



www.BioBoost.eu

Biomass based energy intermediates boosting biofuel production

This project has received funding from the European Union's Seventh Programme for research, technological development and demonstration under grant agreement No 282873

Deliverable

D 4.3 BioBoost Logistic Model

Workpackage:	WP4
Deliverable N°:	D4.3
Due date of deliverable:	31/12/14
Actual date of delivery:	28/05/15
Version:	2
Responsible:	FH OÖ
Authors:	Gabriel Kronberger, Erik Pitzer
Contact:	gabriel.kronberger@fh-hagenberg.at
Dissemination level:	PU-public
Attachments	



Table of Contents

Executive Summary	5
Changes to the Logistic Model	7
1 Automatic Regional Subdivision	7
2 Lower Bound on Logistic Cost to avoid Spurious Transports	9
3 Dynamic Solution Space Reduction	9
4 Delayed Onset of Utilization Mutation	10
5 Semantic Operators	11
6 Fixed Locations.....	12
7 Refinement of Price Curve	16
8 Multi-objective Optimization	18
9 Feedback/Cooperation with TNO	19
10 Full Pathway Results.....	19
11 Post-processing of Transports.....	23
12 Soft Constraints – Penalties	24
13 Analysis Tools	24
14 Appendices/Enclosures	24

Executive Summary

The BioBoost project focusses on de-central conversion of biomass to intermediate energy carriers which are subsequently processed in central plants. For the evaluation and optimization of the logistic processes in this scenario, a logistics model has been developed and implemented in software. The logistic model is part of a full model of the BioBoost value chain and facilitates calculation of costs for logistics and handling on a regional level (NUTS3) on an EU-wide scale.

A first prototype of the simulation and optimization software tool for BioBoost has been released in early 2014. Using this tool it is possible to optimize locations and capacities of de-central plants and to calculate optimized transport modes and routes for these scenarios using for instance OpenStreetMap data for routing. Since the initial release, several improvements and extensions have been developed and incorporated in the simulation software which are documented in this document.

For large NUTS3 regions with big feedstock potentials, a pessimistic calculation of logistic costs had been used in the previous version, as the model assumed that only one plant with the necessary capacity would be built in the center of the region. However, for such regions it is feasible to build multiple plants of smaller capacity scattered within the region. Thus, the logistics costs could be made much smaller in these regions. Therefore, a finer grained subdivision has been artificially created to limit the maximum size and hence provide a fairer and more realistic simulation.

The initial model did scale logistic costs linearly based on the transported amounts. However, in the simulation we often saw transports of very small amounts, sometimes over large distances. When transporting small amounts of feedstock, the costs do not scale in the same way, as we must assume that there is a certain of minimal costs (e.g. for at least one full truck or one pallet every once in a while) even if this unit is not fully loaded. Therefore, a minimum threshold for logistic costs has been added to the model.

The optimization of logistics is difficult from a technical perspective, when the overall costs of an optimized scenario is very close to zero (not clearly profitable & not clearly a loss). In this case the optimization routine must balance two opposing factors. Costs should be reduced by optimizing the logistics; however costs can also be reduced easily by reducing the number of plants and it becomes necessary to prioritize the optimization of the logistic network in the initial phase of optimization to make sure that first a sufficiently optimized logistic network is developed to prevent prematurely scaling down plant capacities. The optimization algorithm (an adapted evolution strategy) has been changed to support this delayed optimization of plant capacities.

The pathway for catalytic pyrolysis foresees co-processing of the energy carrier with fossil crudes in an existing refinery with a surplus of hydrogen. Accordingly, capacities and sites of European refineries were mapped to the refined NUTS3 regions and included in the simulation model.

The logistics model has been extended to support simulation of the full pathways (including central processing) and results for all pathways (fast pyrolysis, catalytic pyrolysis and hydro-thermal combustion) have been calculated.

Finally, the function for the estimation of feedstock purchasing prices, based on the relative amount of used feedstock, has been tuned. Previously, it has been assumed that the price is almost constant up to a level where 50% of the available feedstock is used, above which a linearly and sharply rising price curve has been assumed. This lead to an issue in the optimization routine. Therefore a smoother curve (quadratic with continues transient) has

been fitted to model the feedstock prices more accurately and provide better leverage for the optimization. Finally, an option to override the default prices, imported from Del. 1.1, was included to allow the study of regional scaling effects.

Changes to the Logistic Model

In the following we document the changes that have been made to the logistic model and to the simulation and optimization software tool since the prototype version has been released in D4.2 (MS4).

1 Automatic Regional Subdivision

For larger regions with a large feedstock potential, aggregation of all feedstock to a single de-central plant can create scenarios that are sub-optimal and not realistic. The reason is the amount of feedstock which would suffice to drive several de-central plants, which could be located at more favorable locations in terms of transport distance inside a single region. Of course, this has to be counterbalanced with the economy of scale for larger plants.

Accordingly, we have created a geographical model that subdivides regions into a certain number of sub-regions. As the calculation of all contingencies and appropriate split points would be too expensive during the optimization, all regions were tested and split into sub-regions to hold a maximum of 7500 km². Depending on the initial size of the regions this meant splitting one region in up to 15 regions in one case.




Figure 1 Region Splits

Figure 1 shows some examples for extreme splits. Depending on the number of regions, the shapes and, hence, the distances can vary. The regions have been divided in a way to give the sub-regions approximately equal area. Currently, this sub-division does also not consider the underlying road network. Therefore, it is possible that other sub-divisions might provide even better reduction of transport costs. For the current level of granularity, however, it should suffice to only consider geographic distances for the scaling of routes.

A simple formula for the estimation of the route length reduction would be desirable; it does not seem to be accurate enough, without requiring further information. We have tried to consider simpler pieces of information such as region size, region shape, elongation, or compactness to derive a formula for the calculation of arbitrary regional subdivisions. However, the results were not satisfactory.




Figure 2 Distribution of Relative Distance Ratio Depending on the Number of Subdivisions

In Figure 2 the distribution of distance reductions is shown for a split into up to six sub-regions. While the probability of an around average ratio is relatively high the bottom spread is still too large to warrant a generalization. Therefore, a simple spreadsheet has been created with up to seven sub-region scaling factors for each region that will be incorporated into the simulator.

The second aspect that would have to be considered is the tradeoff between economy of scale and route length reduction. This can be formulated as a single variable optimization problem.

