

Continuous Trading Matching Algorithm

Public Description

Table of contents

- 1. Introduction..... 3**
 - 1.1. Continues trading matching algorithm 3**
 - 1.2. Single intra-day coupling algorithm..... 3**
 - 1.2.1. SOB 3
 - 1.2.2. CMM 3
 - 1.2.3. Market coupling and implicit capacity requests..... 3
 - 1.2.4. Explicit capacity requests 3
- 2. Power transmission network..... 4**
 - 2.1. Delivery areas and market areas 4**
 - 2.2. NEMO hubs 4**
 - 2.3. Borders and interconnectors..... 4**
 - 2.4. Capacity information 4**
 - 2.5. Ramping constraints 5**
- 3. Power markets 7**
 - 3.1. Products and contracts 7**
 - 3.2. Regular orders..... 7**
 - 3.3. Iceberg orders 7**
 - 3.4. Block orders 8**
 - 3.5. Execution restrictions 8**
 - 3.5.1. NON 8
 - 3.5.2. IOC 8
 - 3.5.3. FOK..... 8
 - 3.5.4. AON..... 8
 - 3.6. Linked orders 9**
 - 3.7. Explicit capacity requests..... 9**
- 4. Order book creation..... 10**
 - 4.1. Consolidated order books..... 10**
 - 4.2. Cross-Border trading 10**
 - 4.3. Timestamp..... 11**
 - 4.4. Price-Time-Priority principle 12**
- 5. Order matching..... 13**
 - 5.1.1. Price determination 13
 - 5.1.2. Iceberg orders in regular matching 13
 - 5.1.3. Matching against multiple orders 13
 - 5.1.4. Unmatchable orders..... 14
 - 5.2. Trade creation..... 15**
 - 5.3. Trade cancellation 15**
 - 5.4. Priority setting between implicit and explicit allocation..... 16**

6.	Routing	17
6.1.	Need for routing calculation	17
6.2.	Terminology	17
6.2.1.	Interconnector	17
6.2.2.	Path	17
6.2.3.	Available transfer capacity (ATC)	17
6.2.4.	Flow	18
6.2.5.	Allocated capacity	18
6.2.6.	Ramp rate	18
6.2.7.	Routing plan	19
6.2.8.	Capacity allocation	19
6.3.	Routing model basics	19
6.3.1.	Balance condition	19
6.3.2.	Capacity condition	19
6.3.3.	Ramping condition	19
6.4.	Routing model complications	20
6.4.1.	Complexity introduced by the ramping constraint	20
6.4.2.	Single-period routing with ramping	20
6.4.3.	Error! Bookmark not defined.	

1. Introduction

1.1. Intra-Day market coupling

The intra-day power market is a continuous market allowing 24/7 power trading. Coupling of the national intra-day markets is provided by the single intra-day coupling (SIDC) algorithm.

1.2. Continues trading matching algorithm

The continues trading matching algorithm, called hereafter single intra-day coupling algorithm is incorporated in the XBID (Cross-Border Intra-Day) solution. The XBID solution comprises, among other components, two modules which each perform part of the algorithm tasks: the shared order book (SOB) module and the capacity management module (CMM).

1.2.1. SOB

The SOB module contains the basic functionality for continuous trading, like order entry, order management and order matching. Tradable electricity contracts are automatically generated in the SOB module and made available for trading based on a predefined trading schedule.

1.2.2. CMM

The CMM provides the functionality for managing and allocating available transmission capacity between all areas in the underlying power transmission network (which sometimes is called delivery grid).

1.2.3. Market coupling and implicit capacity requests

The XBID solution enables multiple power exchanges to connect to the central SOB module. Traders cannot connect to the SOB directly. Instead, they connect to one or more trading solutions of power exchanges (NEMOs). Orders on contracts tradable in XBID which are entered into a trading solution of a power exchange are transmitted to the SOB, where they can be matched against each other. SOB always performs these matches regardless whether the orders were entered in the same bidding zone or in different bidding zones and regardless of the power exchange that transmitted the orders – provided, of course, that a match is possible in the first place.

The SOB module maintains a consolidated order book across all participating bidding zones and NEMOs, based on available transmission capacity between bidding zones. The CMM provides the current capacity information. When cross-border trades are concluded, the required cross-border capacity is implicitly allocated in the CMM.

1.2.4. Explicit capacity requests

Explicit market participants can directly access the CMM to transmit explicit capacity requests.

2. Power transmission network

The relevant characteristics of the power transmission network are made available to the CMM by means of configuration settings (delivery areas, market areas, NEMO hubs, borders, interconnectors, ramping constraints, virtual areas) and input data (capacity information).

2.1. Delivery areas and market areas

A delivery area represents a control zone as part of an energy network operated by one TSO. The XBID term delivery area is a synonym to the CACM term scheduling area.

