TÉCNICO LISBOA

Numerical Modelling of Plastic Pyrolysis in a Solar Cavity Receiver Reactor

Anna Viviane Reeves

Thesis to obtain the Master of Science Degree in

Energy Engineering and Management

Supervisors: Prof. Carlos-David Perez-Segarra
Prof. Francisco Manuel da Silva Lemos

Examination Committee

Chairperson: Prof. Maria de Fátima Grilo da Costa Montemor
Supervisor: Prof. Carlos-David Perez-Segarra
Member of the Committee: Prof. Pedro Jorge Martins Coelho

December 2017

Dedicated to Hedwig Lang, who has always loved and supported me more than anybody else, hoping that I would have the opportunity to get a good education that she did not. May this poor soul finally find some peace now and be able to let go, knowing that I am getting there.

Further dedicated to San Barandiaran, for without his love and support and all the salads he made for us I would not have been well enough to do this thesis well.

Acknowledgments

The author would like to thank Professor Lemos and the research group on CSP at the CTTC in Terrassa for their advice on various issues, and the CTTC for providing a fast computer to run the simulation on.

Resumo

Encontrar novas formas de utilizar os resíduos de plásticos é importante para evitar que estes poluam o meio ambiente e para assegurar a sustentabilidade de recursos. A pirólise solar por absorção direta afigura-se como uma candidata promissora. Esta tese contém a primeira tentativa publicada de modelar numericamente este processo. Foi selecionado um reator simples, cilíndrico de leito fixo, iluminado através de uma janela por um refletor parabólico. Foi assumido que o catalisador seria utilizado após o passo de decomposição térmica, e portanto não foi incluído na simulação. Foram incluídos os três plásticos mais comuns com um modelo empírico para a sua decomposição e dados de calorimetria diferencial de varrimento (DSC) da literatura. Foram aleatoriamente atribuídas probabilidades à existência de aditivos, incluindo os pigmentos mais comuns, bem como filtros de UV e carbonato de cálcio, um enchimento comum. Dois produtos de decomposição foram introduzidos, óxido de cálcio e cobre. O espectro solar padrão de referência foi utilizado como radiação de entrada e a energia foi distribuída através de uma metodologia estocástica "Monte Carlo ray tracing". Os resultados obtidos são prometedores e recomenda-se a continuação da investigação. As perdas de radiação ficaram só entre 1% e 3.5% nos casos normais. Verificou-se que uma granulometria de 3-5mm era suficientemente pequena para que a transferência de calor entre o plástico colorido e transparente fosse suficientemente eficiente. No entanto as limitações computacionais implicaram que os reatores simulados fossem de dimensões bastante reduzidas e é possível que aparecen fenómenos adicionais em reatores de dimensões mais realistas.

Palavras-chave: pirólise solar, reator solar, pirólise de plásticos, utilização de resíduos de plásticos

Abstract

Finding more economical ways of utilising waste plastic is important to keep it from polluting our environment and for resource sustainability. Solar pyrolysis by direct absorption seems like a promising candidate as operational costs should be low. This thesis details the first ever published attempt to numerically model this process.

The literature on plastic pyrolysis was reviewed, and a simple cylindrical fixed bed reactor chosen, illuminated through a window via a parabolic dish. The catalyst was assumed to be placed after the thermal decomposition step, for better contact and to avoid poisoning, hence it was not part of the simulation. The three most common plastics were included, with an empirical model for their thermal decomposition and DSC data from the literature. Additives were randomly assigned with certain probabilities, comprising the most common pigments as well as a UV absorber and calcium carbonate filler. Two additional decomposition products were included, calcium oxide and copper.

The standard solar reference spectrum was used as input and the energy distributed via Monte Carlo ray tracing.

The results look promising and thus further research is recommended. Losses from escaping rays are only between 1% and 3.5% for normal feedstock. A practically feasible granule size of 3-5mm is small enough for the heat transfer from coloured to transparent plastic to be efficient enough. However, limitations on computing power mean that the simulated reactors are quite small, and additional phenomena might appear in realistically sized reactors.

Keywords: Solar Pyrolysis, Solar Reactors, Plastic Pyrolysis, utilisation of plastic waste

Contents

List of Tables

List of Figures

Nomenclature

ABS	Acrylonitrile Butadiene Styrene
DNI	Direct Normal Irradiation
DSC	Differential Scanning Calorimetry
ETP	Expanded thermoplastic
FSR	Forward Scattering Ratio
GHI	Global Horizontal Irradiation
HDPE	High Density Polyethylene
HIPS	High Impact Polystyrene
LDPE	Low Density Polyethylene
LLDPE	Linear Low Density Polyethylene
MSW	Municipal Solid Waste
NaN	Not a Number
PA	Polyamide
PC	Polycarbonate
PLA	Polylactic Acid
PP	Polypropylene
PS	Polystyrene
PUR	Polyurethane
SAN	Styrene Acrylonitrile
TGA	Thermogravimetric Analysis
WEEE	Waste Electrical and Electronic Equipment

Glossary

Additive	Materials added to a pure plastic polymer. Sometimes this term is meant to include pigments and sometimes not, it should be clear from the context though.
Argon Bubble	A lattice element containing only argon gas
Dye	A material that imparts colour to another, and dissolves in it
Element	Entity of the simulation, either a plastic piece, argon bubble or sheet
Gasification	Thermal decomposition in a stoichiometrically oxygen deficient atmosphere
Lattice Constant	Distance between the centres of each cubic lattice volume, or length of the sides of the cubes
Lattice Element	Either a plastic piece or an argon bubble
Participating Element	Either a plastic piece or a sheet
Pigment	A material that imparts colour to another, and is not soluble in it, hence will consist of particles suspended within it. The best quality pigments, that have the best opacity and hiding power, have particle sizes similar to the wavelengths of visible light, so that they scatter it well.
Plastic Piece	Cubic part of a granule that behaves like an independent entity
Plastic	Commercial plastic, pure plastic plus additives
Pure Plastic	The pure plastic part of a plastic that may contain additives
Pyrolysis	Thermal decomposition in the absence of oxygen

Rugate Filter	Optical filter based on a dielectric coating, where the refractive index is varied continuously (rather than in steps). They can be tailored to reflect all radiation within a certain wavelength range.
Solar Thermochemical Reactor	A reactor in which an endothermic reaction takes place, driven only by the incoming solar energy

Chapter 1

Introduction

Polymers like cellulose and rubber have been part of nature for a long time. They are made up of long macromolecular chains of a single unit molecule, with carbon atoms forming the backbone. However with the advent of the industrial revolution, man began to create his own artificial polymers, called plastics. As he learned to tailor them to his needs, they soon became so useful that their use mushroomed after the Second World War, and they have become an ubiquitous and indispensable part of our everyday life. The long macromolecular chains make them lightweight, yet tough and durable, even when moulded into a thin sheet. This together with their impermeability to water makes them an ideal material for packaging and containers. By keeping food and other products fresh, they are thus helping humanity to save resources. Use of lightweight packaging material and plastic parts in automobiles also means that these vehicles have to use less fuel and thus emit less greenhouse gases. So while efforts to use less plastic shopping bags are definitely commendable, they only present a small fraction of the total, and a modern world completely without plastics is not conceivable and also not desirable any more.

Biodegradable alternatives like PLA and other bioplastics need to be developed urgently, and most of the plastic we wear could be replaced by hemp and wool clothing. However, this will take time, because normal plastic is too cheap and well established, and both legislation and consumer awareness need to change.

In the meantime, it is to be expected that global plastic production will continue to increase significantly, as a look at the trend from the past in Figure 1.1 confirms. While production in developed countries has already stabilised, developing countries currently use much less per capita. Assuming no major disasters like nuclear wars, their increasing demand will continue to be the main driver of the increase in global plastic production for some time to come. So, for the medium-term, we are faced with the issue of a large amount of already produced and yet to be produced synthetic plastic waste which does not readily biodegrade.

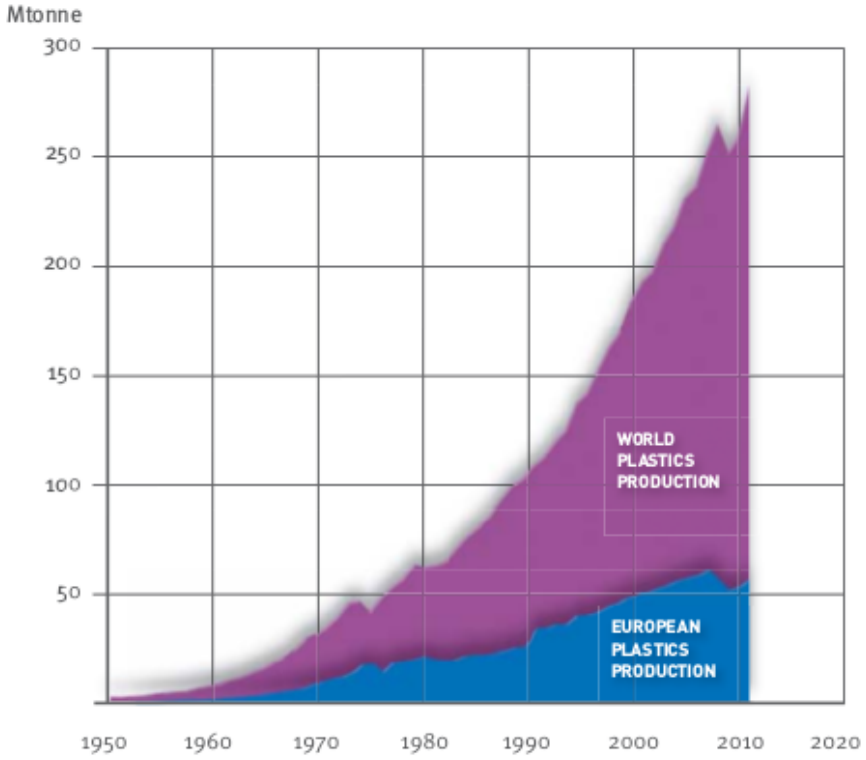

Figure 1.1: Global plastic production over the years [1]

1.1 Motivation

The main problem posed by this is that simply sticking large amounts of plastic waste into some far away landfill and forgetting about it is not such a good idea, even if you happen to live in a country like the US that has enough space available. Besides being a somewhat unsustainable way of dealing with earth's limited resources, plastics also contain a wide array of chemicals, and most of them eventually leach out as they are not bound to the polymer matrix. Many of these contain toxic heavy metals like lead and cadmium, or chemicals like bisphenol A and phthalates that act as endocrine disruptors for animals [2]. Even if there is no groundwater around to contaminate, one can never know if some of those chemicals will not eventually somehow find their way around and enter a food chain. Maybe not in our generation, but that does not mean that it matters less. Also, many of these chemicals have never even been tested for harmful health effects, and in many cases, the nature of the additives is a trade secret of the manufacturer or processor [2]. It seems somewhat irresponsible to be dumping materials containing synthetic chemicals into nature when regulatory authorities do not even know what all those chemicals are and how they could affect all the various life forms that exist on this planet.

And even this is assuming that all countries at least make the effort of bringing all their plastic waste to a far away landfill site and covering it over well. The current reality is that most developing countries simply bring most of their mixed waste to open dumps (see Figure 1.2), which are often very close to cities, where poor people scavenge through them and animals eat from it. In India, there is a whole caste of waste pickers who live on top of the dumps, where they are constantly exposed to the methane emissions and toxins from the fires.

And of course, many people simply throw their rubbish into nature, including the rivers and oceans. In

(a) Deonar dump site in Mumbai. Copyright Dhiraj Singh/Bloomberg

(b) Burning rubbish with children and buffaloes around. Copyright Sriram Kannan

Figure 1.2: Waste management, or rather the lack of such, in India

fact, a whopping eight million tonnes of plastic is estimated to enter our oceans every year [3], where it slowly photodegrades, i.e. breaking up into smaller particles, and releasing some of its toxic additives while soaking up other toxins. There is also evidence that polystyrene actually does degrade in an oceanic environment [4], which is not good since styrene is pretty toxic and many of its additives too. The chemicals cause birth defects and reproductive issues in marine animals, while some, especially birds and turtles, die from eating plastic that they mistake for jellyfish or other types of food, or by getting entangled in it (see Figure 1.3). 44% of all seabirds are known to eat plastic at times [4], and pictures of dead birds found with stomachs full of plastic items have become iconic of their plight, see for example Figure 1.3. Estimates for the amount of deaths are in the millions [5], although it is hard to estimate of course, since most of the plastic sinks to the ocean floor and the ocean depths are still very much unexplored. Our plastic waste might end up wiping out many deep sea species before we even get to know about their existence.

Might the problem improve by itself if developing countries reach a similar level of development and

(a) Albatross Chick that died from being fed plastic. Copyright Chris Jordan

(b) Turtle that got stuck in six pack ring. Copyright Missouri Department of Conservation

Figure 1.3: Examples of marine animals impacted by plastic pollution

environmental awareness as the currently developed nations do? Unfortunately not, unless new technologies emerge that they could adopt. A quick look at Figure 1.4 reveals that even developed countries

struggle to implement plastic recycling, even if they subsidise it like in the EU. The majority still ends up in landfill, while some is incinerated at a low overall efficiency.

Given the immensity of the problem, and the lack of environmental awareness among many citizens in

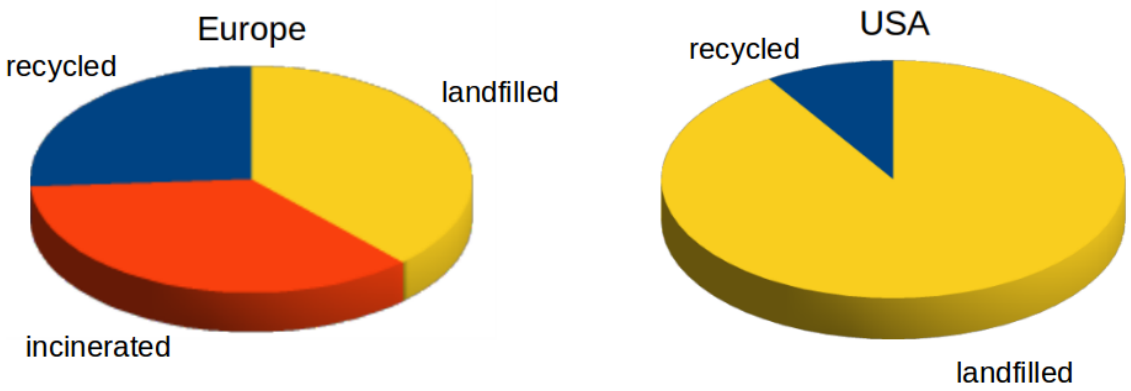

Figure 1.4: Fate of plastic waste in some developed countries [6]

developing countries, it is clear that the only medium term solution to this medium term problem which could have a significant enough impact would be a reasonably profitable business model to utilize most of this waste plastic. Then, people in developed countries could be made to throw their waste plastic into separate bins, while developed countries have enough poor people who would collect and separate waste plastic for a modest fee.

The potential is definitely there, since waste plastic actually still contains a lot of chemical energy [7], as can be seen from Table 1.1. So why are we burning so much coal, ethanol and firewood all over

Fuel	Heating value in MJ/kg
Petrol	45
Diesel	45
Waste Plastic	32
Coal	10 - 25
Ethanol	30
Methanol	22.7
Dry Firewood	16

Table 1.1: Heating values of various common fuels and waste plastic [8]

the world, while polluting our environment with waste plastic, which actually has a higher heating value than those three? The answer of course is that direct burning of plastic creates a lot of pollutant gases, so incinerators that burn plastic along with MSW need to spend more effort and energy on cleaning the flue gas, which offsets the energy gain from the increase in heating value. In fact, burning waste always leads to some polluting gases, so incinerators in countries with stringent environmental regulations like in Europe have to spend so much of their produced energy on cleaning the gas that their overall energy efficiency is only around 15% [9]. The Dutch incinerators are among the most efficient, with 21% [10]. That along with being costly to construct means that they actually end up making a loss despite selling

electricity to the grid, and communities end up paying around three times higher gate fees to an incinerator to take their waste than to a landfill site. [11] Hence, they are mainly used in countries that do not have much space for landfills due to their high population density, like Japan and South Korea. Given the high costs and large amount of public opposition to incinerators, it seems highly unlikely that they will spread fast enough to significantly impact the problem of plastic pollution, especially not in developing countries that can hardly afford them. Besides, it does not seem too desirable from an environmental standpoint anyway, because all the carbon in the plastic chains is converted into carbon dioxide which contributes to global warming, for a meagre return of 15% of its energy content. Of course, if the heat is used for district heating or as process heat for something else then the overall energy efficiency is higher, but many people do not like the idea of having a waste incinerator nearby and heating is not always needed in every part of the world during every season. Maybe such a conversion is not completely unavoidable, but then it should ideally be as efficient as possible. Also, once a community has invested into building an incinerator, it needs to feed it with enough waste every day for many years, even if waste generation fluctuates. This can decrease the incentive to recycle and to turn the organics into biogas and fertiliser.

What other utilisation methods currently exist for plastic waste? There is of course recycling, which is being practiced, though mostly in developed countries and mostly just for PET which is the easiest to recycle. Unfortunately though, recycling plastic is a work and energy intensive process. Since different types of plastic melt at different temperatures, they all have to be separated from each other, which is not a straightforward task. If different types of plastic are blended together, then recycling is not possible at all. Also, the plastics have to be sorted not just by type but also by colour, and then the resultant colour might not have exactly the desired shade. Certain additives can also interact with each other and with the pigments in not so desirable ways.

Then, the heat treatment causes some damage to the polymer chains, meaning that the recycled plastic is only usable in not too demanding applications and does not fetch a high price on the market. This together with the high costs means that recycling facilities usually do not make a profit, and the only reason that they can exist at all here in Europe is because consumers are indirectly subsidising them. Under the European directive on packaging and packaging waste, the packaging producers are responsible for the recycling of their plastic packaging. They organise this by paying a fee per tonne they produce into a fund that subsidises the recyclers so that they can break even. Of course, the added cost is passed on to consumers, as this is how capitalism works and should work. [10]

It also means though that plastic is not recyclable indefinitely, since eventually the quality just becomes too low. So, while recycling can help to mitigate the problem of plastic waste to a certain extent, and should be encouraged for the cases where it is easiest, it cannot by its very nature be the only solution. There will always be the issue of what to do with the plastic that has reached the limit for the number of times that it can be recycled. And there will always be demand for virgin plastics, because there are too many important applications for which they have to be tailor made from scratch to give the desired performance. Who would buy a car made with recycled plastic?

Research is ongoing about mixing waste plastic with cement for making roads, and a 2km road in Ban-

galore has already been built with this mix and is in use. These kind of roads are actually superior to normal ones, lasting more than twice as long, being able to withstand higher loads and being impermeable to water [12]. However, this also merely puts off the problem, albeit into an even more distant future. Concrete cannot be recycled indefinitely, and the toxic additives will still be contained in it and slowly leach out.

Therefore, clearly what should be achieved somehow would be to break apart the carbon chains and separate the fragments from the additives. This clearly calls for some kind of thermal treatment, however preferably not with too much oxygen and high temperatures as in combustion, as this is what causes most of the noxious gases like dioxins to form. One such option is gasification, which is a form of incomplete combustion in an oxygen-reduced environment. However, gasification plants are complex and require a high investment cost, so the economic feasibility depends on the local circumstances [7]. While it is definitely a potential possibility for plastic waste, this type of feedstock lends itself much better to pyrolysis than to gasification, due to its high volatile content. Gasification makes more sense for upgrading a low quality feedstock like petcoke or some general waste [13], but less for plastic which readily releases its valuable oils and gases in an inert atmosphere at only around $400°$C to $500°$C. It would seem somewhat wasteful to heat it to $1100°$C and break apart all of the long chains completely, only to then bring the carbon atoms together again in a later synthesis step to form some organic molecule. The syngas also needs to be cleaned [14] and requires careful handling lest it explode, and, unless burned directly, the conversion step requires extra effort and energy. This together with feed pretreatment means operational costs are high [15], and investment costs are also high since the plant is complex. True, the advantage is of course that any desired chemical can be custom made, and the ones with the currently highest market price can be chosen. With pyrolysis, a whole spectrum of different chemicals is obtained, which need to be separated from each other and then individually marketed. Nevertheless, the plant and the whole process is much simpler, and the type of catalyst can be freely chosen to tailor the product spectrum to a certain degree. The complete absence of oxygen also means that apart from the pollutants already in the plastic, additional ones like dioxins and nitrous oxide cannot be formed at all. Given the global nature of today's markets, finding customers for the many valuable petrochemicals and fuels obtained from plastic pyrolysis should not be a problem. More of a problem should be the separation of these valuable chemicals, since many have boiling points similar to other ones, hence costly solvent extraction often has to be used [14]. The oil can be separated into fractions that can be refined into ultra-low sulphur petrol and diesel, and if needed the petrochemicals shown in Table 1.2 can be extracted and sold separately. The gas has a high calorific value between 42 and 50 MJ/kg [16] and could be burned directly, although it could also be separated. Methane and ethane could be added to the natural gas grid, whereas the propane and butane could be liquefied under pressure and added to LPG fuel containers widely used in developing countries for cooking purposes. Table 1.3 lists the other petrochemicals that could be obtained from the gas.

So, overall pyrolysis seems to be a better option than gasification. It also has the advantage over recycling of not requiring as much pre-sorting. Only certain types like PET and PVC need to be removed as they are unsuitable, but the majority of plastics can simply be pyrolysed together and a small amount

Benzene	Precursor to various chemicals and octane enhancer
Toluene	Solvent, precursor to various chemicals, octane enhancer
Xylene	Solvent, precursor to various chemicals
Naphthalene	Precursor for various chemicals
Indene	Production of indene/coumarone thermoplastic resins
Fluorene	Precursor for pharmaceuticals
Biphenyl	Heat transfer agent and preservative
Styrene	Production of polystyrene and ABS

Table 1.2: Main petrochemicals present in pyrolysis oil and their most common uses [16]

Hydrogen	Use in fuel cells and production of ammonia fertiliser
Ethene	Production of many chemicals and PE plastic
Propene	Precursor to many chemicals, production of PP plastic
Butene	Precursor to chemicals and polymers
Butadiene	Production of synthetic rubber and ABS plastic

Table 1.3: Main petrochemicals present in pyrolysis gas and their most common uses [16] [14]

of contamination by the other types is not a problem. The presence of additives or pigments is not a big problem either, since most of the organic ones decompose too and most of the inorganic ones remain solid and get left behind.

One could thus imagine the two methods coexisting in harmony, with the recyclers picking out the good quality PET, PVC and HDPE and leaving the rest for pyrolysis. This might make more economic sense than trying to force the recyclers to recycle as much as they can. In Germany for example, where mixed plastics are recycled too, the fee paid by the packaging producers was 1430 €/t in 2007, whereas in Belgium and France, where only plastic bottles and flasks are recycled, it was just 220 €/t and 245 €/t in 2010, respectively. [10]

So, why have most people not heard about plastic pyrolysis yet? The problem here is that classical plastic pyrolysis is not very economical due to the high energy input required to break so many carbon bonds, which leads to low profit margins if that energy is purchased. Profitability obviously fluctuates along with the global oil price, which is rather unpredictable as it depends on the politics of many countries. That can make planning ahead for a new plant risky. Several such plants have ended up shutting down again after some time, although in recent times some have emerged that seem to be able to make do by burning the off-gas produced in the process to provide the process heat, while using small amounts of natural gas to initiate a cold start. [17] [18] [19]

They claim to do their own on-site refining so that they can sell ready to use oil. Although they do not mention it and pride themselves on how green they are, this must entail the purchase of hydrogen from the market, which is an extra expense and causes additional extraction of fossil fuel elsewhere, since most commercial hydrogen is currently produced from steam-reforming of natural gas or coal. The same goes for hydrocracking of plastic, which is similar to pyrolysis but with extra hydrogen added from the beginning, rather than afterwards.

