Designing a Heat Distribution System in Combination with an ATES System in Urban Environments using a Slime Mould growth model

By

C.J.J. de Boer

10-07-2025

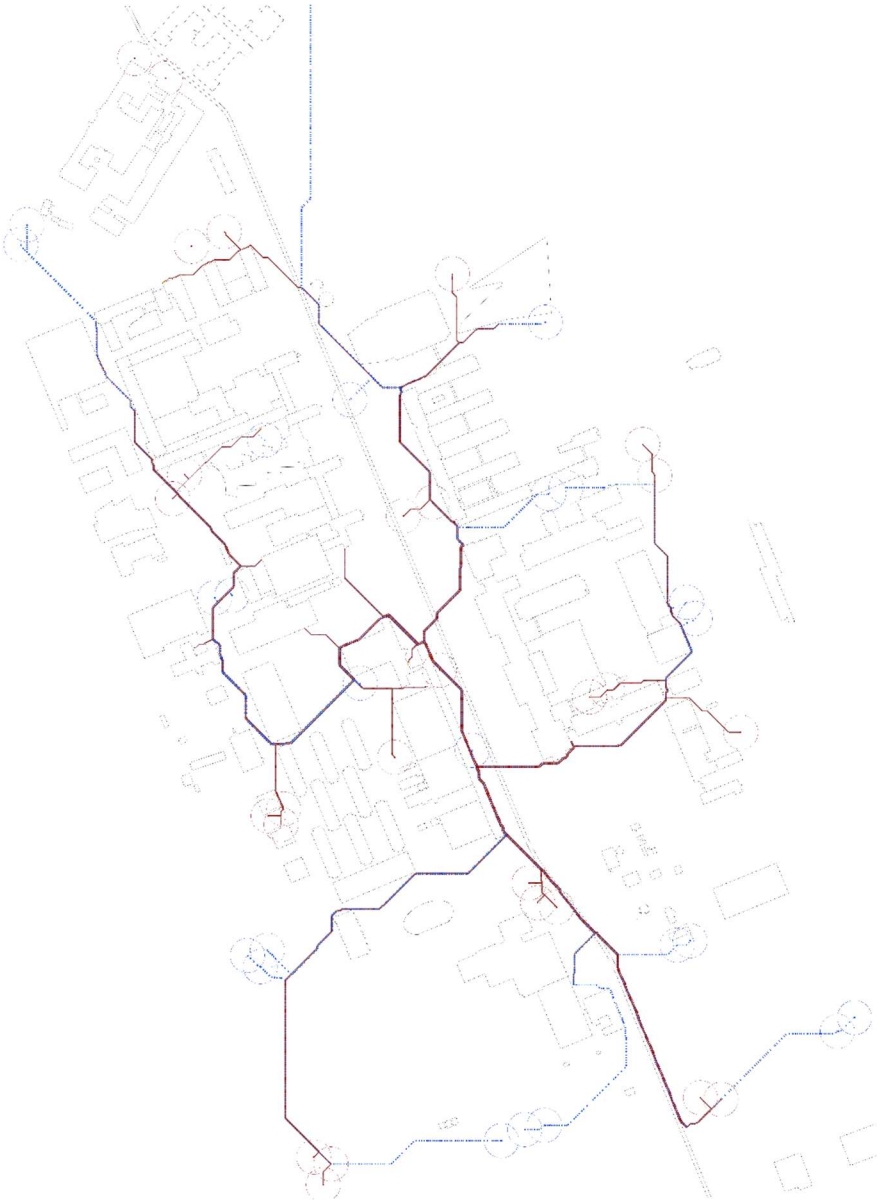

Student number:	5145104
Committee members:	Dr. ir. J.M. Bloemendal
	Dr. O. Nejadseyfi
	Dr. ir A.H.A. Stienen
	Ing. L.A. Medema
Institution:	University of Technology, Delft
Faculty	Faculty of Mechanical Engineering (ME), Delft
Degree:	Master of Science in Mechanical Engineering

Abstract

Climate change is an urgent global challenge. Buildings significantly contribute to energy consumption and CO_2 emissions. To accelerate decarbonization, ATES systems, integrated with 5GDHC networks, offer a sustainable approach. These systems enable seasonal thermal energy storage and bidirectional energy exchange. However, designing these complex combined systems in urban areas poses significant challenges, including integrating dispersed ATES wells, managing phased construction, and navigating dense underground infrastructure. This study investigated a biologically inspired approach, utilizing the network optimization capabilities of slime mould (Physarum polycephalum). A simulation model was developed and verified, incorporating key slime mould behaviours like efficient pathfinding, to generate optimal thermal network configurations. Using the TU Delft Campus as a case study, this research explored the model's capacity to streamline the design of sustainable and efficient heat distribution systems. Key findings indicate that the visual differences in layout between the best results and all results are minimal. In conclusion, this research successfully developed and applied a slime mould growth model for optimal thermal network design, specifically for district heating systems combined with ATES in an urban environment. The model proved effective in generating a network designs that minimize CapEx.

Acknowledgement

I would like to thank everyone who contributed to the successful completion of this thesis. First of all, I am very grateful to my thesis supervisors, Dr. ir. J.M. Bloemendal and Dr. O. Nejadseyfi, for their valuable guidance and collaboration. I've learned a great deal under their mentorship, and I truly appreciated the freedom they gave me, while maintaining a crucial, critical perspective on the project. Their insights were invaluable in shaping this work. I am also very grateful to A. Medema for providing extensive data and for the insightful sparring sessions that helped me navigate through missing datasets. His support was essential for the practical foundation of this thesis. Finally, I would like to thank Dr. ir. A.H.A. Stienen for his invaluable assistance and active brainstorming in finding a suitable thesis topic that aligns with the BME faculty.

Contents

1. Introduction

Climate change is one of the most pressing challenges facing the world today. In 2020, 15,000 scientists signed a declaration warning of a climate emergency [1]. They stated, *"The consequences of global heating are becoming increasingly extreme, and outcomes such as global societal collapse are plausible and dangerously underexplored"* [1]. This underscores the urgent need for the world to decarbonize its energy production.

Buildings play a pivotal role in this transition, as they account for a significant portion of energy consumption and CO_2 emissions. In Europe, buildings constitute 40% of total energy use and 36% of CO_2 emissions [2]. Decarbonizing the building sector is therefore essential for achieving broader climate goals [3].

One promising approach to reduce emissions of buildings is the use of Aquifer Thermal Energy Storage (ATES) systems [4]. These systems sustainably store thermal energy in underground aquifers. During summer, cool groundwater is extracted for cooling building, and the resulting warmed water is then stored in a dedicated warm well. Conversely, in winter, this stored warm water is retrieved to serve as a heat source for a heat pump, supplying the building with heat. After the heat is extracted, the cooled water is returned to the cold well, thereby recharging the cold well for use in the following summer. Such systems are particularly effective in buildings that require both heating and cooling and are already widely implemented in the Netherlands [4], underscoring their highly sustainable nature.

To further enhance the efficiency and sustainability of ATES systems, they can be integrated into a thermal networks. Notably, 5th Generation District Heating and Cooling (5GDHC) systems are designed for this purpose, operating at low, near-ambient temperatures (10–25°C). This design allows for high flexibility and energy efficiency through decentralized heat pumps within each building, enabling bidirectional thermal flow where both heating and cooling demands are served by the same loop [5]. Crucially, this bidirectional capability means that buildings can exchange energy, for instance, a building requiring cooling can supply its excess heat to another building with a heating demand, effectively balancing the network. Importantly, 5GDHC networks are inherently designed to be coupled with seasonal storage solutions like ATES [6], making their combined application particularly effective for optimized energy exchange and overall decarbonization.

However, there are multiple reasons why 5GDHC networks are complex. First, these networks can be complex due to the integration of multiple ATES wells, which are often located far apart and not always situated in the most practical locations [7]. In renovation projects, this complexity is further compounded by the fact that ATES wells and district heating infrastructure are typically constructed in phases. This phased implementation is mainly driven by the high costs associated with drilling the wells, extending the heating network, and preparing buildings to be ATES-ready [8]. Secondly, a 5GDHC network inherently comprises both a cold and a warm network. It is important to note that the warm and cold wells are not connected by the same network, as warm wells are exclusively connected to the warm network, and cold wells are exclusively connected to the cold network. Moreover, there is considerable uncertainty regarding the actual energy demand of buildings. This is particularly evident in university environments, where the cooling demand is difficult to predict due to the unknown nature of future research activities that may take place in these buildings [8]. Connecting an entire residential district to a district heating network can result in a spiderweb-like

system of pipelines [9]. Lastly, the unique characteristics of each building or street further complicate the design, often requiring varying pipe diameters and customized layouts. In addition, the dense underground infrastructure, consisting of existing pipes, cables, and other utilities, makes it even more challenging to optimize the routing and configuration of the network [10].

A biologically inspired approach offers a potential solution. In Japan, researchers conducted an experiment with using actual slime mould (Physarum polycephalum) to organically model network and determine the optimal layout for the city's metro network [9, 10]. Remarkably, the slime mould's design closely mirrored the actual metro system in Japan, which is widely regarded as one of the most efficient and well-structured networks in the world [9, 11]. This experiment highlights the impressive problem-solving capabilities of slime mould, suggesting its potential contribution to designing complex 5GDHC thermal networks [12]. Therefore, the aim of this study was to investigate if an optimal thermal network could be designed using the principles demonstrated by slime mould growth. Using a simulation, it was explored how this biological phenomenon could create an ideal pattern for district heating systems, addressing the inherent complexities of connecting buildings, ATES systems, and other components in an efficient and sustainable way.

Which leads to the following research question: How c*an a slime mould growth model be used for the design of a heat distribution networks in combination with an ATES system in an urban environment?*

1.1. Approach

The core of this thesis involves the development and implementation of an algorithm that mimics the natural behaviour of slime moulds to create efficient networks. Chapter 2 outlines the fundamentals of the slime mould that serve as the basis for the algorithmic design and describes the technical details and specific code, including how cost factors and hydraulic performances are modelled. This chapter also includes the model's verification process, detailing the specific measures taken to ensure the functionality and reliability of the developed model through various tests and comparisons. Chapter 3 includes a full execution of a case study, which involves a complex practical case: the TU Delft Campus. This location was chosen due to its diverse and dynamic energy demands. By testing the model in this realistic and challenging environment, its practical applicability and potential for sustainability improvement was demonstrated. The results of the case study, specifically the networks generated by the algorithm for the TU Delft Campus, are presented in detail in section 4. A comprehensive discussion of the findings, including the performance of the generated networks based on crucial criteria such as economic costs, network topology, and hydraulic performance, will be provided in section 5, offering in-depth insights into the strengths and limitations of the chosen algorithmic approach for designing heat networks and reflects on its broader implications.

2. Method

2.1. Fundamentals of the slime Mould

The fundamentals of the slime mould can give a clear view of how a slime mould moves, behaves, and makes decisions. As a single-celled organism, slime moulds exist as a giant cell with multiple nuclei, enabling coordinated movement and decision-making [13]. They navigate their environment by moving towards food sources, detecting chemical gradients and actively avoiding harmful

substances [14]. This movement is characterized by oscillatory internal fluid flow, which facilitates nutrient distribution and signal processing, allowing for adaptive responses [14]. Furthermore, slime moulds exhibit remarkable energy optimization by initially spreading out randomly, but then reinforcing the most efficient paths while retracting less optimal connections. This process minimizes energy expenditure while maximizing resource acquisition [15]. Their abilities demonstrate decentralized intelligence, as they exhibit complex problem-solving without a brain or nervous system, comparable to biological computing [16]. Finally, slime moulds possess unique characteristics regarding cell splitting and fusing. They can be divided into smaller fragments, each capable of independent movement, although when reunited, they reorganize and function seamlessly as a single coordinated entity once again [17].

2.2. Model criteria

To precisely define the optimization objectives of the model, it is essential to explicate the interests of all involved parties within a heat distribution network. Typically, these parties are the supplier (the operator of the thermal network, including ATES wells and technical rooms) and the user (the user who, for example, wishes to heat or cool their building) [18].

For the user, the certainty of a well-functioning system is crucial, along with the lowest possible costs per Giga Joules (GJ) of heating and cooling, with a maximum of a market-conform price per GJ. The supplier, on the other hand, strives for profitable operation, often expressed as an annual margin of 6% to 12% on the total investment across the complete system [18].

In the development of a thermal network, Capital Expenditure (CapEx) and Operational Expenditures (OpEx) are crucial factors that collectively determine the Net Present Value (NPV) [18]. Within the context of the pipeline network itself, the primary objective is generally to minimize CapEx. The CapEx primarily encompasses the construction costs of the pipelines. These costs are directly related to the total length of the network and the chosen pipe diameters. The goal is to find an optimal balance between minimizing the total pipe length and selecting the most suitable diameters.

2.3. Slime mould algorithm: model implementation

To provide a deeper insight into the precise functioning of the model, this section describes how the slime mould finds the optimal path between a point A and a point B with the use of equations. Subsequently, it explains how it avoids obstacles and grows towards all other food sources. The model is modelled in Grasshopper [19] using the software Rhino8 [20], with the use of Python language.

2.3.1. Movement direction options for the slime mould

The movement capability of the slime mould is entirely determined by the permissible movements from node to node within the grid. In option a) of Figure 1, only horizontal and vertical movement is allowed, with each step having a length of 1. Option (b) also allows diagonal movements, resulting in a step length of √2, thereby significantly increasing path efficiency. Option (c) offers an even more efficient path by permitting steps with a length of $\sqrt{5}$.

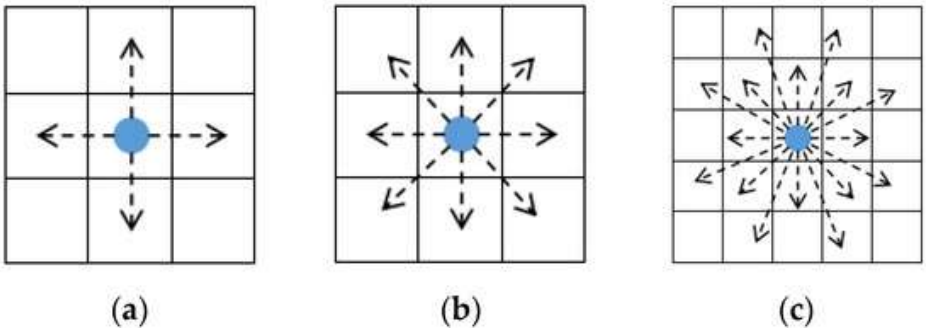

Figure 1: possible direction I a chessboard grid [21]

To evaluate the efficiency of each option, the least favourable path is compared with the Euclidean distance, which represents the direct path [22]:

$$l = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2} \quad (1)$$

For illustrative purposes, in option (*a)*, the least favourable path is the route from point (0,0) to point (1,1). The Euclidean distance in this case is √2, while the path length is 2. This results in a minimum efficiency of $\frac{\sqrt{2}}{2}$. For option (*b)*, the efficiency is $\frac{\sqrt{5}}{1+\sqrt{2}}$, and for option (*c)* $\frac{\sqrt{17}}{2+\sqrt{5}}$.