Each delivery area is assigned to a market area, which represents an uncongested price area. The XBID term market area is a synonym to the CACM term bidding zone. A delivery area can only be assigned to one market area. Any one market area may contain one or more delivery areas.

Orders are always entered into a delivery area. Trade records contain the source and the destination delivery area (which may be the same).

Market areas serve as a container for one or more delivery areas. The main function of market areas is the separation of market zones under the assumption that the transmission capacity between market areas is subject to congestion. Between delivery areas within the same market area, the system will always assume unlimited transmission capacity.

2.2. NEMO hubs

More than one NEMO (power exchange) can be assigned to the same delivery area. NEMO markets within the same delivery area are referred to as NEMO hubs and the solution allowing for several NEMOs to be present in a single delivery area is commonly referred to as a virtual hub solution. There is no technical limitation to the number of NEMO hubs that can be present in the same delivery area.

2.3. Borders and interconnectors

A border is defined as a connection between two market areas. An interconnector is defined as a connection between two delivery areas. Because of these definitions, there can be multiple interconnectors per border.

2.4. Capacity information

Capacity information per interconnector or border is made available per direction by the relevant TSO(s) as daily input information. This information is provided as a combination of NTC and AAC values. The NTC value (Net Transfer Capacity) represents the physical transfer capacity of the power line(s). The AAC value represents the already allocated capacity.

Based on this input, the algorithm calculates the available transmission capacity (ATC) for the border or interconnector at hand according to the below formulas.

$$ATC_{A \rightarrow B} = NTC_{A \rightarrow B} - AAC_{A \rightarrow B} + AAC_{B \rightarrow A} - \text{Intraday Allocations}_{A \rightarrow B} - \text{BM/GenOutage}_{A \rightarrow B} + \text{Intraday Allocations}_{B \rightarrow A}$$

$$ATC_{B \rightarrow A} = NTC_{B \rightarrow A} - AAC_{B \rightarrow A} + AAC_{A \rightarrow B} - \text{Intraday Allocations}_{B \rightarrow A} - \text{BM/GenOutage}_{B \rightarrow A} + \text{Intraday Allocations}_{A \rightarrow B}$$

Where the BM and GenOutage values refer to potential balancing mechanism or GenOutage allocations that should be taken into account when calculating the ATC.

On borders with more than one interconnector, two approaches for capacity information are possible. Either capacity information is assigned to each interconnector separately or capacity information is valid jointly for all interconnectors on the border ('common ATC').

2.5. Ramping constraints

Ramping constraints are restrictions that are applied to interconnectors for grid security purposes. They ensure that the amount of power going down an interconnector does not change too rapidly from one time period to the next and effectively act as a further restriction on the available capacity.

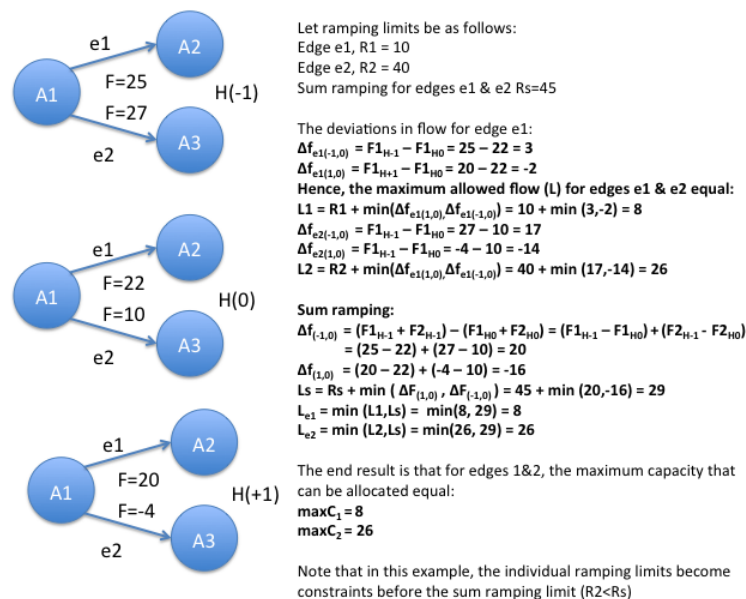
Ramping per border is defined by a value which means that total allocated capacity cannot vary more than this value from one time interval to another (previous and next time interval to be considered).

A ramp rate on an interconnector is always valid in both directions.

Intra-day trading on an interconnector or border with ramp rates is only allowed if the day-ahead flows on the traded contract as well as those on the preceding and following contract are known. As a result, the last contract of the trading day is only made tradable once the day-ahead values for the following day are available to the algorithm. This may be at a moment later than the moment that the border opens for intra-day trading.