This is where the main advantage of utilising solar input energy would lie. Firstly, because some of the

gases are quite valuable as petrochemicals, see Table 1.3. And secondly because after separating out the most valuable constituents, the remainder could then be steam-reformed into pure hydrogen which could be used for the hydrogenation of the oil, which is a necessary step to stabilise it as it will contain many double bonds.
Therefore, utilising solar input seems like the best way of keeping operating expenses low while also being the greenest possible way in which to utilise waste plastic, since it would not cause the extraction of any extra fossil fuels to facilitate it, and the petrochemicals would be utilised by various industries and not necessarily end up as carbon dioxide later. They would also replace some petrochemicals which are currently being custom produced from fossil fuels. Maybe also some of the additives could be recovered and reused, although the current practice seems to be sending the residue to landfill [17].

1.2 Topic Overview

Research on many different types of solar thermochemical reactors for the production of solar fuels has been ongoing for a few decades already, for many different reactions and reactor designs. Many prototypes and pre-commercial pilot plants have been built and tested, though none of them have reached the commercial stage yet. [20] Mainly this is due to the fact that traditional fossil fuels are still cheap and there is not enough will to switch to solar fuels yet, but also because the high temperatures at which most of these reactions proceed (usually between $1000°$C and $2000°$C) make the reactor engineering and materials selection challenging.
And there is of course the issue of unsteady solar energy input which further complicates matters and adds an element of uncertainty. However, this technology in general clearly holds a lot of promise for the future, as storage of intermittently available renewable resources will become more and more of an issue. Also, with fossil fuels becoming ever scarcer and global warming an ever more obvious problem, the switch to solar fuels will have to happen eventually, and the sooner the better.
Electricity generation from renewable sources is of course important too, but solar fuels will be an important complement to that. Firstly because electricity is difficult to store, and some vehicles like planes, drones and ships would be better off with a fuel tank and a fuel cell than with lots of heavy batteries. Secondly because many industrial processes require huge amounts of process heat to drive various chemical reactions, and it simply makes much more sense to supply this heat energy directly from absorbed solar radiation than to capture it with a PV panel that has only around 20% efficiency. And thirdly, the unsustainable way in which humanity is currently exploiting earth's natural resources will clearly have to change as they become ever scarcer. Solar pyrolysis and gasification of waste will become a good way of upgrading it into valuable fuels and chemicals while retrieving some of the other valuable materials, especially metals, from it.
Research on solar pyrolysis has to date mostly focused on biomass. Solar pyrolysis of tyres has already been investigated experimentally by one research group [21], and seems to work. Solar gasification of carbonaceous wastes [22] as well as tyres and PET [23] have also been investigated, and there is a plant operating in Morocco that practises pyrolysis of MSW. [24]

Those mixed wastes may include plastic, however it appears that nobody has seriously investigated the case of pyrolysis of only plastic yet. To date, the only publication on the subject that can be readily found on the internet is the master thesis of a student from the University of Sydney from 2007 [25]. He built a small prototype, did TGA experiments on the most common plastics and reviewed the literature on other solar thermochemical reactors. There is very little detail on the prototype though and the experiments he performed with it, nor any conceptual development like which type of reactor would be the most suitable. He concludes that the idea "seems feasible", however without even mentioning any of the potential difficulties and how they could be overcome. These difficulties, along with the complexity of modelling the process, are probably among the reasons that no professional researchers have taken on the subject yet.

The first obvious difficulty is the fact that many plastics are transparent, not just in the visible but the entire solar spectrum. To make matters worse, a quick look at the plastics around you should reveal that most of the uncoloured ones have a certain milkiness to them. This is because some of the light is scattered, either by the impurities or at the boundaries between crystalline and non-crystalline regions within it. The more light is scattered by the feedstock, the more will escape through the window. However, all is not lost because waste plastic also contains a lot of coloured plastic, and most of the inorganic pigments are stable to well over the temperature at which the plastic pyrolyses. Which brings us to the second difficulty, the low thermal conductivity of plastics. This can however be overcome too, by simply shredding the mixed waste plastic to the millimetre scale. A few simple order of magnitude estimates will reveal that a grain size of around 1 - 5mm should be sufficient for the grains to pass on significant amounts of heat from one to another on timescales of seconds. Although clear statements on the internet on how finely commercial granulators can granulate plastic are not readily found, based on what can be found, 3mm should be possible, and there seems to be no obvious reason for there to be a theoretical limit. Another difficulty can be the stickiness of melted plastic, which is also an issue in conventional plastic pyrolysis. This affects the choice of reactor type, which is addressed in the Theoretical Model section of this report. On the other hand, one aspect that plays in our favour here is the fact that recycling facilities much prefer clear plastic to coloured plastic. If they receive a batch with many different colours, then they first have to sort them all by colour, which requires special equipment and is never 100% accurate. And then the end product might not have exactly the desired shade and may be difficult to adjust. So one could envision a future arrangement where the recycling facilities take most of the clear waste plastic and most of the coloured one goes for solar pyrolysis. Some solar reactors could also specialise in WEEE, which is mostly black with some white. This possibility is discussed later on.

1.3 Objectives

Due to the complexity of this subject, one semester is clearly insufficient for writing a simulation that could capture all the complexities. Besides, many estimates had to be made since a lot of information that would be needed about the content and properties of the additives is impossible to find or may not even be known. Experiments would have to be performed in order to find all those values.

Therefore, the aim of this work was not to get a definitive result on whether the concept is feasible or not, but just an initial idea on whether it seems promising or not. Waste plastic is not the ideal material for solar pyrolysis, as explained in the previous section of this report. Therefore, the focus was on how much of the incoming energy gets lost through the window after being reflected or scattered back upwards, and on whether any of the clear plastic gets left over, isolated from the absorbing materials, or if a layer of white pigment forms at the surface. The pyrolysis reaction will of course continue even if not much more energy comes in, as long as the temperature stays over a certain threshold. But as it consumes its own thermal energy, the temperature drops and this significantly slows down the rate at which the reaction proceeds. Whether such unfortunate situations occur or not depends on how efficiently the energy is transferred from the absorbing additives to the plastic, and a simulation should be able to answer this question, albeit for a very small reactor since a realistically sized reactor with small granules would require a supercomputer with a huge amount of RAM to run on.
Since solar reactors are not commercialized yet, it would not be possible to estimate their cost, hence an economic feasibility analysis was not attempted.

1.4 Thesis Outline

The choice of reactor based on a literature review will be described in the Background section, along with the plastics and additives included. Available solar concentrating technologies and locations are presented along with the reasoning for why only the parabolic dish seems suitable for this application. A sketch of the proposed setup is presented. The approach to the complex problem of plastic decomposition will be explained and the literature sources for the chosen model and data given.
Then, the Implementation section will describe how the reactor was discretized, how the heat transfer was calculated, how the Monte Carlo Ray Tracing was implemented, how corrections to the surface are made and how the plastic pyrolysis formula is numerically integrated. Finally, the last subsection will describe the objects created by the code and how the main function brings together all the different parts and the objects described.
The Results section will give the full details of the base case results and then a sensitivity analysis for several parameters of interest based on this.
The Conclusion section mostly focuses on the many recommendations for further work.

Chapter 2

Background

2.1 Theoretical Overview

2.1.1 Possible reactor types

According to one of the recent review articles [26] on plastic pyrolysis, the issue of the stickiness and high viscosity of plastic causes “serious operational problems” for reactors, hence only three types are commonly deemed suitable at all: Fixed beds, fluidised beds and spouted beds. Additionally, there is the somewhat innovative Cycles Spheres Reactor designed specifically to overcome the high viscosity issue of plastics by [27], though it does not seem that this concept has been implemented anywhere. Nevertheless, all four types will now be briefly assessed regarding their suitability for adaptation to solar plastic pyrolysis.
There is also the possibility of microwave pyrolysis, though it has its issues that have so far prevented commercialization [28] [15]. This clearly cannot be adapted to solar input though, which does not reach the microwave range, so it will not be considered further in this work.

Fluidised Bed Reactors

Fluidised bed reactors are commonly used in many industries for many reactions, due to their advantages of achieving good contact among the reactants and catalyst, if present, of facilitating efficient heat transfer and of keeping the temperature distribution fairly homogeneous. They can thus be a good choice for plastics, since the fluidisation solves the problem of their low thermal conductivity. Care needs to be taken though to avoid potential issues with defluidisation, since the sticky plastic melt can coat the surface of the bed particles and cause them to stick to each other. Some of these agglomerates are broken up again by collisions with the bed particles, however this process obviously depends on temperature. So there is a critical temperature above which the agglomerates do not grow too big. Below this temperature though, defluidisation will occur after a certain time after too many of the agglomerates have sunk to the bottom [29]. While a fluidised bed reactor operating with conventional heating can ensure that the temperature stays well above the critical temperature, this would obviously present a major problem

if the temperature was controlled by intermittent and unpredictable solar energy input. A conventional backup heater would have to be present, thus complicating the design, and it would probably have to be used quite often, because one would never be able to be sure where exactly the critical temperature is when using waste plastic, where each batch contains different types of plastics with different additives. It thus seems unlikely that anyone would ever want to attempt this, especially since there is a similar but better suited alternative, namely the Spouted Bed Reactor.

Spouted Bed Reactors

Spouted bed reactors operate in a similar way to fluidised beds, with gas blowing through and fluidising a bed of sand particles and reactants from below. The only difference is the shape, being conical rather than cylindrical. So whereas in a fluidised bed, the gas blows evenly through the whole circular cross section of the cylinder, in a spouted bed the gas enters through the bottom tip of the cone, taking bed particles with it as it surges right through the middle and emerges through the surface in a spout. Those particles then end up falling back down onto the surface and begin their slow journey downwards, so the whole bed slowly circulates. The high velocity in the central column easily breaks apart any agglomerates that may have formed, thereby ensuring that this type of reactor has no defluidisation issues at all. Therefore, there seems to be no reason why this type of reactor could not be used with solar input energy. The solar flux could be focussed onto the surface, or also the walls, or maybe the fluidising gas could be preheated with it. It might not make for a very homogeneous temperature distribution, which might lead to a bit of a fluctuation in the product spectrum obtained, which is in general not so desirable when operating commercially. It remains to be seen though whether this would be a big issue in practice or not.
In conclusion, the spouted bed seems like a good choice for our purpose, since there are no defluidisation issues to watch out for and the catalyst can be mixed in with the bed, where it will be in good contact with the degrading plastic. The main drawback, like with all fluidising reactors, is the complexity and the associated investment and operational costs. However, the reason it was not chosen for modelling in this work is simply because it would have been way too complicated to do in one semester.

Cycled Spheres Reactor

Whether this type of reactor would work well with solar energy would have to be tested in practice. The metal spheres would surely improve the heat transfer a lot. However, potential difficulties with getting enough heat into the central column can be envisioned. As can be seen from the drawing in the paper [27], the concept includes not just heating at the reactor sides and bottom but also internal heating in the central column. Getting solar heat energy there might be difficult. Without it, one can imagine how it might be difficult to get the screw moving when starting up the reactor, and how it might end up getting stuck for a while after cooling due to a cloudy interlude.
In any case, even if more papers and information was available about this concept, it would have been too complex to model for a master thesis.

Fixed Bed Reactor

Fixed beds are the simplest reactor type, and therefore also the cheapest. They are essentially just a vessel with good heat insulation in which the reactants sit without moving, with an inert gas blowing over the surface to take away the evolved volatiles. Since there is no movement, the contact between reactants and with a potential catalyst can be quite bad if the particle sizes are not very small. And since all the heat has to be transferred from the walls via conduction, it might not reach the centre quickly enough if there is a highly endothermic reaction occurring within it. Because of these issues, variations with a stirrer are common.

When it comes to conventional plastic pyrolysis, fixed beds are obviously not a good choice, at least not for commercial operation, as the low thermal conductivity of plastic severely limits the size of the reactor. Thus, the fixed bed reactors employed by scientists studying plastic pyrolysis have all been very small, with the biggest holding no more than 100g of material [26].

However, with solar energy this would not be a problem at all. With a window at the top and sunlight falling directly onto the surface, there would be no limit on how much power could go into the plastic. Double the input energy and double as much is absorbed. With heat transfer by conduction, the maximum temperature that you can heat the surfaces to sets the limit on how much power is transferred. So, going solar would really make sense here, and actually looks like the only way that this type of reactor could be used for commercial plastic pyrolysis at all. The fact that there may be a high temperature gradient from top to bottom should not be an issue for this kind of process.

Obviously, reflection and reradiation losses will limit the overall efficiency of solar resource utilisation in a case of direct absorption like this one. Therefore, the question of how big these losses would be seems like a very interesting one. While there would be less such losses with spouted beds, fixed beds are much simpler and therefore easier to purchase and maintain, which could facilitate the initial spread of this technology. One can well imagine companies not wanting to make too big an initial investment in a technology that is still new and where the returns are not too certain and buyers for all the various end products still need to be established. Or waste management companies in developing countries that have big problems with plastic pollution and would like to test solar pyrolysis but do not have a huge budget.

So, while on the one hand it is true that this reactor type was chosen for this work because it is the only one that was simple enough to model, it is also true that it does actually seem promising and well worth investigating. Time and experience will tell whether the fixed or spouted bed ends up becoming more popular, or if they might complement each other with the fixed bed aiming at the low end market and the spouted bed at the high end.

2.1.2 Possible feedstock types

The composition of plastic waste does vary somewhat from one country to another. In general though, surely PE and PP are the two most common ones, followed by PVC, PET and PS. Figure 2.1 shows the demand in Europe for 2013 to 2015. This is just to get an idea though. The fractions of demand will not

exactly match the fractions entering the waste stream, as usage changes over the years and different products have different lifetimes.

Most plastics are well suitable for direct pyrolysis. Two notable exceptions are PET and PVC though,

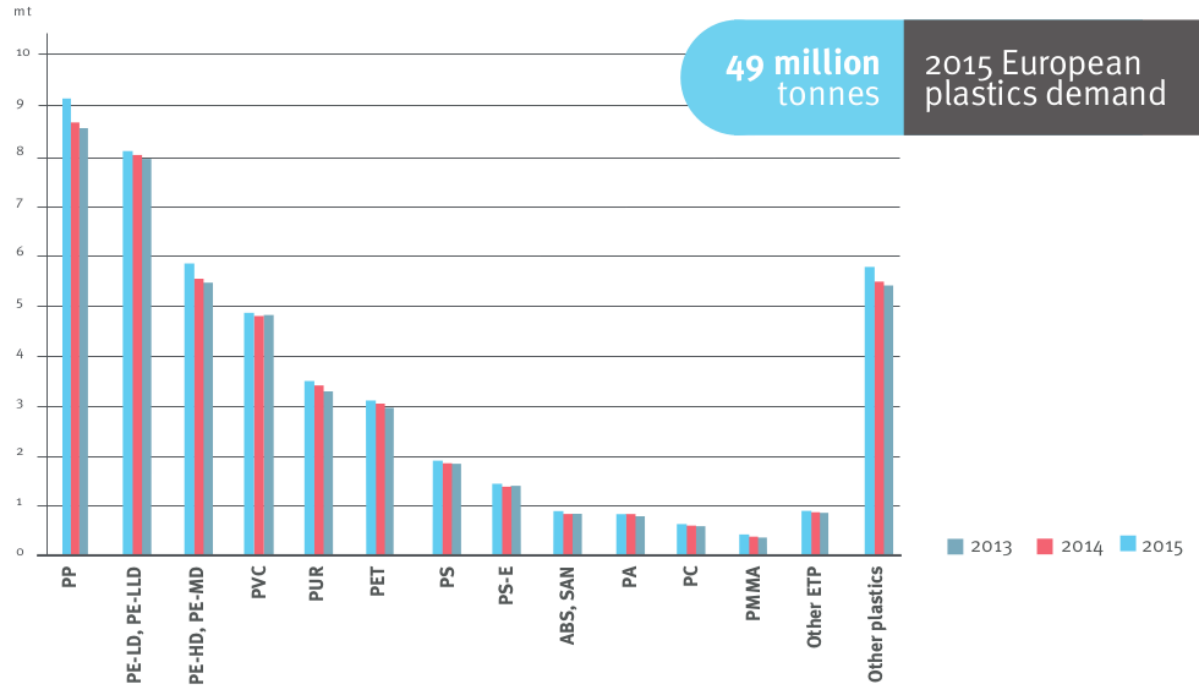

Figure 2.1: Overview of European plastic demand by type [30]

which were not taken into consideration at all. PET contains too much oxygen and thus gives very little oil, and is anyway widely recycled since it is well suited for this. PVC chains have every fourth hydrogen atom replaced by a chlorine atom. Pyrolysing this directly will result in a gas consisting mostly of corrosive HCl, char and other chlorinated products that are of little use. The hydrogen has a higher heating value than carbon and is a clean fuel, so end products with a higher hydrogen to carbon ratio are preferable. There is a lot of research ongoing about PVC pyrolysis and how to remove the chlorine. Adsorbents, catalysts and mechanochemical treatment all help [31] [15], though the removal has to be economical and very effective if damage to the reactor walls is to be avoided and the resulting oil is to be allowed for use as motor fuel. It remains to be seen if this will become feasible. It may make more sense to try and recycle it first though, which is a theoretical possibility. Ultimately, something else will need to be done with it though, whether economical or not. PVC, “the poison plastic”, often also contains a lot of toxic additives, especially plasticizers, so otherwise it will become a big problem in the future, as more and more waste PVC piles up and should not really be thrown into nature. We can only hope that people will become aware of this problem soon enough and stop using it altogether.

In 2015, 5.8% of European plastic demand was for electrical and electronic equipment [30]. Since most of this is black, and most of the rest white, it might be a good choice for solar pyrolysis, since carbon black absorbs well in the entire solar spectrum and titanium dioxide at least absorbs very well in the solar UV part. The majority consists of PP, PE, HIPS, PUR, other ETP, ABS/SAN, PA, PC, PVC and

many other less common ones [30]. So the majority is made of polymers which are suitable for direct pyrolysis,except PUR which contains too much nitrogen, and PA, PC and PVC. PA needs additional treatments, which may be worthwhile though to retrieve the valuable ϵ-caprolactam monomer [32]. The styrene monomer that may be obtained from HIPS and ABS/SAN is also quite valuable [32]. However, there is one caveat, which is the fact that around 30% of them contain flame retardants, of which 41% are halogenated flame retardants [33]. This halogen is usually bromine, although it is slowly falling out of favour due to environmental concerns. These obviously need to be removed somehow, which can be achieved though, either through a catalyst, e.g. zeolites, or mechanochemical treatment, hydrothermal treatment or by adding red mud, which acts both as a catalyst and adsorbent [33].

2.1.3 Possible solar collectors

Four main types of concentrators exist these days: Parabolic Trough, Tower, Dish and Fresnel [34]. The two linear systems, trough and Fresnel, focus the light onto long tubes carrying a heat transfer fluid. Since plastic has a high viscosity, passing it through these tubes would not be possible, and the currently employed HTF oils in parabolic troughs hardly get hot enough to pyrolyse plastic. They would have to use steam or be adapted to work with molten salt, research for which is currently underway. However even so, the temperature difference would not be particularly high, which together with the low thermal conductivity of plastic would lead to rather inefficient heat transfer. So, these two types do not seem very suitable for use with plastic pyrolysis.

Because of the issue of low thermal conductivity, direct absorption of sunlight seems like the way to go. To this end, the two technologies that concentrate all the light towards one small receiver can be employed, namely towers and dishes. However, a cavity receiver ideally has only one small window, to limit losses from escaping rays. This window should ideally be at the top, so that the light falls onto the product first and not the walls. With a huge field full of heliostats like for a tower, it would be difficult to arrange another mirror such that it would reflect all the rays from the heliostats downwards onto the window. This can only practically be achieved with a dish.

2.1.4 Suitable Geographical Locations

Not all countries could practically and economically implement this technology, because some just do not get enough sunshine. However, it is not as simple as “the nearer the equator the better”, since concentrating solar technology can only utilise the radiation that is directly coming from the sun. Any photons that arrive there after being scattered in the atmosphere do not arrive at the correct angle and will thus end up missing the window. So there are actually areas near the equator where the atmosphere scatters so much that they are not suitable at all, probably mostly due to humidity. A good overview of this can be obtained from Figure 2.2, which shows the world map of DNI provided by the German Aerospace Centre. The lighter coloured regions are better than the darker ones.

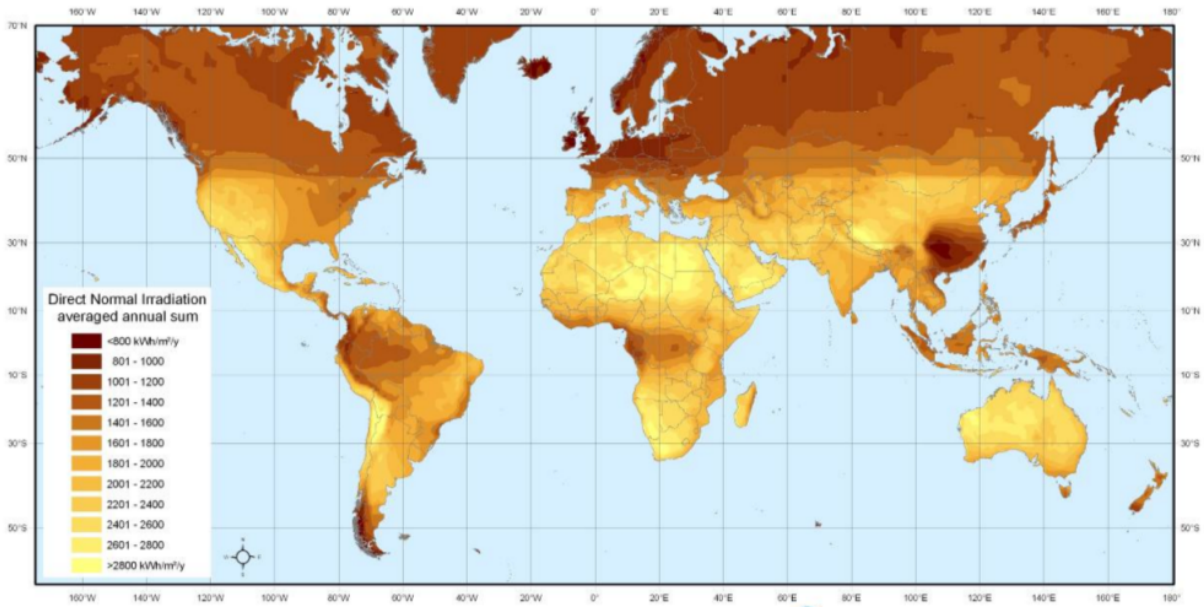

Figure 2.2: World map of DNI [35]