Since option (c) allows larger steps of 2 in the x or y direction, obstacles can be bypassed (see appendix E.2. This issue could be addressed by further modifying the Bresenham algorithm. However, this leads to a substantial increase in both the number of data points and computational complexity. Consequently, option (b) is selected, offering an acceptable path efficiency of $\frac{\sqrt{5}}{1+\sqrt{2}}$.

2.3.2. The Growth of the Slime Mould from Point A to Point B

To determine how a slime mould locates the nearest food source and estimates the associated cost, the A* algorithm, a widely recognized graph search algorithm [23], is employed for efficient pathfinding. This algorithm effectively identifies the least-cost path from a starting node to a goal node by ingeniously combining two cost estimates based on absolute distance [23]. It uses a heuristic function to estimate the cost from a current node (point) to a goal node (desired point). This heuristic, also known as the chessboard distance [24]. The cost estimations is performed using a variant of the Manhattan distance equation:

$$h(n) = |n_x - goal_x| + |n_y - goal_y| \quad (2)$$

Where $n = (n_x, n_y)$ represents the coordinates of the current cell and $goal = (goal_x, goal_y)$ represents the coordinates of the goal cell. Critically, this method ensures that the model never underestimates the true distance between points, which is a key property for the A* algorithm's optimality guarantee.

Subsequently, the A* algorithm utilizes a priority queue as a crucial data structure. This queue efficiently manages a collection of nodes, representing potential locations within the grid that still require examination. Each node in this priority queue is assigned a priority determined by its estimated total cost, denoted as $f(n)$, which represents the total cost to travel from the starting

point to the goal via that particular node [25]. This total cost, $f(n)$, is fundamentally the sum of two distinct components, as expressed by the equation:

$$f(n) = g(n) + h(n) \quad (3)$$

Here, $g(n)$ indicates the actual costs already incurred to move from the starting point to the current node n. In the context of the code, this specifically accounts for the length of the pipe segments and their associated material costs up to node n. Further details on this calculation can be found in sub-section 2.3.3. Conversely, $h(n)$ provides a heuristic estimate of the remaining costs required to travel from the current node n to the final goal. The priority queue is organized to consistently yield the node with the lowest $f(n)$ value first for examination. Through this strategic approach, the A* algorithm consistently selects the paths that appear most promising at any given moment, thereby ensuring the discovery of the optimal route, balancing both shortest path and lowest costs.

In the code, the starting point (point A) for the A* search is the current location of the slime mould. The goal (point B) is a food source (supply or demand node) that the slime mould is trying to reach. The A* algorithm can also examine intermediate steps (follow-up step) to get closer to the food source, where each examined cell in the grid is placed in the priority queue as a potential node and evaluated based on its $f(n)$ score.

2.3.3. Calculation of the Actual Costs g(n)

The costs of an individual segment of a heating pipe between two cells (x_1, y_1) and (x_2, y_2) are calculated by the equation (see Appendix C for further explanation):

$$g(n) = \left(C_m(Q) * F + C_p\right) * l \quad (4)$$

Where, $C_m(Q)$ represents the cost per meter of the selected pipe material, in this case, High-Density Polyethylene (HDPE). This material cost is dependent on Q the flow rate (m^3/h) passing through the segment, with the specific cost values based on various flow rates detailed in Table 1. Next, C_p signifies the basic installation costs per meter. l is defined as the Euclidean distance between the two connected nodes, as further elaborated in Equation 1. Finally, F acts as a weighting factor for any extra costs. Further details about these parameters are presented in Appendix C.

A crucial consideration in the cost calculation is that the costs for bends, branches, constrictions, and the direct connection costs of the district heating network to buildings or ATES wells are not included in this model.

Table 1: Cost per meter (€/m) of High-Density Polyethylene (HDPE) pipes per flow rate (m³/h) and nominal diameter (DN) (mm).

Flow rate (m^3 /h)	5,9	11,9	20,7	37,2	67	108,1	229,8	412	663,5	992	1405
DN (mm)	50	65	80	100	125	150	200	250	300	350	400
Cost per meter ($C_m(Q)$) (€/m)	3	5	11	14	16	26	41	90	140	180	250

2.3.4. Creating the obstacles

To make obstacles for the slime mould, a Bresenham-inspired algorithm is used [26]. While it is based on the original Bresenham line algorithm, which efficiently computes which grid cells best approximate a straight line between two points, this adapted version allows for only horizontally and vertical movement, see Appendix E.1 for further details. This makes it more suitable for use in the context of the slime mould simulation, where flexibility and obstacle avoidance are crucial. Below is a step-by-step explanation, which is inspired on the Bresenham line function [25]. The algorithm takes two grid-based points as input: (x_0, y_0) and (x_1, y_1).

It calculates the horizontal and vertical distances:

$$dx = |x1 - x0| \tag{5}$$

$$dy = |y1 - y0| \tag{6}$$

It determines the step direction for x and y:

$$sx = 1 \; if \; x0 < x1 \; else \; sx = -1 \tag{7}$$

$$sy = 1 \; if \; x0 < y1 \; else \; sy = -1 \tag{8}$$

An error term is initialized:

$$\varepsilon = dx - dy \tag{9}$$

Starting from (x0, y0), the algorithm moves step by step toward (x_1, y_1). At each step, the current point is added to the list of line points. A temporary variable is calculated, to decide the direction of the next step:

$$e2 = 2 * \varepsilon \tag{10}$$

The step direction is chosen via the following formula:

$$\begin{cases} x0 \mathrel{+}= \; sx \; and \; \varepsilon \mathrel{-}= dy \; if \; e2 > \; -dy \\ y0 \mathrel{+}= sy \; and \; \varepsilon \mathrel{+}= dx \;\; if \; e2 < dx \end{cases} \tag{11}$$

This process continues until the end point (x_1, y_1) is reached. The determined grid cells (path points) are then stored in a matrix M, where each cell $M[x][y]$ representing a path receives a specific value.

2.3.5. Avoiding Obstacles

Now that such an effective path can be made between two points, it is useful to avoid obstacles, such as buildings and other objects. This is done by reading the obstacle matrix, created in 2.3.4.

The obstacle matrix is as large as the grid, in this grid each coordinate has a value, which can be found with the function $M[x][y]$, this value can give three options, using the following equation:

$$M[x][y] = \begin{cases} 0 \\ 9999 \\ else \end{cases} \tag{12}$$

Here *0* indicates a free path, incurring no extra cost beyond the normal $f(n)$ value (as explained in section 2.3.2). Conversely, a value of 9999 signifies an impossible path to traverse. This specific

value is chosen to ensure clear differentiation from penalty values applied when the 'difficulty' function for $g(n)$ is intended to be fully utilized. By assigning such a high, distinct value for impassable routes, it was guaranteed that these paths would never be accidentally selected as merely 'very expensive' (i.e. traversable but difficult), thereby preventing unintended route choices. *else* numerical value describes the difficulty of entering the path, where a higher number signifies greater difficulty, by increasing the $g(n)$ value (see section 2.3.3) with the value found in the matrix (note: this function is currently implemented in the code but is not actively used in the simulation).

2.3.6. Expansion of the Network

To efficiently expand the network to all food sources, the created code reuses existing pipe segments (which the code has placed before) where possible. The already installed connections are stored in the dictionary 'Slime', which contains for each occupied cell (x, y) the installation step at moment t installed (x, y) and the pipe diameter $DN(x, y)$.

When the A* algorithm searches for a new path to connect to another food source, the calculation of the actual costs $g(n)$ is adjusted. For a movement of node n to neighbour n', the actual costs $g(n)$ depend on whether the segment already exists in the 'slime' dictionary and whether the required pipe diameter needs to be increased (see Equation 13). This cost-saving strategy influences the shape of the resulting network, where new connections tend to branch off from existing pipes. The segment cost is defined by the following Equation:

$$g(n) = \begin{cases} \left(C_m(Q_{new}) - C_m(Q_{old})\right) * F * l(n,n') \; if \; (n,n') = s_p \\ (C_m(Q_{new}) * F + C_p * l(n,n') \qquad if \; (n,n') \neq s_p \end{cases} \tag{13}$$

Here, $C_m(Q)$ indicates the costs per meter of the pipe with the required capacity for a given flow. The total flow through the segment (n, n') after the new connection is denoted as Q_{new}, while Q_{old} refers to the existing flow through the segment (n, n') if it is already part of the already placed network. The term (n, n') itself refers to the pipe segment between node n and neighbour node n'. Furthermore, l represents the Euclidean distance (length) (Equation 3) of this segment and C_p signifies the basic costs per meter for installing a new pipe. Lastly, S_p refers to the set of all cells (or segments) that are already part of the existing pipe network, which is stored in the 'Slime' dictionary.

The total estimated costs $f(n) = g(n) + h(n)$ (see section 2.3.2) still guide the search of the A* algorithm. However, because the costs of moving over existing, sufficiently dimensioned pipes are lower (no new installation costs need to be incurred), the algorithm will naturally prefer paths that use parts of the already existing network, which results in a potentially more cost-effective and more integrated heating network.

2.3.7. Slime Mould Finds Food Source (Supply or Demand)

When the slime mould model identifies a food source, this source is immediately removed from the lists of potential food sources, specifically slime_data['supply'] and slime_data['demand']. This removal prevents other 'slime moulds' from redundantly attempting to connect to an already served food source, therefore avoiding unnecessary network expansion.

In this model, the 'slime mould' inherently aims to balance flows. Here, a positive value consistently represents energy supply, while a negative value indicates energy demand. For instance, in a cooling scenario, a building would effectively *supply* energy (e.g., waste heat), meaning the associated

(warm) wells would, conversely, have a *demand* for energy. All flow capacities are expressed in m^3/h, as a constant ΔT across the entire system is assumed.

Consider an example: if the 'slime mould' first encounters several sources, each supplying 50 m^3/h, it connects these wells. However, at this initial stage, the diameters of the connecting pipes are not yet determined and remain at a minimal size. If the 'slime mould' subsequently reaches a building with an energy demand of −100 m^3/h, the pipe connecting this building to the wells will then be increased to a minimum capacity of 100 m^3/h to meet the building's energy requirement. If a perfect balance exists between the energy supplied by wells and demanded by buildings, it is possible for isolated pipe networks (or 'islands') to form, resulting in multiple, separate heat distribution networks.

If this simultaneous creation is attempted, two main problems arise. Firstly, it becomes impossible for the 'slime mould' to design both cold and warm networks concurrently, as cold and hot wells would then merely connect amongst themselves, forming isolated 'islands' and because a single building connection point cannot be assigned both an energy supply and a demand simultaneously within the current framework. The proposed solution for this issue involves temporarily designating both hot and cold well as either *suppliers* or *demanders collectively*, making the building assume the opposite role. Furthermore, to maintain energy balance, the capacity of the wells must either be halved or the capacity of the building doubled. This approach will be explored in the case study, as detailed in Section 3.

The second problem is that the 'slime mould' might bypass a demander if it has not yet encountered any suppliers, leading to this demander never being connected to the rest of the network. A potential solution is to develop two variants of the simulation: one where buildings act as suppliers and the warm or cold wells as demanders, and vice versa. While this sometimes might intuitively seem counter-logical (as a warm well is typically a supplier in a heating situation and a cold well a demander), the exact implementation to resolve this in the current code remains unclear. Therefore, the case study in Section 3 investigates both variants.

2.3.8. Fusion of Slime Moulds

The fusion of slime moulds (merge slimes function) occurs when the slime territories of two or more independent slime moulds overlap. When a collision is detected, the data of the colliding slime moulds (including the occupied cells, connections, flow, etc.) are merged into a single slime mould, allowing the network to integrate and potentially more efficient global solutions to emerge. This is because a certain Slime Mould sees another Slime Mould as a food source and vice versa, which ensures that all slime moulds and food sources will be connected to each other at the end of running the code.

2.3.9. Creating both a warm and cold ATES network

The slime mould can create an ATES network that connects different demanders (buildings) with suppliers (ATES wells). However, as discussed in 2.3.7, the network consists of a warm and a cold pipe network. Here, the warm piping will never be connected to a cold well and vice versa. This means that the cold network looks slightly different than the warm network. This is solved by splitting the slime mould code into two sections. First, one network is completely drawn by the code

and when it is ready, the other network is made, but then with an advantage step, the placement costs will be zero for these parts. See formula for explanation:

$$g(n) = \begin{cases} (C_m(Q_{new}) * F * l(n,n') & if\ (n,n') = s_p \\ (C_m(Q_{new}) * F + C_p * l(n,n') & if\ (n,n') \neq s_p \end{cases} \quad (14)$$

For an explanation of the variables, see section 2.3.6.

2.3.10. Model Resolution

Modifying the grid resolution significantly influences both the model's accuracy and computational load. Moreover, it directly determines the step size applied in Section 2.3.1.

A higher resolution increases the level of detail in the environment, which can be particularly beneficial for how the slime mould interacts with obstacles. For example, it allows for more precise placement of obstacles and paths, potentially making it possible for the slime mould to find routes between tightly spaced objects. In contrast, at lower resolutions, fine spatial distinctions may be lost.

Consider Figure 2(a) in this scenario, with a higher resolution, two obstacles are stacked vertically. The top one imposes n penalty points, and the bottom one m. Because of the finer grid, there is a small gap between them that can be navigated. However, in Figure 2(b), due to a lower resolution, the obstacles are merged into a single area with a penalty of $n + m$, blocking the path entirely.