Taking the original formulae for transport, maintenance, construction, and financing (see Deliverable 4.2 “Protocol on Milestone 4”), we can eliminate all quantities that are not affected by the route length or scale choice.

$$\text{Transport Costs } C_{p,r}^t = Q_{p,r}^s \cdot D_{r,t(p,r)} \cdot \frac{P_{p,r}^t + P_{p,t(p,r)}^t}{2}$$

$$\text{Capacity } CP_{p,r}^c = \frac{Q_{p,r}^t}{P_p^{uf}}$$

$$\text{Maintenance Costs } C_{p,r}^m = \left(\frac{CP_{p,r}^c}{P_p^{ds}} \right)^{\sigma^{mc}} \cdot P_{p,r}^{mc}$$

$$\text{Construction Costs } C_{p,r}^c = \left(\frac{CP_{p,r}^c}{P_p^{ds}} \right)^{\sigma^{cc}} \cdot P_{p,r}^{cc}$$

$$\text{Financing } C_{p,r}^i = \lambda \cdot C_{p,r}^c$$

The problem of this task is not the division of a region but the split of a plant if the region receives feedstock from other regions too. In that case it might seem like a split is necessary because the facility is large enough, however, a split does not really reduce transport costs as the sources are external to the region. Therefore, splits were taken as pre-determined based on the size of the original size to provide a fairer view of logistic cost.

2 Lower Bound on Logistic Cost to avoid Spurious Transports

While the overall optimization algorithm chooses only utilization of feedstock and transport targets for each stage of the logistic chain, the optimal means of transport is automatically selected. This optimization step is very simple as it involves only a single comparison. The amount and distances are already known at this stage and all possible means of transport are evaluated for this concrete scenario.

In the past, a frequent and obvious problem with the optimization was the ignorance of irrelevant transports. These were often transports of miniscule amounts over very long distances. Because of the very low cost impact, these long transports could be easily “overlooked” by the optimization ranging in a few thousand Euros in comparison to several billion Euros for the overall scenario.

To mitigate this problem, it was assumed that aggregated yearly transports require a certain minimum volume to justify the establishment of a consistent transport route. Currently this was set to require at least one vehicle, e.g. truck-load, per two weeks, which is about 500 t/a. Every transport edge does now accommodate at least this volume, which reduces the number of these spurious transport links.

The YAML configuration file has been extended with an additional property “min-amount” for every logistic action to specify different minimum values for each type of transport. Therefore, this can still be configured for every configuration of transport means and product. Additionally, an option to specify the maximum transport distances for each transport type (product & mode) has been added. It has been discussed that a maximum transport distance of 100 or 200km can be considered feasible in practice for the de-central routes and up to 500 or 1000km for the central routes that can often use railway transports.

3 Dynamic Solution Space Reduction

While the simulation itself has been refined to minimize per-iteration calculation time, also the optimization scheme has been improved by dynamically reducing the size of the solution space. In other words, the number of different possibilities for varying each scenario is reduced to only meaningful similar scenarios. This is accomplished by analyzing the concrete problem instance and determining which solution candidates are feasible. Using this technique several sets of solution candidates can be excluded from the optimization process. These are solution candidates that can be ruled out up-front as infeasible for the following reasons:

- *No feedstock available:* If a certain type of feedstock is not available in one region, any percentage of feedstock obtained will still result in no obtained feedstock. Therefore, the feedstock utilization vector for this feedstock is dynamically shortened to exclude these regions from manipulation as it would not have any influence on the outcome of the simulation.
- *Nothing to Transport:* In addition to reducing the size of the utilization vectors, the transport vectors can also be shortened. Again, this can be done by examining the current scenario. In particular, for example, for any region that does not have a certain feedstock available, any transport choice for that combination of region and feedstock will not have an effect as nothing can be transported. Therefore, these regions can be removed as viable sources. On the other hand, as described in Section 5 certain regions might be restricted in what types of plants can be erected. Therefore, for these regions it does not make sense to transport and deposit any amount of source product as it will be impossible to perform further processing. In

this case, these regions are removed from the list of viable transport targets which reduces the size of the transport targets vector.

Both of these optimizations are only derived from the base data and do not depend on the current solution candidate as any extra “intelligence” in the mutation and crossover phase of a genetic algorithm could potentially interfere with the inherent strategy of the optimization algorithm itself. Moreover, the intermixing of solution candidate manipulation and evaluation would make the system much more complicated. Therefore, any further “optimization” is left for the actual optimization algorithm.

Early experiments with this technique showed that this solution space reduction accelerates convergence speed of the optimization algorithm by about 10-20% without sacrificing any feasible possibilities or biasing of the search.

4 Delayed Onset of Utilization Mutation

Especially in optimization situations where a profitable outcome is difficult and the optimal solution yields only a small profit in comparison to the turnover, the optimization might be “tempted” to choose “the easy way out” and select no feedstock utilization as the optimal strategy as it does not incur any costs. These cases can usually be considered premature convergence and should be avoided. The idea behind this improvement is to forbid limiting the utilization rate until a sensible routing network has been established.

The problem with premature convergence of utilization is connected to a strategy for initialization of the transport targets. A naïve approach would be to simply assign the originating region also as target region for every transport. On the one hand this limits the transport distance for every region and gives equal chances to each region for the establishment of a de-central or central conversion facility. On the other hand, however, it requires an enormous amount of construction and maintenance costs as the initial solutions have all possible conversion facilities in every single region. Therefore, the optimization chooses the easiest way out and reduces the amount of feedstock utilization which in turn reduces the plant sizes until it completely shuts down all plants, suggesting a no-risk no-profit solution.

The alternative would be to start out with no utilization and let the conversion plant and transport network infrastructure build bottom-up at the same time. However, in our past experiments this did not prove very successful as singular plants without any neighboring regions that participate and especially central plants with very little supply are very far from profitable. Therefore, the optimization is more or less stuck already at the beginning.

These complications led to the development of the described strategy. Initially, all available feedstock is forcefully utilized and the manipulation operators within the optimization is prohibited to reduce it. Only after a certain transport network has established that provides sufficient amounts of feedstock and product to profitably drive de-central and central plants the utilization can be reduced again. At the same time the transport network can still be adjusted which can still relocate or even remove conversion facilities, without, however, falling prey to premature convergence.

This feature was implemented by addition of a parameter to the mutation operator which in turn monitors the progress of the optimization algorithm. This parameter chooses the minimum percentage of iterations that have to be made before any manipulation (i.e. reduction, reduce or shifting) of feedstock utilization is allowed.