2.2 Theoretical Model

2.2.1 Reactor Setup

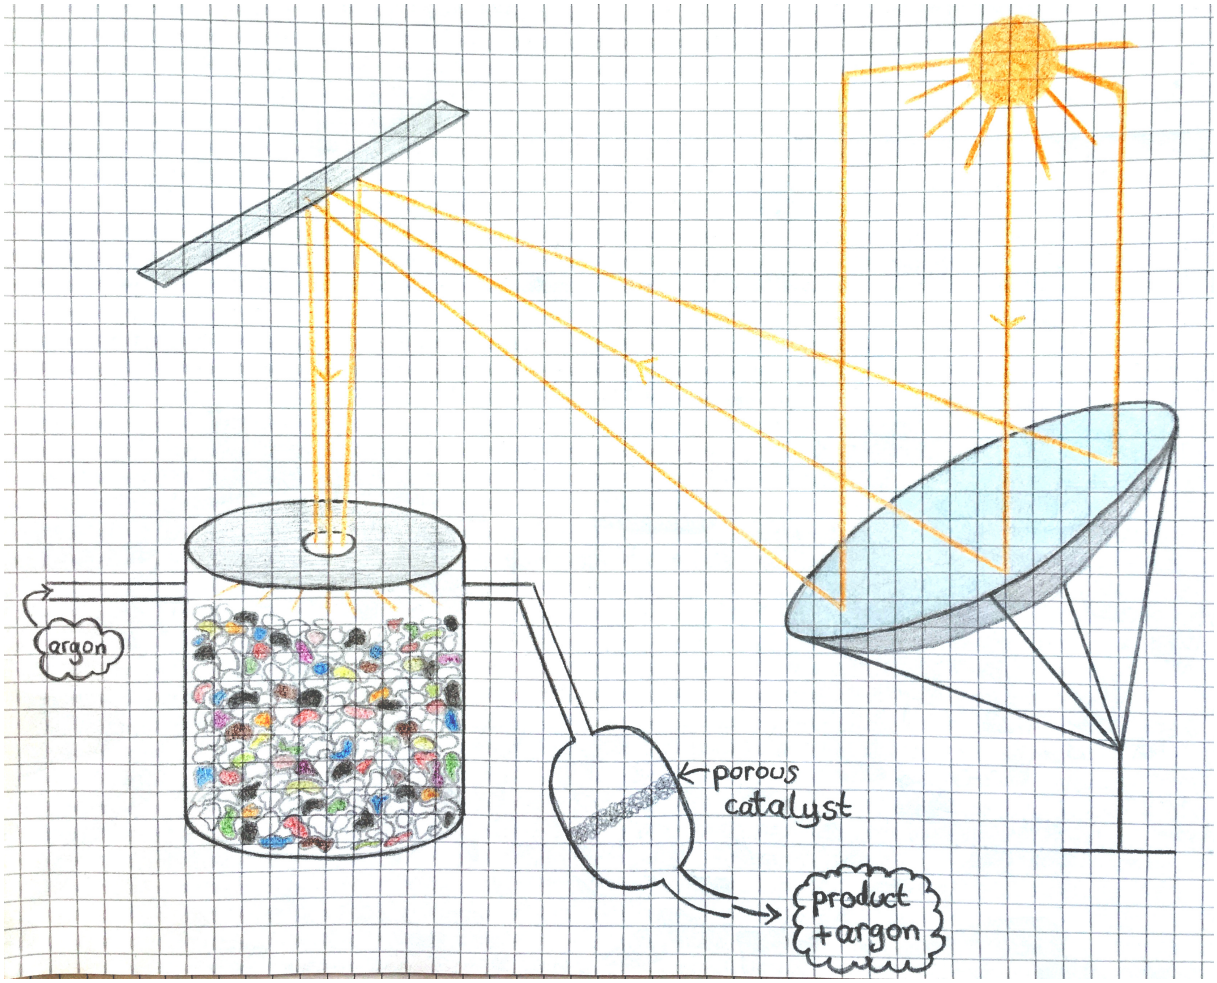

Figure 2.3: Diagram of proposed system

Figure 2.3 gives a basic overview of the assumed system. A parabolic dish focuses and reflects the incoming light upwards onto a mirror that reflects it downwards onto the window of the reactor. This is nothing more than a cylindrical cavity with mirrors on the inside, a window at the top and an inlet and outlet for the argon gas. All the reactor dimensions like radius, height, height of plastic and window radius and thickness are variable and can be set to whatever value is desired before running the simulation.
Fixed beds may be directly or indirectly irradiated. Direct irradiation offers the potential of the highest efficiency, so it is the preferred option unless other reasons make it unsuitable. Usually, the main problem with it is the deposition of carbon and other materials on the window [13]. Some designs avoid this by having no window at all, which however increases the losses from convection and thermal radiation. In the case of plastic, one could imagine that the losses with direct irradiation might be high too due to the light scattering by many of the pigments. With an absorber plate that catches almost all of the incoming radiation and re-radiates it in the infrared, which should be well absorbed by the plastics, this would not be a problem. However, with direct absorption, carbon deposition on the window would probably not be

much of an issue, since the evaporating materials are not too hot and get carried away by the argon gas blowing over the surface. High pressures should not be needed either, which can be a problem for windows too [13]. There also seems to be no big need for an upper cavity to act as a thermal shock absorber for the intermittent solar radiation, which is desirable in solar gasification [13]. Therefore, the obvious first step should be to model direct absorption and see how high the losses really are, and only if they prove to be quite high then the switch to indirect irradiation might be deemed worthwhile.

A non-stirred reactor was chosen, since plastic has such a high viscosity due to the long chains that get entangled that stirring it would use a lot of energy and there would be the risk of the stirrer getting stuck. Commercial plastics tend to have molecular weights far above the critical value above which entanglement becomes an issue. One can calculate from the basic equations [36] that even at 500°C, a typical polyethylene will still have a viscosity similar to that of honey.

Because the contact between plastic and catalyst would not be too good in a non-stirred reactor, and to avoid problems of carbon or metal deposition on the catalyst, most scientists investigating plastic pyrolysis these days have found that it makes the most sense to have the catalytic cracking step after the thermal decomposition step [26], by passing the hot product gas through the catalyst straight afterwards. The same approach was adopted here, with the cracking step being neglected since the choice of catalyst depends strongly on the desired products, and there would not be much to model that would be specific to solar input energy. If there is a dip in solar input, the catalyst may cool but also the reaction in the main chamber will come to a halt, so there will not be much to crack during that time.

Argon gas was assumed to be the inert carrier gas that sweeps away the evolved products and fills any holes between plastic particles. Inlet and outlet for this gas was neglected in the geometry. Unfortunately, there was no time to attempt to estimate the heat losses from convection by this gas blowing over the surface. It would anyway have been no more than an estimate though since the surface roughness is impossible to know. Besides, there seems to be no obvious reason for why the gas would have to blow very quickly, so those losses could be kept low by keeping the speed low. If some of the gaseous volatiles get broken down a bit more by absorbing additional solar radiation, that should not be a problem since smaller volatiles are generally preferred and they get passed through the catalyst at a later stage anyway.

The three most common plastics were chosen to be part of this simulation: HDPE, LDPE and PP. Unfortunately, PS could not be included because the paper from which the reaction models and data for the decomposition are taken does not include it. The relative fractions of these three were calculated from the figures on global non-recycled plastic waste given by [28]. This seemed like the best choice since recycling is preferable to pyrolysis where possible as it requires much less energy, so in practice it will tend to be the non-recycled plastic that will go for solar pyrolysis. The reactor walls were assumed to be adiabatic, though this might not be so realistic. The window was assumed to be made from fused quartz glass, which is standard among such reactors. Infrasil 301 from Heraeus Quartzglas [37] was chosen because it is optimised for good infrared transmission but not for the deep ultraviolet. Most other quartz glasses are manufactured in such a way as to give high transmissivity far into the UV, but since we are only interested in letting through light in the solar spectrum, there is no need for that here. The

transmission calculator on their website was used to obtain the absorption coefficient as a function of wavelength and the Sellmeier equation coefficients for the refractive index as function of wavelength was taken from the datasheet available there.
It was further assumed that the underside of the window has a rugate filter to reflect back thermal radiation, based on the research of [38] which found that this is a promising technology for reducing heat losses through the window. The user can set the minimum and maximum wavelengths of the range for which this filter is designed. The start was set at 2.5μm, as most of the solar spectrum's energy lies before that, and the maximum at 5μm. This should be adequate for preventing the thermal radiation from escaping, which according to Wien's displacement law would have its maximum around 4μm for the hottest plastic. However since thermal radiation could not be considered due to time constraints, the filter should not affect the simulation results presented here much. It is assumed not to affect the incoming radiation.

2.2.2 Solar Energy Input

The standard solar reference spectrum ASTMG173 was used for the energy input as function of wavelength [39]. It assumes an air mass of 1.5 and an average latitude for the USA. Obviously, since we are using concentrating solar technology, the DNI not GHI has to be used. Summing together all the DNI values in the file gives a total of 887.65Wm^{-2}. This value is not as good as some of the really good areas like Chile, South Africa and the Sahara desert, but not too bad either. This value is multiplied by the total area of the solar collector, which can also be set by the user. The value of 56.65m^2 was used, which is the area of the EuroDish Stirling Dish [40].

2.2.3 Included Materials

There are three files for materials other than plastic whose properties are read in by the code. A brief explanation for why each material was chosen is given here.

Additives

Though plastics typically contain many different types of additives like plasticizers, processing aids, flame retardants, antioxidants, UV absorbers, HALS, fillers, clarifying agents, mould release agents, antistatics, nucleating agents, coupling agents, cross-linking agents, antiscratch additives, anti-fogging additives and antimicrobials in addition to the pigments or dyes [41] [2], it was not possible to include many of these due to time constraints. Not only are their chemical properties not always easy or even possible to find, but also the absorption and scattering spectrum for each would be needed, and obviously this is even more difficult to find. Hence, it was decided to focus only on the most important ones for a simulation like this one, i.e. those which seemed likely to have the most effect on the results. UV absorbers clearly seemed like a necessary choice for a simulation involving solar energy absorption. It was also decided to include a filler because of the high weight percentage they may have, which ranges

from 10 - 60%. Some other additives like plasticizers may also have a few percent, but not as much as the fillers and so there was no time for them.

- UV absorber

  Benzophenones are the most commonly used UV absorbers [41], and among these, UV351 is the most common one, as a quick look on alibaba will confirm. Therefore, it seemed sufficient to include only this one.

- Filler

  Calcium carbonate is the most widely used filler [41], and it has a high scattering coefficient (as can be seen from the bright white colour), so it was decided to use this one.

Pigments

The vast majority of plastics contain pigments, not dyes, since most materials will not dissolve well in a molten plastic [41]. Therefore, no dyes were included in this simulation.
Finding figures for the amount of waste plastic that is coloured versus clear and the percentages of each colour is very difficult, unless one has thousands of Euros to spend on market research reports. So the numbers had to be estimated. Looking at pictures of mixed waste plastic, it looks like around half coloured and half clear. Also, the plastic bags from a Chinese landfill had a ratio of coloured to clear slightly over half [42]. A study about plastic recycling found a percentage of 36.5% - 48.5% of coloured plastic in the domestic waste from the UK town of Stockport [43]. So a ratio of one half was adopted, although the simulation should anyway be run with varying ratios to determine the sensitivity of the results to it. And like it was mentioned in the introduction, recycling facilities prefer clear plastic, so in practice the percentage of coloured plastic going for solar pyrolysis is likely to be higher.

- White

  Titanium Dioxide is by far the most commonly used white pigment in general, and for plastics too. In fact, it is the most widely used pigment of all [44]. So the choice here was very clear. It scatters very strongly in the visible and infrared, but has a bandgap in the UV just above the visible. Fortunately for solar pyrolysis, this bandgap is indirect, meaning that the absorbed photons are mostly converted to heat and not reradiated. However only around 6% of the standard solar spectrum lies in the UV, so too much titanium dioxide in the reactor will decrease the efficiency markably. Hence, estimating the fraction of white plastic seemed important, however no definite figures could be found. One vendor is selling HDPE regrinds with a 10-20% white content [45], and the percentage of white plastic bags in the Chinese landfill was 11.34%, though that does not necessarily prove anything.

  So, global figures were used to obtain an estimate. Total worldwide production should be around 6.7 million tonnes this year [46], and around 25% of the total is used for plastic [47], which makes 1.675 million tonnes. The typical percentage by weight of this pigment is 3-5% [47], so taking the mean of that gives 41.875 million tonnes of white plastic. Dividing this by the total amount of worldwide plastic production, 322 million tonnes in 2015 [48], gives a white fraction of 13%. Or,

if one assumes that almost all PET plastic is used for transparent plastic bottles and subtracts 26 million tonnes [49] of this from the worldwide total, it would be 14.15%. So, 14% was used.

It is well known that titanium dioxide can act as a photocatalyst for the degradation of organic material, via the Honda-Fujishima effect. In fact, it is already widely used for that purpose, to help clean waste water and for self-cleaning walls and glass [50]. Scientists are also investigating its potential to make plastic decompose by itself under the influence of sunlight when thrown into nature. However, it seems unlikely that this would significantly influence solar pyrolysis, as this is carried out in an inert atmosphere and the photocatalytic degradation proceeds via the radicals that are created either from water or from air as they take up the electron that crosses the bandgap after absorbing a UV photon. In the absence of such oxygen containing media, most electron-hole pairs recombine within nanoseconds, thereby releasing heat. Besides, pigmentary titanium dioxide consists of particles with mean diameter of around $0.2\mu m$, and only the atoms on the surface of each particle aggregate are even able to participate in a reaction. This is the reason why scientists investigating its potential use for degrading plastic use nano-sized titanium dioxide, not commercial pigment. So, while there seems to be no research to confirm it, even if the plastic has a small amount of moisture or other oxygen-containing additives that could partake in a photocatalytic process, the overall effect would probably be negligible compared to all the other chemical reactions [51]. Also, the paper on solar pyrolysis of tyres included titanium dioxide among the catalysts investigated, and found no such effect [21]. Titanium Dioxide exists in nature in four polymorphic forms, anatase, akaogiite, brookite and rutile. Rutile has the highest refractive index and density, and the other three forms convert to it around 700°C. Hence, for use as a pigment only rutile is used [41], so this is the form assumed by the simulation.

- Black

  Though different types of black exist on the market, carbon black is without doubt the most common one, constituting 10% of worldwide pigment production [44]. So the choice here was also not difficult. Since it is such an efficient light absorber, and does not evaporate under 4000°C, the more of it is in the reactor the better. Hence, some figures on the prevalence of this colour were sought here too. The Stockport recycling trial found a fraction of 5% black plastic [43], though the total also contained 21% PET and 4% PVC which are excluded from pyrolysis, and 29% films which tend to be clear and thus more suited for recycling. The author goes on to mention that the infrared sorting process cannot identify black plastic, which accounts for approximately 8% of the total input. So, this number seems like a better indicator of the real percentage. Assuming that all PET is clear and not included for pyrolysis makes this number go up to around 10%, which was used here, although the fact that the recyclers cannot easily sort black plastic could also mean that they would be more likely to hand it over for solar pyrolysis.

- Yellow, Orange and Red

  Since iron oxides are the second most commonly used pigments of all [44], the choice here was also very clear. Hematite was used for red, goethite for yellow and orange was assumed to consist

of half each of the two. Goethite dehydrates at 180°C and turns into hematite, which is taken into account in the simulation. Additionally, chrome yellow (lead chromate) seems to be quite common too, though its use is declining slowly due to its toxicity. Nevertheless, a small amount was included, since there is still a lot on offer on alibaba and the plastic going for pyrolysis does not necessarily have to have all been produced recently.

- Blue and Green
  Another straightforward choice, since phthalocyanins are very popular due to their low cost and very high tinting strength. Copper Phthalocyanine gives a very deep, clear shade of blue. Substituting most of the hydrogen atoms on the outside of the ring with chlorine atoms gives green, for which 15 out of 16 substitutions were assumed since that is the average amount. [52]

- Violet
  Looking on alibaba and reading around, it seems that both dioxazine violet (aka carbazole violet) and quinacridone violet are very common, and thus it was initially envisaged to include both in equal amounts. However, the paper with all the data on scattering and absorption coefficients unfortunately did not include quinacridone violet, only quinacridone red. Hence, only dioxazine violet was eventually used.

Decomposition Products

The only two decomposition products that could be included in the limited time available were two of the more important ones, copper (from the phthalocyanines) and calcium oxide (from the calcium carbonate). Whenever an organic additive decomposes, it was assumed that each carbon atom flies away with two hydrogen or two chlorine atoms each, and the remaining carbon atoms were calculated based on this. These were assumed to become carbon black, although technically carbon black is not pure carbon, but such subtleties were neglected. The material properties for carbon black used are actually mostly those of pure carbon, since they are easy to find.

2.2.4 Plastic Decomposition

Plastics are made up of macromolecular carbon chains, usually containing tens of thousands of carbon atoms each. The two extra bonds that each carbon atom has are usually taken up by hydrogen atoms, though in some cases there can be chlorine, oxygen, a methyl group or a whole side chain. So, it is not surprising that breaking apart these chains requires a lot of energy and the overall process is very complex and not straightforward to model. Many researchers have attempted to determine kinetic parameters for this process, however there are big variations in their results. This is mostly attributable to the fact that most of them assumed first order kinetics, which is a rather inaccurate simplification. In fact, the process usually starts off being zeroeth order and then slowly turns into first order, in the later stages where any additional bond breakage is likely to lead to evaporation and thus mass loss [53]. Many researchers have thus proposed formulas for the reaction model, or attempted to model the entire

process in more detail, yet none of these models has been able to fit the actual observed TGA data with a high degree of accuracy.

Attempting to model the process on a molecular scale would clearly have been way beyond the scope of a master thesis. Hence, a choice had to be made among the models available in the literature, while keeping in mind that even the best currently available one can be no more than an approximation in this case. Even if it could accurately predict the results of laboratory TGA experiments of the most common plastics, it still would not apply well to a real-life case like this one. Firstly, real waste plastic contains many additives which will all have some effect on the degradation reaction. Some in a more indirect way by simply increasing the distance between adjacent chains and thus inhibiting intermolecular hydrogen transfer, but others in a more direct way. They may decompose before the plastic has fully degraded and release oxygen, chlorine or bromine, which may integrate themselves into the chain fragments. Even worse though, the antioxidants, flame retardants and light stabilisers will slow down the degradation process because that is what they are designed to do. They are added with the precise purpose of preventing damage to the plastic from the propagation of free radicals that may be created by external influences. However the pyrolysis process also proceeds via the propagation of free radicals that are created by the breakage of a single bond. Of course, some of those antioxidants and flame retardants will degrade before the plastic does, but it depends on the type. And virtually all commercial plastics contain antioxidants, often they are added right away by the original manufacturer. Secondly, the UV solar rays have just about enough energy to break a C-C bond in the chain. Plastics that are expected to be exposed to sunlight during their use time usually have either UV absorbers or light stabilisers (HALS) or both added to prevent damage from UV light. However, most UV absorbers will probably decompose before the plastic has finished decomposing (see for example Table A.2 in Appendix A, where the UV absorber used here, UV351, decomposes at 457.9°C), and it is anyway questionable that they would be able to absorb all of the UV rays in the case of intense irradiation in a solar reactor. Further complicating matters, several pigments like carbon black or titanium dioxide also absorb UV rays, which is why the code is set up to not add an extra UV absorber to plastics containing these pigments. So, the question of how much the UV irradiation affects the plastic decomposition is a complicated one, and experiments would need to be performed to obtain at least an estimate.

Thirdly, all the TGA experiments reported in the literature are either performed under isothermal conditions or using a constant heating rate. Most of them do not even perform simultaneous DSC. However in our case, the heating rate will not be constant but is influenced by several factors, like variable solar input, additives undergoing phase transitions and varying conductive heat transfer from the neighbours. But even just deducing what the heating rate would be is not so straightforward, since it depends on how much of the incoming heat energy can be consumed by the reaction and how much is left over to heat the material. Even the few papers which include DSC approach the issue of heat consumption the other way around. They set a fixed temperature programme for their TGA apparatus and then measure how much heat flow is needed to maintain this temperature profile. But in our case, the heat input is given and we need to know the resulting temperature profile. Even if experiments were performed to determine the temperature profile for various fixed heat inputs, those results could probably still not be

extrapolated to a case of randomly varying heat input.

Despite all these issues, a choice had to be made, and given the inadequacy of simple first order models, one with more detail was preferred. It was thus decided to use the empirically derived formulas and values from [54], because the authors' approach of not presupposing any model or constant parameters but rather using the TGA curves to derive the model and the values of activation energy and pre-exponential coefficient as a function of conversion via the advanced isoconversional method seems like the most sensible approach. And after seeing how according to their results, the pre-exponential coefficient varies by several orders of magnitude over the course of the conversion process, one really would not feel comfortable using a simple model where this value is constant any more. Therefore, the reaction models they found for HDPE, LDPE and PP were selected, although they may not apply so well in the case of heating via radiation as opposed to heating a particle solely via conduction heat transfer. The free software im2graph [55] was used to read the values of activation energy and pre-exponential factor from the graphs in the paper. Unfortunately, the lead author did not respond to a request for the full data of the pre-exponential factors, so the values at low conversions had to be estimated since they are several orders of magnitude lower than the maximum values and the scale of the graphs is not logarithmic. The paper unfortunately does not include DSC curves or even any mention of the heat consumption associated with those experiments. Therefore, this information had to be taken from other papers. After some searching, one paper was found with both DSC and TGA of HDPE [56], and one for PP [57]. Although the kinetics should be different for LDPE since the junctions of the sidebranches break more easily than the main chain, the data for HDPE was used for it too since no other suitable data could be found. Two other papers were found which had DSC data, but one did not have any TGA data and one did but at a different heating rate. But both these data from the same experiment were needed in order to be able to convert the DSC curve from one versus temperature to one versus conversion. A simple MATLAB script was employed for this purpose. It reads in the raw data that was read from the graphs with the im2graph software, extrapolates both curves for their values at each 0.5K temperature step, subtracts the estimated amount of heat that must have been for sensible heating and then matches up the values of conversion and heat flow via the shared temperature. Finally, it turns the resulting heat flow data into a column array of 101 values, the first value representing the amount of heat flowing into a 1kg sample before the conversion starts and the remaining hundred the amounts needed to convert the subsequent one percent each. Since not all graphs had the time in addition to temperature as horizontal axis, thus not allowing to convert from power to energy, the heat flow was normalised so that the total matches the values from [58]. Figures 2.4 and 2.5 show the values obtained. The spikes are probably due to inaccuracies caused by the not very detailed input data.

The activation energies and pre-exponential coefficients were also converted to simple csv files with one hundred values for each percentage of conversion. A higher accuracy seems unnecessary given all the other uncertainties and the fact that they were taken from graphs.

It is known that certain plastics may affect the degradation kinetics of other plastics if they are pyrolysed together, via intermolecular transfer of radicals. Especially polystyrene, which begins to pyrolyse at lower temperatures than the common polyolefins, will cause them to also begin pyrolysing sooner [27].