Ramp rates can also be applied to groups of interconnectors and to borders with more than one interconnector where the 'common ATC' principle is applied (i.e. the capacity assigned to the border is distributed evenly over all interconnectors on that border, cf. section 2.4).

Ramping on a group of interconnectors (shared ramping) follows the same principle as for ramping per border (interconnector) but considers the sum of allocated capacities for the group of interconnectors concerned. The ramping limits for the individual interconnector under shared ramping constraint can differ from one another. Furthermore, the shared ramping limit is a value which can differ from the individual interconnector limits.



$$\text{CapUp}_h = \min(\text{ATCUp}_h, \text{flow}_{(h-1)} + \text{ramp}, \text{flow}_{(h+1)} + \text{ramp})$$

$$\text{CapDown}_h = \min(\text{ATCDown}_h, \text{flow}_{(h-1)} - \text{ramp}, \text{flow}_{(h+1)} - \text{ramp})$$

Where CapUp_h is the capacity to be displayed in one direction, ATCUp_h is the updated capacity in one direction (update performed by TSOs), CapDown_h is the capacity to be displayed in the opposite direction, ATCDown_h is the updated capacity in opposite direction, and. Ramp is the ramp rate.

3. Power markets

The relevant characteristics of the power markets are made available to the algorithm by means of configuration settings (products and contracts) and input data (orders and explicit capacity requests).

3.1. Products and contracts

A product can be thought of as a unique set of trading features – the most important of which is the delivery duration (referred to as ‘delivery period’ in the documentation) – or, alternatively, as a template for a contract. A contract identifies the actual object of an order and a trade, and can also be referred to as an instrument.

'Hourly' would be an example of a product; it is the product for which the delivery period lasts exactly one hour. The contracts generated by this product are hour 1, hour 2 ... hour 24 of any given calendar day.

Products have a trading unit (MW will be used, but it could also be kW) and a trading currency (which could be anything, even though only EUR will be used). Products also have a minimum and maximum contract quantity and a minimum and maximum price.

The delivery duration of a product must be chosen in such a way that 24 hours is an exact multiple of these delivery durations. Gaps between the contracts generated from a product are allowed. Thus, a peak contract of 12 hours may be separated by a gap of 12 hours from the peak contract of the next day.

3.2. Regular orders

A regular order (also referred to as limit order) represents an instruction issued by a market participant to a power exchange to either buy or sell a certain quantity of a certain contract at a price with a certain limit (i.e. a minimum price for supply orders and a maximum price for demand orders) in a given delivery area.

3.3. Iceberg orders

Iceberg orders are limit orders which are only visible with part of their total quantity in the market, while their full quantity is exposed to the market for matching, and which may come with a range of limit prices.

For each iceberg order, a peak size and a total quantity are defined. An iceberg order is displayed in the market with a quantity equal to its peak size (if the order is not matched immediately). The quantity visible in the order book, is called the shown quantity. The shown quantity of an iceberg order is always identical to its peak size when it is newly entered and not matched.

When an iceberg order is matched in a trade, its total quantity is reduced by the trade quantity. If the shown quantity before the trade was greater than the subtracted amount, the order remains visible in the market with the remaining shown quantity. If the shown quantity before the trade was less than or equal to the subtracted amount, a new slice of order quantity is made available in the market sized at the peak size of the iceberg order or at the remaining quantity, if this is smaller than the peak size. When the quantity of the last slice has been reduced to zero, the iceberg order is fully executed and removed from the order book.

Iceberg orders can be entered with a peak price delta. Each new slice will then be entered with a new limit price which is reduced by the peak price delta for buy orders and increased by the peak price delta for sell orders.

3.4. Block orders

By definition, a block order is an order with a total delivery period greater than the smallest delivery period of the product for which the order is being entered. A user-defined block order is an order on multiple predefined contracts belonging to the same product and having consecutive delivery periods. It is the market participant that combines these contracts into one user-defined contract by placing a block order on them. A user-defined block order cannot be an iceberg order.

3.5. Execution restrictions

Each order has one of the following execution restrictions: NON (none), IOC (immediate or cancel), FOK (fill or kill), or AON (all or nothing). Only iceberg orders do not have any of these execution restrictions.

3.5.1. NON

An order submitted with the execution restriction NON is either executed immediately, or, if the order cannot be matched right away, entered into the order book. Partial order executions are allowed and NON orders can be executed against multiple other orders and create multiple trades.

If an order is submitted with no value in the field execution restriction, the value "NON" is assumed.

3.5.2. IOC

An order submitted with the execution restriction IOC is either executed immediately, or, if the order cannot be matched, deleted without entry in the order book. Partial executions are allowed and IOC orders can be executed against multiple other orders and create multiple trades.

An IOC order is never displayed in the order book.

3.5.3. FOK

An order submitted with the execution restriction FOK is either executed immediately and with its full quantity or, if the order cannot be matched with its entire quantity, deleted without entry in the order book. FOK orders can be matched against multiple existing orders in the order book and in that case, create multiple trades.