5 Semantic Operators

The simultaneous optimization of transport network and feedstock utilizations (or alternatively plant locations and capacities) that is necessary in the BioBoost project is challenging because the optimization of the transport network is interrelated to the optimization of the feedstock utilization vectors and vice-versa. Standard variants of meta-heuristic optimization problems cannot be used in a straight-forward manner and have to be adapted accordingly.

Experiments with standard operators used in evolutionary algorithms, such as single-point crossover or single-point mutation, showed that these operators are not effective for the optimization of BioBoost scenario. The reason for this is that the chosen representation has the property that there are strong dependencies between separate elements of solution candidates. This has the effect that changing only one element of the transport target vector without considering the values of the other elements most likely has a detrimental effect on the encoded scenario. Frequently, multiple such detrimental changes in a row are necessary to identify possible improvements. Similarly, simply combining parts of two encoded scenarios through single-point crossover has a very small probability of generating an improved scenario because e.g. the neighbourhood relation between regions is completely ignored. We have therefore developed new evolutionary operators specifically tuned for optimizing BioBoost scenarios. The following operators have been implemented:

- Plant-based crossover: This operator combines two scenarios on the level of plants. If both parent solution candidates have a plant in the same region the plant will also exist in the new solution candidate. If the plant has the same supplier regions the plant in the new solution candidate will also have the same suppliers. Otherwise plants are copied with their supplier information randomly from the parents. Conflicts are resolved randomly.
- Manipulation operators:
 - Plant Splitter: Splits an existing plant into two plants of the same type, one plant is kept in the original location the second plant is assigned to one of the neighbouring regions. The set of supplier regions for the original plant is split into two sets randomly, the first set contains the suppliers to the first plant the second set contains the suppliers for the second plant. This operator only changes transport targets
 - Plant Merger: Merges two existing plants of the same type into one larger plant. Two plants are selected and the supplier regions for both plants are merged to one set of suppliers to one of the two plants. This operator only changes transport targets.
 - Plant Mover: Takes a plant and moves it to one of the neighbouring regions. This operator only changes transport targets.
 - Supplier Equalizer: Determines the set of all supplier regions for one plant (de-central or central) and equalizes the utilizations of those suppliers. This operator only changes the utilizations of supplier regions to one plant. The overall capacity is not changed only the distribution of acquired feedstock from regions is changed. This operator only changes utilizations.
 - Supplier Randomizer: Determines the set of all supplier regions for a plant and slightly changes all utilizations randomly (this is the symmetric operation to the supplier equalizer). This operator only changes utilizations.

- Supplier Utilization Exchange: Selects two supplier regions for the plant randomly and exchanges the utilization values. This operator only changes utilizations.
- Supplier Toggler: Selects a random supplier and either sets the utilization to zero if it is non-zero or sets the utilization to 50% if it is zero. This operator only changes utilizations.
- Plant Killer: Selects a plant and removes it from the scenario. This operation also changes all supplier regions for the plant and set the utilizations for those supplier regions to zero. The plant killer is an essential operator when the goal is to identify the most cost-effective plant location. However, it is detrimental when the goal is to maximize produced amounts with minimal costs.

6 Fixed Locations

Another new feature was the introduction of predefined (i.e. free) capacities and maximum capacities for certain regions and plant types. This allows, for example, the co-location of plants together with existing plants at a reduced cost, or, on the other hand, the restriction of plant capacities or even exclusion of certain regions as possible installation sites. In particular for central refineries this was required as these facilities depend on the availability of hydrogen which is currently only economically viable when co-located to a larger refinery, where hydrogen is generated via stream reforming from natural gas.

This information is used, on the one hand, as already mentioned, for the exclusion of certain possibilities of transport targets. On the other hand, of course, for the evaluation of plant sizes, either reducing the cost for co-located usage of available capacities or limiting plant size to the specified maximum size for each region.

The implementation of this feature includes several new properties in the YAML configuration file that can be specified for each conversion technology. In particular the following extra values can be specified now:

- **available-capacities**: This property can contain a list of pairs specifying the available capacity for each region.
- **max-capacities**: Similarly, this property can contain a list of pairs specifying, for each region the maximum possible capacity for a plant of this type. In addition the special region name 'default' can be used to indicate the default maximum capacity, which can be useful in situations where it is generally not possible to build a plant of this type in any region, such as when special requirements or regulations are required for the plant to operate properly, e.g. consumption of hydrogen.
- **available-maintenance-factor**: This factor can indicate a possible reduction in maintenance cost when utilizing existing capacities. For example, when existing personnel can be deployed.

As an important side note, it should be mentioned that in previous versions in case the maximum capacity is reached, any feedstock that has been transported to the plant could not be converted anymore and is, therefore, was completely lost. To better steer the optimization algorithm, in newer versions the simulation allows plants with exceeding capacities but adds penalties to the total cost which are not considered economically but

only by the optimization algorithm to select different scenarios with no or negligible excess of capacity.

On the following page an example conversion description as found in the YAML configuration file is shown with the new properties. In this case, there are capacities available in region FR614. At the same time the default maximum capacity is limited to zero, i.e. no plants can be built except in designated regions. Finally, the “available maintenance factor” is set to one, which indicates that there is no cost benefit by using existing personnel or resources.

```

conversions:
- label: FastPyrolysis
  feedstock: Straw
  safety-stock: 365 # days
  dry-matter-loss: 0.025 #percentage
  storage: # EUR/t/a
    investment: 1.15
    labor: 2.75
    other: 2.85
  products: # [t/t] names and mass ratios
    Biosyncrude: 0.675676
    CO2: 0.324324
    WaterVapor: 0.108108
    CoolingWater: -0.344595
    ElectricityIn: -0.087838
  cost: 0 # [EUR/t] variable cost without feedstock cost
  design-capacity: 219123.38028 # [t/a]
  construction: 11003716.52 # [EUR/a]
  min-construction: 1743968 # [EUR/a] (1/10)^0.8 of design size const. cost
  maintenance: 7278442.59 # [EUR/a]
  min-maintenance: 727 844.259 # [EUR/a]
  construction-scaling-exponent: 0.8 # factor
  maintenance-scaling-exponent: 1 # factor
  utilization-factor: 0.9

  available-capacities: { FR614: 2000000 } # [t/a]no construction cost
  available-maintenance-factor: 1 # = same maintenance as non-allocated
  max-capacities: { Default: 0, FR614: 20000000000 } # [t/a]

```

A list of existing European refineries can be found on Wikipedia¹. The given crude oil daily processing capacity may be used as additional information to which extent an existing refinery can reduce the costs of a co-located refinery for a BioBoost pathway. The list was used to map existing capacities onto NUTS3 regions. Subsequently, we filtered the full list of plants to only keep plants that have a crude oil processing capacity of around 10Mt/a or more because we assume that such large plants have the necessary facilities (complex hydrocrackers) for co-processing of catalytic pyrolysis oil. The mapping was done manually, using Google Maps² and Eurostat Statistical Atlas³. Additionally, processing capacities were converted into million tons per anno (Mt/a), where a barrel is equivalent to approximately 0.137 tons⁴. In the simulation model we approximate that 2% of the total capacity can be used for co-processing CP oil (mainly determined by hydrogen availability).