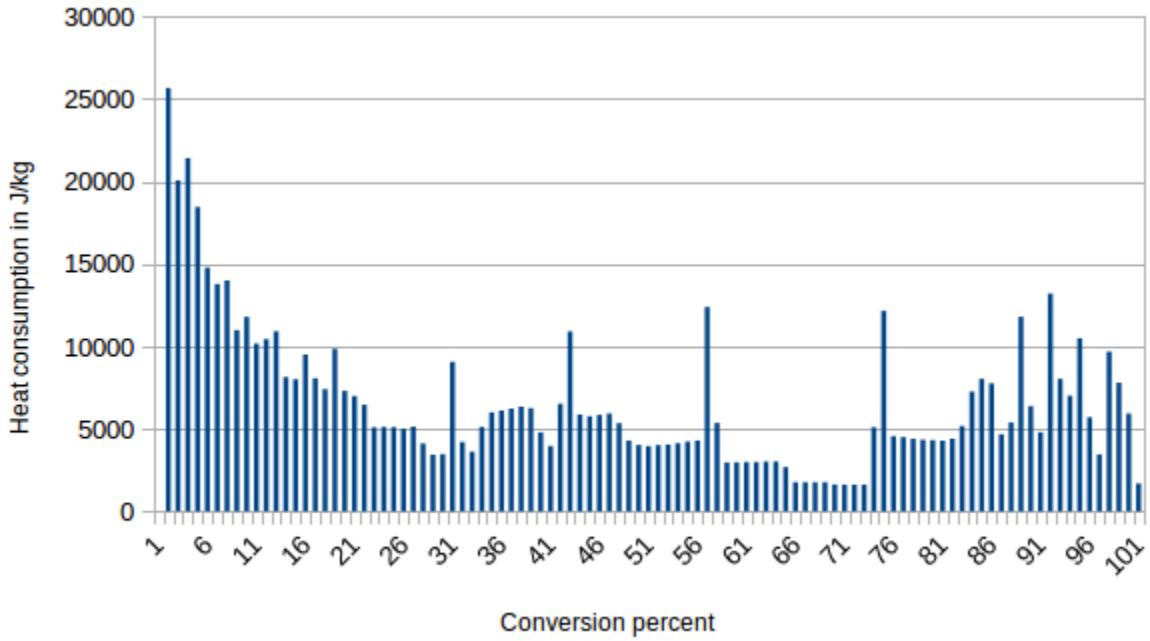

Figure 2.4: Inferred energy consumption during pyrolysis of HDPE at each conversion percent

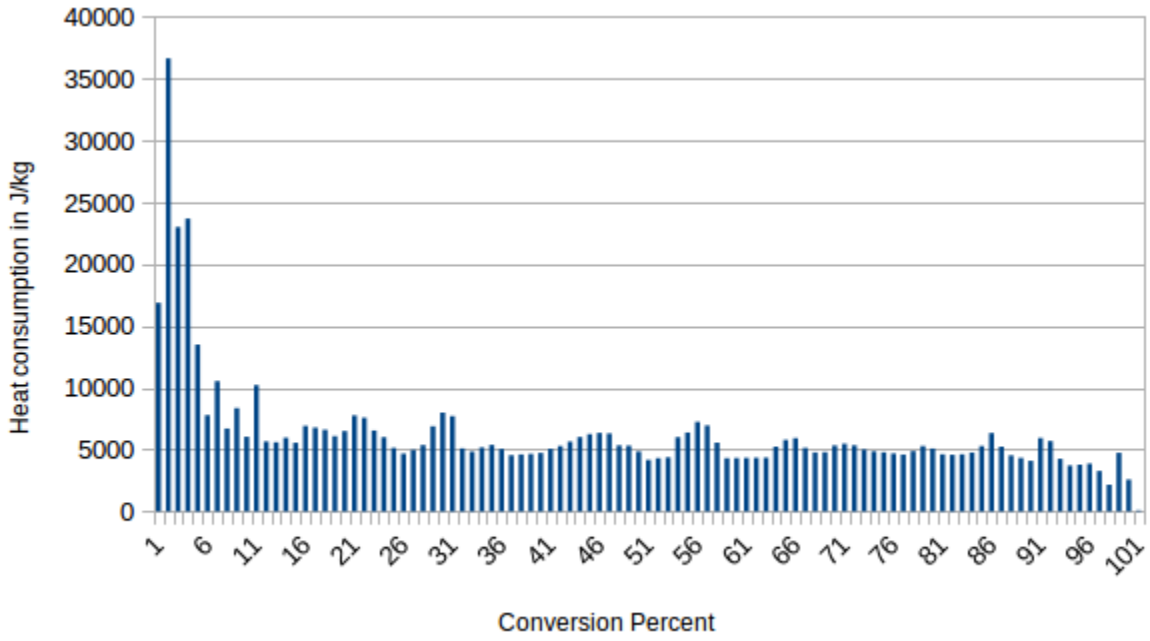

Figure 2.5: Inferred energy consumption during pyrolysis of PP at each conversion percent

However, even if PS were included in this simulation, this effect should be negligible since the mixing needs to be on a molecular scale for it to be significant, as concluded by [59], and particle diameters of a few millimeters are assumed in this simulation, with no stirring. A look at the self diffusion coefficients of various polymers [60] [61] shows that they are very low, of the order of $10^{-13}\frac{m^2}{s}$ or even less. From this, it can be estimated that it would take a few thousand hours or more for different types of plastic to diffuse into each other, of course assuming that the cross-diffusion coefficients are of a similar order of

magnitue than the self diffusion coefficients, which does not necessarily have to be the case. This result is not surprising, given that plastics consist of very long macromolecular chains, and it takes a while for such chains to move via reptation.

Chapter 3

Implementation

Due to the complete novelty of this kind of simulation, and the fact that it combines the three fields of ray tracing, heat transfer and chemistry, no existing software could be used. Instead, the whole simulation had to be built from scratch using a full programming language, and C++ was chosen as it is the main programming language used by the research group at the CTTC supervising this thesis. The C++11 standard had to be used since the good quality random number generators needed for the ray tracing are not available in earlier versions.

3.1 Numerical Model

3.1.1 Discretization of Reactor

After being run through a shredder and then a granulator, the waste plastic will consist of many granules of all kinds of shapes. Since random shapes would have been too complicated to model though, they were assumed to be cubic in shape, all with the same size. This makes the heat transfer more efficient than in reality. To partially compensate for this and make it a bit more realistic, it was decided to subdivide each granule into several pieces, all of the same size and also cubes, with the exact same initial properties, and to have a certain fixed probability for each piece on the outside of the granules to be empty. This empty spot will be filled by the argon gas blowing through the reactor, so they are called argon bubbles. The remaining pieces of the granules will be called plastic pieces, for lack of a better word.

The argon bubbles do not rise upwards like they would in real life, although future versions of the code could include this aspect. In reality, they will rise slowly due to the high viscosity of plastic. One can estimate the upper speed limit from the standard formula $v = \frac{1}{3}\frac{r^2 \cdot g}{\nu}$ to be around $\frac{1}{12}\frac{mm}{s}$ for the hottest plastic, which has the lowest viscosity. Since the whole simulation with the parameters chosen in this work takes a bit more than three hours, this effect may not be negligible.

Each plastic piece and each argon bubble is represented by an object in the code, and each object has its own properties, including temperature. The objects, being cubic in shape, are stacked tightly face to face in an imaginary three-dimensional cubic lattice. The physical equivalent of these objects will be

called lattice element in this work. The lattice remains fixed in space throughout the simulation, as does the volume of the elements within it. That is, the density of the materials anyway does not change with temperature in this simulation, but even if all that data could be found and used, it would only affect the volume fractions and molar volumes of the constituent materials of a plastic piece. The volume it is assumed to occupy in the lattice is always the same, the lattice constant to the power of three. Even if its mass decreases because an additive evaporates or decomposes or it is in the middle of the plastic pyrolysis process it will still occupy the same volume. In order to maintain consistency with the heat transfer formulae, an effective density is then calculated instead for use in those, by dividing the new mass by the constant volume.

The element at each lattice position can be accessed via the index values of the three-dimensional array that holds all the objects. The elements can only move by "jumping" from one lattice position to another, which happens instantly. Although the plastic pieces are imaginary subdivisions of a real granule, they may move totally independently of each other, even when the granule is still solid. Keeping solid granules together would have introduced further complications for which there was no time.

The origin of the Cartesian coordinate system through which all positions in space are expressed is located at the bottom of the reactor, outside of it. The circular reactor wall just touches the x and y axes, so the centre of the circular bottom wall is located at (r,r,0), where r stands for its radius. The transitions from one lattice element to the next are located at integer multiples of the lattice constant. Since the reactor is cylindrical but the coordinate system Cartesian, some of the element columns will take part in the simulation and some will not. Those that do not always remain empty, and for those that do, there always needs to be an element at every position in the column. At the end of the simulation, all of those elements will end up being argon gas. Whether a given column takes part is determined in the following way: The reactor bottom is imagined to be filled with squares of the same dimension as the granule dimension. Whether or not the granules above each square take part or not depends on whether the granule centres lie within the reactor or not. However, if a granule takes part then it does not necessarily mean that all of its constituent pieces take part too, since some of them may not fully lie within the reactor. So the code computes the distance to the corner which is furthest away from the centre, and only lets an element column within an existing granule take part if this distance is less than or equal to the reactor radius. All of the area between the partaking lattice elements and reactor wall is assumed to be filled with argon gas, though it does not take part in the heat transfer, does not carry away heat by convection and does not absorb or scatter. Its only effect is via the refractive index of argon gas, which is used in Snell's law and the Fresnel equations when a ray is at the boundary between this area and the lattice. Figure 3.1 shows an example of what this may end up looking like, in a case with four subdivisions of each granule in each spatial dimension. Argon gas has no colour of course, but it is coloured light blue here just to show it more clearly.

Regarding the vertical direction, the lattice and its elements only extend up to the initial height of the plastic. The space between lattice and reactor ceiling is assumed to be filled with argon gas, which again only has an effect via its refractive index. The height of the plastic mass can be set by the user, but if this value does not evenly divide by the granule size, then the height is adjusted slightly to the

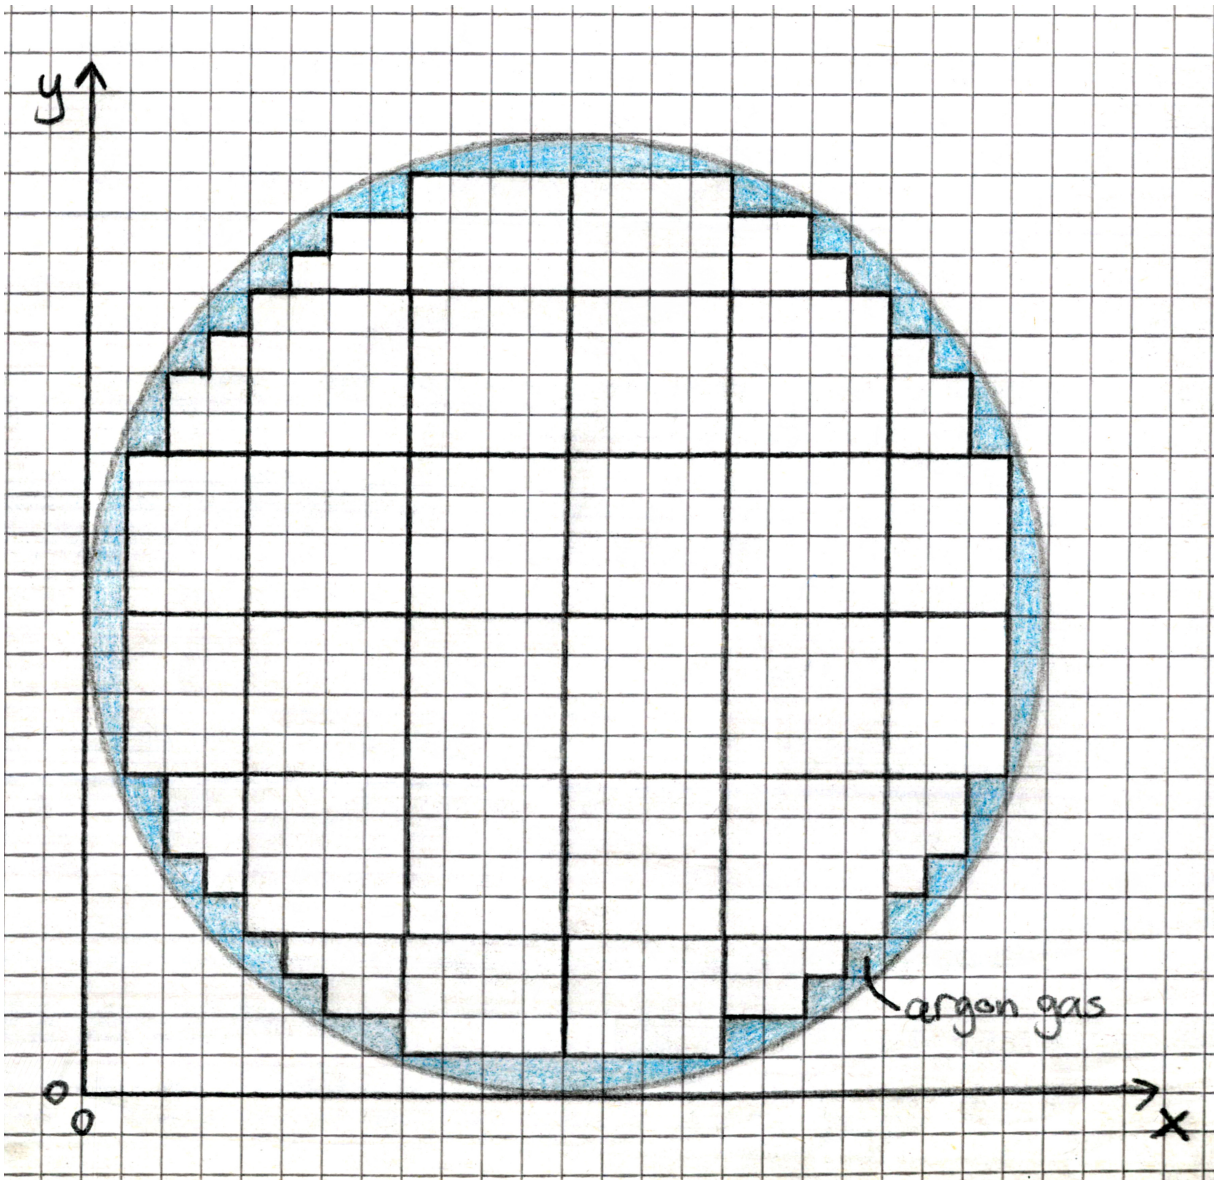

Figure 3.1: Example of element columns participating in simulation

nearest value that does.

Above this gas filled area, there is the window. Its radius and thickness can be set by the user, but apart from that all its properties are fixed and do not change in time. It does not have a temperature, although in practice it will heat up and needs to be kept from overheating, so it would be good to include this aspect in future work. Since its properties do not vary in space or time, it is not discretized at all, but simply treated as a fixed part of the geometry.

When talking about “plastic pieces”, the word “plastic” is assumed to mean a real waste plastic, that is including additives. To avoid confusion with the plastic minus additives, this will be called pure plastic in this work, as in the code. It could have also been defined the other way around, with the real waste plastic being called composite plastic, but it is all a matter of definition. Now, as the simulation progresses, it will happen that the pure plastic of a plastic piece finishes decomposing, but some additives remain because their evaporation or decomposition temperatures are much higher than the pure plastic

reaches while decomposing. This presents a problem because it is not realistic to have the remaining mass occupy a full cube of the lattice, when it can be as low as 0.1% of the original mass in the case of some organic pigments. Or even less if the organic pigment has already decomposed and left behind a few decomposition products. So, it was decided to instead convert what is left of the plastic piece into a thin sheet which has the same area as the cube surfaces and lies sandwiched between the element that replaces the decomposed one and the one below. Or, in case there already is a sheet or a sheet stack at this boundary, then the new sheet goes to the top of this stack. A sheet stack is simply a collection of sheets all perfectly stacked one on top of another, like the pages of a closed book. Or, in case there happens to be an argon bubble in the lattice position underneath, then it falls down through it.

Each sheet has a thickness that is calculated from its mass, density and area of its horizontal face (the square of the lattice constant). However, if this thickness was added to the total height of the column it is located in, and the lattice elements above it were shifted upwards correspondingly, it would greatly complicate the already complicated ray tracing geometry. Therefore, it was decided to keep the lattice fixed in space regardless of how many sheets there may be. They thus do not occupy any volume in this lattice nor in space, although when a light ray reaches that boundary it does have to go through them, and may be absorbed or lose energy or be scattered within them according to their thickness value and other properties. Like the lattice elements, the sheets can only move by “jumping” instantaneously from one allowed position to another. Unlike the lattice elements, they cannot move sideways though, only downwards.

Figure 3.2 in Subsection 3.1.3 gives examples of ray paths, however it may also serve as a visualization of the simulation elements described here, albeit in two dimensions only. Sheets there are drawn with their real thickness, because zero thickness cannot be drawn, so this should be kept in mind.

All these simplifications do introduce some inaccuracy, although they do not make the simulation totally unrealistic; based on the data for the thermal expansion of PE in [36], it can be estimated that PE will expand its volume by 41% from an increase in temperature of 500K from 25°C. And, as can be seen from the data tables in Appendix A, most additives make up only a few percent of the total mass of their plastic. The only exception are fillers, like the most common one calcium carbonate which is included in this simulation, because in this case the producers try to add as much of it as they can to save costs. It can be up to 60% and is stable up to around 825°C [41], so the plastic will always pyrolyse first and leave behind sheets containing it. Neglecting the thickness of these sheets is not too realistic. Also, in a big simulation the sheet stacks could easily become quite high, although in this work it was anyway not possible to run big simulations.

The word ”element” will be used when referring to lattice elements as well as sheets. Since the argon objects do not really do much except to occupy empty space, the plastic pieces and sheets together will be called ”participating elements”.

3.1.2 Heat Transfer

The heat transfer part works in an explicit way, meaning that the heat flow during the time period between the present and the next instant is calculated from the present temperatures, not the future ones or at some point in between. This necessitates restricting the time step to avoid temperature over- or undershoots.

The main thing it does is to calculate the amount of heat incoming or outgoing by conduction according to Fourier's law 3.1 during the minimum time step for all participating elements and to store the results in arrays for later use. This is achieved by calculating the heat flow across each boundary once and then adding and subtracting this amount to and from the two participating elements.

$$\dot{q} = -k\frac{dT}{dx} \tag{3.1}$$

where $\cdot q$ is the amount of heat flowing across a boundary per area and per unit time, k is the thermal conductivity, and $\frac{dT}{dx}$ the temperature gradient in space at this boundary. This equation was discretised in the following way:

$$\Delta E_{cond} = \dot{q} \cdot A \cdot \Delta t = -k_{int} \cdot \frac{T_1 - T_2}{d_{1-2}} \cdot \Delta x^2 \cdot \Delta t \tag{3.2}$$

where ΔE_{cond} stands for the conduction energy received or lost by participating element 1, A stands for the area of the two surfaces touching each other (which in this simulation is always the square of the lattice constant, Δx), Δt the time step, k_{int} the interface conductivity of the two participating elements, T_1 the current temperature of participating element 1, T_2 the current temperature of participating element 2 and d_{1-2} the distance between the centres of the two participating elements. It is thus implicitly assumed that the mass of both objects is concentrated in their centres, although this does not affect the fact that they touch each other at the boundary. The maximum time step that can be applied for each participating element is derived from this formula by considering all the heat flows across its boundaries and requiring that a hot object cannot lose all of its thermal energy to its colder neighbours in one step (neglecting phase transitions and assuming that $\Delta E = mC_p\Delta T$, where m is mass and C_p heat capacity at constant pressure), and then multiplying the resulting expressions by a factor of 0.25 so that the temperatures do not fluctuate too much.

Among all the calculated maximum allowed time steps for each participating element, the lowest one is chosen and applied to all the heat transfer between all the participating elements. However, this value is also restricted by a maximum that can be set by the user, as the heat transfer part and solar input use the same time step and it would not be realistic to allow too much solar energy to enter before checking how the elements respond to it. However in case the time step becomes very small, which can happen due to very thin sheets, there is a safeguard to prevent rays with too little energy from being distributed, as this would be inefficient. In that case, the solar energy accumulates in a separate variable until it reaches a set threshold.

The interface conductivity k_{int} used in equation 3.2 is calculated from the standard harmonic mean formula from Patankar [62] for each boundary, and stored in three arrays for the three directions in

space and another one for the boundaries involving sheets:

$$k_{int} = \frac{1}{\frac{f}{k_1} + \frac{1-f}{k_2}} \tag{3.3}$$

where f is the distance from the centre of element 1 to the interface divided by the distance between the centres of both elements.
Sheets only transfer heat vertically, not horizontally.

3.1.3 Distribution of Solar Energy Input

The solar input is assumed to be constant in time and always comes from the focal point of the lens. It is distributed one ray at a time. The amount of energy per ray varies because the time step is determined by the heat transfer part, so the total amount of energy to be distributed each time may vary, thus requiring different amounts of rays in order to get an acceptable value for the energy per ray. The maximum and minimum amount that this value can take can be set in the Dimensions file, along with the threshold energy value below which a ray is considered to have been absorbed and the code stops tracking it further. Currently, the maximum is set as a function of the lattice constant, by estimating how much energy it would take to raise the temperature of a plastic piece containing only pure plastic by 20K. This value is obviously a somewhat arbitrary choice, but some kind of choice has to be made between accuracy and computational effort. 20K may seem like a lot, but with the small lattice constant of 3mm length that we are considering, it already leads to rays with a maximum of just 0.89J. How many are then to distribute per unit time will depend on the scale factor, since the incoming power is scaled along with the volume of the reactor. The minimum is set at one tenth the maximum, and the threshold at one hundredth of the minimum.
The dish area can also be set there, and this value is simply multiplied by the sum of DNIs per area obtained from the solar spectrum supplied in order to get the total power entering the reactor. Losses due to imperfections of the mirror, tracking, spillage or absorption by the window are not taken into account at all, since no technical assumptions are made about the concentrating system.
The total incoming power is then multiplied by the minimum time step obtained from the heat transfer part to determine the amount of energy entering in that period. If this energy is higher than the required minimum, or the sum of this energy with that which has accumulated from previous steps where there wasn't enough, then the minimum number of rays needed so that each of them will have no more than the maximum energy allowed per ray is calculated, and the total energy to be distributed is divided by that number to obtain the actual amount per ray at this time step.
For each of these rays, a wavelength is then chosen via the Monte Carlo technique [63]. The probability that a ray will have any particular wavelength among the included ones is simply:

$$p(\lambda) = \frac{DNI(\lambda)}{DNI_{tot}} \tag{3.4}$$

where $DNI(\lambda)$ is the amount of direct normal irradiation at wavelength λ and DNI_{tot} the sum of these values over all wavelengths. From this, the cumulative probability distribution is calculated:

$$P(\lambda) = \sum_{l=\lambda_{min}}^{l=\lambda} p(l) \tag{3.5}$$

Then, the code draws a random number between 0 and 1 and goes through the participating wavelengths in ascending order, starting at the lowest (λ_{min}), and as soon as the random number is less than the value of $P(\lambda)$ it stops and sets the wavelength of the ray equal to the wavelength it has reached.

A ray is further characterised by its current position in the Cartesian coordinate system, which is stored in an array with three values, and two angles in radians that determine its direction. One is the angle from the vertical which is measured from the positive z axis and ranges from 0 to π and the other the angle in the x-y plane, measured counterclockwise from the positive x axis and ranging from 0 to 2π. Initial values are determined as follows. It is assumed that the window is a diverging lens which spreads out the rays so that they just reach all of the plastic surface. Based on this assumption, the position of the focal point of the lens is determined. The maximum angle from the vertical that a ray can have is then calculated such that all rays appear to be coming from this focal point. Each ray is then given an angle between exactly downwards and the maximum deviation from that, with equal probability. It is also given an angle in the x-y plane with uniform probability.

The path of each of these rays is then tracked until it either hits a plastic piece with carbon black pigment, or its remaining energy falls under the threshold or it is lost by either escaping through the window or being absorbed while in the window. This is achieved by always finding the next position in space at which the ray leaves the element it has just entered, as well as the entry surface of the next element. The entrysurface variable is of type char (letter) and can take only one of six values - e, w, n, s, t or b, standing for east, west, north, south, top and bottom respectively.

Three integer variables i, j, and k keep track of the index values of the array that corresponds to the lattice element through which the ray is currently going, so that the properties needed can be easily accessed from the object at this array position. When the ray goes through a sheet stack, another index variable s keeps track of which sheet of the stack it is currently in.

The only exception to this tracking mechanism is the case when the ray height exceeds the lattice height, or if it gets reflected right away before even entering the lattice. There it is tracked by a different function that determines whether it hits the wall or ceiling and finds those positions on them, until it reaches the lattice surface again or gets lost through the window. If it then gets reflected by the surface due to Fresnel reflection, the function is called again and again until it either enters or gets lost. If it enters the window, then the code simply keeps track of whether it is at the top or bottom surface and updates its x and y coordinates only.

It can also happen that a ray leaves the participating lattice columns sideways, entering the argon filled area next to the wall. In that case there is another function that tracks it until it either reenters a participating element or its height exceeds the lattice height, in which case it is handed over to the function described in the previous paragraph. No Fresnel reflection is applied for reentry in this case, only Snell's

law. This function uses the same function as the main function does to update the ray position to the exit point of a cubic lattice volume it is entering (there is still a hypothetical lattice there, even though it does not contain elements). It uses a second coordinate vector though to first project the exit coordinates, and then checks if the ray would first hit the reactor wall instead. If so, perfect reflection from the tangent to the wall at this position is assumed and a different function has to determine the new exit point from the cubic volume. The possibility of multiple reflections from the wall within one cubic volume had to be taken into account. The possibility of absorption by the walls was neglected.

The rays never split like they may do in reality when going from a medium with one refractive index to one with another. Instead, the fractions that would be reflected and transmitted are calculated from the Fresnel equations and used as probabilities:

First, in case

$$\theta_i > asin(\frac{n_2}{n_1}) \tag{3.6}$$

then total internal reflection occurs, so no random number needs to be drawn. θ_i is the incidence angle of the ray (the angle it makes with the normal to the surface it impinges on, always the one less than $\frac{pi}{2}$), n_1 the refractive index of the element it is currently in and n_2 the refractive index of the element it is trying to enter.

