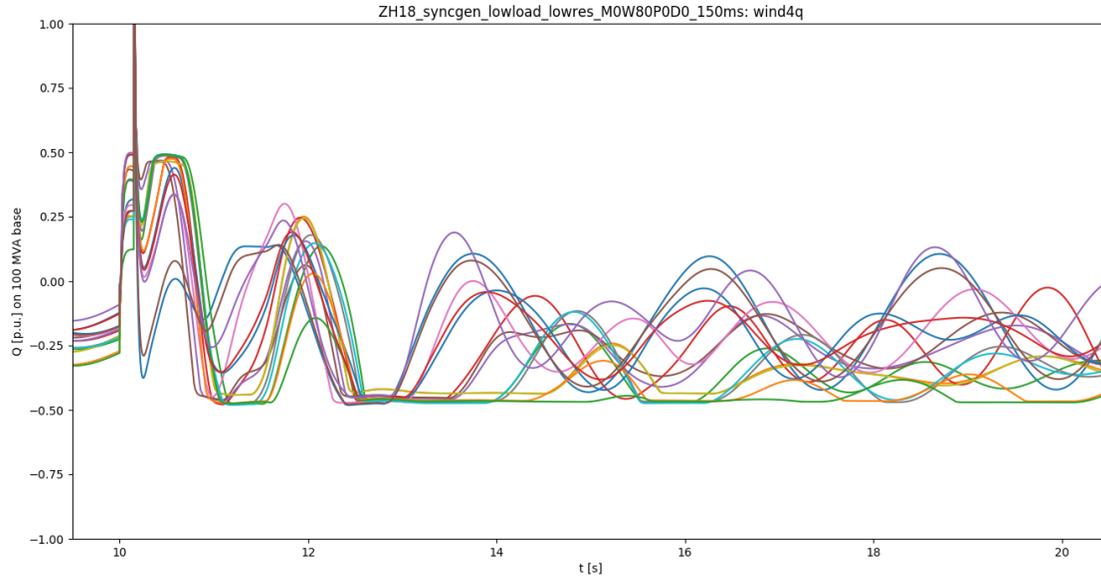
Figure 79: Reactive power output of Type 4 wind turbines for a 150 ms clearing time. Scenario: South Holland 2018, low load, low RES; 80% of wind generation modelled as Type 3, 20% as Type 4

The resulting voltage profiles, as well as the reactive-power profiles of both wind turbine types, were all practically identical. Compare, for example, figures 74, 76 and 78 to figures 75, 77 and 79, respectively.

So, although there is a clear difference between 0% and 10% Type 3 modelling, and between 90% and 100%, apparently no step in between has any discernible effect on the system's fault response. The explanation is that, when using the WECC default parameters, *wind turbines' reactive-power support capabilities are unconstrained by the rating of the turbine.* A 10 MW wind farm will provide the same Q injection as a 100 MW wind farm at the same location. Hence, although there is a clear difference between having one (Type 3 or Type 4) model at a certain bus and having two (Type 3 *and* Type 4), their respective ratings make no difference. This is consistent with what we saw in the main sweep.

Such behaviour is, of course, questionably realistic. We attempted to rectify this with targeted changes to `WT3E1` and `WT4E1` model parameters, such as voltage-to-Mvar gains and reactive-current command limits. However, the changes we applied produced such unexpected results, that we decided not to pursue these attempts further, and to leave a more detailed investigation to researchers with a deeper background in control theory.

Chapter 6

Conclusion

Looking at our results from chapter 5, we can draw the following conclusions about the effect of component modelling on the large-disturbance stability of regional grids.

- Motor load modelling degrades dynamic performance, decreasing critical clearing times and increasing the risk of voltage-envelope violations, because the increased reactive-power consumption of reaccelerating motors hinders voltage recovery. Even a small variation in the percentage of load modelled as motor load can make a significant difference in system fault response. See sections 5.2.2 and 5.4.1.

- Wind turbine modelling, whether Type 3 or Type 4, provides a significant boost to dynamic performance even at low wind penetration – at least when using the default parameters suggested by the WECC. Critical clearing times are increased, sometimes drastically, and voltage dips after fault clearing are mitigated or even eliminated entirely. This is due to the reactive power injection ordered by fast-response voltage controllers in both Type 3 and Type 4 wind turbine models; a less desirable effect of these control settings are sustained oscillations in the turbines' reactive power output, which, however, are barely reflected in voltage profiles. Both Type 3 and Type 4 models can provide reactive-power support unconstrained by turbine ratings; one consequence is that when installed wind power is divided over Type 3 and Type 4 models, they improve dynamic performance even more than either one on their own, independent of the ratio of Type 3 to Type 4 generation.