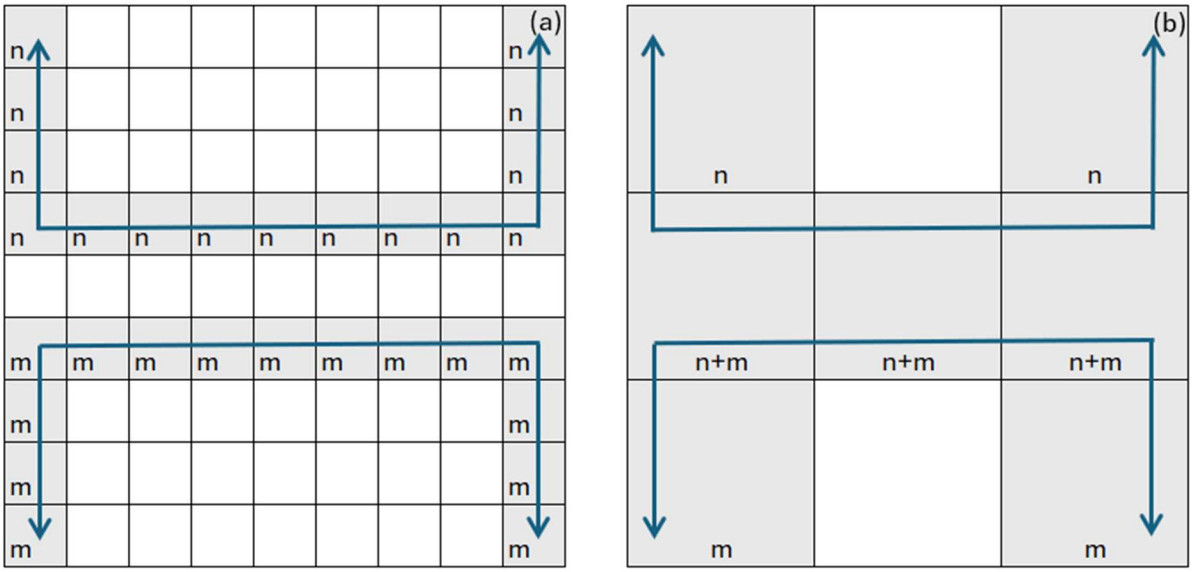

Figure 2: (a) an obstacle with a high resolution, (b) an obstacle with a low resolution

However, increasing resolution comes at a cost. A key challenge is the large volume of data required. For instance, modelling a 1000 by 2000 meter area at a resolution of 1x1 meter results in 2 million data points. Increasing the resolution to 0.1x0.1 meters increases this number to 200 million points.

Such a high level of detail may not always be necessary and can significantly burden the system. Higher resolution increases memory usage, slows down processing, and demands greater computational resources. Therefore, it is important to find a balance between a resolution that provides sufficient detail for the results, without overwhelming the system with excessive data.

For example, a space of approximately one meter between two buildings is typically required for the installation of a district heating network. This means that a resolution of 1 meter is sufficient for the current model. However, when including additional obstacles, such as electrical cables or other underground infrastructure, it may be beneficial to increase the resolution for improved accuracy.

2.4. Model verification

The model is developed in multiple phases to allow for verification and early detection of potential issues. This step-by-step approach improves both robustness and traceability of errors in the model.

2.4.1. Basic model, Single slime mould

The process starts with a basic version of the model (for earlier versions, see Appendix F.1). This version focuses purely on the slime mould's core behaviour, as described in Sections 2.3.1, 2.3.2 and 2.3.6, where a single mould seeks to connect its food sources as efficiently as possible.

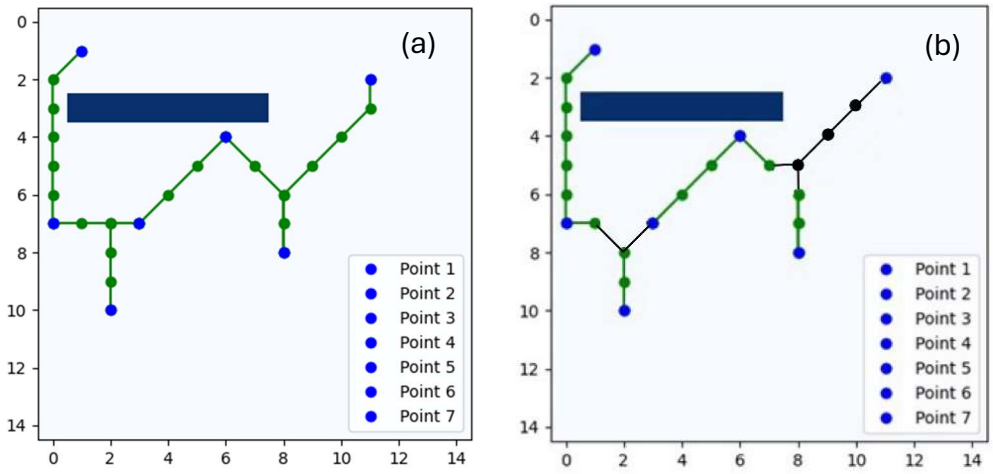

Figure 3: (a): Result of the slime mould code, used length is 26,7, (b): Optimization of the result made in (a) by a human, used length is 26,1

Figure 3 reveals an algorithm-generated path with a total length of 26,7 units. This result demonstrates that the model establishes a solid and efficient baseline. However, through subsequent manual revision and optimization of this path, it was possible to further reduce the length to 26,1 units, as shown in black in Figure 3(b). While this indicates there is still room for human refinement after automatic generation, the overall finding remains that the model produces a cost-effective route.

2.4.2. Basic model, multiple slime moulds

The next step was to introduce multiple slime moulds, expanding the model as described in section 2.3.8. This allowed for the simulation of multiple slime moulds cooperating for optimal pathfinding.

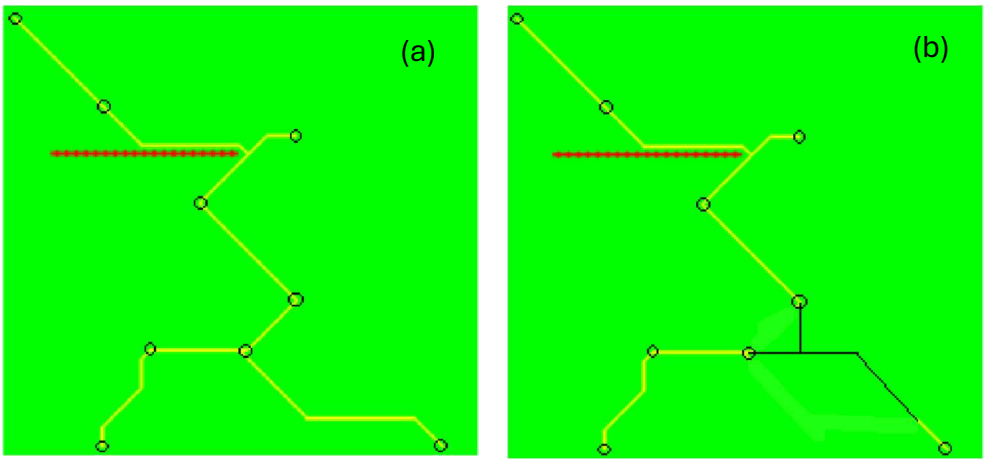

Figure 4: (a): Results V2 of the slime mould code, (b) optimization by human of the results made in (a)

Upon reviewing the results in Figure 4(a), it became apparent that there is a slightly better route than the ones chosen by the model as seen in Figure 4(b). This can be explained by the optimization strategy: the model builds paths incrementally, point-to-point, and does not re-evaluate previously established segments. A real slime mould is capable of doing this, but due to computational limitations, re-evaluation was not implemented.

2.4.3. Slime mould

Once the base functionality was verified, the model was tested using actual DWG files from the TU Delft Campus. These files were read via Rhino8 and then converted into coordinates of the line vertices (see appendix B.1). As described in Section 2.3.4, these vertices were transformed into matrix M, which assigns a weight factor of 9999 to them. Additionally, a random selection of start and end points was tested to see if the slime mould avoids the obstacles using the equation from Section 2.3.5. Initially, with the base Bresenham algorithm from Section 2.3.4, this was not the case. However, after algorithmic adjustments, the slime mould successfully avoids all obstacles, as shown in Figure 5.

Figure 5: slime mould avoiding all obstacles

2.4.4. Implementation of pipe sizing

Subsequently, functionality is added to the code to enable automatic pipe diameter selection, as described in Section 2.3.3. This addition marks a fundamental shift in the optimization strategy: moving from a primary objective of minimizing total pipe length to one focused on minimizing overall network cost. Up to that point, the model's core optimization objective is solely based on the total physical length of the network. However, as comprehensively explained in Section 2.3.3, the introduction of automated diameter selection means that segments carrying higher energy flows are now proportionally considered more expensive than those accommodating lower flows. This nuanced change elevates cost to become the explicit and paramount objective of the entire optimization process. Consequently, the model now dynamically determines the appropriate pipe diameter for each individual network segment based directly on its calculated energy flow.

As a direct result of this enhanced complexity, analysing the model's output has become significantly complex**,** given that not every segment contributes equally to the total network cost. Therefore, to ensure the robust reliability and absolute accuracy of the results generated by this modified optimization approach, it is deemed essential to subject the updated model to a rigorous analysis of a controlled, simplified scenario. For this purpose, a representative, small-scale network design is utilized, as illustrated in Figure *6*. This specific network features a limited number of defined supply and demand nodes and clearly delineated pipe segments, allowing for a manageable and transparent assessment. This specific case is strategically chosen because its inherent simplicity allowed for a thorough and transparent verification of one critical aspect: the precision of the automated pipe diameter selections.

Figure 6: (a) pipe sizing growing after each well, (b) impact of pipe sizing because of supply and demand, (c) solved pipe sizing problem

In Figure *6*(a) three wells (yellow arrows) are shown, each supplying $50\ m^3/h$. According to the table in Section 2.3.3, the selected pipe for each of these individual flows should be $DN\ 125$ $(DN = 125mm)$, which aligns with expectations. The combined pipe carrying a total flow of $150\ m^3/h$ correctly results in a DN200 selection.

In Figure *6*(b), two cold wells (yellow arrows) are visible, each supplying $50\ m^3/h$, along with a building connection (red arrow) demanding $82\ m^3/h$. The pipe exiting the top of the figure connects to additional wells and building connections outside the view. Logically, one would expect the pipe going to the wells to be able to handle $100\ m^3/h$, and the one leading to the building to carry

$82\ m^3/h$. The remaining $18\ m^3/h$ should exit the system via the top pipe, which would suggest a relatively small diameter. However, this is not the case, indicating the system is not functioning as intended.

After further adjustments to the code, the root of the issue was found: the absolute values of all supply and demand flows were being used, which caused the model to lose track of flow direction and misassign pipe segments.

In Figure *6*(c), two red arrows represent building connections with energy demands of $183\ m^3/h$ and $86\ m^3/h$. The pipe diameter increases step by step along this path, which is logical. Further down, two yellow arrows represent wells delivering $25\ m^3/h$ and $80\ m^3/h$. After these wells, the main pipe narrows again, as expected these two wells must meet the energy demand of the buildings above. Additionally, the pipe continues downward, since the combined supply from these two wells is still insufficient. A flow shortage of $164\ m^3/h$ remains, which must be supplemented from further downstream.

2.4.5. Final verification

Finally, the model was extended to generate two separate networks instead of a single one, as described in Section 2.3.9. In this setup, it is essential that all cold wells are connected to the buildings via one network, and all warm wells via another. Additionally, the slime mould is designed to optimize the network layout by maximizing the overlap of shared paths between the two systems.

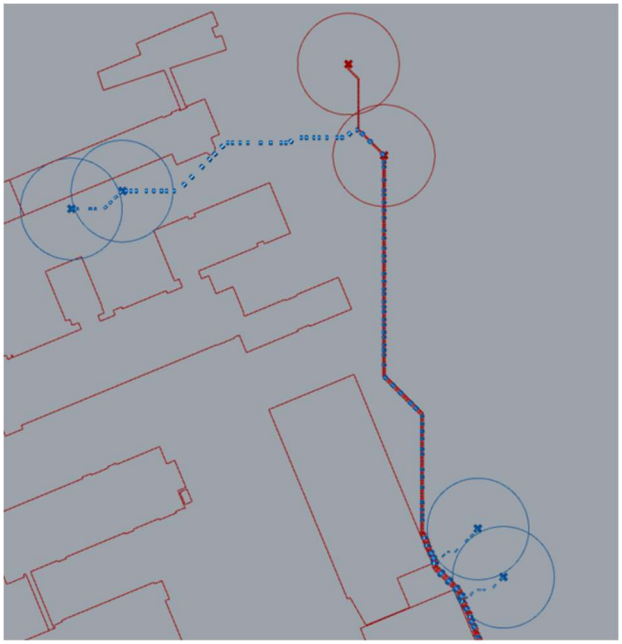

Figure 7: Two generated networks, one to the cold wells and one to the warm wells

In Figure 7, it is clear that the blue (cold) wells are only connected through blue pipelines, while the red pipelines exclusively connect the red (warm) wells. Furthermore, a significant portion of the path is shared between the two networks, which is the intended goal. This shared path approach minimizes excavation work, as only a single trench needs to be dug instead of two.

3. Case Study

3.1. Site description

To answer the research question, the developed slime mould model is applied to the TU Delft Campus. The campus data, provided in DWG format by Anne Medema [8], is first filtered and transformed into a set of coordinates (see appendix B.1). These coordinates are then used by Section 2.3.4 to create obstacles. As the Campus is currently not yet equipped for an ATES system, certain assumptions are necessary. These assumptions, along with the relevant variables and input parameters that need to be specifically defined for this application, are discussed in this section. Figure 8 visually represents the Campus context for this modelling, displaying the various assumptions and input parameters that are explained in the subsections.

Figure 8: The visualisation of all assumption made on the TU Delft Campus

3.2. Energy Demand of Buildings

For a properly functioning ATES system, the thermal energy injected in summer should roughly equal the heating demand during winter see Appendix H. However, this is currently not the case. For example, assuming the buildings have an energy label of D (see Appendix A for the energy label classification), there is no cooling demand, while the heating demand is $363{,}9 MJ/m^2/year$, this is clearly unbalanced. To reflect a more sustainable scenario, all buildings are assumed to have energy label A and to serve as a university function. Floor areas for each building were provided by Anne Medema [8].

3.3. Building Connections

Each building is assumed to be connected to the new ATES network in the same way it is currently connected to the existing district heating network. For buildings with multiple connections, the total energy demand is assumed to be equally distributed across all connections. For newer buildings such as building 29 and 33, which already have individual ATES systems, the connection points to those buildings are manually chosen. These selected locations do not necessarily reflect technical feasibility in terms of proximity to mechanical rooms or infrastructure.