In case this expression does not evaluate to true, then the Fresnel equations are invoked. This requires knowledge of the angle with the normal that the transmitted ray would have, the cosine of which is used directly:

$$cos(\theta_t) = \sqrt{1 - (\frac{n_1}{n_2} \cdot sin(\theta_i))^2} \tag{3.7}$$

The probability of reflection is then given by:

$$p(reflection) = 0.5 \cdot [(\frac{n_1 \cdot cos(\theta_i) - n_2 \cdot cos(\theta_t)}{n_1 \cdot cos(\theta_i) + n_2 \cdot cos(\theta_t)})^2 + (\frac{n_1 \cdot cos(\theta_t) - n_2 \cdot cos(\theta_i)}{n_1 \cdot cos(\theta_t) + n_2 \cdot cos(\theta_i)})^2] \tag{3.8}$$

A random number between zero and one is drawn to decide which of the two occurs. If it is bigger than this value, then the ray is transmitted and Snell's law is applied to adjust the angles of the ray:

$$\theta_t = asin(\frac{n_1}{n_2} \cdot sin(\theta_i)) \tag{3.9}$$

In the case of reflection, it is assumed to be strictly in the plane of incidence, however a small random angle is added to the reflection angle to account for the fact that in reality, the surfaces will not be exactly straight. This is drawn from a Gaussian distribution, the standard deviation of which is set at the beginning of the code ($\frac{\pi}{10}$ was used).

When the ray traverses a plastic piece or sheet with non-zero absorption coefficient, its energy drops by a factor of

$$f = e^{-\kappa \cdot pl} \tag{3.10}$$

where κ is the wavelength dependent absorption coefficient and pl means path length within this element. This value is calculated by the attenuation function of CVcontent and Sheetcontent described in

the previous section. Of course, the array element of the counter array for the solar energy absorbed that corresponds to this lattice element receives the energy difference:

$$E_{abs}(i,j,k)_{new} = E_{abs}(i,j,k)_{old} + E_{ray} \cdot (1-f) \tag{3.11}$$

Or in the case of a sheet, this equation would have to be written differently, since the array is four-dimensional:

$$E_{abs,sheets}(i,j,k,s)_{new} = E_{abs,sheets}(i,j,k,s)_{old} + E_{ray} \cdot (1-f) \tag{3.12}$$

In the case of a ray entering a participating element with non-zero scattering coefficient, a function is called to perform this propagation. This function begins by drawing a random number to determine the first path length before the first scattering event using the inverse of the scattering coefficient, which is the average path length:

$$pl = -\frac{ln(rn)}{S(\lambda)} \tag{3.13}$$

where ln is the natural logarithm, rn a random number between 0 and 1 and S the Kubelka-Munk scattering coefficient at the wavelength λ that this particular ray has.

Then it checks if the projected coordinates of this scattering event lie within this participating element or not. If they do not, then in the case of a lattice element it calls the function that updates the ray position to the exit point of this element assuming it travels in a straight line through it. In the case of a sheet, it just does some basic trigonometry using the angle from the vertical, since the x and y coordinates of a ray are assumed not to change while traversing a sheet.

If the projected coordinates do lie within it, then it first calculates how much of the ray's energy is absorbed along the way there using equations 3.10 to 3.12, and checks if the ray energy has fallen underneath the threshold. If not, it calls another function that simulates the scattering event. However, the full Mie equations would have been too complicated to implement here, especially combining them with the forward scattering ratio values. Besides, pigment particles are far from perfect spheres, and there are enough papers out there pointing out how inaccurate Mie theory results are when it comes to scattering from irregular shapes (e.g. [64] or [65]). In reality, not only will different types of pigments have different shapes, but there can be variations in shape as well as particle size distribution even from one batch to the next of the same manufacturer. So it seems questionable whether a rigorous mathematical treatment of the scattering in this case would be worth the effort.

Therefore, it was decided to simply assume isotropic scattering in the direction perpendicular to the impinging ray, and to find a function to use as probability density function for the angle of deflection from the original direction θ. This function needs to have two parameters so that they can be calculated by imposing the two conditions it needs to fulfill, namely that the total area underneath the curve should be exactly one and that the area under the first half be equal to the forward scattering ratio. Several functions were investigated as possible candidates, but many were not suitable because they would not be non-zero over the full range of forward scattering ratios. $\frac{1}{x}$ has an infinite area underneath the curve

and in the case of an exponential function, it seemed impossible to solve for the coefficients. So in the end, the function

$$p(\theta) = a \cdot \theta^{-b} \tag{3.14}$$

was chosen since it can be used for any FSR. The coefficients can be calculated as:

$$a = -\frac{log_2 FSR}{\pi^{-log_2 FSR}} \tag{3.15}$$

$$b = 1 + log_2 FSR \tag{3.16}$$

Diffraction or partial absorption by the scattering particle are neglected. Only the ray angles are changed by the function that performs the scattering. Then, the function that propagates the ray along in the scattering element begins anew by randomly determining another path length and checking whether the new projected coordinates lie within it or without, and so on until it reaches a boundary of the element, or one hundred scattering events have occurred within it, in which case the ray gets fully absorbed by the element.

Whenever a ray reaches a vertical boundary between lattice elements, the code first checks if there are sheet objects in between them. If so, Fresnel reflection, Snell's law and scattering within the sheets are applied in the same way as for lattice elements. One difference here though is that since the sheets are assumed to occupy no space at all, the ray position remains fixed in space while the ray travels around within however many sheets there may be at this boundary. Of course, the thickness variable of each sheet is used to calculate the path length and the absorption from that, and if the sheet scatters, the scattering function keeps track of the height of the ray within the given sheet. On the one hand, this is to keep things simple by avoiding the possibility that the ray may reach the side of the sheet, but also to avoid the question of the height at which it would enter the neighbouring lattice element, which has no clear answer.

Scattering by the pure plastic itself is neglected, as well as absorption in the solar spectrum by it. While the latter assumption is very reasonable (only impurities absorb solar radiation, not pure perfect plastic itself), the former is not, as a quick look at the non-coloured plastics around you should confirm. However, the scattering takes place mostly at the boundaries between crystalline and non-crystalline regions, hence depends strongly on the degree of crystallinity, which will vary from one plastic batch to another. Due to the complexity and lack of clear literature data, this aspect was neglected due to lack of time to go into it. It would also cause the code to take even longer to run. Nevertheless, future simulations should include this aspect, as it might well affect the results, by preventing rays from penetrating far down.

All scattering and absorption spectra for the pigments were taken from [65], and digitized via the free software im2graph [55]. The data obtained in this research is from paints, so the pigment particles are embedded in a medium with a similar refractive index to plastic. However, there will obviously be some differences, especially regarding the much higher pigment volume concentration in paints, which affects the scattering as it can be influenced by interparticle distance. The authors also use a simple two-flux

model for calculating the coefficients from the measurement data and not a full four-flux model. They are aware that this may lead to some error, but their stated aim is only a general characterisation of each pigment to get an idea of which might be suitable choices for cool roofs. Also, in some of the cases where they tested different paints containing the same type of pigment from different manufacturers, there can be significant differences between the graphs. Nevertheless, given all the other uncertainties and estimates in this simulation, their data should be good enough, and it is anyway in most cases the only or best available data that can be found. Especially scattering coefficients are not widely reported, and especially not for the entire solar spectrum. Reflection spectra are more often published, but they can give only a general idea of where the wavelengths with the highest scattering coefficients are. More information than is given is usually required to calculate the exact scattering coefficients from those.

Figure 3.2 shows an example of what some possible ray paths might look like. The dots along their paths show the positions in space at which the ray position is determined. Argon gas is again shown in light blue, though it does not have a colour in reality. The sheets are drawn with a certain thickness, however it should be kept in mind that in the actual simulation they do not occupy any space in the lattice and the ray positions do not change when a ray is going through them, it is just that it is impossible to draw this.

- Ray 1 has almost the maximum initial angle that is allowed from the vertical, so it ends up first going into the empty area next to the participating lattice, where it gets reflected from the reactor wall. Then, it enters a sheet stack and goes through the first, red sheet. It is reflected from the next sheet via Fresnel reflection and re-emerges again from the top of the stack. Next, it enters a plastic granule which was initially iron oxide yellow (goethite) but some of its pieces have already converted to iron oxide red (hematite) by dehydration. Inside this, it gets scattered a few times and keeps losing energy via absorption until its remaining energy falls under the threshold value.

- Ray 2 is first reflected via Fresnel reflection from a plastic piece containing chrome yellow, then leaves the lattice area and is reflected from the reactor wall, reactor ceiling and again via Fresnel from a red plastic piece. It manages to go through a thin black sheet without losing all of its energy, where it gets scattered slightly. Then it enters a granule containing white pigment and gets scattered inside it several times, before escaping from it and entering a red granule, where it is scattered once before its remaining energy falls under the threshold.

- Ray 3 is very unlucky. It is reflected right away by a white granule at the top of the lattice, enters the window, gets reflected twice by Fresnel reflection at the top and bottom surfaces of the window before finally escaping and so its energy is counted as lost.

- Ray 4 is drawn to show how a ray can be scattered multiple times even within a sheet. In this case, it enters a sheet stack which has a thick sheet of calcium carbonate filler mixed with a little carbon black in the middle, where it is scattered twice and ends up reemerging from the top of the stack. It should be kept in mind though that in the simulation, only its height within the current sheet and direction angles change during this process, not the actual Cartesian coordinates of the ray.

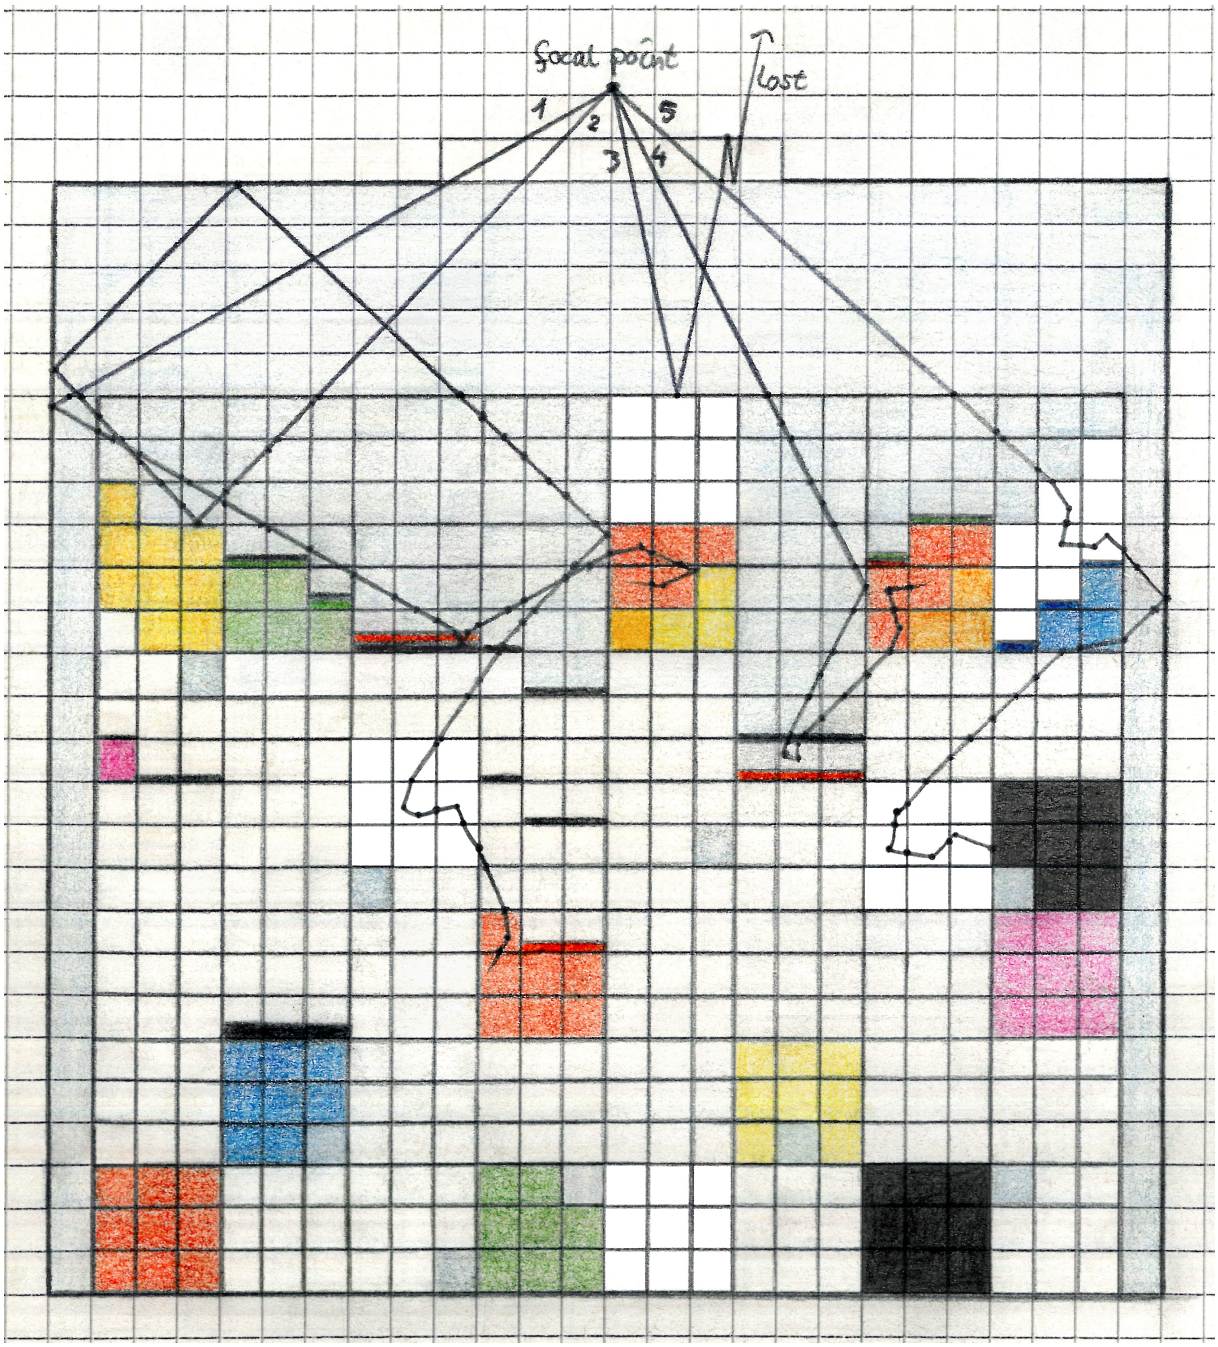

Figure 3.2: Cross sectional view of a simulation at some time instant, with examples of possible ray paths

- Ray 5 goes twice through white granules where it is scattered many times. It is supposed to show also how a ray may enter the argon filled side area and be reflected by the wall before reentering the lattice. Also, it ends up being absorbed instantly by a black plastic piece. Note that the same does not happen to rays 2 and 4, which manage to traverse thin black sheets. This is because black sheets may have been created from the decomposition of organic molecules and thus might be very thin (see for example the granule of dioxazine violet at the left side which has already almost completely decomposed). Or they were initially carbon black pigment but the sheet may have already partially evaporated. Or they might be mixed with another pigment that has a high

backscattering ratio. So, full absorption is quite certain in the case of a plastic piece starting off with carbon black pigment, but is not guaranteed in the case of sheets, hence they do not have the 'isblack' boolean that plastic pieces have.

3.1.4 Surface Corrections

Treating the surface correctly would require considerable effort, since its movement depends on the viscosity as well as the surface tension, which vary from one participating element to the next. For pure polymers, formulas for the variation of their viscosity with temperature can be found in the literature, however even just finding figures for how much the viscosity is affected by plasticizers is a difficult task. Not to mention figures for the effect of all the pigments and other additives. Experiments would have to be performed to determine all the necessary parameters to then perform a rather complex fluid mechanics simulation, which would have been well beyond the scope of this work. Besides, the way in which this simulation is set up with cubic lattice elements and sheets means there is no smooth surface that could gradually move. The only allowed movement is jumps from one lattice position to another. Integrating real fluid flow into this setup would be very complicated.
However, something had to be done to prevent the surface from becoming too unrealistically uneven. One can well imagine a granule with carbon black at the surface, which absorbs much more radiation than its neighbours and thus dissolves quickly, leaving behind a sheet of carbon which keeps on absorbing a lot of radiation and thus burns itself a deeper and deeper hole. In reality, at some point the neighbouring liquid will close the hole over when it becomes too deep. This may actually become a big problem, because the pigments that absorb the most radiation will all eventually sink down and leave a layer of titanium dioxide and/or calcium carbonate at the surface, which will make it difficult for non-UV rays to enter and thus lead to high losses through the window. However, there will be some carbon left over from the decomposing plastic and other additives, and if the surface temperature becomes high enough, this might react with the titanium dioxide to form titanium carbide. This would be fortuitous since titanium carbide is a very efficient light absorber, as well as an economically valuable material, though the purity would be questionable. According to [66], formation of TiC begins at around $1050°$C in the case of microwave synthesis and $1200°$C in case of conventional heating. It is not clear from the paper whether the lowering of reaction temperature due to enhanced sintering from the use of microwaves would also apply in the case of solar radiation. Experiments would have to be performed to determine the necessary temperature as well as reaction kinetics, since the published research assumes differing conditions to this case.
Therefore, since something had to be done to not allow unrealistically steep surface gradients, it was decided to allow the user to set two parameters that control this behaviour. One is the time step with which the surface is periodically checked for necessary corrections, and the other is the minimum gradient that triggers a correction. Gradient here means simply the height difference in lattice constants between direct neighbour columns. Every time the next flow correction time is surpassed, the gradients in both horizontal directions are calculated from the array that keeps track of the surface height. The

steepest gradient is determined and then only one of the gradients which has the same value as the steepest is corrected by moving one plastic piece horizontally over from a direct neighbour, to fill the lowest empty spot of the hole. Of course, that position will not be completely empty, there will be an argon object there. This object is moved to the top of the column from which the plastic piece was taken, after this column has been shifted downwards to fill the vacant lattice position. The whole process is thus a circular swapping operation involving three or more lattice elements. In case there are sheet objects inside that column, they move down with it, but in case there are some at the bottom, they do not move sideways.

After each shifting operation, the function that checks the surface for exposed bubbles runs. This is necessary because if an argon bubble is exposed by the shifting and it happens to be right next to a hole, then in the next iteration the code could end up trying to fill one hole with another hole.

Then, the surface gradients are recalculated for the entire surface, and if there are still gradients as steep as the value set in the Dimensions file, then one of the steepest is corrected again and so on, until no more gradients exceed that value.

Apart from the gradient corrections, the function that checks for exposed bubbles also runs straight after every occasion where a plastic piece has finished its plastic pyrolysis process and an argon object is created. If it finds an argon bubble horizontally right next to the bottom of a hole with a plastic piece somewhere above it, then it shifts the column above the bubbles downwards, placing the displaced argon bubbles above the column to fill the lattice elements left vacant there.

3.1.5 Plastic Decomposition

The numerical implementation of the models and data for this process described in the previous chapter was performed as follows. Each plastic object has a variable called alpha which represents the degree of conversion in percent, and is always initialised to zero in the constructor. The threshold temperature at which the plastic begins to consume heat for pyrolysis was estimated from the DSC graphs, and came out as 340°C in both cases. As soon as this temperature is exceeded, the conversion begins, albeit rather slowly at first. The first thing the code does in such a case is to check whether the input energy supplied would raise the temperature by more than 10K if the plastic was not consuming any of it and no other phase transitions occur. If so, then it sets a boolean to true that tells it that this is a case requiring integration of the formula with respect to temperature as well as conversion, because the temperature, being in the exponential, influences the conversion rate significantly. In this case, it calculates the amount of energy that would be needed to raise the temperature by 10K and uses the ratio of this to the total input energy to calculate the time it would take from the time step supplied to the function from the main function, assuming the energy input stays constant in time during this time:

$$\Delta t_{10K} = 10K \cdot m \cdot C_p \cdot \frac{\Delta t}{\Delta E} \tag{3.17}$$

where m is the total mass of this plastic piece (including additives), C_p its overall heat capacity, Δt the time step given to the function from the main function and ΔE the remaining energy still to be used.

Then, it determines the rate of conversion from the general formula using the present values:

$$\frac{d\alpha}{dt} = A(\alpha) \cdot e^{-E_a/RT} \cdot f(\alpha) \tag{3.18}$$

where $f(\alpha)$ is the reaction model, R the universal gas constant and T the current temperature. Each PurePlastic object knows its values for A, E_a, heat consumption and which reaction model it should use, so the function performing the conversion then calls functions of its object with α and temperature as input to receive the values it needs.
From this, it determines the time it would take to convert 1% at the current rate:

$$\Delta t_{1\%} = (\frac{d\alpha}{dt})^{-1} \tag{3.19}$$

Then, the lowest of all three is chosen as actual integration time step:

$$\Delta t_{min} = min(\Delta t_{10K}, \Delta t_{1\%}, \Delta t) \tag{3.20}$$

The current rate of change is then multiplied by the integration time step to obtain the amount of conversion during this time, which is added to α:

$$\alpha_{t+\Delta t_{min}} = \alpha_t + \frac{d\alpha}{dt} \cdot \Delta t_{min} \tag{3.21}$$

unless of course α would exceed 100% conversion, in which case it is set to exactly 100.
The new mass of the pure plastic is calculated from the new conversion:

$$m_{new} = m_0 \cdot (1 - \frac{\alpha}{100}) \tag{3.22}$$

where m_0 is the initial mass of pure plastic before the start of conversion, which is stored in a separate variable from the current mass.
The remaining time is calculated simply:

$$\Delta t_{new} = \Delta t_{old} - \Delta t_{min} \tag{3.23}$$

Also, the remaining available energy and new temperature are determined. For this, it is necessary to know how much energy the plastic consumed in going from α_t to $\alpha_{t+\Delta t_{min}}$. As mentioned previously, each PurePlastic object knows the specific heat consumption for each percentage point. So, at this point the Plastic object calls a function of its constituent PurePlastic object which interpolates between the values it holds to return the specific heat consumption, which is then multiplied by the initial mass m_0 to get the energy consumption E_{cons}.
In the case of no temperature integration, the adjustment is quite simple:
If E_{cons} is less than ΔE:

$$\Delta E_{new} = \Delta E_{old} - E_{cons} \tag{3.24}$$

Or, in case the energy consumption is more than what is available, the temperature drops:

$$T_{new} = T_{old} - \frac{E_{cons} - \Delta E_{old}}{m \cdot C_p} \tag{3.25}$$

and ΔE is set to zero and the function concludes.
In the case of temperature integration, the adjustment is handled differently. Before the integration starts, the value of $\frac{\Delta E}{\Delta t}$ is stored in a variable. This represents the amount of available energy per unit time, assuming that it is uniformly distributed over the available time. It will be called AER, for available energy rate, since it needs a name. Then,

$$\Delta E_{new} = \Delta E_{old} - AER \cdot \Delta t_{min} \tag{3.26}$$

$$T_{new} = T_{old} + \frac{AER \cdot \Delta t_{min} - E_{cons}}{m \cdot C_p} \tag{3.27}$$

In both cases, the new ΔE is also used to update the temperature projection, in case the pyrolysis reaction stops due to lack of time or completes without using up all the available energy.

If there is still remaining time, energy and plastic left, the loop will execute again and perform the same calculations with the new values and new rate of conversion for α.
Ideally, it should be also taken into account that some additives might decompose or evaporate while the plastic is decomposing, but due to time constraints this further complication had to be neglected. The inaccuracy should not be too big, given that it is only dioxazine violet that would be affected.
Char residue is also neglected for the same reason, which is no problem in the case of PP since it does not char [57], and for PE it is very small [16].