- PV modelling has a similar effect to Type 4 wind modelling, and for the same reasons.

- The effect of HVDC modelling on system dynamic performance is very small and not consistently positive or negative. Different voltage and current ramping rates have no discernible effect.

As mentioned in section 1.3, this thesis has been a first step in a larger process. That means the above conclusions are not set in stone; rather they are indicators of what might be expected in further studies, and of the areas on which such studies could most profitably focus. That brings us to our recommendations for future research.

Chapter 7

Discussion and recommendations

A crucial weakness in our research is the unavailability of manufacturer-specific parameters for the generic wind and PV models used. The choice of these parameters could impact results considerably; and, as we saw in section 5.4.3, the default parameters we used led to results whose plausibility can be questioned. We recommend that future researchers approach the issue from a more sophisticated control-theory perspective, and examine the effects of variations in the modelling parameters of WT3, WT4, and PV models.

With regards to HVDC modelling, it must be remarked that we studied only one HVDC terminal in one grid; presumably the importance of modelling would increase if more HVDC links are added to the grid, as appears to be the policy of many European TSOs for the medium term.

The numerical instability we encountered in several simulation runs often confounded us. In our attempts to understand and solve these problems, we were frustrated by our lack of insight into the numerical mathematics underlying PSSe's internal processes. Hence, we feel that a reference paper on numerical mathematics, as applied to large-disturbance power system dynamic simulations, would be a valuable contribution to the field.

In many scenarios, configuration **M0W0P0D0** (i.e. the "static base case," with no dynamic modelling of any components except synchronous generators) was very unstable, with a critical clearing time in the 100-130 ms range. This means the wind and PV controllers, when modelled, were "boosting" the dynamic performance of an otherwise unstable grid. Whether these controllers can be *relied on* to provide stability is a difficult question – and a dangerous one to answer prematurely. As mentioned in section 1.3, this thesis has only been a first step; we advise against any update to TenneT's operational or planning policy until more detailed studies have been performed.

Bibliography

[1] IEEE/CIGRÉ Joint Task Force on Stability Terms and Definitions. "Definition and Classification of Power System Stability". In: *IEEE Trans. Power Syst.* 19.2 (2004).

[2] T. van Cutsem and C. Vournas. *Voltage Stability of Electric Power Systems*. Springer, 1998.

[3] C.W. Taylor. *Power System Voltage Stability*. McGraw-Hill, 1993.

[4] J.G. Slootweg. "Wind Power: Modelling and Impact on Power System Dynamics". PhD thesis. TU Delft, 2003.

[5] P. Kundur. *Power System Stability and Control*. McGraw-Hill, 1994.

[6] J.J. Grainger and W.D. Stevenson, Jr. *Power System Analysis*. McGraw-Hill, 1994.

[7] C.P.J. Jansen et al. *Dynamic study of Elia and TenneT power systems in large import situations*. Technical report. TenneT & Elia, 2013.

[8] F.J.C.M. Spaan. *External Constraints Flow Based*. Technical report. TenneT, 2016.

[9] R.A.C.T. de Groot. *Modellen van het Nederlandse hoogspanningsnet en de omliggende landen ten bate van stabiliteitsstudies aan het 220 en 380 kV net*. Technical report. KEMA, 2013.

[10] P. Schavemaker and L. van der Sluis. *Electrical Power System Essentials*. Wiley, 2008.

[11] A.E. Fitzgerald, C. Kingsley, Jr., and S.D. Umans. *Electric Machinery*. 6th. McGraw-Hill, 2003.

[12] P.C. Sen. *Principles of Electric Machines and Power Electronics*. 2nd. Wiley, 1997.

[13] W.W. Price et al. "Load Representation for Dynamic Performance Analysis". In: *IEEE Trans. Power Syst.* 8.2 (1993).

[14] *Recommended Practice for Protection and Coordination of Industrial and Commercial Power Systems*. IEEE Standard 242-1986. IEEE. 1986.