3.4. Obstacles

Only buildings, the tramline and two fields behind building 37 and the parking garage are considered as physical obstacles see Figure 8 for all the obstacles shown on the map. For the tramline, jacket pipes are present to allow the ATES network to pass through (see Appendix G for jacket pipe placements). It is assumed that none of the obstacles can be crossed, not even using underground drilling methods.

3.5. ATES Wells

All existing ATES wells on Campus are assumed to remain in use. As significant demand and available space remains, additional ATES wells are placed based on locations suggested in the IF report (Appendix D.1), which guide the positioning of the remaining wells.

3.6. Scenario's

For the scenarios, there are two variable inputs, the input of different starting points for the slime mould, and what the slime mould perceives as a positive food source and what as a negative food source, as discussed in 2.3.7.

3.6.1. Starting points

Six starting points, which are shown in Figure 8 and detailed with their coordinates in Table 2, are used in every possible non-empty combination to create a comprehensive set of simulation scenarios. The total number of such combinations is determined using the formula $2^n - 1$, where n is the number of unique start points. This formula calculates all possible subsets excluding the empty set, which corresponds to a simulation without any start points an invalid case, as the algorithm requires at least one point to operate. With six start points, this results in $2^6 - 1 = 63$ valid combinations. Accordingly, 63 unique scenarios are generated, covering all possible configurations involving one or more of the selected start points.

Table 2: the coordinates of 6 randoms starting points

x	840	616	307	225	261	627
y	64	1018	767	1014	1233	886

3.6.2. Positive and negative food sources

As described in Section 2.3.7, the slime mould model can interpret different food sources as either positive or negative, depending on their definition in the input data. For this study, four different scenarios are examined, each consisting of 63 variants, as shown in Table 3.

Table 3: Overview of the different scenario's used in the case study.

Scenario's	Negative (+)	Positive (-)
1 (B+, W-)	Buildings	Warm wells
2 (W+, B-)	Warm wells	Buildings
3 B+, WC-)	Buildings	Warm and cold wells
4 (WC+, B-)	Warm and cold wells	Buildings

In all these cases, an initial optimization is performed on the primary network. Subsequently, if applicable, the complementary network (for instance, the cold network in scenarios primarily focused on warm wells) is generated. For scenarios involving both warm and cold wells, optimization is performed on each network independently. Due to computational constraints and the need to manage the extensive number of results, the specific combination of cold wells with buildings is not included in the optimization analysis. From each scenario, 63 variants are generated, as seen in 3.6.1.

3.7. Assessment framework

The main goal of this assessment framework is to systematically evaluate the usability and effectiveness of the Slime Mould Growth Model for the design of heat distribution networks in combination with an ATES system in an urban environment. This evaluation consists of a thorough analysis of both the model's performance and the intrinsic characteristics of the resulting network designs. The evaluation of the network solutions generated by the Slime Mould Model are conducted based on the following criteria.

3.7.1. Model Robustness & Solution Completeness

The primary objective is to identify the scenarios that are most robust in generating a high number of complete network solutions, which is essential for the model's practical applicability. The focus here is on analysing how the role assignment of buildings (source/consumer) influences the completion rate in both individual and combined configurations.

3.7.2. Cost-Length Relationship & Efficiency

Building upon the completeness of the solutions, the aim here is to investigate how the generated network solutions correlate with economic efficiency (costs) and physical extent (length), and to gain insights into potential trade-offs between these factors for different variants. The focus is on analysing the correlation, spread, and the influence of other optimization parameters on the cost-length relationship per strategy.

3.7.3. Flow Rate Distribution & Pipe Diameter Selection

To further deepen the understanding of the observed cost-length relationships, the goal here is to analyse the underlying technical factors that determine the cost and length of the networks. This specifically targets the distribution of required flow rates and the resulting choices for pipe diameters, which provides insight into the origin of the observed efficiency patterns. The focus is on analysing average distribution patterns and variability with mean and standard deviation (STDev) of pipe choices across different flow ranges per variant.

3.7.4. Geographical Paths and Utilization Frequency

Finally, this evaluation serves as an integrating visualization of the previously analysed aspects. The goal is to make the spatial outcomes of the model generation visible by visually presenting the frequency of path selection and the relationship between the most cost-effective paths and the most frequently used routes, as a comprehensive overview of the model's performance. The focus is on examining the frequency of path selection and the correspondence between the most cost-effective paths and the most frequently used routes on the Campus map.

4. Results

4.1. Completion different variants

To assess the model's robustness and its capacity to find complete network solutions across different configurations, the number of completed results per scenario is presented in Table 4.

Table 4: Overview of how many simulations were completed for each scenario.

Scenario's	**Completed results**
1 (B+, W-)	24
2 (W+, B-)	5
3 B+, WC-)	0
4 (WC+, B-)	23

Table 4 illustrates how the role assignment (source or consumer) of buildings influences the completeness of model solutions. In individual scenario's (1 and 2), the 'positive' building role performed better, while in combined scenarios (3 and 4), the 'negative' building lead to the most complete results.

4.2. Cost-length relationship

The relationship between the total network length and its corresponding costs is presented in Figure 9. This fundamental correlation provides essential insight into the cost-efficiency of the developed network configurations.

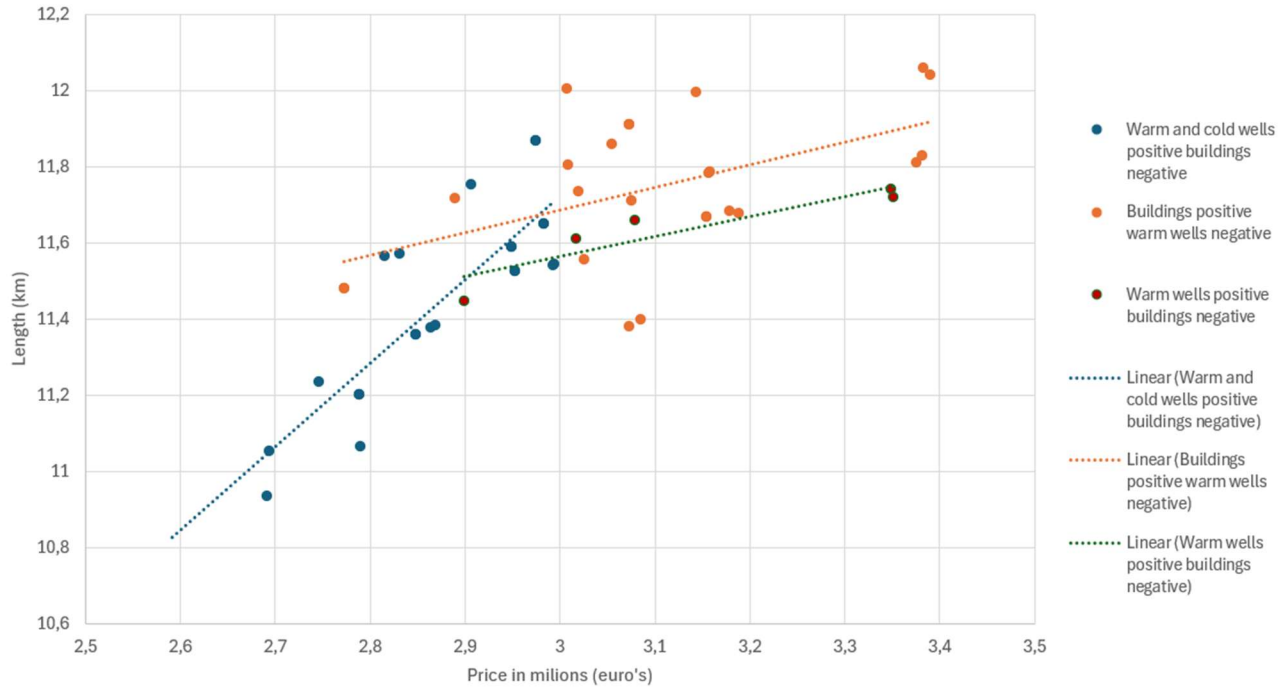

Figure 9: Visualisation of cost-length relationship of the different simulations.

A clear positive linear correlation is observable. For the blue line, representing the integration of both warm and cold networks, there is relatively little spread among the variants. Although a cost difference of €2.8 million to €3 million is visible at approximately 11.6 km length, and a length spread of 11.5 km to 11.9 km at around €3 million in costs, these differences are considered not significant in this context. However, this pattern differs for the orange line, representing a phased optimization approach where warm wells are prioritized. Here, the spread among the variants is considerably greater. For instance, at a length of approximately 11.7 km, there are network variants ranging from €2.9 million to €3.35 million, which represents a difference of about 15 percent in costs. It is also notable that some networks in this group, despite being considerably longer in total length, can still be significantly cheaper than other, shorter networks within the same category. Furthermore, the red line is also present in Figure 9. As shown in Table 4, this line contains only 5 data points. Nevertheless, it is observable that this line exhibits a similar slope to the orange line. Moreover, the two lines are not far apart.

4.3. Flow Rate Distribution and Pipe Diameter Selection Across Scenarios

Figure 10 presents the distribution of required flow rates (and the price) across different network variants. This distribution directly influences the selection of appropriate pipe diameters for the network segments. Given the challenge of directly comparing the individual results a combined mean and STDev graph has been generated to synthesize the data from the different network variants.

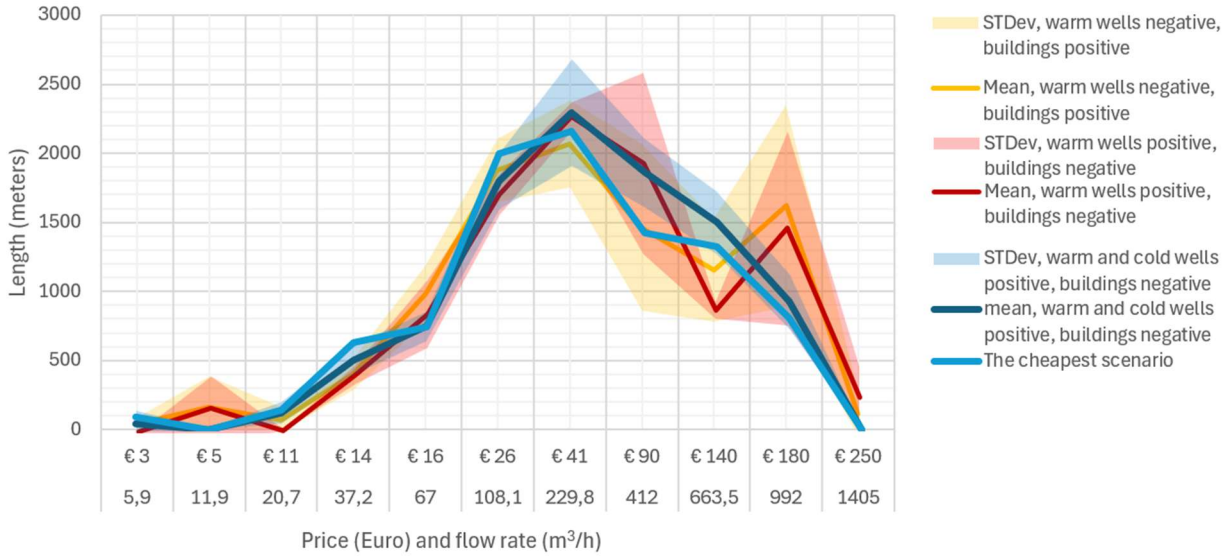

Figure 10: Mean and STDev graph of the relationship between pipe length and flow rate of the pipe.

Figure 10 displays the mean and STDev of the network length distribution across three valid optimization variants, represented by shaded error bands around the mean.

Upon inspection of Figure 10, it is evident that the average choice of network length shows consistent behaviour across all variants for flows ranging approximately from 20.7 m^3/h to 108.1 m^3/h. At a flow of 11.9 m^3/h, the 'warm wells (+), buildings (-)' and 'warm wells (-), buildings (+)' variants utilize considerably more network length compared to the 'cold and warm wells (-), building (+)' variant.

Moving towards relatively larger flows, a divergence in patterns becomes apparent. Between 229.8 m^3/h and 663.5 m^3/h, the average network length chosen for the single well variants (warm wells + / -) is notably lower than for the combined (both cold and warm wells) variant. This trend reverses for flows between 992 m^3/h and 1405 m^3/h, where total length becomes significantly more expensive. Furthermore, a pronounced peak in network length is observable for at least one of the single-source variants around 992 m^3/h.

Regarding the standard deviation, it is noteworthy that for the combined variant (both cold and warm wells), the STDev band initially widens and then narrows, exhibiting a somewhat parabolic shape, indicating greater variability in the middle flow ranges. In contrast, for the 'warm wells (+), buildings (-)' and 'warm wells (-), buildings (+)' variants, this pattern is not observed.

4.4. Geographical Visualization of Heating Network Solutions

In Figure 11, a visual trendline of the heating network is presented. Additionally, the most cost-effective network is also shown. These trendlines are visual representations on the Campus DWG map, with different shades of green showing the frequency with which a particular path is selected or utilized across multiple generated network solutions by the slime mould model. A darker shade of green indicates a higher incidence of path selection among the various model outputs.