3.1.6 Code Structure

The code consists of a main function and several additional functions in separate files that it uses for the ray geometry, moving around of objects within the reactor and the reading in of the input data. All simulation constants can be set by the user in Dimensions.hpp, and all other parts of the code include this file and thus know about these values.
Additionally, there are files for the various objects: Material and its extension PurePlastic, CVcontent and its extensions Argon and Plastic, and Sheetcontent.
The code requires quite a lot of input files in order to run, although the amount depends on how many additives, pigments and their decomposition products the user wants to include. He has to set the number of each of these in the separate Dimensions file which contains all the adjustable simulation constants. The code then reads that many rows of the Excel files which contain the material properties: PropertiesPlastics.xls, PropertiesAdditives.xls, PropertiesPigments.xls and PropertiesDecompProducts.xls. The first column of those files contains the name, which the code then adds to the file names from which it tries to read the absorption and scattering spectra. So if for example carbon black is called ”Carbon_Black_7” in the Pigments.xls file, then there needs to be a corresponding file called ”Absorption_Spectrum_Carbon_Black_7.csv” and another called ”Scattering_Spectrum_Carbon_Black_7.csv”

in the source folder. Since reading of Excel files is not part of standard C++, the code also requires the free version of the libxl library from libxl.com. It was done this way so that the user has a good overview of the materials and their properties and can modify them easily. Putting them all in csv files instead would have been messy. Additionally, it requires csv files with the solar spectrum and absorption and reflection coefficients of the window as function of wavelength. All units are, if need be, converted to SI units immediately after being read in, to avoid possible errors due to unit conversion.

Material

The properties read in are then used to instantiate one object each of the class Material, which has the fields shown in Table 3.1.

None of these properties change in time, so each plastic object simply holds copies of the pointer to

Qualifier	Type	Variable name	Description
const	string	MaterialName	Name of this material as in input file
const	double	M	molar mass
const	double	k	conduction heat transfer coefficient
const	double	density	density at room temperature
const	double	Cp	heat capacity at constant pressure
const	double	FSR	forward scattering ratio, ratio of rays scattered into the forward hemisphere to total rays scattered by a particle
const	double	Tm	melting temperature
const	double	Em	melting energy (latent heat of fusion)
const	double	Te	temperature of evaporation, if applicable
const	double	Ee	energy of evaporation, if applicable
const	double	Tdecomp	temperature of decomposition, if applicable
const	double	Edecomp	energy of decomposition, if applicable
const	array<double, lambdaMax>	n	refractive index as function of wavelength
const	array<double, lambdaMax>	scattering	scattering coefficient as function of wavelength
const	array<double, lambdaMax>	absorption	absorption coefficient as function of wavelength
const	string	Pdecomp1	name of decomposition product 1, if it exists
const	string	Pdecomp2	name of decomposition product 2, if it exists
const	double	nPdecomp1	number of moles of decomposition product 1
const	string	nPdecomp2	number of moles of decomposition product 2
	bool	isblack	is this object black or not, if so it absorbs all incoming rays immediately

Table 3.1: Variables of class material

the material objects contained in it. In future versions of the code, the temperature dependence of those

properties could be taken into account, although often this information is not easy to come by.
Each material object contains all the refractive index values for each wavelength that is included in the solar reference spectrum. If a value is given in the properties file, then this one value is put into all the wavelengths by the constructor. If instead the file has a value of “-1” in place of the refractive index, then the code will notice and try to read the wavelength dependent refractive index values from a csv file called “Refractive_Index_MaterialName.csv”. It then calls a different constructor which takes an array of values for this property.
If a material sublimates, a “-1” has to be put in the properties file for its melting temperature. The code checks for such values when reading in those temperatures, and if it encounters one, it puts a NaN instead of a value. Similarly for evaporation or decomposition temperatures, since most materials can have only one of the two (unless it decomposes over a range of temperatures and evaporates during the process, but gradual decomposition is not included in this code, except for plastic, which is handled separately). The function that determines the temperature change of a plastic object, when searching for phase transitions always checks that their value is not NaN, and ignores those.
PurePlastic includes the same basic material properties and overloaded constructor, but also holds arrays for the data specific to the plastic degradation reaction: The pre-exponential coefficient, activation energy and heat consumption as function of conversion at each percentage point. Its constructor has a body, which sets the value of an integer variable that determines which function is used for the reaction model based on the name of the plastic. This function has to be manually inserted by the user if a new plastic type is added, or the code will throw an exception. There seems to be no obvious user friendlier way of doing this, since in this case the actual functions are given, not tabulated data that could be read in from a file. Theoretically the functions could be converted to tabulated data, but it would lead to needless loss of accuracy.
Each Plastic Object has an array with the pointers to the Material objects that it consists of, as well as the mass of each and the resulting mass and mol fractions and all the combined properties resulting from them.

CVcontent

CVcontent stands for cubic volume content and is an abstract superclass, or pure virtual base class, whatever you prefer to call it. It is not possible to instantiate this class, however it is needed so that a 3D array called PCV can be created with pointers to this base class type, which in reality are objects of its two derived classes. Since the reactor is cylindrical but the content arrays rectangular, those array elements which do not take part in the simulation remain as null pointers. A 2D array of booleans is created right in the beginning that keeps the information about which columns are included and which are not.
It also has another boolean overlay called isCVargon that keeps track of which array element is which. Initially, it was envisaged to have the argon gas bubbles participating in the heat transfer, but after realising how this leads to a very small time step for it, they had to be excluded from doing so. So the current code still contains artefacts from this original setup, like heat capacity and thermal conductivity

of the argon gas and functions within the Argon class to recalculate the gas properties. Most of these could be removed, however at present, CPU time is much more of an issue than RAM, so they were left there in case in future there may be a desire to include them again. If the argon objects were completely removed and replaced with null pointers, it would be necessary to modify the main function though, since it uses the refractive indices of the elements in the old and new volume when applying Snell's law, and for this it calls the base class function getn() of each element. The argon implementation of this function always returns 1.000281. If this was removed, the main function would always first have to check if a lattice element is argon or not and then set n itself.

The variables of the superclass are shown in Table 3.2. None of them has any const qualifier, since all the properties change in time.

It also has the function stubs shown in Table 3.3.

Type	Variable name	Description
double	density	density at room temperature
double	k	conduction heat transfer coefficient
double	Cp	heat capacity at constant pressure
double	m	total mass
double	T	temperature in Kelvin
bool	isblack	is this object black or not, if so it absorbs all incoming rays immediately
uint	NMs	number of materials present in this object

Table 3.2: Variables of class CVcontent

Probably most of these are no longer necessary as stubs in the superclass and could be moved to

Qualifier	Return Type	Function Name	Input variables	Description
virtual	double	getn	int lambda	Returns value of refractive index of this object
virtual	double	getAbsorption	int lambda	Returns value of absorption coefficient
virtual	double	getScattering	int lambda	Returns value of scattering coefficient
virtual	double	getFSR	int lambda	Returns value of forward scattering ratio
virtual	bool	givedEtoObject	double &dE, double dt, const array< Material*, nAdditives > &Adds, const array< Material*, nPigments > &Pigments, const array< Material*, nDecompProds > &DPs, array<double, 660> &GasT	Adds energy dE to this object
virtual	double	attenuation	double lambda, double pathlength	Calculates value of attenuation by absorption of a ray travelling distance pathlength through this object

Table 3.3: Functions of class CVcontent

the Plastic subclass. This subclass currently has the additional variables shown in Table 3.4, and the additional two functions shown in Table 3.5.

calculateProperties() is called once by the constructor, and then by recalculateProperties(). It takes the

Type	Variable name	Description
array<double, lambdaMax>	n	refractive index
array<double, lambdaMax>	absorption	absorption coefficient as function of wavelength
array<double, lambdaMax>	scattering	scattering coefficient as function of wavelength
array<double, lambdaMax>	FSR	forward scattering ratio as function of wavelength
vector< Material* >	mat	vector of material objects, plastic always first, then pigment if applicable, then other additives
vector< double >	compmass	mass of each component of the composition, plastic always first
vector< double >	compmass0	copy of above vector, with initial mass so that progress of an ongoing decomposition/evaporation can be tracked
vector< double >	massfrac	mass fraction of each component of the composition, plastic always first
vector< double >	molfrac	molar fraction
vector< double >	fracunmelted	fraction of each material that has not yet melted
vector< double >	xV	product of mol fraction and molar volume, used in calculation of refractive index
double	xVtot	Sum of xVs
double	alpha	fraction of total pure plastic that has been already converted (pyrolysed and vaporised)

Table 3.4: Variables of class Plastic

Qualifier	Return Type	Function Name	Input	Description
	void	calculateProperties		function to calculate overall properties of the composition
	void	recalculateProperties		function to recalculate total mass/mols and mass/mol fractions after a change in mass

Table 3.5: Functions of class Plastic

values of each constituent material to put together the composite value of the plastic object, depending on their respective mass and molar fractions. The molar fractions are only needed in the calculation of the composite refractive index n, for which the Newton formula is used:

$$n_{tot}^2 = \frac{\sum_{i=1}^{m} x_i V_i n_i^2}{\sum_{i=1}^{m} x_i V_i} \tag{3.28}$$

where x is the molar fraction and V the molar volume:

$$V = \frac{M}{\rho} \tag{3.29}$$

where M is the molar mass and ρ the density. This formula applies to well mixed liquids, so it may not actually apply so well in the case of big pigment particles suspended in a plastic matrix.
The heat capacity and thermal conductivity are combined based on their mass fractions via a simple arithmetic average. The absorption and scattering coefficients are combined based on their respective volume fractions though, and the FSR based on how much each of the materials scatters at each given wavelength. Each material provides only one value for its FSR, since the paper from which all those properties are taken only estimates one, but the composite depends on wavelength because the scattering coefficients do.
recalculateProperties() is called whenever one of the components of the compmass vector changes due to phase transitions. It calculates the new sum total and the new fractions, and then calls calculateProperties(). However, since four of those properties depend on wavelength, the latter involves a lot of calculations, so the whole simulation would take forever if it was called after every tiny change. This is one reason that there is a copy of the compmass array called compmass0, so that the current mass can be compared with its original value and the function only gets called when a multiple of 5% relative change has been exceeded. The other reason is that most DSC data is normalized with respect to the initial sample mass.
The implementation of givedEtoObject of the Plastic subclass is called every time that there is energy to add to the plastic pieces of the main simulation, for each plastic piece even if the change in energy is zero. This is because it might be in the plastic decomposition stage, and the pyrolysis reaction continues even without further energy input as long as the temperature is over the decomposition temperature. Of course, the reaction consumes heat so then the temperature drops accordingly, which means that the reaction rate goes down and it might stop if the temperature falls below the threshold.
The function works by projecting the temperature that the object would reach with the given energy input, using the standard formula

$$\Delta E = m \cdot C_p \cdot \Delta T \tag{3.30}$$

where m is the total mass and C_p the heat capacity at constant pressure. Then, it goes through the materials checking all their phase transitions to see whether they would occur before this projection. If there are more than one in this temperature range, then the lowest is chosen first, irrespective of type. If necessary, it then brings the temperature up to the transition value using the inverse of equation 3.30 to decrease ΔE before using what is left of ΔE to convert material. The special case of plastic pyrolysis was described previously in section 3.1.5. For all other cases, it first checks how much energy would be needed to completely melt, evaporate or decompose the material i:

$$E_{m,needed} = fracunmelted(i) \cdot E_m(i) \cdot m_i \tag{3.31}$$

$$E_{e,needed} = E_e(i) \cdot m_i \tag{3.32}$$

$$E_{decomp,needed} = E_{decomp}(i) \cdot \frac{m_i}{M_i} \tag{3.33}$$

where $E_m(i)$ is the specific melting energy of material i in $\frac{J}{kg}$ or latent heat of fusion, $E_e(i)$ is the specific energy of evaporation/vaporization of material i in the same units and $E_{decomp}(i)$ is the specific decomposition energy in $\frac{J}{mol}$. If not enough is available, it calculates the fraction that would be melted or evaporated or decomposed, and in the latter case also checks for possible decomposition products. If this material has decomposition products, it first checks if they are already present in this plastic, and if not adds them. Then, the masses of these decomposition products go up correspondingly. For the first decomposition product DP1 the formula would be

$$\Delta m(DP1) = nPdecomp1(i) \cdot \frac{\Delta E}{E_{decomp}(i)} \cdot M(DP1) \tag{3.34}$$

If there is enough available, then all the remaining material is converted, ΔE is adjusted correspondingly and the loop starts again. It uses what remains of ΔE to project a new temperature and checks again for other transitions, and so on until it is either consumed by a phase transition or no more transitions are found, in which case the temperature is set to its projection and the function is finished. In case the incoming energy is negative and the plastic is not in the decomposition process, the same happens but in the opposite way. The projected temperature will be lower than the current one and it checks only for possible solidification, not for any unlikely condensation or reverse decomposition events. The function usually returns false, except if the plastic has fully decomposed, then the main function knows that it has to delete the object and possibly create a sheet object from the remaining additives.
The function attenuation is a very simple one, implementing equation 3.10. Its name is actually a bit of a misnomer, since attenuation is defined as the sum of absorption and scattering, so it should be called absorption instead. But then the array of absorption coefficients would have to be renamed too, and the variable that takes its return value in the code.

Sheetcontent

Sheetcontent is very similar to CVcontent, with just a few differences: It has only one constructor, which takes a plastic object whose pure plastic has finished pyrolysing and leaves behind some additives, and copies its values into the equivalent variables of its own. It has an additional thickness variable, which is calculated from its surface area, mass and density. Its density variable is the real density, as calculated from the densities of its constituent materials, whereas for CVcontent it is the remaining mass divided by the fixed volume. It obviously lacks the extra code for plastic decomposition, and the givedEtoObject function returns true if all materials have disappeared. It also takes a boolean by reference that it sets to true if a material has just disappeared and only one is left, so that the main function knows to check if there happen to be any other sheets just above or below which have the exactly same material, so they can be merged. This is important because the thinnest sheets determine the time step for the whole heat transfer part, and this can slow down the whole simulation significantly.

Main Function

The main function begins by declaring all its working variables. Then it reads in the values it needs from the input files as described previously. The solar reference spectrum is given as tabulated data for various wavelengths. It reads in the two columns with wavelength and DNI value, and then transcribes them into a single column array which has the wavelength in nm as index value. Those index values for which no values are available remain as NaN. The spectrum goes up to 4000nm, with only every fifth value given towards the end, so this does mean that a lot of NaNs have to be carried around. It makes it much easier to access wavelength dependent values like refractive index, absorption, scattering coefficient and forwards scattering ratio though. Otherwise, an array with two columns would have to be carried around and the code would have to go through it and search for the index position of the required wavelength every time. Writing n[lambda] instead is not only more convenient, but also saves computation time. The other variables may be given for different wavelengths in their input files, but the function that reads them in converts them to the same one-column array and interpolates the values so that all wavelengths present in the DNI array are also present in the other types. The only requirement for the input spectra is that the wavelength values have to be integers. If the values do not cover all of the wavelength range of the simulation, then the first and last given ones are inserted for the missing ones.

If a required file is not present, the code throws an exception and terminates. Else, it proceeds to set up the reactor. It begins by determining which lattice columns take part and which do not. Then it iterates through all the partaking granules and draws random numbers to give them a plastic type and maybe some additives, all depending on the prevalence values set in the input files. Each granule can only have one pigment, but any amount of additives up to the maximum. There is an additional restriction, namely that in case the pigment chosen is Carbon Black or Titanium Dioxide, then no UV absorber will be included since those pigments already absorb the UV well enough. A simple linear interpolation is performed between the minimum and maximum additive content and a random number. For each plastic piece of this granule that does not happen to be randomly determined as Argon, it then passes to the Plastic constructor the pointer to the PurePlastic Object, a vector containing the pointers to the included materials and another with the determined mass fractions. The constructor calculates the plastic fraction from the sum of the additive fractions and then determines the combined properties of the resulting object.

There is a 2D array that keeps track of the surface height of each column, via the index value of the highest plastic piece, and this surface height is adjusted after the initial setup in case there are argon bubbles at the surface.

Once the setup is in place, the progression in time begins. It is contained within a while loop that executes as long as there are still plastic pieces left. A counter variable keeps track of their amount. The heat transfer part described in section 3.1.2 comes first, since it determines the minimum time step. It is followed by the Monte Carlo ray tracing, as described in section 3.1.3 but only once enough energy has accumulated for there to be at least one ray with the minimum energy set in the Dimensions file. If not,

it is skipped and the conduction heat transfer energy is added straight to the objects, using the function givedEtoObject described in section 3.1.6.
Else, the energy from conduction and solar input is added together before calling this function with it. If in this process the plastic finishes decomposing, the main function deletes the object. If no additives are left over, but some energy is, then that energy is added to the sheet or plastic piece below, depending on what is there. If there are leftover additives, then the plastic object is first handed to the constructor of Sheetcontent before being deleted. The resulting sheet goes to the top of the sheet stack in between this cubic volume and the one below, if there are already sheets there. Else it becomes the first. However, in case there are bubbles of argon gas right below, then the sheets cannot be left suspended in the air, so they fall down through the bubble onto the top of the lowest plastic object, or the reactor floor. In case there are already sheets there, then they just stack up on top of each other.
If the decomposing plastic object was at the surface, an object of argon gas is created in its place. If it is below the surface, then the entire column above it is shifted downwards, both plastic and sheet objects. An argon gas object is then created at the surface. In both cases, the surface check for exposed bubbles runs afterwards, as described in the last paragraph of section 3.1.4.
The sheet objects are stored in a 4D structure: Three dimensions of std::arrays and one dimension of a vector. The vectors all start off with zero elements as there are no sheet objects in the beginning. Whenever a new one is created, the vector function push_back() is called to add an element to the vector at this position in space. Whenever a sheet object decomposes fully, erase() is called to delete the object and also the vector element that held it, and any remaining sheets automatically shift to fill the gap. A major difference between the arrays holding the plastic objects and those holding the sheets is that the former stores pointers, but the latter the actual objects since there is no need for polymorphism in that case. Again, leftover energy is passed on to whatever is below if possible. Whenever a new sheet is created, a sheet is deleted or its second material has just disappeared, the code checks if this has caused two identical sheets with only one material to lie on top of each other, and if so merges them together. This could be theoretically also done for cases where two sheet objects happen to have two identical materials, but it is much more unlikely to occur, would involve a much higher coding effort and computational time for checking, and would not significantly improve the time step issue because sheet objects with two materials are unlikely to be very thin. This issue is alleviated much more in a different way: Whenever their thickness falls below a certain value, currently one thousandth the lattice constant, they get merged into the sheet or plastic below them. In case there is an argon bubble or the reactor bottom below, then they are merged upwards into the sheet or plastic above, if present. In the case of an isolated sheet at the reactor bottom, nothing can be done, except to hope that it will not last long.
Once all the energy has been distributed, the surface flow correction part described in section 3.1.4 may run, if the time set for it has been exceeded.
The user also needs to set a time interval with which summary information is written to a file. Writing this information for every time step would result in huge data files with needlessly precise values, and would slow down the code too since it calculates the maximum, minimum and average temperature of plastic and sheet objects each time. Once the loop over time has finished, it checks if the last values were not

written to file, and if so does that. Then it goes through the remaining sheet objects and calculates the total remaining mass and average temperature of each additive, outputting this information to the screen and writing it to another file with overall simulation information. To this file, it also writes the information about how much energy it would take to fully evaporate each additive in the beginning and at the end, and their sum totals.

3.2 Verification and Validation

Since no similar research or models exist, it was not possible to perform much verification or validation. Obviously, many debug statements were inserted in the code to check that certain things that should not happen do not happen, and to inspect the variables in the Eclipse debugger. Many initial flaws in the code led to the ray position going out of sync with the lattice array indices, or the boolean overlay going out of sync with the actual position of argon objects, and this quickly leads to runtime errors. So the fact that the code does run now without runtime errors means that these aspects work correctly now, and since it is a huge stochastic simulation, even the most unlikely but possible events had to be accounted for since they will all happen eventually.

To check that no energy is going missing anywhere in the code, it projects the amount of energy it would take to evaporate all the plastic including additives and their decomposition products in the beginning, and recalculates it again at the end. This figure compares well with the total incoming solar radiation minus losses. Not exactly of course, but that is because the plastic fragments evaporate at different temperatures, so it is impossible to predict exactly how much energy will be needed to heat up the plastic.

Chapter 4

Results

The maximum necessary window size to let in almost all the necessary flux was read from Figure 2 of [67]. From the intercept of the lines with the vertical axis, it can be seen that at 7cm and 8cm aperture radius, pretty much all of the incoming radiation is let through, while it drops to 95% at 6cm and 85% at 5cm. The concentrating system used for the solar furnace at the Paul Scherrer Institute catches sunlight with a heliostat of 51.8m^2 surface area, which is not too different from the surface area of 56.65m^2 used in this simulation, so the results from this paper should apply fairly well to this simulation. Therefore, based on this information it seemed unnecessary to perform a sensitivity analysis for the window size, since the losses from radiation escaping through the window are usually only around 1-4% with the maximum necessary window radius of 7cm. Even just going down to 6cm would mean losing 5% of the solar input, which clearly could not possibly be worthwhile. From the same figure it is also possible to get an idea of how insignificant the losses from thermal radiation should be at the comparatively low temperatures at which plastic pyrolysis proceeds.
The window thickness was taken from [38].

4.1 Base Case Results

A reactor of 1m diameter, 1m height and 0.8m plastic height is assumed to correspond with the full dish area of 56.75m^2. The standard granule size is always 3mm, regardless of the reactor dimensions, since this determines the effectiveness of heat transfer between coloured and clear plastic. Unfortunately, simulating a realistically sized reactor like this with such tiny granules would require a huge amount of RAM, way more than any normal computer has. So, the reactor dimensions are scaled down by a scale factor set in the Dimensions file. The CTTC computer on which this simulation was run has enough RAM to start a simulation with a scale factor of 0.06, however it would probably take weeks to finish. So, 0.025 was chosen since it runs in a few hours. 0.03 was run once too, but took a long time. 0.035 was attempted as well, but after running for a full day and night, the internet connection temporarily cut out and the session was lost, so only partial results are available for this scenario.
Window radius is 7cm times scale factor, and window thickness 5mm times scale factor.

Although the code is set up to handle granules made up of several plastic pieces, given the small amount of total lattice elements that could be simulated in a realistic time, subdividing the granules did not seem worthwhile, so the parameter for amount of pieces per granule is always set to 1.
The surface is checked for gradients every 10s, and corrects those with five or more difference in surface height.
The wavelength range goes up to 4000nm, since that is the maximum included in the reference spectrum used.
The probability for an element to be a hole is set to 3%, although in the case of only one plastic piece per granule the actual probability ends up being higher, which should be corrected in future versions.
The maximum timestep was set to 0.001s. Obviously, even if the heat transfer part allows it, not too much radiation should be distributed at once before checking what happens to the objects that absorb that radiation. Since a large amount of radiation enters the reactor and the plastic pieces are quite small, even this seemingly small timestep could lead to large jumps in temperature of some of them. However, simulations with 0.0001s take quite a long time to run, whereas 0.001s is still reasonably fast. These two cases are compared later on. Besides, as soon as the first sheet objects begin to appear, the time step anyway becomes smaller than this so the maximum no longer matters.
Progress information is outputted to the console every 30s and written to file every 5s.