[15] *Distribution Automation Handbook*. ABB. 2010.

[16] J.E. DeCastro et al. "Stall Protection of Large Induction Motors". In: *IEEE Trans. Ind. Appl.* 31.5 (1995).

[17] F. Nozari, M.D. Kankam and W.W. Price. "Aggregation of Induction Motors for Transient Stability Load Modelling". In: *IEEE Trans. Power Syst.* 2.4 (1987).

[18] *PSS®E 34.2 Program Application Guide*. Siemens PTI. 2017.

[19] L.M. Vargas et al. "Load Modeling of an Induction Motor Operated with a Variable Frequency Drive". In: *IEEE Elect. Power & Energy Conf.* 2008.

[20] X. Liang and W. Xu. "Modelling Variable Frequency Drive and Motor Systems in Power Systems Dynamic Studies". In: *Ind. Applicat. Soc. Annu. Meeting*. 2013.

[21] T. Petrů. "Modelling of Wind Turbines for Power System Studies". PhD thesis. Chalmers tekniska högskola, 2003.

[22] E.H. Camm et al. "Characteristics of Wind Turbine Generators for Wind Power Plants". In: *Power and Energy Soc. General Meeting*. 2009.

[23] J.A. Bos. "Connection of large-scale wind power generation to the Dutch electrical power system and its impact on dynamic behaviour". Master's thesis. TU Delft, 2008.

[24] L. Lindgren, J. Svensson and L. Gertmar. *Generic models for wind power plants: Needs and previous work*. Technical report. Elforsk, 2012.

[25] P. Pourbeik. *Proposed Changes to the WECC WT4 Generic Model for Type 4 Wind Turbine Generators.* Technical report. EPRI, 2011.

[26] *Wind Power Plant Dynamic Modelling Guide.* WECC. 2010.

[27] A. Ellis et al. "Generic Models for Simulation of Wind Power Plants in Bulk System Planning Studies". In: *Power and Energy Soc. General Meeting.* 2011.

[28] A. Honrubia-Escribano et al. "Generic dynamic wind turbine models for power system stability analysis: A comprehensive review". In: *Renewable and Sustainable Energy Reviews* (2017).

[29] M. Seyedi. "Evaluation of the DFIG Wind Turbine Built-in Model in PSS/E". Master's thesis. Chalmers tekniska högskola, 2009.

[30] V. Vittal and R. Ayyanar. *Grid Integration and Dynamic Impact of Wind Energy.* Springer, 2013.

[31] *PV Power Plant Dynamic Modelling Guide.* WECC. 2014.

[32] A. Ellis, M. Behnke and R. Elliott. *Generic Solar Photovoltaic System Dynamic Simulation Model Specification.* Technical report. Sandia, 2013.

[33] G. Chaspierre. "Dynamic Modeling of PV Units in response to Voltage Disturbance". PhD thesis. Université de Liège, 2016.

[34] N. Mohan, T.M. Undeland and W.P. Robbins. *Power Electronics: Converters, Applications, and Design.* 2nd. Wiley, 1995.

[35] E.W. Kimbark. *Direct Current Transmission.* Volume 1. Wiley, 1971.

[36] B.K. Johnson. "HVDC Models Used in Stability Studies". In: *IEEE Trans. Power Delivery* 4.2 (1989).

[37] L. Vanfretti et al. "iTesla Power Systems Library (iPSL): A Modelica library for phasor time-domain simulations". In: *SoftwareX* 5 (2016).

[38] H. Krishnappa. "Model validation and feasibility analysis of Modelica-based dynamic simulations using OpenIPSL and CGMES". Master's thesis. TU Delft, 2017.

[39] *PSS®E 34.2 Program Operation Manual.* Siemens PTI. 2017.

[40] F. Troalen. *3D diagrams of Eurostag generator models for improved dynamic simulations.* Technical report. RTE, 2017.

[41] N. Nenadovic. *Kortsluitberekeningen 2010-2016.* Technical report. TenneT, 2010.

[42] J. Rittiger. *PSS/E HVDC model 'SIEBNC'.* Technical report. Siemens, 2011.

[43] R. Dimitrovski et al. "Dynamic Network Security Assessment Model for the Interconnected Power System of Continental Europe". In: *CIGRÉ Int. Symp.* 2017.

[44] A. Banerjee et al. "Motor Protection Principles". In: *Annu. Conf. Protective Relay Engineers.* 2008.