RegionID	Refinery	Capacity (Mt/a crude oil)	Capacity (kt/a CP oil)
AT127	Schwechat Refinery (OMV)	8.80	176
BE211	Total Antwerp Refinery (Total)	18.00	360
BG341-0	LUKOIL Neftochim Burgas (LUKOIL)	10.40	208
DE122	Mineraloil Refinery Upper-Rhine (Karlsruhe)	14.25	285
DE211	Ingolstadt Refinery (Bayernoil(OMV/Agip/Rosneft/BP))	13.10	262
DE211	Ingolstadt Refinery (Gunvor)	5.50	110
DE40I	Schwedt Refinery (PCK Raffinerie(Shell/Rosneft/BP/AET))	10.50	210

¹ http://en.wikipedia.org/wiki/List_of_oil_refineries#Europe

² <https://maps.google.com/>

³ <http://ec.europa.eu/eurostat/statistical-atlas/gis/viewer/>

⁴ http://nesteoil-webannualreport.com/sanasto_en_EN.htm

DEA23	Rheinland Werke Godorf & Wesseling (Royal Dutch Shell)	17.50	350
DEA32	Ruhr Öl Refinery (Rosneft/BP)	13.30	266
DEE0B	TOTAL Refinery Mitteldeutschland Spargau (Total)	11.35	227
ES213	Bilbao Refinery (Repsol YPF)	11.00	220
ES612	Gibraltar-San Roque Refinery (CEPSA)	12.00	240
ES620-1	Cartagena Refinery, (Repsol YPF)	11.00	220
FI1B1-1	Porvoo Refinery (Neste Oil Oyj)	10.30	206
FR232	Normandy Refinery (Total)	17.50	350
FR232	Port Jérôme-Gravenchon Refinery (ExxonMobil)	13.50	270
FR511	Donges Refinery (Total)	11.55	231
ITG19	Esso Augusta Refinery (ExxonMobil)	9.50	190
ITG19	Impianti Nord Refinery (ISAB ERG)	8.00	160
ITG19	Impianti Sud Refinery (ISAB ERG)	10.70	214
ITG27	Sarroch Refinery, Sardegna (Saras S.p.A.)	17.00	340
LT008	Mazeikiu Refinery (Mazeikiu Nafta - PKN Orlen)	13.15	263
NL339	BP Rotterdam Refinery (BP)	20.00	400
NL339	Shell Pernis Refinery (Royal Dutch Shell)	20.80	416
PL121-1	Plock Refinery (PKN Orlen)	13.80	276
PL633	Gdansk Refinery (Grupa LOTOS)	10.50	210
PT181	Sines Refinery (Galp Energia)	10.00	200
SE232-1	Lysekil Refinery	11.00	220
UKD63	Stanlow Refinery (Essar Oil)	13.60	272
UKE13	Humber Refinery (Phillips 66)	11.05	221
UKE13	Lindsey Oil Refinery (Total)	11.15	223
UKJ33	England Fawley Refinery (ExxonMobil)	16.50	330
UKL14	Pembroke Refinery (Valero)	10.75	215
UKM26	Scotland Grangemouth Refinery (Ineos and PetroChina)	10.25	205

The BioBoost simulation software, allows displaying available potentials in a map view, as shown in Figure 3. The full list of refineries mapped to regions is contained in the appendix.