The complete output written to SummaryResults.txt is shown here. The values for the pigments are displayed in Table 4.1 and Table 4.2 for greater clarity.
Total number of partaking elements are 259, of which 237 contain plastic, giving a fill fraction of 91.5058%.
Initial plastic mass is 0.0066056kg, of which 0.00611667kg are pure plastic, giving an overall additive fraction of 7.4018%.
It should take around 3208.26J to pyrolyse all of the 0.00165931kg of pure HDPE.
It should take around 3879.81J to pyrolyse all of the 0.00287983kg of pure LDPE.
It should take around 2197.22J to pyrolyse all of the 0.00157753kg of pure PP.
...
Finished with simulation at time t = 12575.8, which corresponds to 209.597min.
Total incoming energy was 9898.34J, of which 3.50402% was lost.
Total remaining mass in reactor is 0.00046999kg.
It would have taken a further 1411.13J to fully evaporate all the remaining mass.
9551.5J of solar radiation was absorbed, and the difference between initial and final remaining energy needed is 9527.42J.
The time varying results from the file SimulationResults.csv are now displayed in graphical form, as plotted in Matlab.
Figure 4.1 shows the time evolution of the number of plastic pieces. Note that this is not equivalent to the total mass of plastic, though future versions of the code could calculate and output that too. The number total is also important though since it gives an idea of whether the differently coloured granules pyrolyse together or separately. As Figure 4.1 shows, the number is constant until 150 minutes, where it begins to decrease rapidly and quite smoothly, eventually becoming very steep and finishing at around

Material	Initial Mass [kg]	Initial energy needed [J]	Percent of additive mass
UV531	$5.31568 \cdot 10^{-6}$	7.25002	1.0872
Calcium Carbonate	$3.98758 \cdot 10^{-4}$	1439.49	81.5568
Iron Oxide Yellow	$2.97402 \cdot 10^{-6}$	6.19687	0.608266
Chrome Yellow	$5.26084 \cdot 10^{-6}$	1.84228	1.07598
Iron Oxide Orange	$7.29369 \cdot 10^{-6}$	12.3522	1.49176
Iron Oxide Red	$1.75534 \cdot 10^{-5}$	27.7556	3.59015
Phthalo Blue	$2.71085 \cdot 10^{-7}$	0.147654	0.0554443
Phthalo Green	$7.252 \cdot 10^{-7}$	1.06981	0.148323
Dioxazine Violet	$4.8838 \cdot 10^{-7}$	0.934975	0.0998868
Carbon Black	$1.46176 \cdot 10^{-5}$	53.5021	2.9897
Titanium Dioxide	$3.56751 \cdot 10^{-5}$	102.708	7.29651

Table 4.1: Initial values of additives and their decomposition products

Material	Final Mass [kg]	Final energy needed [J]	Average T [°C]
UV531	0	0	N/A
Calcium Carbonate	$3.98758 \cdot 10^{-4}$	1265.95	532.901
Iron Oxide Yellow	0	0	N/A
Chrome Yellow	0	0	N/A
Iron Oxide Orange	0	0	N/A
Iron Oxide Red	$2.44064 \cdot 10^{-5}$	24.9364	812.196
Phthalo Blue	0	0	N/A
Phthalo Green	$2.9499 \cdot 10^{-7}$	0.313236	513.136
Dioxazine Violet	0	0	N/A
Carbon Black	$1.51609 \cdot 10^{-5}$	47.17	799.457
Titanium Dioxide	$3.13519 \cdot 10^{-5}$	72.744	753.54
Calcium Oxide	0	0	N/A
Copper	$1.76767 \cdot 10^{-8}$	0.0176764	1119.4

Table 4.2: Final values of additives and their decomposition products

210 minutes. This is a good result, because it means that the heat transfer is efficient enough so that no transparent plastic gets left over for long.

Figure 4.2 shows the time evolution of the maximum, minimum and number average temperature of the plastic pieces. It confirms the efficiency of the heat transfer, since the maximum and minimum remain close together. It also shows how the maximum does not fluctuate greatly, which means the maximum time step is not unreasonably large. Since the incoming solar energy is constant, unsurprisingly the average temperature goes up linearly in the beginning. Then, there is a small dent when the melting of the plastics occurs, and the minimum clearly shows the three melting temperatures of the three plastic types and their degree of crystallinity. HDPE takes longest since it has the highest crystallinity.

At around 100 minutes, the temperatures reach 340°C and the plastic pyrolysis reaction begins. As this consumes a lot of energy, the slope of the curves changes noticeably. It is lowest at the beginning of this process because the plastics consume the most energy per percent conversion in the beginning, as Figures 2.4 and 2.5 in Section 2.2.4 show.

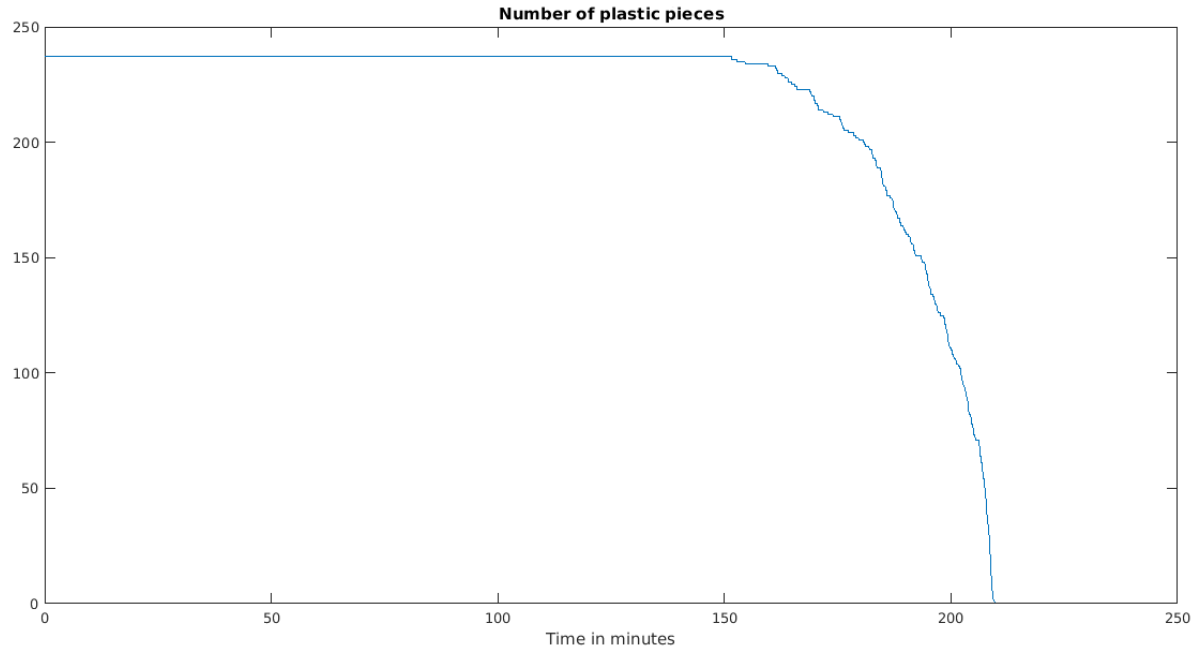

Figure 4.1: Time evolution of number of plastic pieces

Figure 4.3 shows the total number of sheet objects, whose creation unsurprisingly correlates with the

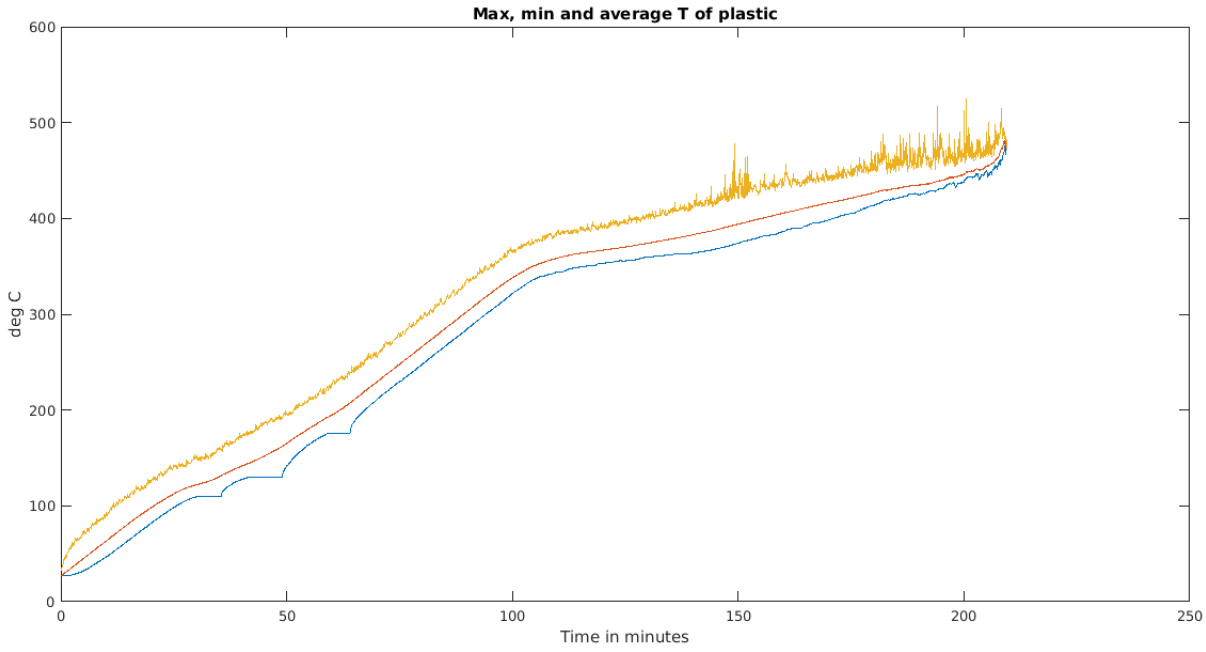

Figure 4.2: Time evolution of temperature of plastic pieces

disappearance of the plastic pieces.

Figure 4.4 shows their temperatures, which for most of the time are constrained by the remaining plastic objects. This is good because they absorb a lot of light, but then pass on most of that energy via conduction to the neighbouring plastic. Only towards the end, when there is almost no plastic left can the average and maximum become very high. One sheet of carbon black just reaches its evaporation point at just over 4000°C right at the end of the simulation. In practice though, the remaining plastic might be

more uniformly distributed than in this simulation where it remains confined to a cube, so this might not happen. Also, it should be possible to turn off the solar input energy a little before the end, to prevent potential damage to the inside mirrors from very hot pigment particles, because there will still be enough heat energy left for the remaining plastic to finish its decomposition.

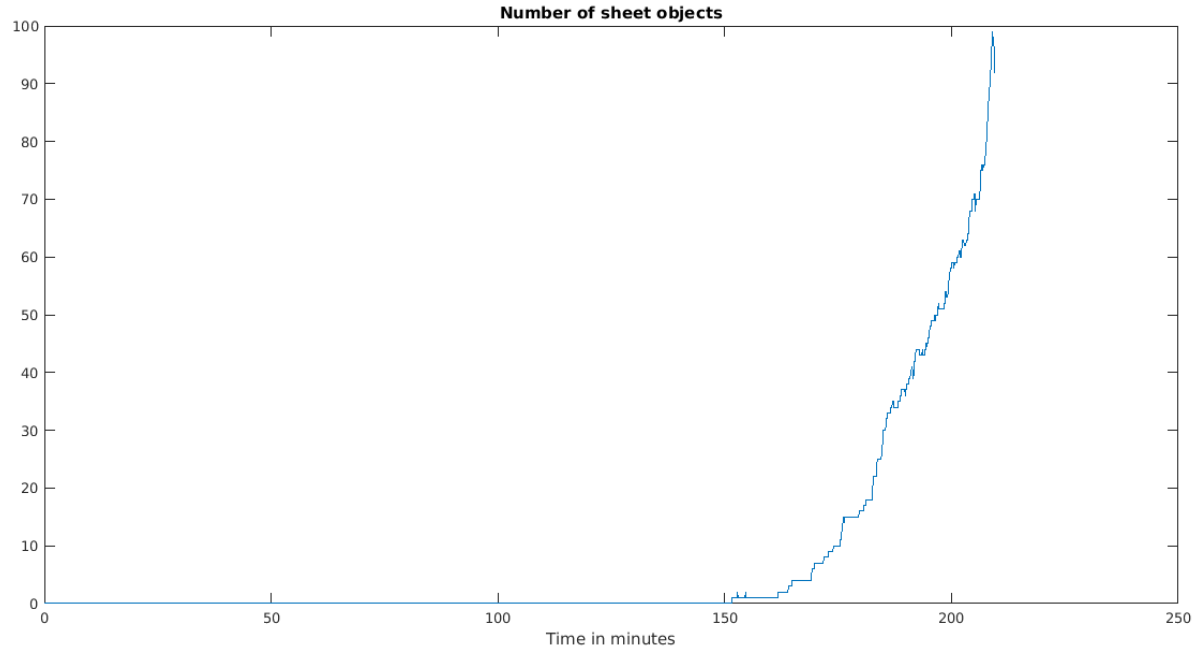

Figure 4.3: Time evolution of number of sheet objects

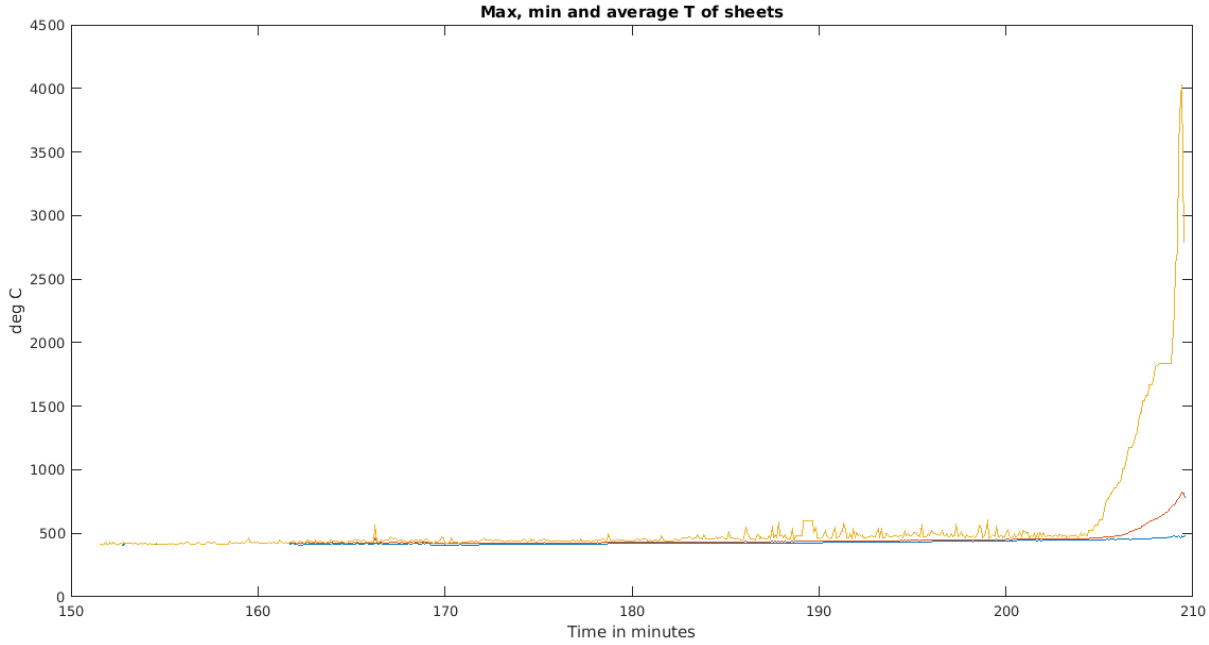

Figure 4.4: Time evolution of temperature of sheet objects

Figure 4.5 shows the calculated cumulative fraction of energy that is lost by escaping through the window. Being cumulative, it fluctuates a little in the beginning before stabilising. As the first sheet objects begin to appear, this increases the absorption efficiency near the surface a lot and so the losses go down quite a bit, from just over 4% to 3.5%. This is a good result, since it means that accumulation of leftover additives at the surface does not seem to present a problem, in fact it would be a good thing. The losses are also manageably low, and lower than one might fear, considering that 48% of the plastic contains no pigment at all, 14% contains highly scattering Titanium Dioxide and 10% contains highly scattering calcium carbonate filler.

Finally, Figure 4.6 shows the time step at which the simulation is progressing. In the beginning, the

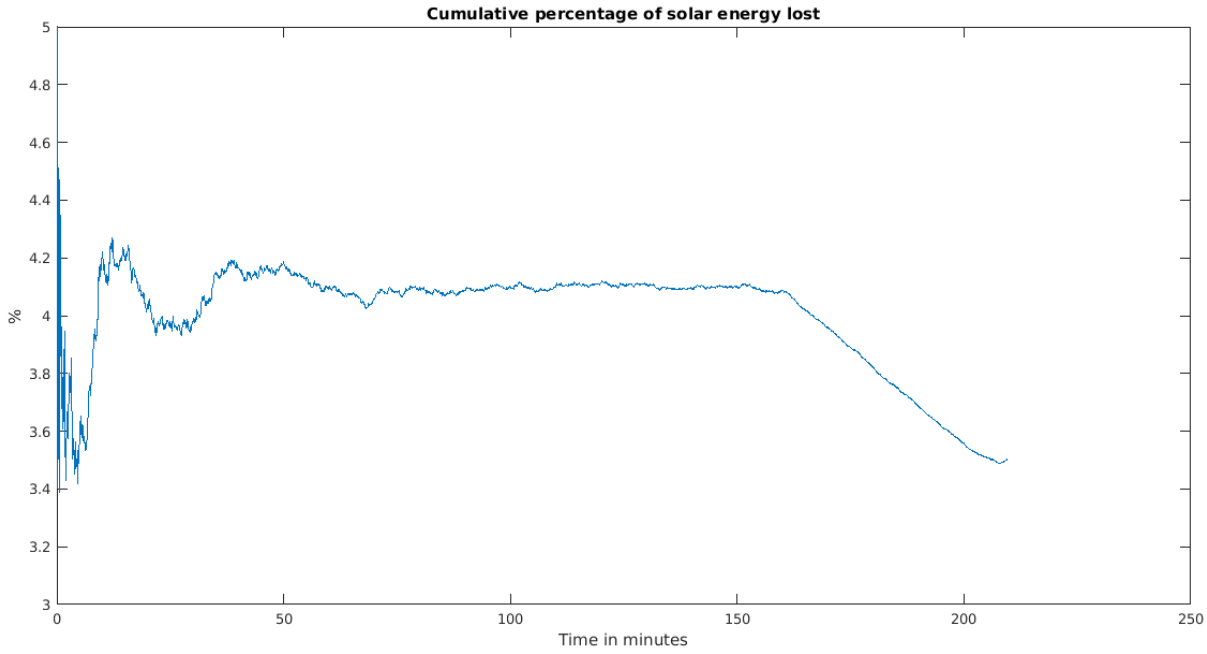

Figure 4.5: Time evolution of cumulative fraction of energy lost

heat transfer could be much faster but the amount of radiation per step has to be restricted, so it stays equal to that maximum. Here, it can be seen that all the safeguards that merge sheets to prevent the time step from becoming even lower due to very thin sheets are mostly working well. Without them, it would quickly go down to the order of 10^{-8}s or even less, and the whole simulation would take forever. The later stages are of course slow even so and take up most of the simulation time, but this cannot be avoided.

At some point during the debugging process, additional code was inserted to check the temperature at which the pyrolysis fragments evaporate, to see how much energy they carry away. Since those results are available, they are shown here in Figure 4.7 in case they are of interest.

4.2 Sensitivity analyses

Due to the many assumptions that had to be made, and the fact that in reality feedstocks will always vary from one batch to another, several simulations were run with varying input parameters to see how much

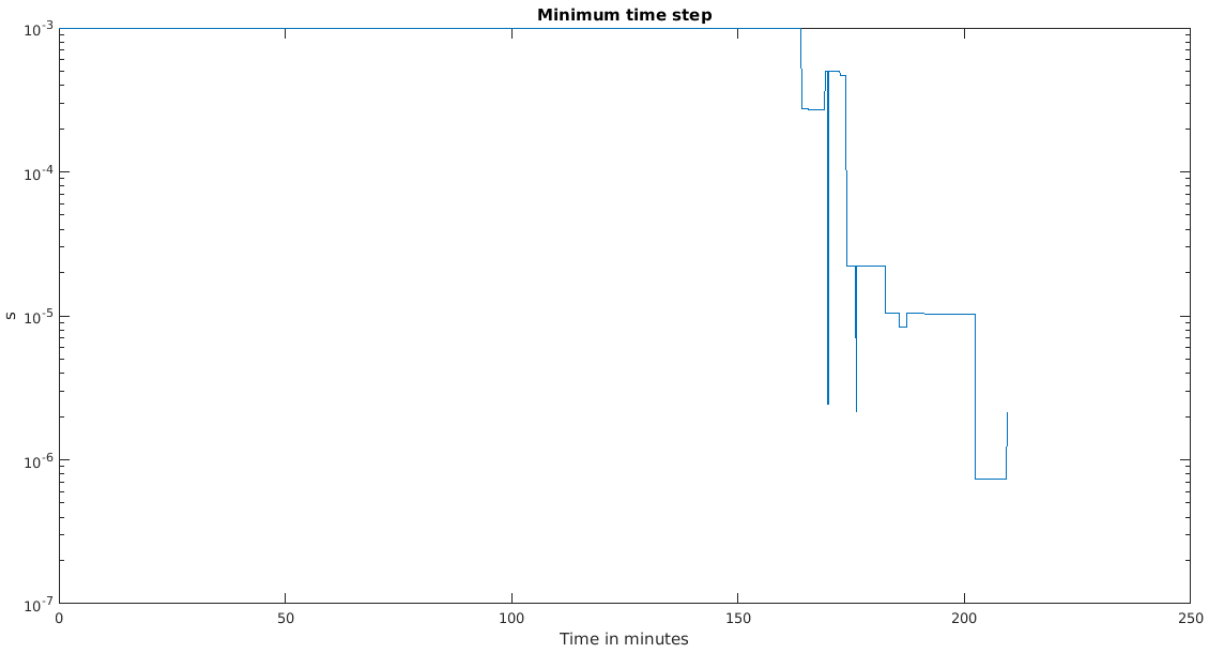

Figure 4.6: Time evolution of simulation time step

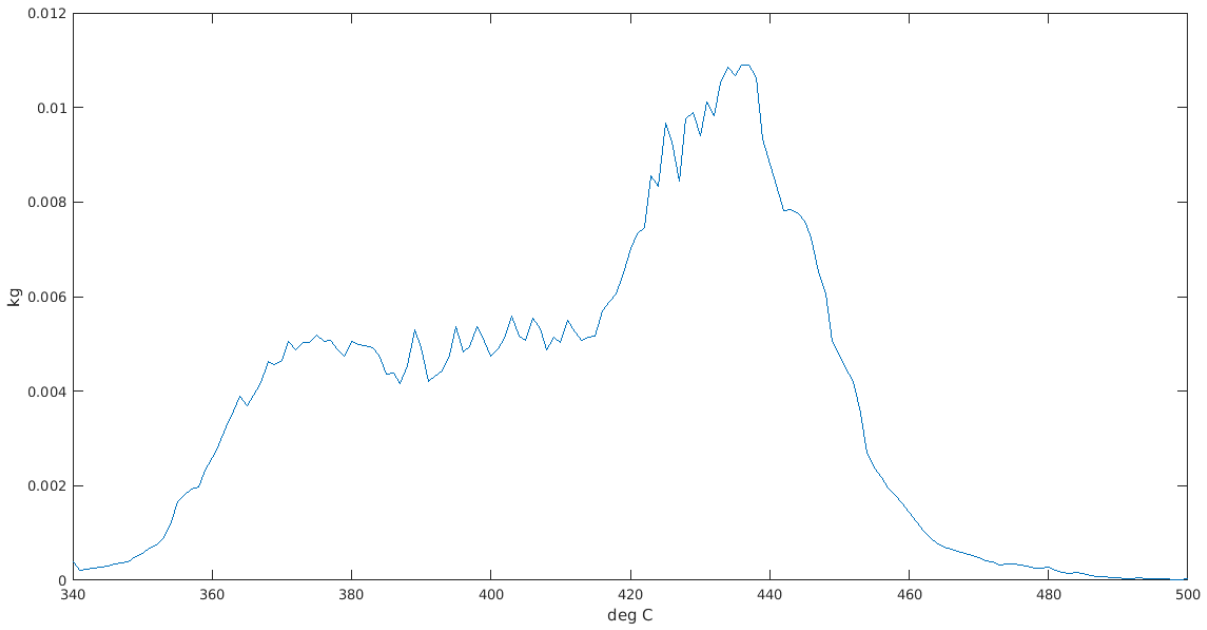

Figure 4.7: Temperature at which the pyrolysis fragments evaporate

that would affect the results, and if there are feedstocks for which it would no longer work well. The base case results are shown again here in Figure 4.8 in the smaller format, so that they can be more easily compared to the other scenarios.