[45] W.W. Price et al. "Standard Load Models for Power Flow and Dynamic Performance Simulation". In: *IEEE Trans. Power Syst.* 10.3 (1995).

[46] *Reliability Guideline: Distributed Energy Resource Modelling.* NERC. 2017.

[47] K. Stray. "Ny HVDC-kapasitet og vindkraft i Sør-Norge". Master's thesis. NTNU, 2009.

[48] M. El Chehaly. *Lower Churchill Project – Harmonic Impedance Studies.* Technical report. SNC-Lavalin, 2012.

[49] *Commission Regulation (EU) 2016/631 establishing a network code on requirements for grid connection of generators.* 2016.

[50] TenneT. *Netkaart Nederland.* Online, accessed 16-01-2017. Available: https://www.tennet.eu/fileadmin/user_upload/Company/Publications/Gridmaps/NL/Gridmap_Netherlands_NL.pdf.

[51] TenneT. *Nederlands transportnet, voorjaar 2014.* 2014.

Appendix A

Structure of the Dutch power grid

First of all, a word of thanks to Loe von Berg, who provided the lion's share of information for this appendix.

A.1 Overview

Originally, each province in the Netherlands constructed and operated its own electrical transmission grid. The grids of the northeastern provinces (Frisia, Groningen, Drenthe and Overijssel) had an operating voltage of 110 kV, the other provincial grids used 150 kV. In the 1950s and 1960s, an extra-high-voltage (EHV, meaning here 220 and 380 kV) grid was constructed to interconnect the different provincial grids; this was called the *koppelnet*, Dutch for "connecting grid". Although the EHV grid was originally meant as an add-on, allowing different provinces to support each other in emergencies, the provincial grids gradually became more and more connected physically, and eventually lost all operational autonomy as well. As of today, all electrical infrastructure at voltage levels of 110 kV or higher is operated and maintained by TenneT, the Dutch transmission system operator (TSO).

Nevertheless, the historical provincial grids are still very clearly visible in the structure of the Dutch power system today. TenneT operates an EHV grid at 220 kV (in the northeast of the country) and 380 kV (everywhere else), as well as seven *regional grids* at 110 or 150 kV. These regional grids are not connected to each other directly at HV (110 or 150 kV) level, only through the EHV grid.

Physical transmission lines connecting certain regional grids at HV level do exist, but they are disconnected in normal operation. They can be used in emergencies to "shift" the stations at either end from one grid to another. For example, the Zealand/Brabant grid is connected to the Limburg grid by the line Maarheeze-Nederweert; in case of a calamity in Brabant, Maarheeze can be connected to the Limburg grid (and disconnected from Zealand/Brabant) and vice versa for Nederweert in case of a calamity in Limburg.

Table 34 lists the seven regional grids with their main characteristics. The difference between *Petersen grounding* and *solid grounding* is only of interest for transient studies; however, it has operational consequences in that Petersen-grounded and solidly-grounded equipment cannot be connected. The possible shifting of stations from Petersen-grounded Limburg to the solidly-grounded FGU grid through the Haps-Teersdijk line seems to contradict this – however, in this case only ungrounded transformers are shifted, avoiding compatibility problems. (In the past, when South Holland was Petersen-grounded, the same applied to the Sassenheim-Haarlemmermeer link.)

The boundaries of the different regional grids do not always line up with the provincial borders; most notably, the northern part of Flevoland (the *Noordoostpolder*) is connected to the GDO rather than the FGU grid, because this area was part of Overijssel at one point.

In this appendix, we shall have a brief look at each of the regional grids in sections A.2 through A.8, but first let us consider the EHV "backbone" of the Dutch power system.