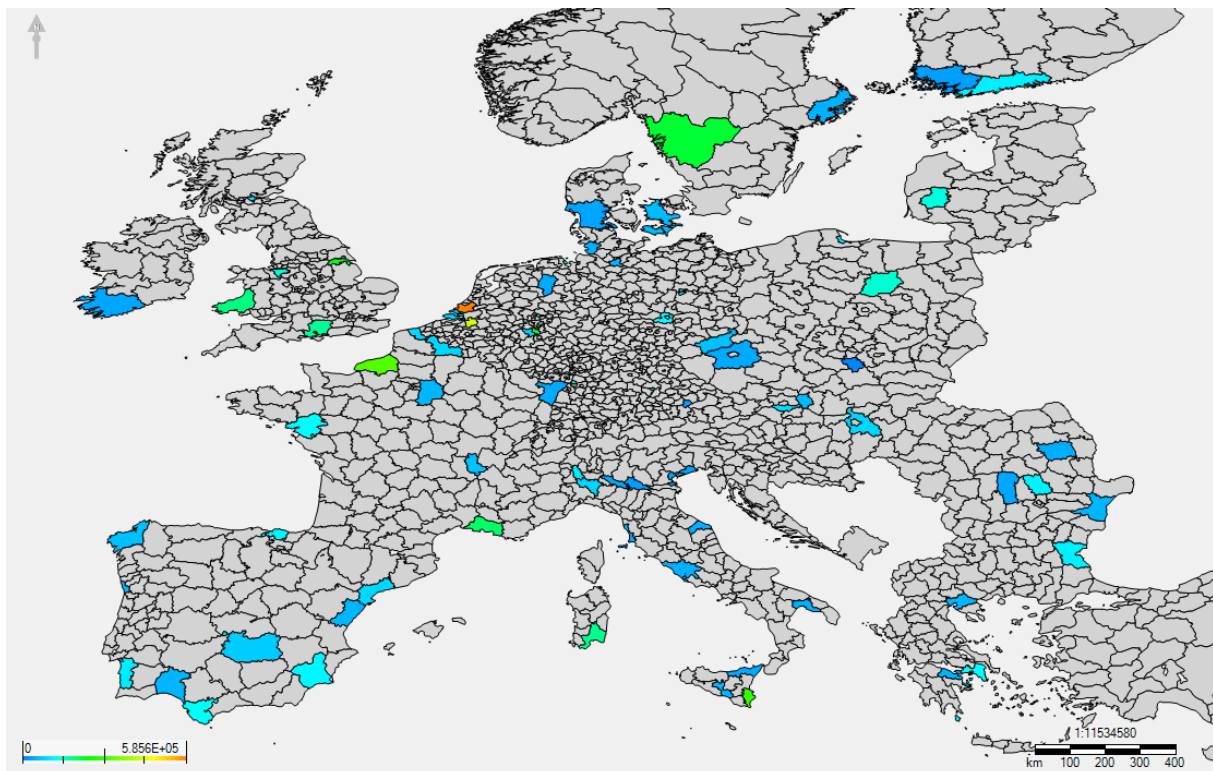


Figure 3: Capacities of refineries in EU mapped to NUTS3 regions. Plants for central processing of catalytic pyrolysis oil are limited to a subset of these locations.

7 Refinement of Price Curve

The first version of the simulation software already provided the possibility to consider increasing feedstock prices due to market saturation. At first we proposed an exponential function, where a saturation factor was calculated for any percentage of feedstock utilization in a region to increase the price up to a maximum penalty. It turned out that this approach pushed the optimization algorithm to favor feedstock utilizations close to the given exponent.

To allow for a lighter gradient of the price-supply function, SYNCOM suggested to linearly increase the feedstock price, once a certain rate (e.g. 50%) of the available feedstock is bought. Below that rate a fixed price is used. As additional parameter, a maximum price at 100% utilization can be defined. This function was incorporated into the simulation model and validated. As a result, the optimization always found solutions where the feedstock utilization in a region is either zero or close to the defined rate, where the feedstock price starts to raise. Only when either the feedstock price was lowered significantly, or the product price after conversion was raised to an unrealistically high value, the optimization algorithm yielded utilization rates that are higher than the threshold value.

To overcome this problem a quadratic function was later implemented to remove the “artificial” threshold that was brought in by the step-function idea before. The quadratic function was designed to have its minimum at 0% utilization and has two fixed values at 50% and 100% utilization, respectively. A disadvantage of this approach is that the feedstock price at low utilization rates can get too low. The optimization algorithm tended to minimize feedstock acquisition and as a consequence thereof a very limited number of suggested conversion facilities.

In Figure 4 the different price/supply curve versions described above are shown graphically.




Figure 4: Preliminary versions of feedstock price/supply curves

For the most recent version of the simulation software a more accurate price function was used which considers both, the proposed price/supply relation from Deliverable 1.1 – Feedstock costs and the smoothness of the function which is required to avoid “artificial” optimization traps. Therefore, a fixed price is used for utilizations up to 50%. Starting from that point, a quadratic function is applied that starts with a gradient of zero and goes through a defined point at 100% utilization. Necessary parameters can be adjusted in the YAML configuration file. For each feedstock entry, the tag “price” sets the base price and the tag “max-price” defines the maximum price at full utilization. An example is shown in Figure 5.

```
products:
- label: Straw
  price: 46.212 #EUR/t
  max-price: 100.138 #EUR/t
  price-scaling: { AT11: 1.022, AT12: 1.019,... }
  max-price-scaling: { AT11: 1.031, AT12: 1.029... }
```

Figure 5: Price curve configuration example

As a new feature, the tags “price-scaling” and “max-price-scaling” can be used to override the default prices with region specific prices. This can be achieved by setting a scaling factor for the default price and the maximum price, for a specific region or a group of regions. For example, according to the values shown in Figure 5, the straw price for all regions having a name that starts with “AT11” can be calculated as follows:

- base-price (0% - 50% utilization) = $46.212 * 1.022 = 47.23$
- max-price (100% utilization) = $100.138 * 1.031 = 103.24$

According to Deliverable 1.1 – Feedstock costs, the average prices were calculated for the BioBoost reference biomasses straw, forestry residues and organic municipal waste. The resulting default price/supply curves for those are shown in Figure 6.




Figure 6: Default price/supply curves

8 Multi-objective Optimization

A very important extension for the optimization software module has been the support for different objective functions and even for multi-objective optimization. Previously, the only objective value that we used in the optimization routine had been the total profits. We have however come to the conclusion that this objective value does not really help for answering central BioBoost questions such as: “Where would it be most beneficial to invest into the first plant?”, “How much would it cost to produce transport fuel at that plant?”, or “How much transport fuel can be produced?”. Therefore, we added new objectives such as minimization of relative overall costs for the production of final product (e.g. transport fuel) per ton, and the maximization of total amount of transport fuel produced. We also implemented a multi-objective optimization algorithm that supports optimization of both objectives (maximum produced amount at minimum costs) based on the well-known non-dominated sorting genetic algorithm (NSGA-II) multi-objective GA variant. This algorithm produces a set of Pareto-optimal solutions which can be analysed individually after the optimization run. Figure 7 shows an example for a possible result for this multi-objective approach; each dot represents one possible solution.




Figure 7: Pareto-optimal set of solutions with respect to total amount of produced bio-fuel and production costs per ton of bio-fuel.

9 Feedback/Cooperation with TNO

The versions of the simulation and optimization software environment which have been released in the consortium have been tested by TNO and SYNCOM. Software bugs and usage of incorrect data have been reported by TNO to FH OÖ in order to be fixed. These items have been communicated verbally or per e-mail.

10 Full Pathway Results

In this section result examples for the full value chain on each of the three reference pathways are illustrated. The results shown below have been calculated on the available data as of December 2014. As the data are not yet complete and we expect that some details might still change in the course of the BioBoost project we also expect to recalculate new optimized scenarios for refined data later on. These updated results will be reported in a later deliverable.

The maps indicate biomass transports with blue arrows and intermediate energy carrier transports with red arrows. The colour of the NUTS3 regions indicates how much of the available potential is utilized; see Figure 8 for the corresponding colour scale. A utilization of 1 would mean that 100% of the available feedstock should be purchased. The distribution of costs and revenues is shown in corresponding pie charts.