An FOK order is never displayed in the order book.

3.5.4. AON

An order submitted with the execution restriction AON is either executed against exactly one other order with its full quantity or entered into the order book. Partial executions are not allowed.

The execution restriction AON is only allowed for user-defined block orders and user-defined blocks are always AON.

3.6. Linked orders

Multiple orders can be submitted at once as linked orders (also called basket orders with a 'linked' execution instruction). These orders can be on the same or different contracts, products, delivery areas and market areas. All orders in a linked basket must be executed immediately and with their full quantity (i.e., they always have a FOK execution restriction). If this is not possible, the linked orders will all be deleted.

3.7. Explicit capacity requests

An explicit capacity request is a request to allocate a certain amount of capacity on a given interconnector in a given direction for a given time period.

4. Order book creation

4.1. Consolidated order books

Orders that are submitted into the SOB are either executed immediately, deleted immediately or entered into the order book. Some order types can also be executed partially and be either deleted or entered into the order book for the remaining part. Entry into the order book means that the order is stored in the system and made visible to other users until either the order is deleted or deactivated or a matching order is submitted.

Per contract, the SOB maintains a consolidated order book. This order book contains all orders on its contract entered in all delivery areas where this contract is tradable.

While there only is one consolidated order book for each contract, buy and sell orders inside the same order book can only be matched against each other if either both orders were entered in the same market area or they were entered in market areas that are connected for cross-border trading.

Which orders of a given order book are visible in any given delivery area may differ in function of the available capacity between the delivery areas. Hence, every delivery area has its own local view on an order book, which only displays the orders that can be matched in that delivery area.

Orders on a contract that is tradable in two or more delivery areas in the same market area are displayed in all local order book views belonging to this market area. As a result, also without cross-border trading, the local view of an order book may contain more orders than were entered in the delivery area it belongs to (but only in market areas that contain more than one delivery area).

A new order book is created whenever a new contract becomes tradable. The order book is updated whenever orders are added to it, deleted from it, modified (volume, price, execution restriction, activation of a new iceberg slice), or matched (either partially or fully), or the underlying contract expires or trading on it is halted. Whenever there is a relevant cross-border capacity update or change, the local views of an order book are updated. After each order book update or local view update, the algorithm validates if new matching opportunities have arisen and, if so, the resulting trades are concluded.

4.2. Cross-Border trading

Cross-border trading is trading between different market areas. The borders between market areas are subject to congestion. Trading between market areas is only possible if there is available transmission capacity. When orders from different market areas are matched against each other, the required transmission capacity for the transaction is implicitly allocated by the CMM.

The CMM provides the SOB with ATCs (and in case ramping is involved: offered capacities (OCs)) between all connected market areas. If positive transmission capacity is available between two delivery areas, the orders entered in one of these delivery areas will be displayed in the local order book of the other delivery area.

Local views will be enriched with cross-border orders if sufficient capacity is available. Depending on available transmission capacity, the same order can be displayed in multiple

local views. Orders that allow for partial matching (i.e. orders that do not have the AON execution restriction) will be displayed in the local views of other market areas with part of their quantity if the available transmission capacity is smaller than their full quantity. Orders with the AON execution restriction, which do not allow for partial matching, are always either displayed with their full quantity or not displayed at all remotely. For block orders, the (time unit with the) most restrictive ATC/offered capacity is leading. An order that was visible in multiple local views is removed from all local views after full matching, deactivation or deletion.

For each border, the CMM maintains two ATC values: one for each direction. Based on these two values the SOB calculates for each pair of market areas the maximum volume (in MW) of buy and sell orders belonging one market area that can be shown in the other market area. The total volume of these external orders shown in a market area's local view cannot be greater than the corresponding available capacity in either direction between these two market areas.

Buy orders require the flow of power to the delivery area of the buy order. So, the total quantity of buy orders that are displayed from connected delivery areas cannot exceed the ATC *directed towards* the delivery area of the buy orders.

Sell orders require the flow of power out of the delivery area of the sell order. So, the total quantity of all sell orders that are displayed from connected delivery areas cannot exceed the ATC *directed away from* the delivery area of the sell orders.

If the outgoing ATC value from market area 1 to market area 2 (sum of all possible routes) is a positive value x :

- (1) buy orders from all delivery areas of market area 2 are displayed in the local view of the order books of all delivery areas of market area 1. The maximum volume of all external buy orders displayed in the local views of market area 1 is x .
- (2) sell orders from market area 1 can be displayed in the local view of the order books of all delivery areas of market area 2. The maximum volume of all external sell orders displayed in the local views of market area 2 is x .