Grid	Voltage [kV]	Grounding	Connection points to EHV (220/380 kV) grid	Direct connections to other HV grids
Zealand/Brabant	150	Petersen	Borssele, Geertruidenberg, Eindhoven	Limburg (Maarheeze-Nederweert)
Limburg	150	Petersen	Boxmeer, Maasbracht	Zealand/Brabant (Nederweert-Maarheeze); FGU (Haps-Teersdijk)
Flevoland/Gelderland/ Utrecht (FGU)	150	solid	Dodewaard, Doetinchem, Lelystad	Limburg (Teersdijk-Haps); North Holland (Breukelen-Diemen & 's Graveland-Diemen)
South Holland	150	solid	Krimpen, Crayestein, Simonshaven, Maasvlakte, Westerlee, Wateringen, Bleiswijk	North Holland (Sassenheim-Haarlemmermeer)
North Holland	150	solid	Diemen, Oostzaan, Beverwijk	South Holland (Haarlemmermeer-Sassenheim)
Groningen/Drenthe/ Overijssel (GDO)	110	solid	Hengelo, Hessenweg, Meeden, Weiwerd, Vierverlaten, Zeyerveen	Frisia (Luttelgeest-Lemmer)
Frisia	110	solid	Bergum, Louwsmeer, Oudehaske	GDO (Lemmer-Luttelgeest)

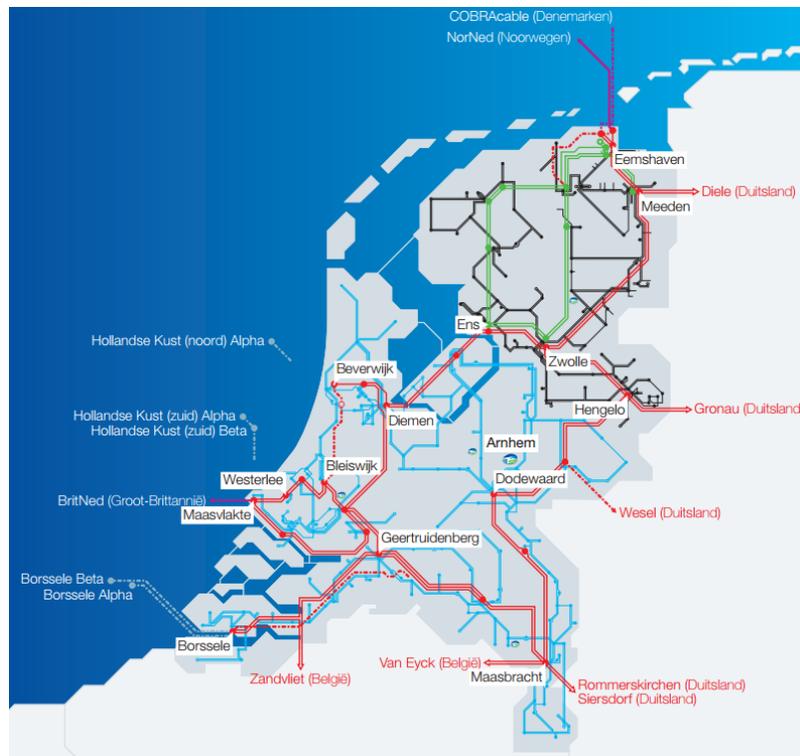Table 34: The regional grids in the Dutch power system

Figure 80: Simplified map of the Dutch power system; dotted lines indicate planned connections [50]

Figure 80 shows a highly simplified map of the Dutch EHV and HV grid. (As in all figures in this text, 380 kV lines are drawn in red, 220 kV lines in green, 150 kV in blue and 110 kV in black; purple indicates HVDC connections.) At its heart is the *380 kV ring*, which includes the stations Krimpen aan den IJssel (or Krimpen for short), Diemen, Lelystad, Ens, Zwolle, Hengelo, Doetinchem, Dodewaard, Boxmeer, Maasbracht, Eindhoven and Geertruidenberg. The loop structure is very robust, as an outage on any line will still leave a path between any two stations in the loop. Hence, it has been applied in other parts of the grid as well: a second 380 kV ring exists around the Rijnmond area (Rotterdam and surroundings); in the northeast of the country, there is a 220 kV ring (Hessenweg-Zeyerveen-Vierverlaten-Bergum-Louwsmeer-Oudehaske-Ens).

International interconnections – known as *tie lines* – to Belgium, Germany, Great Britain (BritNed) and Norway (NorNed) exist at the EHV level; at the HV level there are no interconnections between different TSOs' control areas as this would complicate system operations considerably. An undersea cable to Denmark (COBRA) is under construction.

A.2 Zealand/Brabant

Figure 81: Detail of the TenneT grid map for 2014, showing Zealand within the Zealand/Brabant grid [51]

Figure 82: Detail of the TenneT grid map for 2014, showing Brabant within the Zealand/Brabant grid [51]

Although Zealand and Brabant are currently operated as a single grid, they can be seen as two quite distinct "sub-grids" *(deelnetten)* with their own characteristics, and the line Rilland-Woensdrecht as the connection between them.