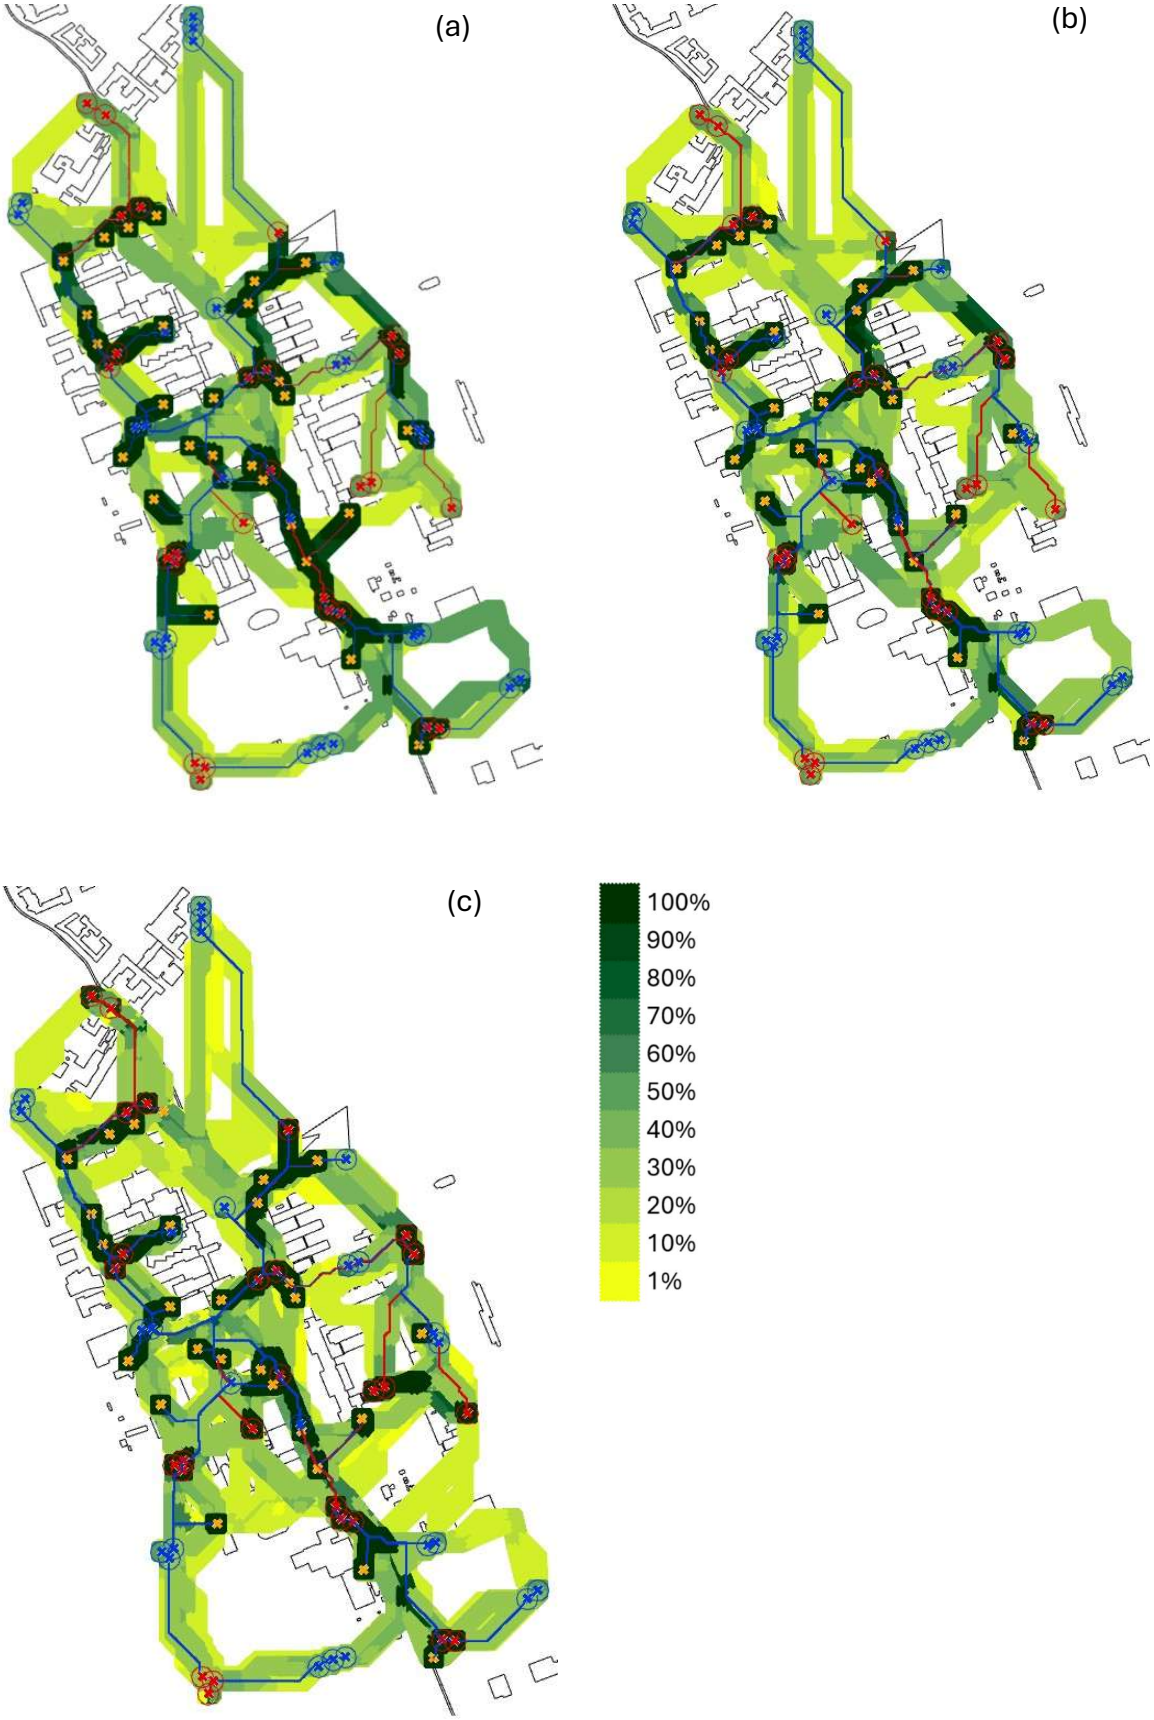

Figure 11: (a), top 10% cheapest scenarios of the warm and cold network, with a price range from 2.69 to 2.96 million, showing the best scenario in red and blue. (b) all completed scenarios, with a price range from 2.69 to 3.39 million, showing the best scenario in red and blue. (c) all scenarios, with a price range from 2.69 to 3.39 million, showing the best scenario in red and blue.

Figure 11(a) shows the option which are within 10% above the cost of the cheapest option. This shows that the best scenario almost always aligns with the most popular route between the points. An exception appears in the middle-right area of the figure, where a more popular route is bypassed.

In Figure 11(b), all full scenarios are presented. The best scenario generally follows the darker (more popular) paths, similar to Figure 11(a). However, a deviation is observed near the Spar supermarket on Stieltjesweg, where the best scenario does not follow the most popular route. The overall pattern in Figure 11(b) closely matches that of Figure 11(a). However, Figure 11(b) explores a significantly larger number of routes compared to Figure 11(a). Figure 11(c), which includes all scenarios and the best scenario, also shows that the best scenario mostly follows darker (i.e., more frequently used) paths, as seen in Figure 11(b). Compared to Figure 11(b), Figure 11(c) contains several additional paths, such as near Balpop 3 and Parking. Figure 11(a) demonstrates that even a cost difference of only 10% results in a wide variety of scenarios. In some areas, the differences are minor, while in others, the outcomes diverge significantly.

5. Discussion

The integration of findings from the results offers profound insights into the behaviour and performance of the Slime Mould Model under various scenarios. A crucial starting point is that in section 4.1, the positive building variant generates significantly more solutions in the individual scenarios, while the negative building variant dominates in the combined scenarios. This suggests that the interaction between the building definition (source versus consumer) and the optimization strategy (individual versus combined) fundamentally influences the model's outcome.

The efficiency analysis in section 4.2 support this thought. Although the cost differences between the single (+) and (-) variants are relatively small, indicating comparable efficiency when a network is optimized in isolation, the large differences between the combined and individual variants are significant. The fact that the combined variant yields the most cost-effective results, shows that it is most favourable to only perform optimizations of this variant in the future. This efficiency is further supported by the findings from section 4.3, which indicate that the distribution of the combined variant exhibits an increasing and then decreasing pattern. While the two individual variants also show an increasing and then decreasing pattern, they subsequently peak. This suggests that relatively large pipe diameters are used, resulting in nodes and pipe sections with high flow rates, where the failure of these specific sections could lead to a large part of the network failing. However, during the model's creation, the failure risk of a 5DGHC network and how this could be resolved was not concretely examined. The combined variant suffers less from this due to fewer choices of large pipes.

Section 4.4 shows that for the different results, diverse network routes are proposed to connect the points. This even applies to the solutions within the top 10% of results. This suggests that there is no single unique, optimal network configuration that consistently yields the lowest costs. Instead, there appears to be a spectrum of highly cost-effective solutions that achieve their optimal price through different combinations of pipe segments with varying prices and flows. This phenomenon is further illustrated by the peaks and valleys and the wide spread of data points in the results.

Therefore, the term 'best result' can be better interpreted as a collection of robust, cost-efficient solutions, each with its own unique distribution of pipes and corresponding physical layout. This

offers designers considerable flexibility, as additional, non-financial criteria, such as constructability, minimal environmental disruption, or future expansion possibilities, can be decisive in choosing a definitive design from this set of nearly optimal solutions.

However, the comparison between Figure 11(a-c) reveals that an analysis purely based on the pattern is limited. The visual differences between these two types of figures appear small, which can be misleading. This is because these figures do not account for pipe diameters, which are crucial for the actual costs and functional properties of the network. As a result, the differences in the visual representations may seem small, while the actual, quantifiable differences, as visible in section 4.3, are much larger.

For finding the most suitable network, this implies, on the one hand, that there is no single unique 'best' network. On the other hand, it means that although the trend graph in figure 4.4 can provide useful information about the distribution of pipe segments per variant, it is not directly possible to 'draw' a specific pattern based on these results to create one of the best outcomes. From this, it can be concluded that the most practical approach is to look at, for example, the top 10% cheapest variants and then make a manual selection based on specific preferences or additional criteria. Nevertheless, in certain cases, the absolute cheapest option from the model can simply be chosen.

Although the developed model is based on the principles of slime mould growth, there are inherent differences with the biological *Physarum polycephalum*. The real slime mould continuously revises and improves its paths through dynamic internal flows, whereby efficient connections are reinforced and less optimal ones are retracted. This reflects an adaptive and continuous optimization process. In contrast, the computational model lacks this dynamic post-hoc revision of individual network paths within a single simulation. However, this difference is functionally bridged by the applied methodology. By using multiple, different starting points, the model generates a series of diverse network solutions. The ability to select the most optimal route from these generated variants mimics the collective "trial-and-error" of the natural slime mould. This process enables the model, through aggregation of results, to produce a robust set of network designs that minimize CapEx, thereby functionally approaching the efficiency of the actual slime mould.

5.1. Model limitations

Model’s limitations hinder the model’s ability to reach its full potential. First, the model does not generate alternative locations for building connections. In some cases, relocating a building connection (slightly) could be the most cost-effective solution, but the model does not suggest this. Second, the model does not analyse variable energy flows, but solely the peak demands of all buildings. Furthermore, it is assumed that the peak demands of all buildings occur simultaneously. Moreover, the model disregards the imbalance between winter energy demand and summer energy surplus in ATES wells. It is assumed that this is resolved at the building level through dry coolers or other measures. Lastly, elevation differences are not considered. While this is often not a significant issue in the Netherlands, it can cause complications in cities with more elevation differences.

It is important to acknowledge that the actual costs for constructing a heat network in practice often turn out to be higher than the calculations presented in Appendix C.1., which means the cost estimations of this model cannot be directly compared one to one with the reality of implementation. This mismatch is further supported by insights from Rick Molenaar of TechniPlan Adviseurs, who indicated that a source network typically uses a pressure drop significantly smaller

than 150 Pa/m. This lower pressure drop is necessary because the pipes must function as a type of open distributor, suggesting that the model has often selected pipes that are too small when designing the network, and consequently, the actual price will likely be even higher than currently predicted.

5.2. Future potential

5.2.1. Integration of Variable Installation Costs and Existing Infrastructure

The model should be extended to account for variable installation costs under different soil conditions (e.g., installation under a road is more expensive than under a sidewalk or a green area). Crucially, this involves integrating information from the Kadaster [27], which provides essential insight into the location of existing underground infrastructure like pipelines and sewage systems. While the model currently only uses impassable obstacles (as described in section 3), it could be advanced by introducing difficulty factors. For instance, a sidewalk could represent a factor of 1, a building a factor of 3 (a hypothetical value, based on the higher cost of trenchless drilling), and crossing a district heating pipeline a factor of 4 (due to complexity, costs, and necessary coordination with the operator, often resulting in an underpass at 2.5 meters depth). These factors could impact excavation costs or total pipeline costs, allowing for the exact quantification of specific passage costs for a given trajectory. The NEN7171-1, NEN7171-2, and NEN7244 standards are of further relevance in this context.

5.2.2. Potential of Existing District Heating Tunnels (TU Delft Case Study)

Specifically for the TU Delft case study, further investigation into the potential of existing district heating tunnels is warranted. TU Delft already possesses an extensive network of underground tunnels for its current district heating system, which could potentially be utilized for the new network. However, the suitability and potential conversion costs of these tunnels have not yet been evaluated. Furthermore, no consideration has been given to the obstruction these tunnels might pose if they cannot be utilized.

6. Conclusion

This research focused on the development and application of a biologically inspired model, based on the principles of slime mould growth, in the context of designing an optimal thermal network. Through simulation, it was investigated how this model can create an efficient pattern for district heating systems, in response to the stated research question: "Can a slime mould growth model be used for the design of a heat distribution network in combination with an ATES system in an urban environment?" The results demonstrate that the Slime Mould Model has indeed proven capable of generating a robust list of network designs, specifically aimed at minimizing the CapEx of the network.

References:

[1] W. J. Ripple *et al.*, “World Scientists’ Warning of a Climate Emergency 2022,” *Bioscience*, vol. 72, pp. 1149–1155, 2022, doi: 10.1093/biosci/biac083.

[2] A. P. Jansen, “Forecasting the space heating demand of Dutch households,” 2020.

[3] T. Fotiou, P. Fragkos, and E. Zisarou, “Decarbonising the EU Buildings|Model-Based Insights from European Countries,” *Climate*, vol. 12, 2024, doi: 10.3390/cli12060085.

[4] M. D. Jackson, G. Regnier, and I. Staffell, “Aquifer Thermal Energy Storage for low carbon heating and cooling in the United Kingdom: Current status and future prospects,” *Appl Energy*, vol. 376, 2024, doi: 10.1016/j.apenergy.2024.124096.

[5] “Network evolution: from early generations to today’s systems | Dyneo - AI Technologies for thermal networks,” *https://en.dyneo.ch/blog/levolution-des-reseaux-thermiques-des-premieres-generations-aux-systemes-avances-daujourdhui*, 2025.

[6] T. E. van den Akerboom, “Clustering buildings with ATES systems to improve Amsterdam’s Fifth Generation Heating and Cooling Network,” *https://repository.tudelft.nl/record/uuid:7bbaa640-b659-43c5-b031-a902376c91f6*, 2024.