If the incoming ATC value to market area 1 from market area 2 (sum of all possible routes) is a positive value y :

- (1) sell orders from all delivery areas of market area 2 are displayed in the local view of the order books of all delivery areas of market area 1. The maximum volume of all external sell orders displayed in the local views of market area 1 is y .
- (2) buy orders from market area 1 can be displayed in the local view of the order books of all delivery areas of market area 2. The maximum volume of all external buy orders displayed in the local views of market area 2 is y .

4.3. Timestamp

When an order or an explicit capacity request is entered into the SOB (orders) or CMM (explicit requests), it receives a timestamp. An order receives a new timestamp whenever its volume, price or execution restriction is modified or its state is changed from inactive to active. Iceberg orders also receive a new timestamp whenever a new slice is activated. All timestamps are assigned sequentially

4.4. Price-Time-Priority principle

Orders are ranked in two processes: the order book calculation process and the order matching process. Order ranking is a deterministic process, which applies the price-time-priority principle. Orders are ranked per side (buy/sell) best price first, meaning that buy orders are ranked from highest to lowest price and sell orders are ranked from lowest to highest price. Orders with the same price limit are prioritized by their time stamp oldest first.

4.5. Local view calculation

Orders from other market areas are selected based on available capacity and their ranking according to the price-time-priority principle. Iceberg orders are displayed with their visible quantity and not with their total quantity. AON orders can only be displayed with their full quantity; other orders can also be displayed with a fraction of their quantity.

The local view for each delivery area is calculated independently. Consequently, the display of an order in one delivery area does not influence the display of the same order in another delivery area. Local views are per delivery area and contract. The local views of the delivery areas within the same market area are all the same.

The views for buy and sell side of a local order book are calculated independently. Local views will be published in parallel, and buy and sell will be merged per area.

5. Order matching

Order matching is a deterministic process in the XBID system, which results in the conclusion of a trade.

For orders to match, they must be of different sides (buy/sell) and on the same contract. There is no cross-matching between orders on different contracts. In addition, the price limit of the sell order must be equal to or below that of the buy order (i.e. the intersection of the two order execution ranges may not be empty).

When an order is matched in a trade, its quantity is reduced by the trade quantity. Orders with the execution restriction FOK or AON can only be matched with their full quantity. Orders with the execution restriction NON or IOC can also be matched partially.

Matching of orders is based on their ranking, which follows the price-time-priority principle (cf. section 4.4). Orders are always executed at the best possible price, meaning that the best (i.e. highest-price) buy order is always matched with the best (i.e. lowest-price) sell order first. Orders with the same price limit are prioritized by their time stamp oldest first.

The ID algorithm supports two different matching processes: regular matching and batch matching. Regular matching (section 5.1) is triggered by the entry of an order with a new time stamp. Batch matching (section 5.2) is triggered by an update of the cross-border capacity.

5.1. Regular matching

Regular matching is triggered by the entry of an order with a new time stamp. An order with a new time stamp may be a newly entered order, a modified order, a (re)activated order that was inactive before, or a new slice of an iceberg order.

5.1.1. Price determination

When two orders are matched in a regular match, one of these orders must be an order with a new timestamp and the other one must be an order already present in the order book.

The price at which two orders are matched becomes the price of the trade that is concluded. Two orders are matched at the limit price of the order that was already in the order book. If a buy order with a new timestamp is matched against an existing sell order, the limit price of the sell order becomes the trade execution price. If a sell order with a new timestamp is matched against an existing buy order, the limit price of the buy order becomes the trade execution price.

5.1.2. Iceberg orders in regular matching

In a matching process where a single order with a new timestamp is matched against more than one slice of an iceberg order already in the order book, the price is always determined by the iceberg order already in the order book and never by the order with the new timestamp, even if the timestamp of the iceberg order is renewed during the matching process.

5.1.3. Matching against multiple orders

If an order with a new timestamp can be executed, it is not necessarily executed at a single price (except orders with the execution restriction AON), but may sequentially generate

multiple transactions at different prices against multiple different orders that already existed in the order book. As soon as the order has been executed against all orders at a certain price limit, the next best price level becomes best and the order continues to be matched against orders entered at this price level. This process continues as long as the incoming order remains executable and has a positive order quantity. Subsequently, the order is deleted if the order quantity has reached zero or if it has the IOC execution restriction. In all other cases, the order is entered into the order book with its remaining quantity.

5.1.4. Unmatchable orders

If an order with a new timestamp cannot be executed against any existing order, it is entered into the order book, unless it has the execution instruction IOC or FOK. If it has the execution restriction IOC or FOK, it is deleted.

5.2. Batch matching (or intra-day auction)

Batch matching is triggered by an increase of the cross-border capacity. Batch matching rounds are also referred to as intra-day auctions, because the matching displays certain auction-like characteristics.