Figure 8: Utilization color scale ranges from 0 to 1

10.1 Fast Pyrolysis Pathway Example (Straw)


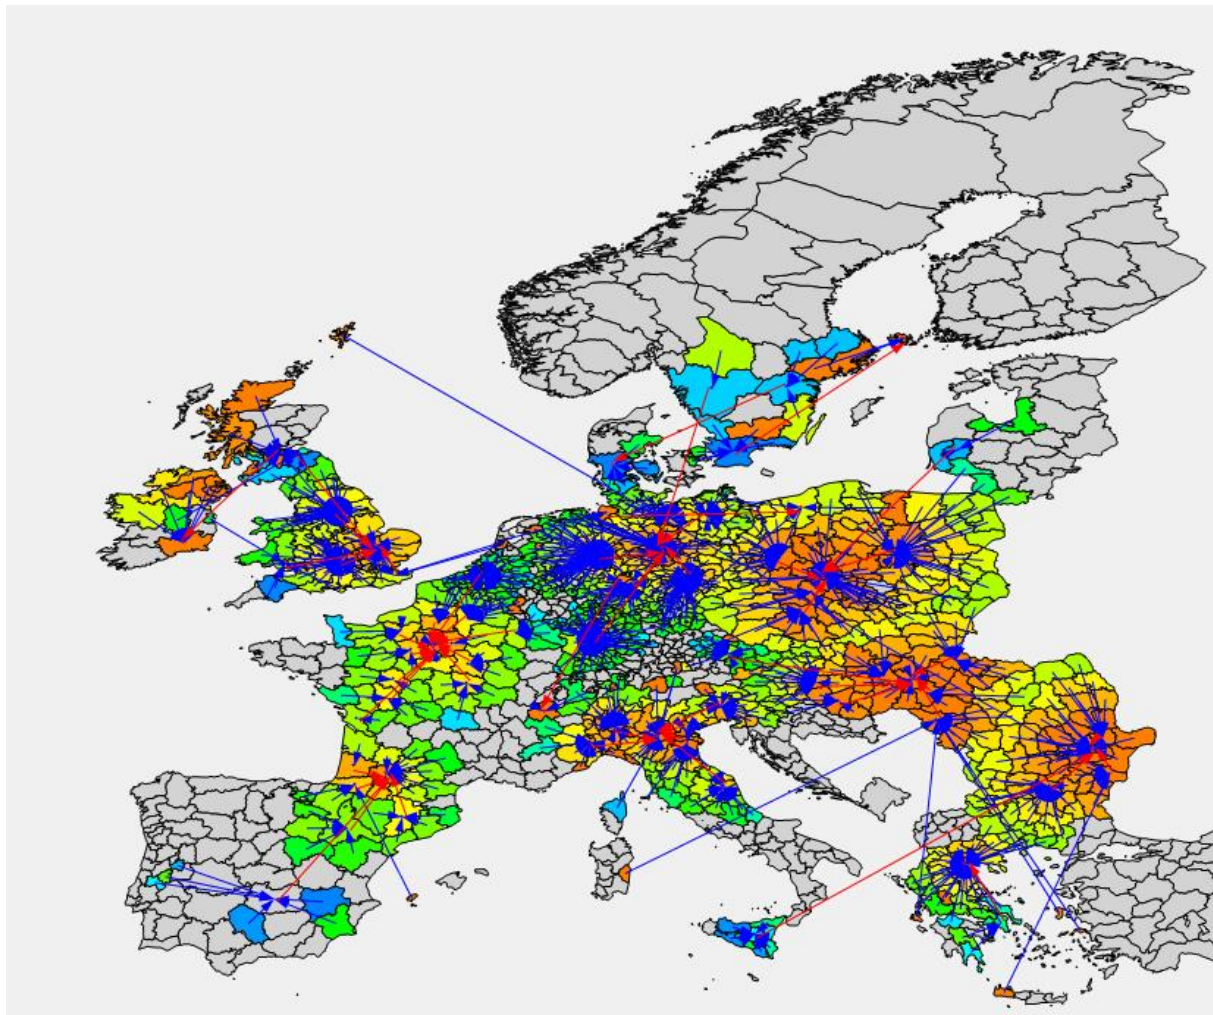


Figure 9: Map of transports and feedstock utilizations as well as cost and revenue break-down for the fast pyrolysis pathway