[7] Prem, “Wko,” *https://duurzaammbo.nl/warntekoudeopslag-wko*, 2025.

[8] L.A. Medema, “Campus and Real Estate – Energy,” 2025.

[9] K. Marlein, “Three levels of district energy network design maturity,” *https://www.iqgeo.com/blog/three-levels-of-district-energy-network-design-maturity*, 2025.

[10] T. Lu, “The Vital Role of Underground Utilities in Modern Infrastructure,” *https://hastingsutilities.ca/the-vital-role-of-underground-utilities-in-modern-infrastructure/*, 2025.

[11] F. Jabr, “How brainless slime molds redefine intelligence,” *Nature*, 2012, doi: 10.1038/nature.2012.11811.

[12] “Ride the Slime Mold Express!,” *AAAS Articles DO Group*, 2021, doi: 10.1126/article.30953.

[13] “Introduction to the ‘Slime Molds,’” *https://ucmp.berkeley.edu/protista/slimemolds.html*, 2025.

[14] A. Boussard, A. Fessel, C. Oettmeier, L. Briard, H.-G. Döbereiner, and A. Dussutour, “Adaptive behaviour and learning in slime moulds: the role of oscillations,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 376, p. 20190757, 2021, doi: 10.1098/rstb.2019.0757.

[15] AskNature, “Brainless Slime Mold Creates Smart Networks — Biological Strategy — AskNature,” *https://asknature.org/strategy/cytoplasm-creates-most-efficient-routes/*, 2020.

[16] S. Savić and S. Grant, “Slime Mold and Network Imaginaries An Experimental Approach to Communication,” *Leonardo*, vol. 55, pp. 462–467, 2022, doi: 10.1162/leon_a_02248.

[17] A. Tero *et al.*, “Rules for biologically inspired adaptive network design,” *Science (1979)*, vol. 327, pp. 439–442, 2010, doi: 10.1126/science.1177894.

[18] D. Van der Kooij, “TechniPlan Adviseurs,” 2025.

[19] R. McNeel, “Grasshopper,” 2025.

[20] R. McNeel, “Rhino8,” 2025.

[21] H. Zhang, Y. Tao, and W. Zhu, “Global Path Planning of Unmanned Surface Vehicle Based on Improved A-Star Algorithm,” *Sensors*, vol. 23, 2023, doi: 10.3390/s23146647.

[22] “Euclidean Distance | Formula, Derivation & Solved Examples - GeeksforGeeks,” *https://www.geeksforgeeks.org/euclidean-distance/*, 2025.

[23] Wang S., “Graph Search Algorithms: A Practical Overview,” *https://www.puppygraph.com/blog/graph-search-algorithms*, 2025.

[24] “Chebyshev distance - Wikipedia,” *https://en.wikipedia.org/wiki/Chebyshev_distance*, 2025.

[25] P. E. Hart, N. J. Nilsson, and B. Raphael, “A Formal Basis for the Heuristic Determination of Minimum Cost Paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, pp. 100–107, 1968, doi: 10.1109/TSSC.1968.300136.

[26] “Bresenham’s Line Generation Algorithm - GeeksforGeeks,” *https://www.geeksforgeeks.org/bresenhams-line-generation-algorithm/*, 2024.

[27] “Oriëntatieverzoek kabels en leidingen,” *https://www.kadaster.nl/producten/woning/orientatieverzoek?gad_source=1&gad_campaignid=15791772891&gbraid=0AAAADqvNe9SDwPVeNumiQzGsUBo3WBXg&gclid=Cj0KCQjwjo7DBhCrARIsACWauSkgs60TbFLvS5y5otUYPYtwkrtNxc4cGePi3sRJrhfAOTeTyymMkV0aAiadEALw_wcB*, 2025.

[28] M. Dr. ir. Bloemendal, “Course: Building heating and cooling requirements,” 2025.

[29] “Tyleenslang HDPE 50 mm x 3,0 mm rol 100 m,” *https://pvcvoordeel.nl/tyleenslang-50mm-rol-100-meter/?gad_source=1&gad_campaignid=17182275280&gbraid=0AAAADsAsf0bITA95Cs2u_LBMuMoqbWbw&gclid=Cj0KCQjw2tHABhCiARIsANZzDWpzY2OBpRjgKBQMpV4MfRNBuvak7DX3rWDcHZ3-F0MgneVx42r1BuwaAuhFEALw_wcB*, 2025.

[30] “Drukbuis Tyleen PE 100,” *https://smoldersbv.nl/Winkel/tyleen/tyleen-slangen-buizen/drukbuis-tyleen-pe-100-6mtr/*, 2025.

[31] Shane, “DN, De en Φ uitgelegd: De verschillen begrijpen | MachineMFG,” *https://www.machinemfg.com/nl/dn-de-and-%CF%86/*, 2023.

[32] “NPS (Nominal Pipe Size) Guide: Dimensions, Charts & Specifications,” *https://www.engineeringtoolbox.com/nps-nominal-pipe-sizes-d_45.html*, 2025.

[33] Ing. H.C. Roel *et al.*, “ISSO-publicatie 39 Energiecentrale met warmte en koude opslag (WKO),” 2017. Accessed: May 02, 2025. [Online]. Available:

https://bouwzo.nl/reader/publicatie/isso-publicatie-39-energiecentrale-met-warmte-en-koude-opslag-wko/2017

[34] "Geothermal energy: proven way of sustainable energy production | Deltares," *https://www.deltares.nl/en/expertise/areas-of-expertise/energy-transition/aquifer-thermal-energy-storage*, 2018.

[35] S. Fabozzi, G. De Luca, and L. Vanoli, "Fourth generation district heating and cooling," in *Polygeneration Systems: Design, Processes and Technologies*, Elsevier, 2021, pp. 323–350. doi: 10.1016/B978-0-12-820625-6.00003-7.

[36] S. Boesten, W. Ivens, S. C. Dekker, and H. Eijdems, "5th generation district heating and cooling systems as a solution for renewable urban thermal energy supply," *Advances in Geosciences*, vol. 49, pp. 129–136, 2019, doi: 10.5194/adgeo-49-129-2019.

Appendix

A

A.1 Energy labels

Label			Small office (<5000m2)	Large Office (>5000 m2)	shop	primary school	sec. school / university	Sport facility	appart-ment	individual house
A										
	Space heating	[MJ/m2/y]	52.1	44.7	100.7	89.9	76.8	93.2	56.7	37.8
	Cooling	[MJ/m2/y]	23.7	40.8	25.8	58.3	35.1	22.3	14.4	22.2
	DHW	[MJ/m2/y]	5.0	5.0	5.0	5.0	5.0	45.0	51.1	47.5
	max heating power	[W/m2]	34.8	29.8	67.2	60.0	51.2	62.1	37.8	25.2
	max cooling power	[W/m2]	28.7	45.8	30.8	63.3	40.1	67.3	65.5	69.7
B										
	Space heating	[MJ/m2/y]	52.1	44.7	64.4	276.4	76.8	123.3	61.8	37.8
	Cooling	[MJ/m2/y]	23.7	40.8	58.3	40.8	35.1	13.8	13.4	22.2
	DHW	[MJ/m2/y]	5.0	5.0	5.0	5.0	5.0	45.0	51.1	47.5
	max heating power	[W/m2]	34.8	29.8	42.9	184.2	51.2	82.2	41.2	25.2
	max cooling power	[W/m2]	28.7	45.8	63.3	45.8	40.1	58.8	64.5	69.7
C										
	Space heating	[MJ/m2/y]	150.2	128.1	60.0	225.0	123.4	134.6	108.7	67.0
	Cooling	[MJ/m2/y]	28.7	32.3	112.9	54.2	36.4	18.8	12.8	17.2
	DHW	[MJ/m2/y]	5.0	5.0	5.0	5.0	5.0	45.0	74.5	74.2
	max heating power	[W/m2]	100.1	85.4	40.0	150.0	82.3	89.7	72.5	44.7
	max cooling power	[W/m2]	33.7	37.3	117.9	59.2	41.4	63.8	87.3	91.5
D										
	Space heating	[MJ/m2/y]	303.9	227.9	148.8	363.9	300.4	311.1	251.0	205.9
	Cooling	[MJ/m2/y]	33.4	33.0	19.3	0.0	0.0	0.0	6.8	9.0
	DHW	[MJ/m2/y]	5.0	5.0	5.0	5.0	5.0	45.0	75.9	76.3
	max heating power	[W/m2]	202.6	151.9	99.2	242.6	200.2	207.4	167.3	137.2
	max cooling power	[W/m2]	38.4	38.0	24.3	0.0	0.0	0.0	0.0	0.0
E										
	Space heating	[MJ/m2/y]	370.3	364.0	174.0	360.7	281.9	363.6	277.4	188.7
	Cooling	[MJ/m2/y]	28.9	33.9	20.7	0.0	0.0	0.0	0.0	0.0
	DHW	[MJ/m2/y]	5.0	5.0	5.0	5.0	5.0	45.0	38.0	55.8
	max heating power	[W/m2]	246.9	242.7	116.0	240.5	188.0	242.4	185.0	125.8
	max cooling power	[W/m2]	33.9	38.9	25.7	0.0	0.0	0.0	0.0	0.0
F										
	Space heating	[MJ/m2/y]	364.0	284.9	195.8	431.6	330.3	387.3	526.1	344.4
	Cooling	[MJ/m2/y]	28.3	29.7	20.8	0.0	0.0	0.0	0.0	0.0
	DHW	[MJ/m2/y]	5.0	5.0	5.0	5.0	5.0	45.0	43.8	62.5
	max heating power	[W/m2]	242.7	189.9	130.5	287.7	220.2	258.2	350.7	229.6
	max cooling power	[W/m2]	33.3	34.7	25.8	0.0	0.0	0.0	0.0	0.0
G										
	Space heating	[MJ/m2/y]	367.2	104.6	66.6	134.1	107.1	125.9	277.4	163.9
	Cooling	[MJ/m2/y]	27.0	12.4	5.0	0.0	0.0	0.0	0.0	0.0
	DHW	[MJ/m2/y]	5.0	1.4	1.4	1.4	1.4	12.5	15.1	23.1
	max heating power	[W/m2]	244.8	251.1	159.8	321.7	257.1	302.1	665.8	393.3
	max cooling power	[W/m2]	32.0	49.6	23.0	0.0	0.0	0.0	0.0	0.0

Figure 12: An overview of an estimated energy demand for a building, given its specific function and energy label [28].

B

B.1 Data Acquisition and Processing

The data for the TU Delft campus was provided in the form of a DWG file. To read this DWG file and make it usable for Section 2.3.4, several steps are required. After the DWG file is imported into an active Rhino document, the geometric objects (such as lines and polylines) are extracted and processed.

The processing involves systematically retrieving all layers from the recently imported DWG file via the function layers = sc.doc.Layers. The code then iterates over each of these layers (for layer in layers:) to identify the objects within them. For each layer, all present objects are retrieved using objects = rs.ObjectsByLayer(layer.FullPath). A crucial step is checking whether an object is a curve and, more importantly, if it is a polyline, which is verified with if curve and rs.IsPolyline(curve):. This is essential because polylines, as linear elements with vertices, represent the desired geometry for the obstacles. If an object turns out to be a polyline, the function polyline = rs.PolylineVertices(curve) retrieves all individual corner points (vertices) of that polyline. Each vertex is a 3D point with X, Y, and Z coordinates, where the Z-coordinates are consistently set to 0, as the DWG file represents a 2D map.

Subsequently, all these processed points are saved in a TXT file, structured in the following format:

===layername1===
Figuur1: x1,y1,z1 x2,y2,z2 …. xn,yn,zn
FiguurN: x1,y1,z1 x2,y2,z2 …. xn,yn,zn
===layernameN===
Figuur1: x1,y1,z1 x2,y2,z2 …. xn,yn,zn
FiguurN: x1,y1,z1 x2,y2,z2 …. xn,yn,zn

In addition to the TXT file, a CSV file is also generated. This CSV file makes it possible to select only the desired layers for further processing. Furthermore, each layer can be assigned a weighted 'weight factor'; for now, these are all set to 9999. This functionality is essential for filtering the data. As can be seen in Figure 13(b), a large part of the DWG file is excluded by this function, in contrast to the unfiltered imported DWG file visible in Figure 13(a). Table 5 shows the layer selection menu, where the 'weight factor' can be set and layers can be made visible or invisible.

Table 5: Layer selection menu, for assigning a weight factor and making layers visible or invisible

layername	Weight	visibility
Default	9999	No
geoBGT_scheiding-hek	9999	No
geoBGT_spoor-tram	9999	Yes
geo_kunstwerkdeel-duiker	9999	No
geoBGT_straatnamen	9999	No

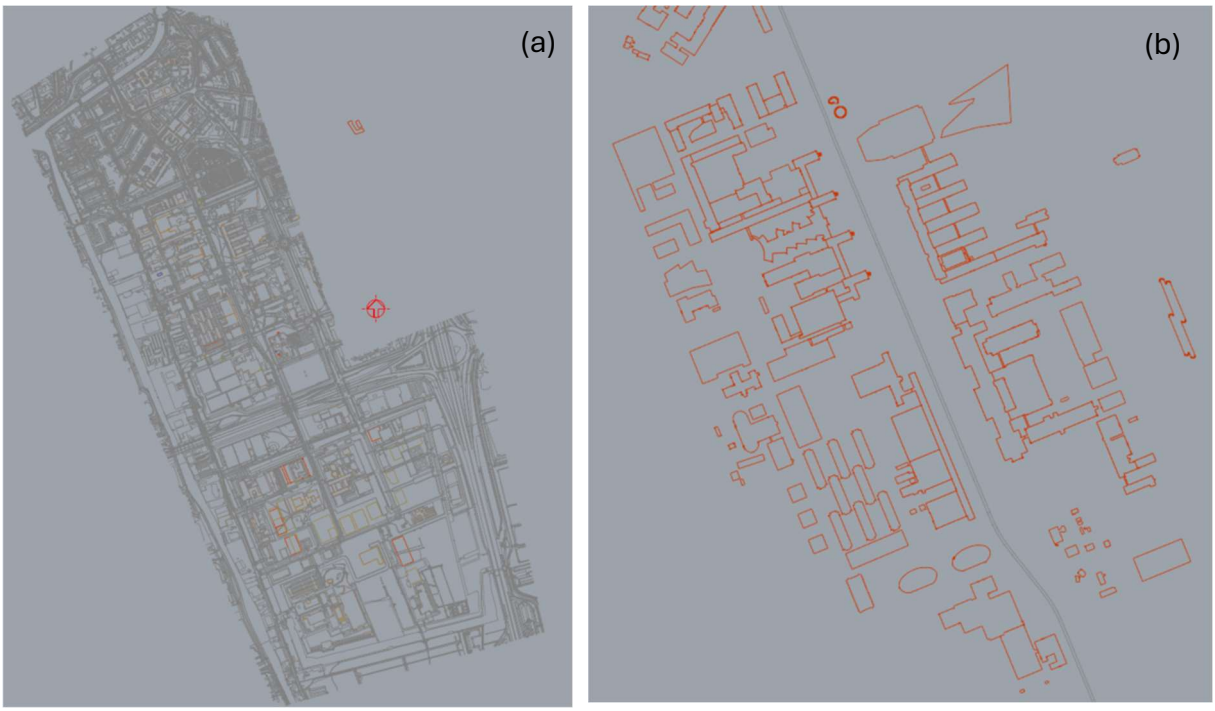

Figure 13: (a): displays the imported DWG file in rhino, (b): highlights the selected layers

C

C.1 Heat distribution network costs

The installation of a district heating network involves various costs, as shown in the table below.