An increase of cross-border capacity may lead to a crossed order book, in other words to sets of orders that were not matchable before due to insufficient cross-border capacity becoming matchable. For user-defined block orders, which cannot be partially matched (because they have the AON execution restriction), this may happen by any capacity increase. For orders that do allow for partial matching, this may only happen if an original zero value is increased, as any non-zero value would already have been used up by partial matching.

In such cases, the regular matching process cannot be used, as there is no order with a new timestamp. All orders that have become matchable were already present in the order book and it is not obvious which one(s) of these should set the price of the trade(s). Therefore, a different matching process is followed, which is referred to as batch matching.

The regular matching of orders, as well as explicit capacity allocation, is suspended for the duration of the batch matching.

5.2.1. Iceberg orders in batch matching

In batch matching, iceberg orders with a peak price delta of zero participate with their total remaining quantity. Iceberg orders with a non-zero peak price delta participate with each slice individually. After the first slice is executed completely the next slice with a new price limit and timestamp will participate until either all slices are fully executed or no further execution is possible anymore.

5.2.2. Contract (or order book) sequence in batch matching

A capacity increase may enable the immediate matching of orders for multiple contracts and for different delivery periods. All orders in the affected order books will participate in the batch matching. The batch matching is performed per contract, as long as capacity is available.

5.2.3. Price determination

All pairs that were matched in the same batch matching round get the same trade price. This trade price is the arithmetic mean of the price limits of the last pair matched in that round.

5.3. Trade creation

A trade represents the obligation to transfer a certain amount of energy from seller to buyer in exchange for the financial trade value (or: at the trade price). A trade also contains information on the one or two delivery areas between which the energy is transferred, the path(s) along which it is transferred, as well as the delivery period of the energy.

Whenever two orders are matched, a trade is created. Even though one order can match with multiple other orders, a trade is always between exactly two orders. The attributes *price*, *quantity*, and *value* of a trade as well as its *timestamp* are derived from the characteristics of the order matching event.

How the price of a trade is established is explained in sections 5.1 (regular matching) and 5.2 (batch matching). How the quantity of a trade is established is explained in the same sections. The financial value of a trade is calculated according to the below formula:

$$V (EUR) = Q (MW) * p (EUR/MWh) * d (h)$$

Where:

V is the value of the trade in Euro.

Q is the quantity of the power traded in Megawatt.


p is the price per energy unit agreed between buyer and seller in Euro per Megawatt hour.

d is the duration of the delivery period of the traded contract in hours.

The timestamp of a trade reflects the moment that one of the orders that will enter into the trade is found to be executable (against the other order) by the algorithm.

5.4. Trade cancellation

A trade may be cancelled subject to the applicable market rules. Trade cancellation is technically only feasible if counter-allocation of the capacity/capacities allocated implicitly in the trade still is possible. It is not necessary that this counter allocation follows the same path that the original allocation followed. However, the allowed counter allocation is restricted to the interconnectors used in original trade routing plan. In the example below, a trade may have been concluded on the paths DA1 – DA2 – DA4 and DA1 – DA3 – DA4, and it may be cancelled by counter-allocation on the path DA4 – DA3 – DA1.



5.5. Priority setting between implicit and explicit allocation

Cross-border trading, which involves the implicit allocation of interconnector capacity, and explicit capacity allocation are processes that, in principle, compete for the same interconnector capacity. (Obviously, this is only true for those interconnectors that allow for explicit capacity allocation). If there is competition for the same capacity between regular matching and explicit allocation, it is the time stamp of the order that triggers regular matching and the time stamp of the explicit allocation request (cf. section 4.3) that determine which of these processes gets the capacity allocated. Both order entry (and modification, activation, and activation of a new slice of an iceberg order) events and explicit capacity requests enter the XBID system through a single queue and get their timestamp assigned in this single queue. As a result, all orders and all explicit capacity requests have unique timestamps, so that it is always possible to sort them by time stamp unequivocally.

Batch matching is triggered by a capacity update. As soon as a batch matching round starts, both regular matching and explicit capacity allocation are suspended. This means that batch matching always takes precedence over both regular matching and explicit allocation. Regular matching and explicit allocation are only enabled again after all batch matching rounds belonging to a single capacity update event have been completed.

6. Routing

The flow of a certain quantity of power between delivery areas may be routable via different routes. The selection of the actual route applied is a deterministic process following certain rules. This process is called routing calculation.

6.1. Need for routing calculation

A routing calculation is performed either in the context of an order book (re)calculation or in that of a trade flow calculation.

When the local view of an order book is updated, the SOB calculates if sufficient transport capacity is available to display orders entered for other delivery areas in the area that the local view belongs to. The order book recalculation is triggered by a change in ATC or by a change in order book contents. An order book recalculation starts no earlier than the previous calculation has terminated. Therefore, not every trigger necessarily produces an order book recalculation, but all triggering data changes that occurred during an order book recalculation will be taken into account in the subsequent recalculation.