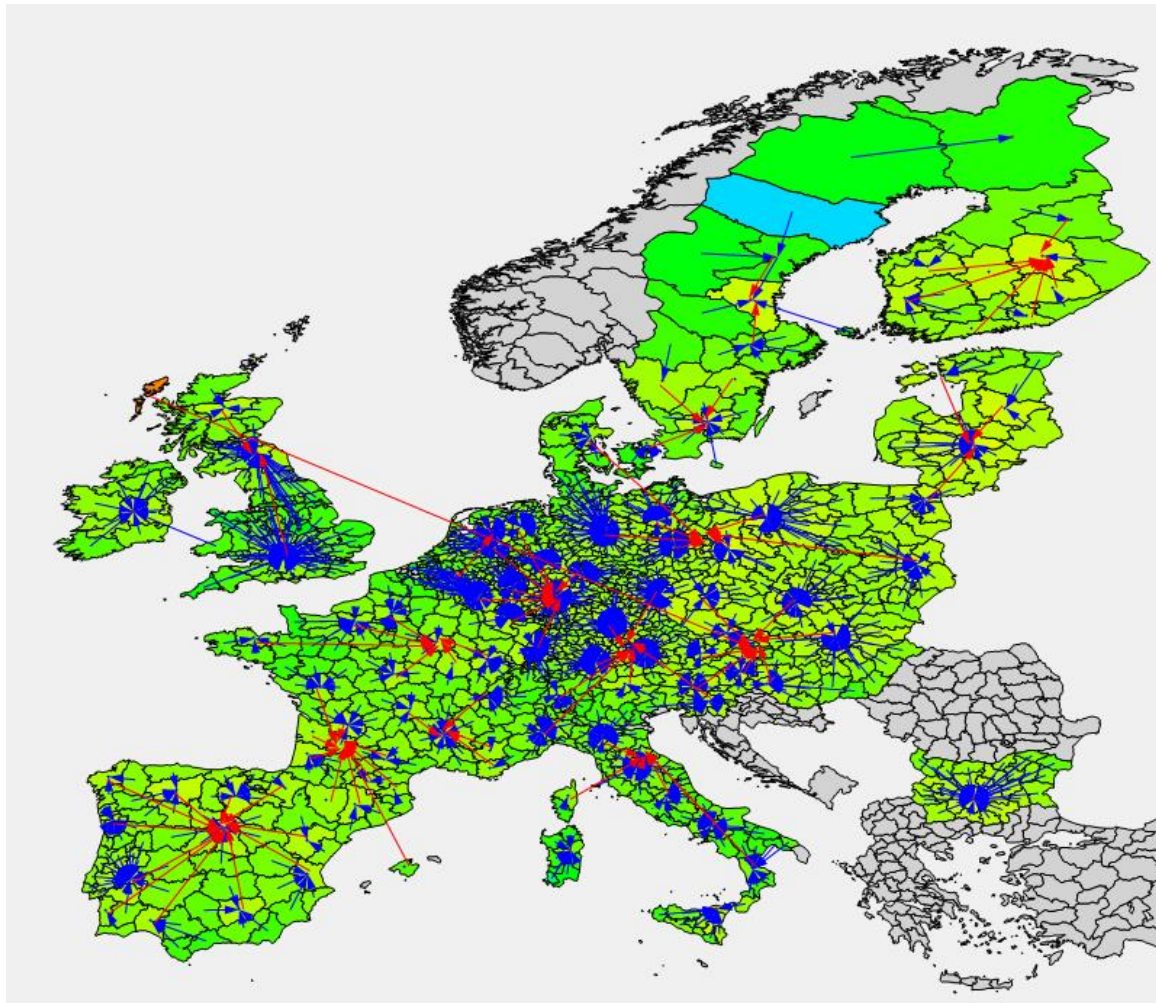


Figure 10: Map of transports and feedstock utilizations as well as cost and revenue break-down for the catalytic pyrolysis pathway

10.2 Hydrothermal Carbonization Pathway (Organic Municipal Waste)


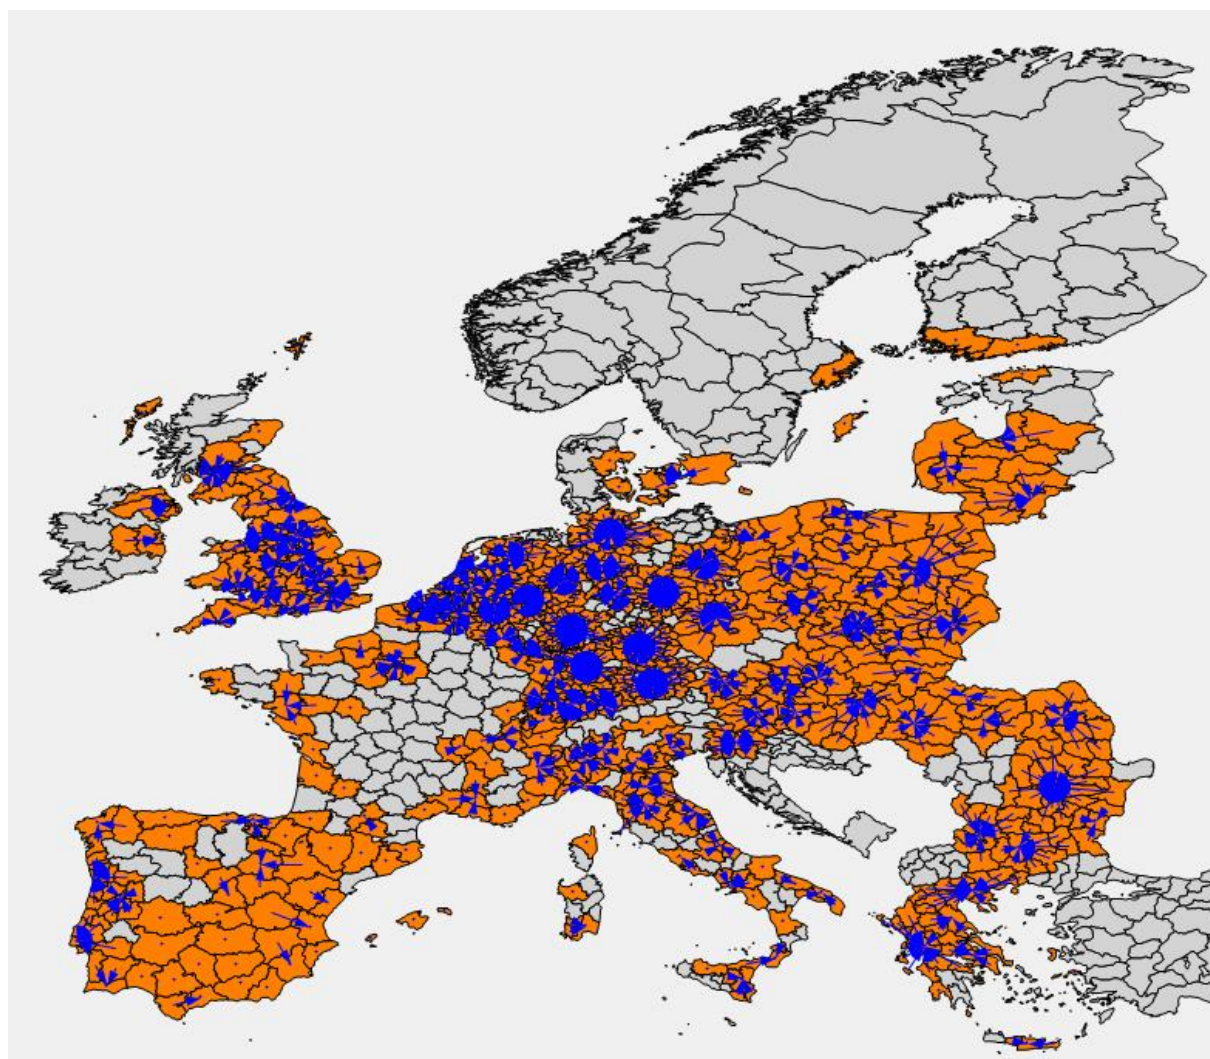


Figure 11: Map of transports and feedstock utilizations as well as cost and revenue break-down for the HTC pathway

11 Post-processing of Transports

Preliminary results showed that optimized scenarios often contain long distance transports of relatively small amounts which are visualized in the map view as very long arrows. These transports are an artifact of the meta-heuristic optimization algorithm which finds good solutions in a relatively small runtime, but is not guaranteed to find a globally optimal solution. Even though some material is transported over long distances these transports do not have a strong negative effect on the overall costs of the solution because only relatively small amounts of material are transported and generally the costs of logistics are a small factor in the overall costs. As a possible solution to this problem we suggest to apply a post-optimization step to the generated results to remove these overly long transports. This post-processing step might have a detrimental effect on the overall costs of the optimized solution but we ignore and accept this reduction in solution quality in exchange for a more realistic result.