Zealand is a small province with few inhabitants, and thus quite little load, although there is considerable industrial activity around Terneuzen, and between Borssele and Vlissingen. A large amount of generation is connected at Borssele. The town's name is synonymous with the only nuclear power plant in the Netherlands, and a large coal-fired plant is located there as well, but the latter is on its way to the exit. Large offshore wind farms – which will also be connected at Borssele – are moving to replace them, however. Thus we are in the inconvenient situation that a lot of power is generated in the "remote corner" Borssele, with no large loads nearby and few connections to the wider grid.

To alleviate this, TenneT is working on the *Southwest 380 kV* project, which will include a 380 kV connection at the Rilland station, a doubling of the 380 kV circuit Borssele-Rilland, and a new 380 kV line from Rilland to somewhere near Tilburg (the exact route has not been determined yet). Once Southwest 380 kV is complete, the 150 kV line Rilland-Woensdrecht will be disconnected, so that Zealand and Brabant become separate 150 kV grids.

Brabant is quite a different story. It is a larger province with a handful of large cities (Eindhoven, Tilburg, Breda, Den Bosch) and several thermal power plants, such as the Amer plant – however, most of them have been switched off for economic reasons.

A.3 Limburg

Figure 83: Detail of the TenneT grid map for 2014, showing the Limburg grid [51]

The Limburg grid contains a decent amount of residential and industrial load, but hardly any generation – only the Swentibold plant in Geleen (near Maasbracht). This means a lot of power has to be "imported" from the EHV grid. Therefore Limburg's two connection points to the EHV grid, Boxmeer and Maasbracht, are serious bottlenecks; especially Boxmeer is vulnerable to outages as it has only one transformer.

The Haps-Teersdijk line can be used to connect Limburg to the FGU grid. (Cuyk is connected to Haps, and thus part of Limburg, in normal operation.)

A.4   FGU

Figure 84: Detail of the TenneT grid map for 2014, showing the FGU grid [51]

The FGU grid is spread out quite widely, with a relatively even distribution of load. Some wind power is connected in Flevoland, but thermal generation – although physically available – is in decline for economic reasons. A known issue in the FGU grid is loading of the north-south connections between Lelystad and Dodewaard with transit flows.
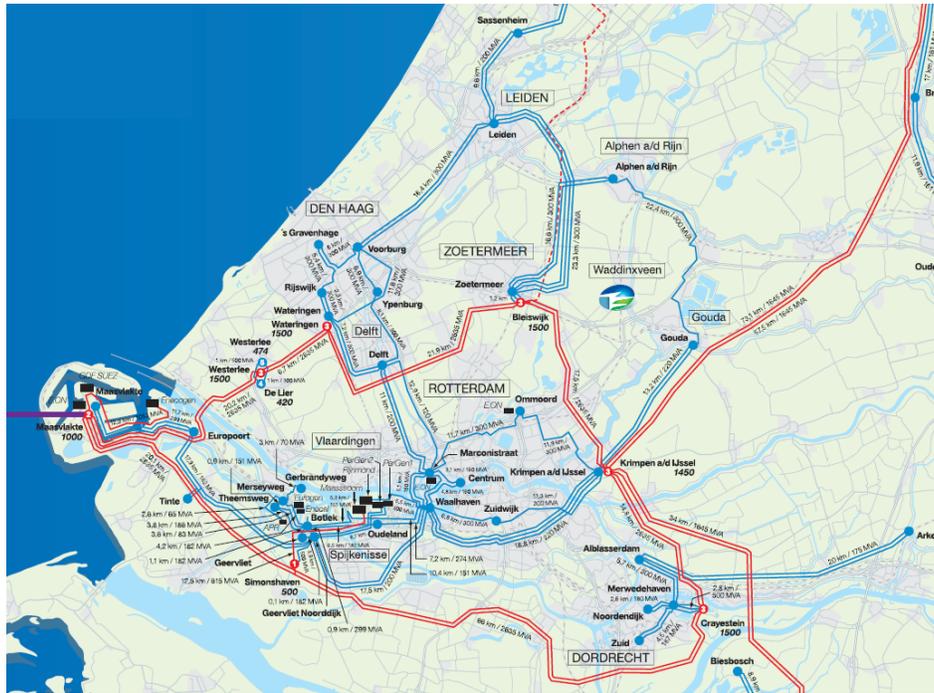
A.5   South Holland