Whenever a cross-border trade is created or cancelled, the SOB calculates the necessary flows in the network to transport power between the source and destination area. The result of this process is a routing plan with netted and aggregated power flows, which in turn result in one or more capacity allocations in the CMM. The process is explained in more detail in the remaining part of this chapter.

6.2. Terminology

6.2.1. Interconnector

An *interconnector* $A \rightarrow B$ is a directional connection between source delivery area A and destination delivery area B. Only one interconnector per direction and pair of delivery areas is allowed. If there is an interconnector in one direction, the interconnector in the opposite direction must also exist; it is called the reverse interconnector.

6.2.2. Path

A *path* $A \rightarrow B \rightarrow \dots \rightarrow Z$ is a sequence of distinct delivery areas in the grid where subsequent areas of the path are connected by interconnectors. The first area of a path is called its source; the last area of a path is called its destination. A path cannot contain the same area more than once (no loops).

6.2.3. Border resolution

The *border resolution* is the smallest allowable duration of a power transmission time period on a border.

6.2.4. Available transfer capacity (ATC)

Available Transfer Capacity (ATC) is an interconnector attribute indicating the maximum available amount of power that can be transported in the direction of the interconnector. The quantity varies per period and changes after each capacity allocation.

6.2.5. Flow

A *flow* is an interconnector attribute indicating the flow of power in the direction of the interconnector that is used in the routing calculation. The value cannot be negative and must be smaller than or equal to the ATC. The quantity depends on the period.

6.2.6. Allocated capacity

Allocated capacity is an interconnector attribute indicating the netted and aggregated sum of all capacity allocations on an interconnector and its reverse. This quantity depends on the period.

6.2.7. Cost coefficient

A *cost coefficient* is an interconnector attribute accounting for the mathematical cost of a flow. It is always a positive value. The cost coefficient is independent of the period and determines over which path power should be routed preferably. It has no financial bearing, but is merely a mathematical construct to make distinctions between interconnectors in terms of routing priority. Interconnectors with a lower cost coefficient will be prioritized over interconnectors with a higher cost coefficient. The cost coefficient is direction-independent, i.e. it is the same for an interconnector and its reverse. The default value of a cost coefficient is 1.

6.2.8. Ramp rate

A *ramp rate* is an interconnector attribute indicating the maximum allowable flow change between consecutive time periods. If ramping conditions apply to an interconnector, a ramp rate must be defined. Ramp rates are non-negative, direction-independent and period-independent.

6.2.9. Transport

A *transport* is the transfer of power through the grid, determined by a path and a quantity. A cross-border trade may require the transportation of power on several paths, i.e., a set of transports.

6.2.10. Network flow

A set of transports starting at the same source and ending at the same destination can be assembled into a *network flow*. Network flows can be obtained by merging multiple transports where parallel and opposing flows are combined per interconnector and reverse interconnector. If two transports have an interconnector in common, the resulting network flow on this interconnector is the sum of the flows of the individual transports.

6.2.11. Internal netting

Internal netting applies if the paths of two transports share the same pair of adjacent areas, but have flows in opposite directions, i.e. if a certain interconnector is part of the first transport and its reverse is part of the second transport, then internal netting applies on that interconnector and its reverse.

If the path of one transport includes an interconnector that is the reverse of an interconnector in a second transport, the magnitude of the resulting network flow is the absolute value of

the difference of the two individual transport flows. The direction of the network flow is entirely along the interconnector with the flow value that was larger initially. There is a zero netted flow in the opposite direction. Netting in network flows is allowed since the amount of allocated capacity on an interconnector is immediately made available as additional available capacity on the reverse interconnector.

6.2.12. Routing plan

A collection of network flows for a set of successive periods is called a *routing plan*.

6.2.13. Cost of a power flow

The *cost of a power flow* on an interconnector is defined by the product of the flow and the cost coefficient. The cost of a transport is defined as the sum of the cost of the flows on the interconnectors in the path of the transport. The cost of a network flow is defined as the sum of the products of the netted and aggregated flows on the involved interconnectors and the corresponding cost coefficients. Due to flow netting, the cost of a network flow can be smaller than the sum of the cost of the individual transports.

6.2.14. Capacity allocation

The allocation of network flows for cross-border trades in the CMM is called *capacity allocation* and the respective flows are converted into allocated capacity. Allocations of capacity are netted in the CMM.

6.3. Routing model basics

Every routing model representing the transport of power through a grid must satisfy the following flow constraints: the balance condition, the capacity condition and the ramping condition.

6.3.1. Balance condition

The total influx and efflux of power must be equal in all areas except for the source and the destination area. In case there is one source and one destination area, the total efflux of the source area is equal to the total influx of the destination area, which, by definition, is the transported quantity from source to destination in the period under consideration. This condition must hold for all periods independently.