Table 6: A price overview of all the costs for building a 4th generation heat distribution network.

Description	Quantity	Unit	Unit Cost	Total Cost
Engineering and Preparation of Route				
Preliminary research/test trenches/clash detection	1.0	Euro	€ 2,640.00	€ 2,640.00
Engineering	1.0	Euro	€ 4,400.00	€ 4,400.00
Preparation	1.0	Euro	€ 1,200.00	€ 1,200.00
Supply				
Pipes and accessories	1.0	Euro	€ 11,933.00	€ 11,933.00
Bends and appendages	1.0	Euro	€ 7,046.00	€ 7,046.00
Labour				
Digging and closing trench for heat network installation	115.0	m1	€ 40.49	€ 4,656.35
Installing heat network	230.0	m1	€ 6.28	€ 1,444.40
Removing and temporarily restoring pavement	230.0	m2	€ 21.75	€ 5,002.50
Digging and closing work pit for routing	2.0	Euro	€ 809.75	€ 1,619.50
Creating opening for pipe entries	4.0	Euro	€ 175.00	€ 700.00
Installing Roxtec or equivalent for sealing	4.0	Euro	€ 363.00	€ 1,452.00
Installing, insulating and welding pipes	115.0	m1	€ 289.00	€ 33,235.00
Inspection of welded joints	2.0	times	€ 1,275.00	€ 2,550.00
Supervision				
Site supervisor	16.0	hours	€ 80.00	€ 1,280.00
Work preparer	12.0	hours	€ 75.00	€ 900.00
Project manager	4.0	hours	€ 110.00	€ 440.00
			Subtotal:	**€ 80,498.75**

Based on the cost breakdown presented in Table 4, several assumptions can be made to simplify the integration of pipe diameter (DN) into the financial model.

For the entire 'Labour' section, excluding the line item for 'installing, insulating, and welding pipes,' it is assumed that the cost per meter is largely independent of pipe diameter. For example, trench excavation work is not significantly more labour-intensive for a DN200 pipe compared to a DN20 pipe, as the trench must be excavated regardless of the specific pipe diameter. These activities are therefore treated as having a constant cost per linear meter.

Similarly, the 'Engineering and Route Preparation' costs are considered linear per meter and are assumed not to vary with DN size. These include activities such as initial surveys and design work, which are typically influenced by total project length rather than pipe dimensions.

In the case of 'Supervision,' a dependency on pipe size may exist; larger diameters generally require more attention and planning. However, this relationship is expected to follow a linear cost model.

Since 'Placement Costs' are likely to dominate small DN size networks, these costs are also treated as diameter-independent for simplicity.

For the categories 'Supply' and the labour item 'installing, insulating, and welding pipes,' a strong correlation with pipe diameter is expected. These costs scale with material usage and complexity, which increase with larger DN sizes. As such, these items are modelled as pipe diameter-dependent costs and will be dynamically adjusted in the financial model based on the selected DN. This is the 'Material Cost' in the formula. Furthermore, a factor will be used since the welding costs, as well as bends etc., are strongly related to the DN size of the pipe.

Price Sizing Formula

For the installation of 230 meters of piping over a distance of 115 meters (supply and return), the total cost can be divided into fixed and variable components. The fixed costs, in this case the 'Placement Costs,' amount to €28,284.75 for the 115 meters, which equals €245.95 per meter.

The variable costs depend largely on the choice of pipe material. Nowadays, HDPE is a common material for district heating networks [8], as it is significantly more cost-effective than traditional options. The table assumes a DN50 pipe, which costs €52 per meter (or €104 per meter for both supply and return). If HDPE is used instead, the cost drops to only €3 per meter for the same DN size [29]. This price will account for the 'Material Cost'.

Larger DN sizes result in exponentially higher costs (see Table 7). Welding costs add only about 2% to the material costs, though this depends on the number of bends and branches required. For a €10/meter HDPE pipe, welding costs are roughly €1/meter, making the total €22 per meter for both supply and return. However, since bends and branches are not directly included in the cost, this will be accounted for in the factor, which is estimated to be 1.5, so 50% of the cost comes from welding and branches etc.

Since HDPE pipes are almost always used in ATES networks, these costs will be included in the model.

The cost formula is as follows:

$$g(n) = \left(C_m(Q) * F + C_p\right) * l \tag{15}$$

Filled in it will look like this:

$$g(n) = (C_m(Q) * 1{,}5 + 245{,}95) * l \tag{16}$$

Table 7: Data provided by [8, 30]

Flow rate	5,9	11,9	20,7	37,2	67	108,1	229,8	412	663,5	992	1405
DN	50	65	80	100	125	150	200	250	300	350	400
Price/m	3	5	11	14	16	26	41	90	140	180	250

C.2 Appendix C2. Pipe sizing in a heat distribution network

To ensure efficient hydraulic performance while minimizing pumping energy consumption, a Python-based model was developed to automate the selection of appropriate pipe diameters for low-temperature district heating networks. The model is designed to work exclusively with DN size, allowing for direct applicability in real-world systems.

Design Guidelines

In ATES-based heating networks, a maximum pressure loss of approximately 150 Pa/m is commonly used as a design guideline [8]. This value strikes a balance between minimizing pump energy consumption and avoiding oversizing of the pipe network.

Choosing a pipe diameter that is too large can increase the thermal inertia of the system and lead to higher heat losses, due to the increased surface area of the pipe. On the other hand, selecting a diameter that is too small can result in very high pressure losses, causing the required pump power to increase significantly. The model addresses this trade-off by automatically selecting the smallest pipe that keeps pressure losses within acceptable limits.

Input Parameters

The primary user-defined input to the model is the required thermal power output P (in watts). Based on this, the model calculates the necessary volumetric flow rate using the following relation:

$$\dot{V} = \frac{P}{C_p \Delta T \rho} \tag{17}$$

In this formula, V represents the volumetric flow rate in cubic meters per second ($\frac{m_3}{s}$), P stands for the desired power output in watts (W), Cp is the heat capacity of water (a constant value of $\frac{4186J}{kg}K$), ΔT denotes the temperature difference between supply and return (6.6 Kelvin), and ρ is the density of water (a constant value of $1000\frac{kg}{m_3}$). This temperature difference is predefined based on the design of the 5th-generation district heating system.

Calculation Methodology

For each available DN pipe size, the model performs the following steps: The computed flow rate is used to determine the average velocity in the pipe, using the internal cross-sectional area of each standard DN diameter.The Reynolds number is calculated using the standard formula:

$$Re = \frac{\rho v D}{\mu} \tag{18}$$

with $\mu = 0{,}000355\ Pa \cdot s$ as the dynamic viscosity of water at 15°C. This value indicates whether the flow is laminar or turbulent.

Friction factor calculation:

$$\begin{cases} f = \dfrac{64}{Re} & if\ Re < 2300 \\ f = \dfrac{0.25}{\log_{10}(\frac{\frac{e}{D}}{3.7} + \frac{5.74}{Re^{0.9}}))^2} & if\ Re > 2300 \end{cases} \tag{19}$$

where ϵ is the pipe wall roughness (0.045 mm)
The pressure loss per meter is calculated using the Darcy–Weisbach equation:

$$\Delta p = f \frac{L}{D} \frac{1}{2} \rho v^2 \tag{20}$$

The length L is set to 1 meter to obtain the pressure drop per meter of pipe.
For a given pressure loss limit in this 150 Pa. The model iterates through all DN sizes and selects the smallest pipe diameter that results in a pressure loss below the specified maximum threshold. This ensures compact system design and minimizes material use, while avoiding excessive pumping energy.

The result includes the selected DN size, the corresponding flow velocity, pressure loss per meter, and Reynolds number. This information supports hydraulic validation and energy efficiency assessments [31, 32].

D

D.1 ATES wells placement

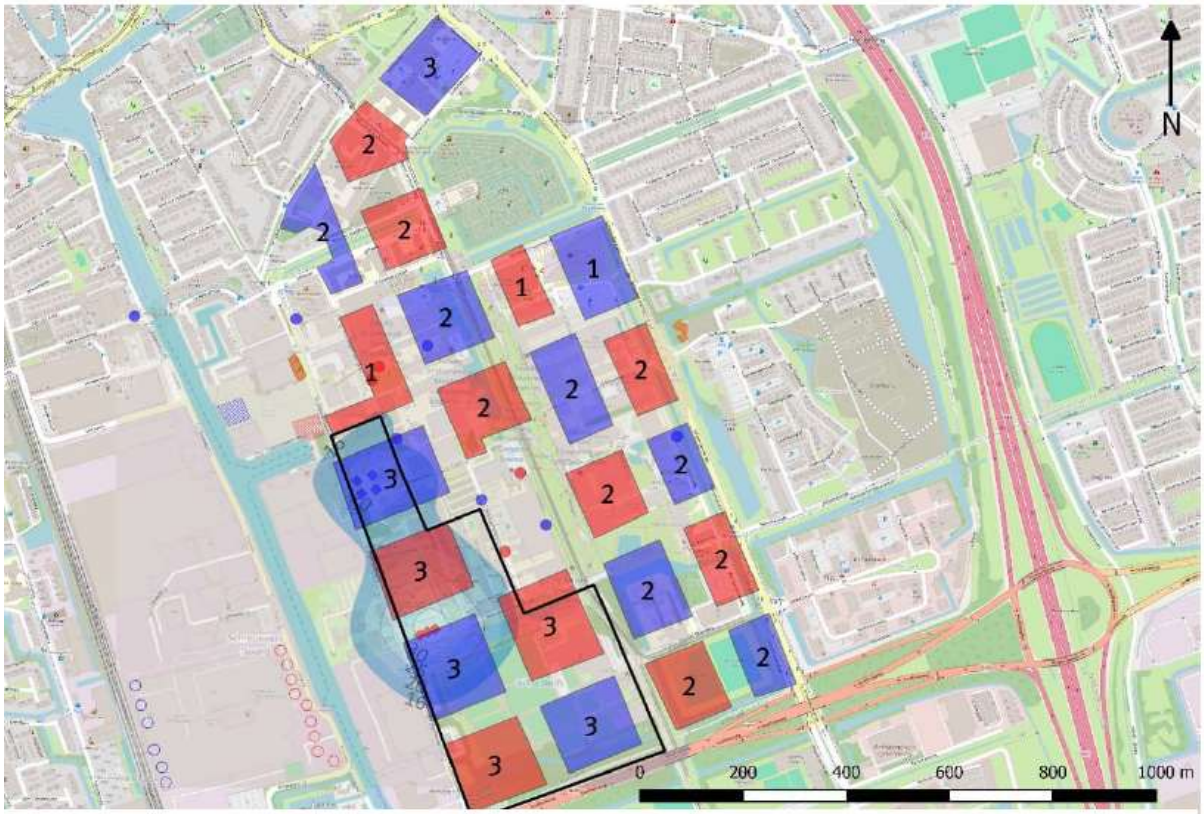

Figuur 5 Kader HTO gebied met het thermisch invloedsgebied tot 16°C uit de HTO effectenstudie van IF Technology.

Figure 14: Suggested locations of ATES sources by IF Technology[8].

E

E.1 Bresenham-inspired reason

The reason a modified Bresenham's algorithm [26] is used becomes evident when comparing its output to the standard version. When the standard Bresenham's algorithm is applied, it results in Figure 15(a), clearly showing the algorithm's tendency to make diagonal movements. Figure 15(b) presents the modified version, where the code now executes only horizontal and vertical steps.

The implications of these movement types for the slime mould's pathfinding are illustrated in the subsequent diagrams. In Figure 15(c), it can be observed that the slime mould, due to its ability to make diagonal movements, is able to successfully avoid the gray obstacles. Conversely, Figure 15(d) demonstrates that with only horizontal and vertical movements, avoiding the obstacle is no longer possible, forcing the slime mould to traverse through it.

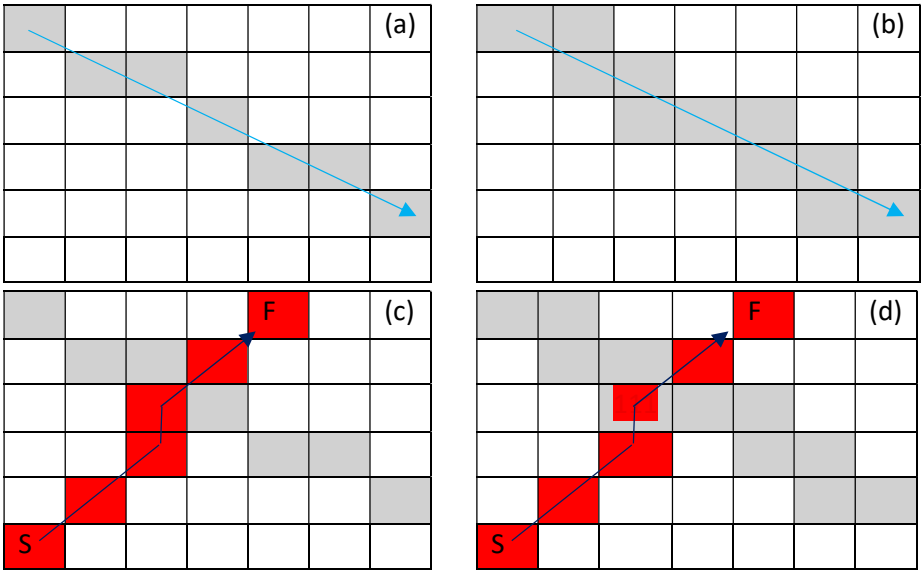

Figure 15:

E.2 Bresenham algorithm skipping an obstacle

Figure 16 provides an example where the slime mould can bypass an obstacle by utilizing option C, as discussed in the section 'Movement direction options for the slime mould'. This is exemplified by a movement from node point (0,0) to node point (2,1).