Figure 85: Detail of the TenneT grid map for 2014, showing the South Holland grid [51]

South Holland is easily the most heavily meshed grid in the Netherlands, with more EHV connection points than any of the other regional grids. With Rotterdam and The Hague it contains two of the largest cities in the country, and it is heavily urbanised in general, with several medium-sized cities and towns such as Leiden, Delft and Gouda. Moreover, the port of Rotterdam is the largest in Europe and contains a massive amount of industrial load. There is also ample generation, concentrated in the industrial areas of the Rotterdam port, particularly the Maasvlakte. The Westland region (around the stations Westerlee and De Lier) contains many greenhouses with CHP units.

A disadvantage of the heavy meshing and many EHV connections is that the short-circuit power levels can become unsettlingly high. As beneficial as this is to stability, it can lead to unacceptably high short-circuit currents in the event of an actual fault. To counteract this, the Botlek-Waalhaven line is disconnected in normal operation. Thus the South Holland grid effectively consists of four unconnected pockets: the port of Rotterdam, the "regular" South Holland grid (which includes Rotterdam itself), the small Westland pocket (Westerlee and De Lier) and the Crayestein pocket.

The line Sassenheim-Haarlemmermeer, which connects South Holland to North Holland, is mostly used to shift offshore wind farms on the North Sea from one grid to the other in response to their changing infeed.

A.6   North Holland

Figure 86: Detail of the TenneT grid map for 2014, showing the North Holland grid [51]

North Holland lies on the boundary of two different worlds. The southern part of the province, around Amsterdam and Haarlem, is part of the extremely densely populated *Randstad*, and thus presents a huge residential load. (Amsterdam has a port as well, but it is insignificant compared to the port of Rotterdam.) The steelworks at Velsen forms a large industrial load, but also contributes to generation, with a plant driven by its waste heat. Another significant concentration of generation is connected to the Hemweg station in Amsterdam.

By contrast, the northern part of North Holland is sparsely populated, except for the medium-sized cities of Alkmaar and Den Helder. A significant amount of onshore wind is installed in this region. The north-south divide is clearly reflected in figure 86, showing a clutter of stations and lines around Amsterdam, and large stretches of "empty" land in the north.

Transporting all the generated power from Velsen can be an issue. TenneT studies have shown that if one of the three routes Velsen-Oostzaan (150 kV), Velsen-Beverwijk-Oostzaan (150 $\rightarrow$ 380 kV) and Velsen-(Waarderpolder-)Vijfhuizen-Nieuwe Meer (150 kV) is out of service, an additional outage on Velsen-Oterleek can lead to instability.
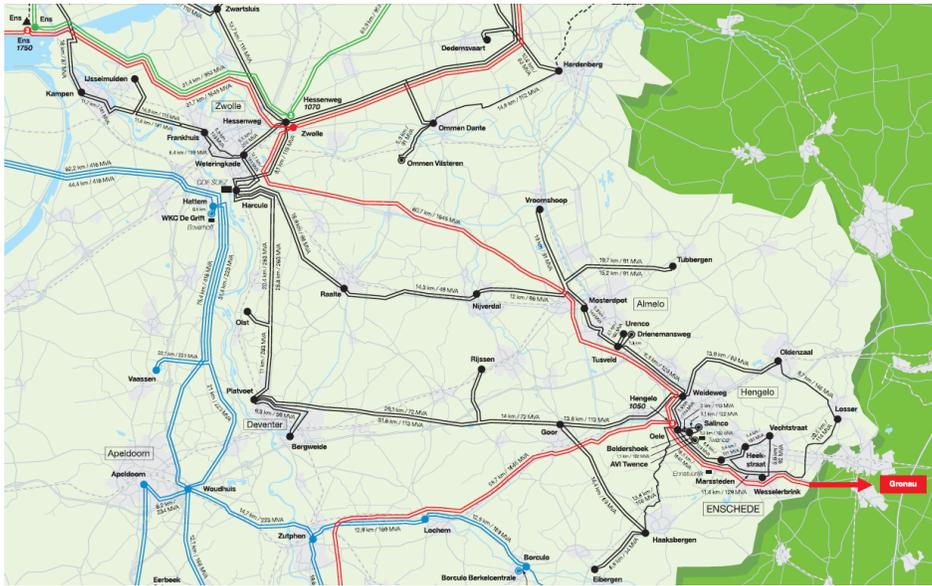
A.7 GDO