6.3.2. Capacity condition

A valid flow across an interconnector has a non-negative magnitude and must not exceed the ATC of the interconnector. This condition must hold for all periods independently.

6.3.3. Ramping condition

The netted flows on an interconnector must conform to the ramping condition. Ramping conditions are restrictions for the flows in successive periods on a given interconnector and render the routing of power for successive periods interdependent.

6.4. Routing model complications

A routing plan satisfying all flows constraints is called feasible. It consists of a collection of network flows each satisfying the balance and capacity conditions. Additionally, the netted network flows are compatible with the ramping condition.

A feasible routing plan is a potential candidate for the solution of a routing problem. However, the rules do not uniquely determine the choice of the flows and additional criteria are required to determine the preferred routing plan. The first criterion is the minimum cost flow principle.

6.4.1. Minimum cost flow principle

The XBID routing model applies the minimum cost flow principle to select preferred routing plans. The minimum cost flow principle selects routing plans with minimum cost among all feasible routing plans satisfying the flow constraints.

If the minimum cost flow rule is not sufficient to define a unique solution, the selection of one flow plan among equivalent flow plans is deterministic and repeatable, but not based on business rules.

6.4.2. Complexity introduced by the ramping constraint

Ramping constraints impose restrictions between allocated capacities in subsequent periods on an interconnector and its reverse, so that the routing plans of adjacent periods become interdependent. The uncoupled minimum cost flow routing problem (MCF problem), i.e. routing the individual network flows independently without ramping constraints, is a well-studied mathematical problem of linear optimisation, for which several efficient algorithms are known. The coupled minimum cost flow routing problem, i.e. the same problem after inclusion of the ramping condition, introduces a new level of complexity, which requires the application of more general optimisation techniques. These techniques require significantly more computational effort and could lead to unacceptable performance of the routing process. In order to benefit from faster MCF routing algorithms, routing tasks covering a single period product (hourly, half-hourly and quarterly products) are treated differently to routing tasks where the ramping condition couples the flow of two or more periods (block bid products).

6.4.3. Single-period routing with ramping

The routing tasks originating from the order book calculation and trade flow calculation of hourly, half-hourly and quarterly orders involve contract periods spanning one single time period. These single-period routing problems can be reduced to MCF problems by introducing ramping-reduced available capacity bounds on the flows.

Only one network flow needs to be calculated since the routing plan is only for a single period p . The inclusion of the ramping constraints boils down to additional restrictions on the routable flow in period p . Altogether, denoting the ATC for period p by $u(A \rightarrow B, p)$, the ramping-reduced available capacity $u'(A \rightarrow B, p)$ is constrained by three upper bounds,

$$u'(A \rightarrow B, p) \leq u(A \rightarrow B, p) \quad (\text{ATC})$$

$$u'(A \rightarrow B, p) \leq r(A, B) - f(A \rightarrow B, p) + f(B \rightarrow A, p) + f(A \rightarrow B, p+1) - f(B \rightarrow A, p+1) \\ (\text{backward ramping})$$

$$u'(A \rightarrow B, p) \leq r(A, B) - f(A \rightarrow B, p) + f(B \rightarrow A, p) + f(A \rightarrow B, p-1) - f(B \rightarrow A, p-1)$$

(forward ramping)

Forward ramping is the adjustment of the ATC for the ramping condition with the flow of the preceding period, and backward ramping is the adjustment for the ramping condition with the flow of the next period.

The quantities $f(A \rightarrow B, p)$ refer to allocated capacities, hence the upper bounds at the right-hand side are flow-independent. For single-period routing tasks, a MCF routing calculation is performed with the ramping-reduced available capacity.

The routable quantity is the sum over all paths. Potentially, the calculated routable quantity is smaller than the requested transport. For AON routing problems, this would imply that no solution is found; else, the routable quantity will be used for allocation.

6.4.4. Multi-period routing with ramping

The routing tasks originating from order book and trade flow calculations of block bids involve contract periods that span two or more time periods. A heuristic methodology is applied to approximate the optimal solution, as the problem is too complex to be solved in the minimum time available (trade-off between accuracy and speed).

In case the process does not result in a solution of the cheapest routing plan, no other routing plan will be explored. This implies that a solution found may be suboptimal. It also implies that an existing solution may not be found.

6.4.5. Correction for border resolution

As the duration of power contracts is independent of the underlying grid resolution (the border resolution of the interconnectors in the grid, which may differ from interconnector to interconnector), the time period of a network flow may not be identical to the existing border resolutions. If the duration of the routing time period is larger than the border resolution, the available capacity for routing is taken to be the minimum of the available capacities in the time sub-periods.

The determination of the available capacity for the routing time periods is a preparatory step before entering the actual routing calculation.