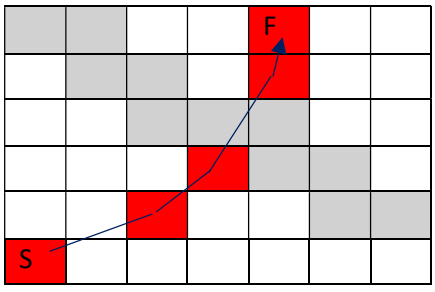

Figure 16:

F

F.1 Basic version for the verification

In Figure 17, a visualizations of the initial model variants were created to illustrate the model's behaviour.

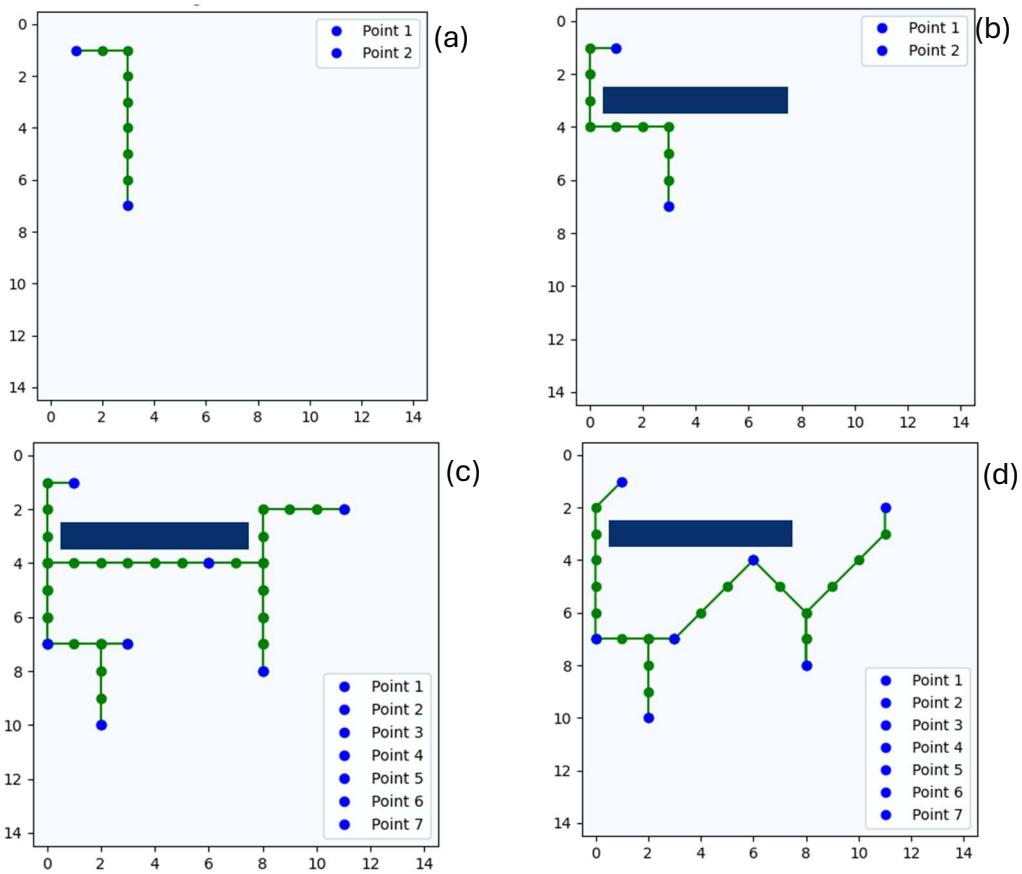

Figure 17: (a): slime mould that goes from point 1 to point 2, (b): slime mould that goes from point 1 to point 2, while skipping an obstacle, (c): slime mould that connects 7 different points, (d): slime mould that connect 7 different points, with horizontal, vertical and diagonal movements.

G

G.1 Jacket pipe placement

Two types of jacket pipes have been installed beneath the tramline on the campus. One consists of 3x160mm pipes, laid at a depth of 970mm, and the other consists of a single 500mm pipe, laid at a depth of 2400mm.

The 1x500mm variant has been installed in four locations under the tramline. The first is located approximately 40 meters north of the front side of Building X. See table below:

Table 8: The location of the 1x500mm jacket pipes respectively to X placed under the tramline

Point	1	2	3	4
distance	40	150	360	207

The 3x160mm variant has been installed in 16 locations. The first point is situated 63 meters north of the front side of Building X under the tramline, followed by additional points as shown in the table:

Table 9: The location of the 3x160mm jacket pipes respectively to X placed under the tramline

Point	1	2	3	4	5	6	7	8	9	10	11	11	12	13	14	15	16
distance	63	97	24	58	47	52	50	55	29	80	37	50	49	49	58	48	100

In Rhino 8, these coordinates are represented as follows:

Table 10: The location of the 1x500mm jacket pipes placed under the tramline in coordinates for Rhino8

Point	1	2	3	4
X	728	642	503	421
Y	298	420	753	947

Table 11: The location of the 3x160mm jacket pipes placed under the tramline in coordinates for Rhino8

Point	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
x	717	658	645	622	604	585	565	545	534	503	489	470	450	430	411	373
y	316	390	410	464	506	554	600	650	677	752	786	830	877	930	975	1065

Note**:** The locations of the jacket pipes are estimated based on data extracted from PDF files. The exact positions may vary.

H

H.1 Background: main components in a 5GDHC network

To fully grasp the functioning of a district heating network, it is essential to analyse several key components. While the complete Aquifer Thermal Energy Storage (ATES) system will be discussed at a general level, the focus will primarily be on the elements necessary for a fundamental understanding of the heat and cold wells.

An ATES system can broadly be divided into three main sections: the warm and cold wells, the heat distribution network, and the technical room. The technical room houses various components, such as heat pumps, heat exchangers, buffer tanks, regeneration systems, after heaters, and other auxiliary equipment. Although a detailed understanding of the technical room is not strictly required, a basic knowledge of heat pump operation is necessary to comprehend the principles behind the wells in an ATES system. Subsequently, the ATES system itself will be explained in more detail and additional extra information on the heat distribution network.

Heat pump

Heat pumps transfer thermal energy between a heat source (air, water, or ground (well)) and a heating system using refrigerant phase changes. A compressor increases the refrigerant's pressure and temperature, the compressed gas then releases heat to the heating circuit as it condenses. An expansion valve reduces pressure, causing the refrigerant to cool and absorb heat from the source. Heat pumps can provide both heating and cooling, requiring an external heat source or sink. Aquifer Thermal Energy Storage (ATES) systems can be more efficient heat sources than ambient air. The theoretical efficiency of a heat pump is calculated as:

$$COP_{theoretical} = \frac{T_{hot}}{T_{ho} - T_{cold}} \quad (21)$$

Where temperatures are in Kelvin. Actual COP is lower than the theoretical COP (approximately 62% for water-water and 55% for air-water heat pumps).

Heat Exchangers

Heat exchangers transfer heat between two media (e.g., water-to-water in ATES) without direct contact. This separation can prevent corrosive water from the ATES side from circulating through the system and allows for system isolation [7]. Heat exchangers have thermal losses related to temperature differences, effective heat exchangers typically have a 1 Kelvin temperature difference (ΔT). For example, if medium A is at 10°C and medium B is at 20°C, medium A can be heated to a maximum of 19°C, and medium B can be cooled to a minimum of 11°C [33].

ATES system

Aquifer Thermal Energy Storage (ATES) systems offer a sustainable approach to heating and cooling in buildings by utilizing underground aquifers to store thermal energy.

ATES System Operation

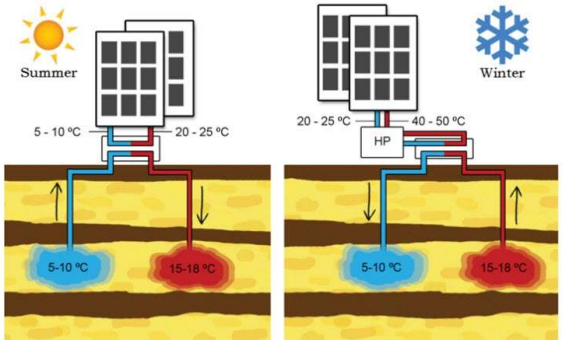

Figure 18: A simple overview of how an ATES system works [34].

Summer Cooling Cycle

During the summer, groundwater typically around 9°C, though this can vary by location can be used to cool buildings. The water is pumped up and flows through heat exchangers where it absorbs heat from the building. After passing through the exchanger, the water temperature rises to approximately 10°C, and continues to warm up to around 18°C as it cools the building. This heated water is then stored in the warm well at a temperature of about 17°C. The cooling process primarily requires energy for pumping, which accounts for roughly 2.5% of the system's total cooling and heating capacity[33].

Winter Heating Cycle

During the winter, water stored in the warm well initially around 17°C cools slightly over time but generally remains at approximately 14°C. This water, typically about 13°C after passing through a heat exchanger (0), is extracted and used as a heat source for a heat pump. The heat pump draws energy from the water, reducing its temperature to around 6°C on the evaporator side, after which the cooled water is pumped into the cold well. On the condenser side, the heat pump increases the temperature of the building's central heating water, for example from 38°C to 45°C.

Storage Capacity and Maximum Flow Rate

The storage capacity of an Aquifer Thermal Energy Storage (ATES) system is primarily determined by the properties of the aquifer. However, since the system will only consider the balance between supply and demand, storage capacity won't be a direct focus. It is noted that storage capacity can influence the number of wells that need to be connected to the heat distribution network.

The maximum achievable flow rate in an ATES system depends on the hydraulic conductivity of the aquifer and the well design. For this paper, the maximum flow rates were provided by Anne Medema [8]. Furthermore, to address insufficient energy supply during peak capacity, additional wells were identified using a report from IF Technology and guidance from Anne Medema [8] see appendix D. While the exact placement of these additional wells could have some uncertainty, it is acknowledged that for this case study, well placement can be adjusted as needed. The maximum flow rate of a single well is in Delft typically up to 50 m^3/h but can be higher (e.g., 75 m^3/h) for deeper wells [8].

Heat distribution network

A district heating system (DHS) is a centralized infrastructure designed to supply thermal energy for space heating and domestic hot water to multiple buildings from a shared heat source. Instead of individual gas boilers or electric heaters, buildings connected to the network receive hot water via insulated underground pipelines. The water is heated at a central plant typically using fossil fuels, biomass, waste incineration, or combined heat and power (CHP) plants and circulated to consumers [5].

The main components of a conventional district heating system include: a central heat production unit; one or multiple supply and return pipelines within the network; substations in individual buildings; and end-user heating systems (radiators, floor heating, etc.).

Heat is delivered at temperatures typically ranging between 70–200°C, the temperature depends on the generation [5, 35] and after the heat is extracted, the cooler return water is sent back to the central plant to be reheated.

Some 4th generation heat distribution networks have also 2 extra pipes for cooling as can be seen in Figure 19.

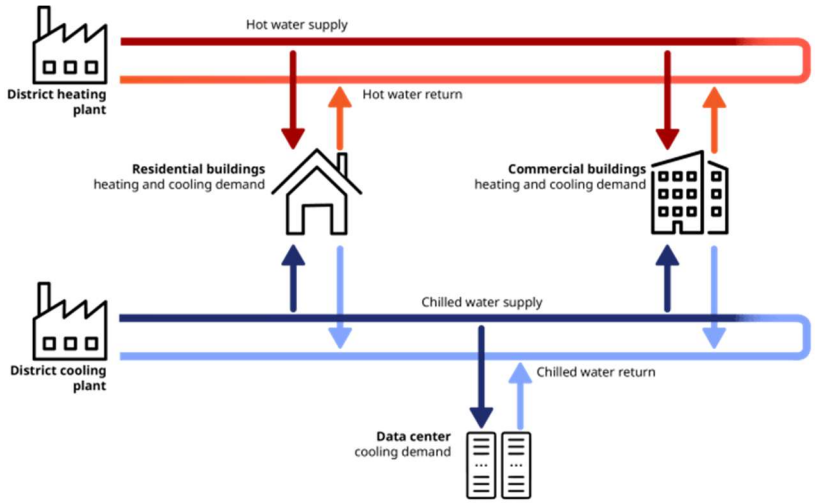

Figure 19:4th generation heat distribution network [4]

5GDHC systems

5GDHC systems (fifth generation district heating and cooling) mark a paradigm shift by operating at low, near-ambient temperatures (10–25°C) and relying on decentralized heat pumps within each building for heating. This design offers high flexibility, energy efficiency, and excellent potential for integrating renewable and residual heat sources [5, 35]. Key features include: an ambient temperature network where the water is not heated centrally but circulates at the natural temperature of the ATES wells; decentralized heat pumps, allowing buildings to extract heat via an individual heat pump or inject heat by directly cooling the building; bidirectional thermal flow, serving both heating and cooling demands through the same loop and enabling energy exchange between buildings; and integration with storage systems, such as seasonal storage using underground solutions like ATES.

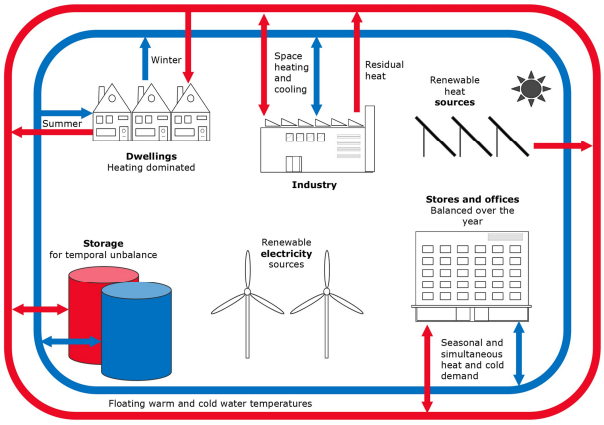

Figure 20: 5th generation heat distribution network [36]

In this paper, the focus will be on 5th generation district heating and cooling systems as seen in Figure 20. While it is technically possible to implement older generation DHS using the same code, such configurations will not be discussed further in this study.