We implemented a so-called 'solution editor' that allows to perform automatic and manual changes to optimized scenarios. This solution editor allows to directly manipulate (delete or change) transportation vectors or to change feedstock utilization values for each region. It also provides a way to automatically remove transports along the Pareto-front (longest-distance vs. smallest amounts) to improve the solutions. The algorithm iteratively removes the highlighted points from the Pareto-Front until either (1) a maximal number of points have been removed, or (2) the overall costs of the solution have increased over a certain threshold. A screenshot of the solution editor is shown in Figure 12.

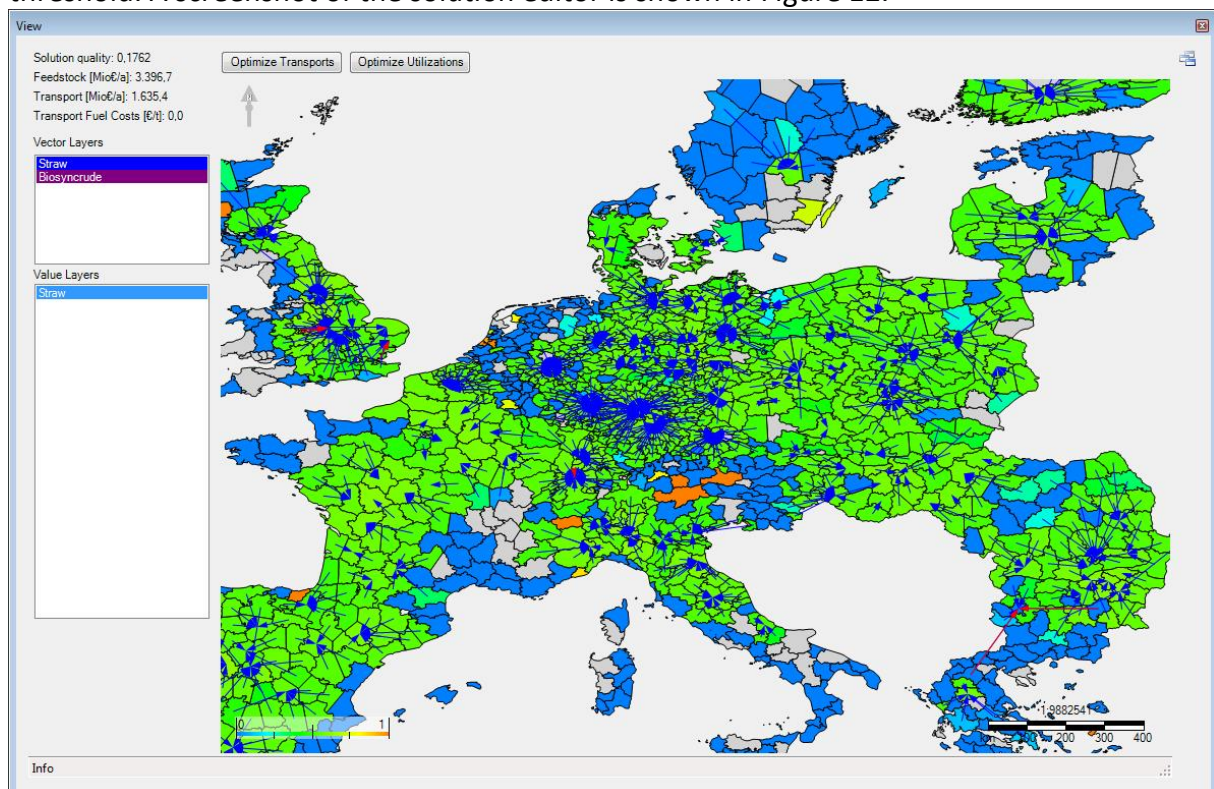


Figure 12: Screenshot of the solution editor that allows to post-process optimized solutions

12 Soft Constraints – Penalties

Through discussion of the results with other partners additional soft requirements were stated. One major concern was that transport distances were still quite large despite the assumed decentral nature. This concern was addressed by adding a limit for transport distances. However, this limit is, in principle, allowed to be exceeded in the simulation but is affixed with a strong penalty that drives the optimization towards shorter routes. These maximum distances can be adjusted in the YAML file for each combination of feedstock and transport mode. For example, farm tractors have a much smaller maximum distance than trucks or trains.

Another concern was that through the economy of scale and the comparatively small logistic costs plants would be scaled up to any limit that was allowed. We did not want to blindly limit the plants' capacities but based these maxima on estimated logistic network saturation, such that a scale exceeding 700 truckloads per day would incur an additional penalty subsuming the anticipated costs and problems concerning infrastructure overload and social aspects.

Finally we added a similar soft constraint for the exhaustion of the maximum possible capacity in a certain region, e.g. limited by the amount of available hydrogen for refineries. Again, we chose not to pose a hard limit that could have caused difficulties for the optimization to steer clear of these scenarios as everything above maximum capacity would seem equally bad. Instead we added a penalty that starts at the maximum capacity and increases to infinity as larger and larger plants are built to force the optimization towards capacities within reasonable limits, however, counterbalanced with other factors such as very large local feedstock availability that can now be iteratively resolved.

13 Analysis Tools

In collaboration with other partners who use the simulation and optimization tool for analysis we have collected ideas for further analysis and incorporated the following tools: As the final fuel production cost is a very interesting key figure, it seems only natural to also track its progress during the individual steps of the logistic chain. For this reason we have created many different map layers that visualize the distribution of absolute and relative cost factors such as logistics costs, conversion costs, supplies, also in relation to feedstock or final product. This makes it very easy to arrive at subtotals or identify profitable plant locations even while the optimization is in progress. In part this was possible by creating two separate versions of the simulator: One very fast implementation that only calculated final fitness values composed of rate of return and fuel amount, and second simulator version that produces all kinds of intermediate results and aggregations e.g. total logistic costs for all supplying decentral plants of a certain central plants.

14 Appendices/Enclosures

In separate appendix file "D4.3 Appendices.zip" we have included the YAML file, the distance matrices, the split shape files in ESRI format and an MS Excel document with the filtered list of refineries that were used to estimate free hydrogen capacities.