Figure 87: Detail of the TenneT grid map for 2014, showing Overijssel within the GDO grid [51]

Figure 88: Detail of the TenneT grid map for 2014, showing Groningen and Drenthe within the GDO grid [51]

The load on the GDO grid is concentrated in the larger cities of Overijssel (Zwolle, Deventer, Enschede); the city of Groningen; and the natural gas wells of the NAM around Slochteren. There is a massive concentration of generation (thermal and wind) at the Eemshaven, which is also where the NorNed and (in the future) COBRA cables are connected; transporting all this power from the relatively remote Eemshaven node causes TenneT considerable headaches. However, these are specifically EHV grid issues, and thus outside the scope of this description.

A small amount of generation is dotted around the 110 kV GDO grid as well, mostly in southern Drenthe.

A.8   Frisia

The Frisian grid is a rather small and simple grid. Frisia being a sparsely populated province, the load is not very large. There is only one thermal power plant of any significant size in Frisia, but onshore wind power is on the rise in this mostly rural province; a planned offshore wind farm in the IJsselmeer will also be connected to the Frisian grid.

Figure 89: Detail of the TenneT grid map for 2014, showing the Frisian grid

Appendix B

Abbreviated names of substations

Abbreviation	Full name	Highest voltage level [kV]
BSL	Borssele	380
CBN	Cambron	50
EWD	Ellewoutsdijk	150
GSE	Goes Evertsenstraat	50
GSP	Goes de Poel	150
HAS	Hydro Agri Sluiskil	150
HST	Thermphos (Hoechst)	150
KNG	Kruiningen	150
KRK	Kreekrak	380
MDB	Middelburg	150
OTL	Oosterland	50
PCN	Pechiney	150
RLL	Rilland	150
SLO	Sloe	150
SVG	Sas van Gent	50
TLN	Tholen	50
TNZ	Terneuzen	150
TZE	Terneuzen Elsta	50
TZO	Terneuzen Oost	150
TZZ	Terneuzen Zuid	50
VSG	Vlissingen	150
WAP	Willem Anna Polder	150
WDO	Westdorpe	150
ZRZ	Zierikzee	50

Table 35: Names of substations in Zealand. Oosterland is not to be confused with Oterleek, a 150 kV station in North Holland also abbreviated OTL. Terneuzen Oost is also denoted TNO instead of TZO.

Abbreviation	Full name	Highest voltage level [kV]
AB	Alblasserdam	150
AK	Arkel	150
AP	Alphen aan den Rijn	150
BTL	Botlek	150
BWK	Bleiswijk	380
CST	Crayestein	380
DDM	Dordrecht Merwedehaven	150
DDN	Dordrecht Noordendijk	150
DDZ	Dordrecht Zuid	150
DE	Delft	150
DLR	De Lier	380
ERP	Europoort	150
GBW	Gerbrandyweg	150
GD	Gouda	150
GV	's-Gravenhage (Den Haag)	150
GVN	Geervliet Noorddijk	150
GVT	Geervliet	150
KIJ	Krimpen aan den IJssel	380
LD	Leiden	150
LYO	Lyondell	11
MDH	Middelharnis	150
MSW	Merseyweg	150
MVL	Maasvlakte	380
ODL	Oudeland	150
OM	Ommoord	150
REC	Rijnmond Energie CV	150
RTC	Rotterdam Centrum	150
RTM	Rotterdam Marconistraat	150
RTN	Rotterdam Noord	150
RTW	Rotterdam Waalhaven	150
RTZ	Rotterdam Zuidwijk	150
RW	Rijswijk	150
SMH	Simonshaven	380
SS	Sassenheim	150
TNT	Tinte	150
TWG	Theemsweg	150
VB	Voorburg	150
VLW	Vondelingenweg	150
WL	Westerlee	380
WTR	Wateringen	380
YP	Ypenburg	150
ZT	Zoetermeer	150

Table 36: Names of substations in South Holland. Geervliet is also denoted GVL instead of GVT. Krimpen aan den IJssel is also denoted KP instead of KIJ.