4.2.1 Sensitivity to random number generator seed

The random number generator is seeded with a fixed integer, for reproducibility reasons. This means that those simulations with the same reactor dimensions and same input materials will end up having

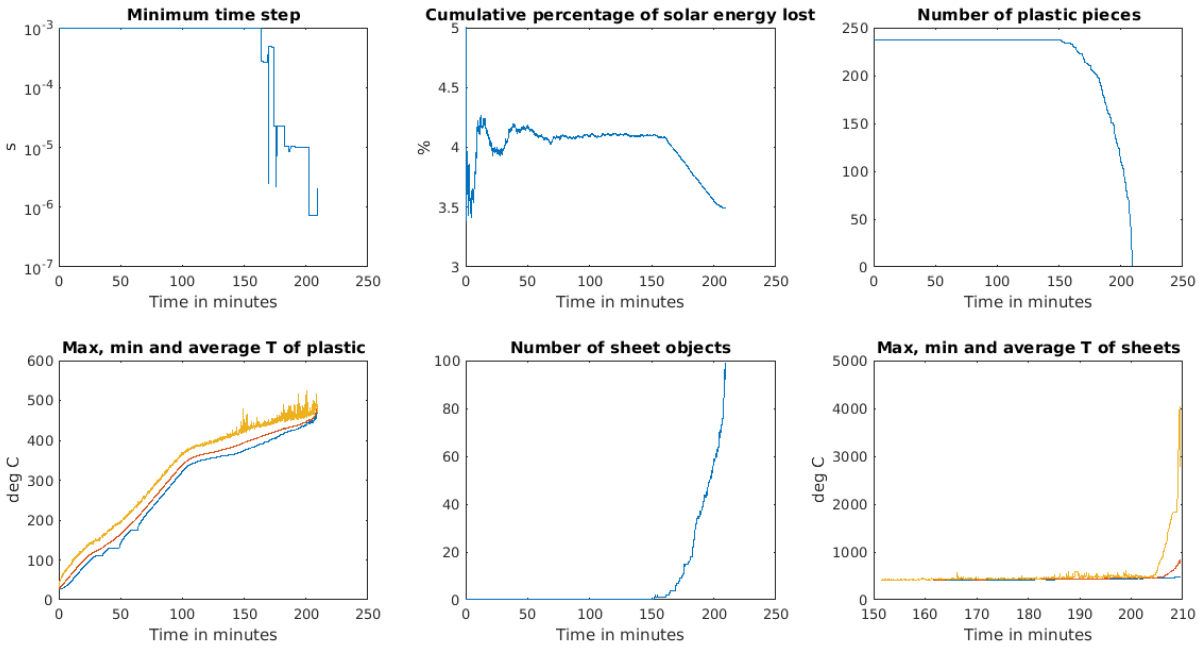

Figure 4.8: Overview of results for base case

exactly the same granules in the same positions. This is desirable, so that those cases where only the maximum time step or flow correction parameters are different can be better compared. Those other differences will then affect the random numbers drawn in the ray tracing.

In order to be better able to compare the base case with those cases where the reactor dimensions are the same but input materials are different, another simulation was run in which the only difference to the base case is the value of the seed. With only 237 plastic granules, the actual percentages of coloured ones might not be too close to the given ones, and might vary significantly from one randomly determined batch to another. Figure 4.9 confirms this suspicion. It seems that in this new setup, more granules with carbon black happen to be near the surface, as the losses have gone down to just over 1% and the maximum plastic temperature is higher above the average. The shape of the curves for number of plastic and for number of sheets has changed slightly, and two sheets of carbon black manage to evaporate just before the end, which causes the losses to go up again slightly right at the end.

Therefore, caution needs to be exercised when making comparisons to the base case, especially when it comes to the magnitude of losses. Also, it might be a good idea to keep a small amount of black granules separately and throw them onto the surface of each batch.

4.2.2 Sensitivity to maximum time step

Would the results be much more accurate if the maximum time step was much lower? Hard to say, but probably not by much. The results of a simulation with 0.0001s instead of 0.001s are displayed in Figure 4.10, and the curves are very similar, just slightly varying in shape. However, since this whole simulation is of a stochastic nature, one would not expect identical results.

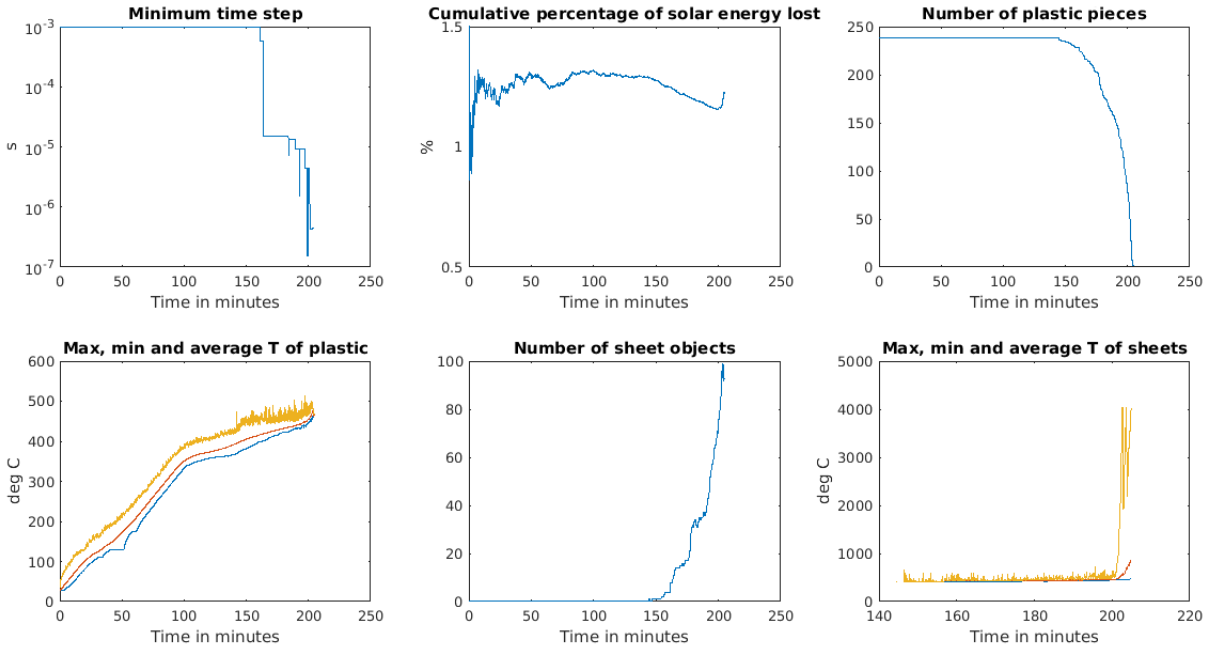

Figure 4.9: Overview of results for case with different random number generator seed

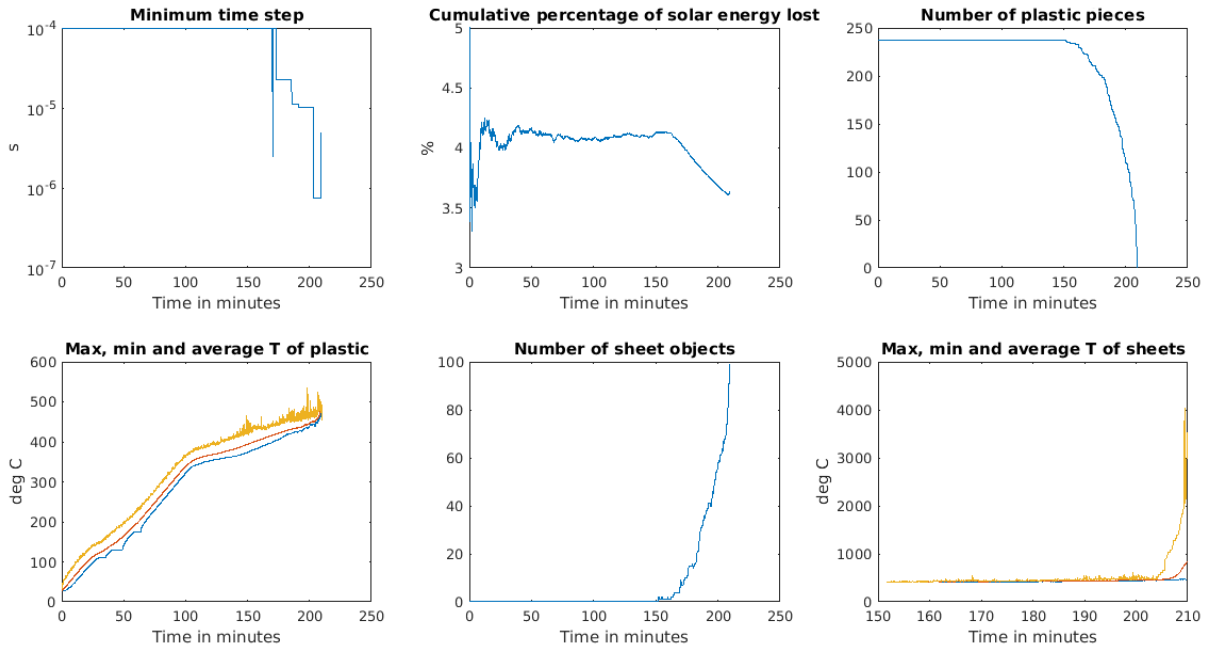

Figure 4.10: Overview of results for case with dtmax = 0.0001s

4.2.3 Sensitivity to flow correction

Checking the surface every 10s for gradient corrections is maybe a bit too frequent, given the high viscosity of plastic. This parameter should have been changed to a more realistic one, but it was overlooked. Figure 4.11 shows the results for a simulation where it is checked every 1min, and like in the previous case, the curves are slightly different but not by a lot. So again it is hard to say whether these results are more accurate or just different because it is a stochastic simulation.

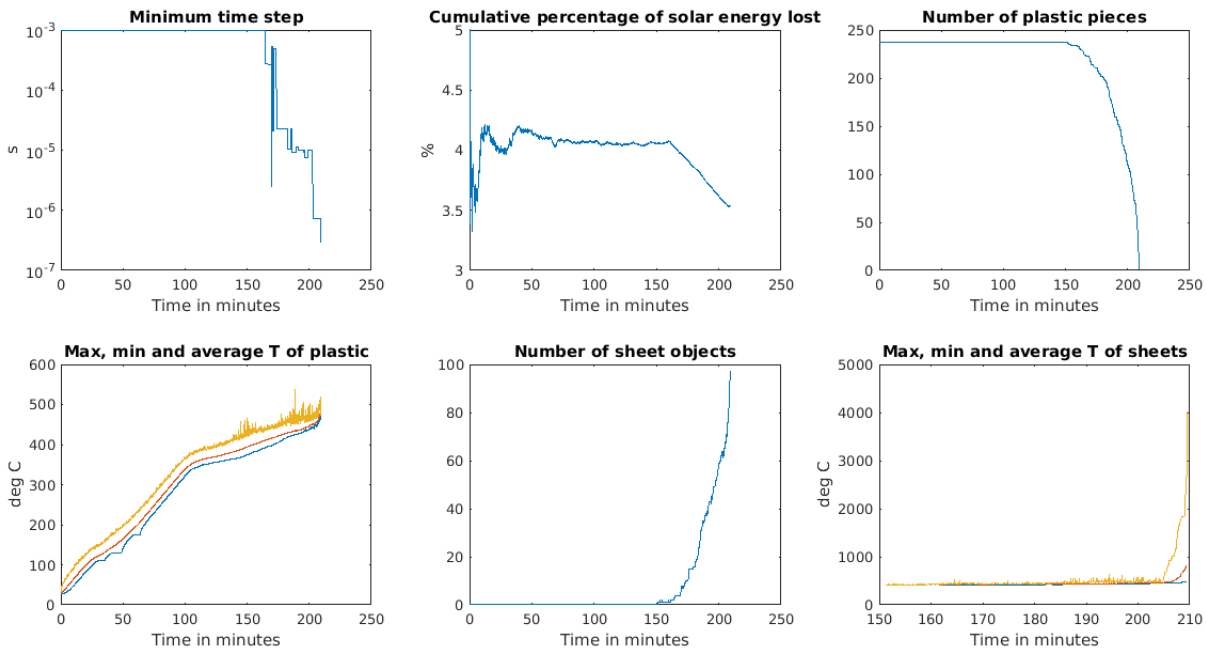

Figure 4.11: Overview of results for case with flow correction every 60s

4.2.4 Sensitivity to scale factor

The scale factor is obviously very important, since it determines the amount of granules (whose size is fixed) that can be in the reactor, which influences the general accuracy and applicability to real life of the simulation. Figure 4.12 shows how the results change if the scale factor is lowered from 0.025 to 0.02. With only 100 plastic granules, a stochastic simulation could probably produce quite variable results for differing random number generator seeds. This should be kept in mind when looking at details like whether sheet objects evaporate or not, as these might be due to pure coincidence. Again, this case seems to have some granules with carbon black at the surface, so losses are low and some of them evaporate at the end.

Figure 4.13 shows the results when going instead to a higher scale factor of 0.03, and again all the curves are different.

Since there are now different plastic objects taking part in the different simulations, and at different locations, it is not surprising that the curves for number of objects and the amount of losses all vary considerably in shape. Not much can be concluded from this. Much bigger simulations would have to be run for the statistical differences to become small enough that general conclusions can be drawn. One such conclusion can be drawn already though, namely that the temperature stratification from top to bottom increases significantly with size. This is obvious from the fact that the vertical difference between average, maximum and minimum temperatures becomes much larger the bigger the reactor is, which means that the surface gets hotter than the average because less and less rays manage to penetrate far down. There is also a more subtle difference that shows this, which is that the gradient of the temperature curves from the point at which the pyrolysis reaction begins onwards becomes straighter with increasing scale factor, because at each point in time the variability in degree of conversion reached

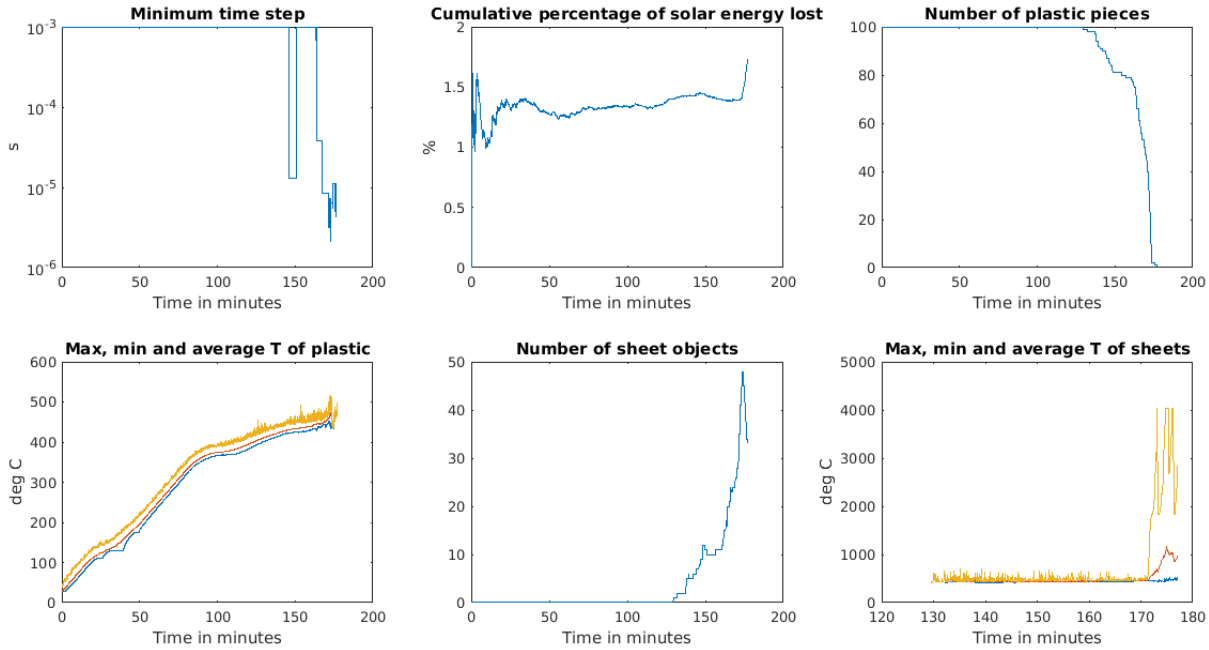

Figure 4.12: Overview of results for case with scale factor 0.02

and thus the specific energy consumption is greater. The tiny reactor with only 100 granules probably has hardly any temperature difference between top and bottom, but based on this trend one can say that the temperature gradient in an industrial scale reactor should be significant. Probably the plastic at the top will already be pyrolysing before the bottom has even melted.

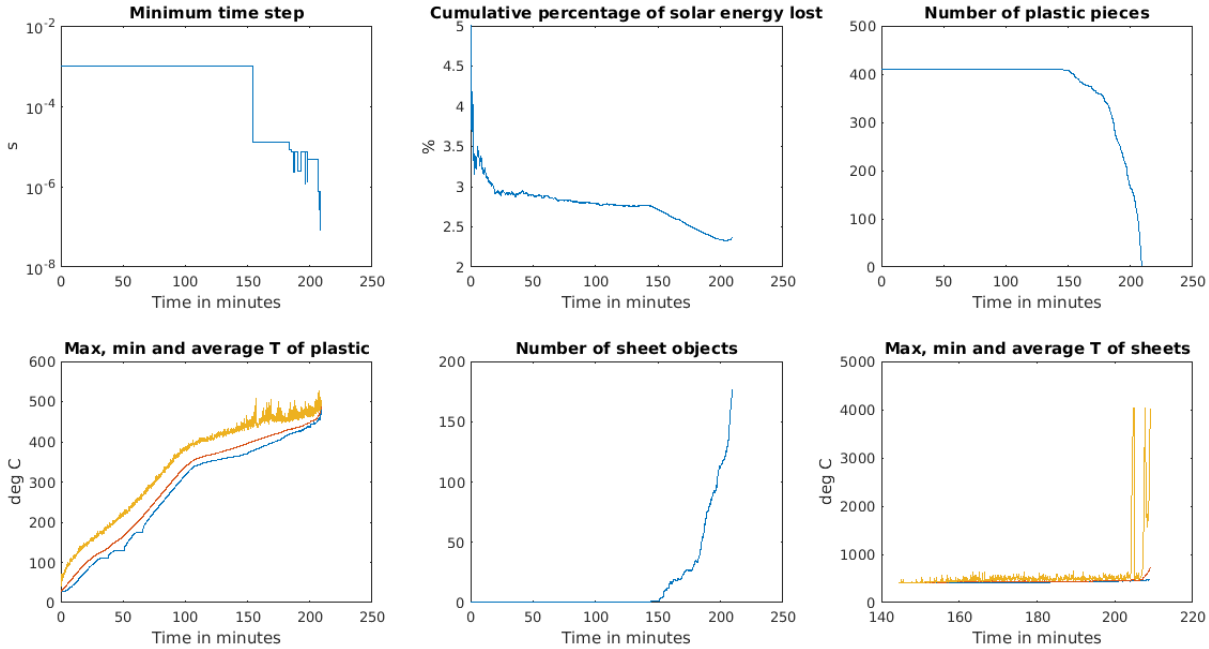

Figure 4.13: Overview of results for case with scale factor 0.03

4.2.5 Sensitivity to colour content

The base case has an average of 52% pigmented plastic, while the rest does not absorb any radiation at all, unless it has a UV absorber or a calcium carbonate filler. The filler is only very weakly absorbing in the solar spectrum though. These fractions are the guess at the real fractions of all plastic waste, although as mentioned before, recycling facilities prefer clear plastic so in practice the pigmented fraction may be higher. This would only affect the results in a positive way though, so it was deemed more interesting to check what would happen in a case where by chance the clear fraction is very high. Figure 4.14 shows the results for a case which has only half as much pigmented plastic, because the prevalence value of each pigment is multiplied by 0.5 by the main function after reading it in from the file.

One might think that with so little pigmented plastic, solar pyrolysis should not work well any more at all, however the results do not look that bad. Well, losses go up to 7.5%, before the sheets appear, then drop to 7%, but that is still manageable, assuming that such feedstock compositions will be the exception and not the norm. Apart from that, the overall results do not look too different to the base case, and no clear plastic gets left over at the end. There are two noticeable dips in the minimum temperature at around 170 and 210 minutes, which must mean that clear granules ended up getting isolated from all pigment containing objects. However, at that point they already contained enough heat energy to complete the pyrolysis reaction by themselves without dropping below the threshold temperature. In a real life case, this probably would not happen though.

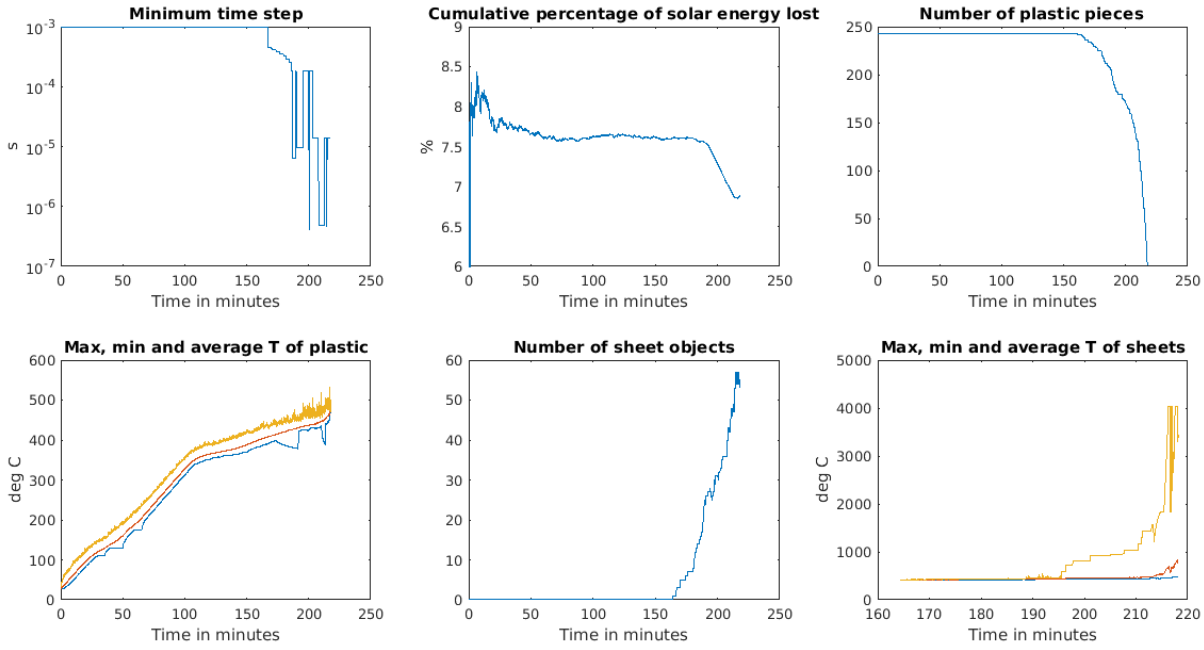

Figure 4.14: Overview of results for case with only half as much pigmented plastic

4.2.6 Sensitivity to light scattering

Since scattering by the pure plastic itself is neglected, one concern might be that the losses would be much higher in real life as more rays might end up being scattered upwards near the surface. Or that if

a batch contains a lot of white plastic, the losses might be very high. To address this possible concern, a simulation was run with 30% prevalence of calcium carbonate filler instead of just 10%. The 10% is anyway just a guess.

Figure 4.15 shows the results and they do not look too bad. The losses are much higher than in the base case, fluctuating around 6.5% in the beginning and then going down to 5.3% at the end. But again, that is still manageable, and the good news is that no non-absorbing granules get left over at the end, in fact the curve is quite steep at the end. This should not be surprising actually, because calcium carbonate has a much higher thermal conductivity than plastics, 2.7 versus 0.2-0.5 $\frac{W}{mK}$, so the overall heat transfer is more efficient. This is probably also the reason that the minimum temperature curve stays slightly closer to the average.

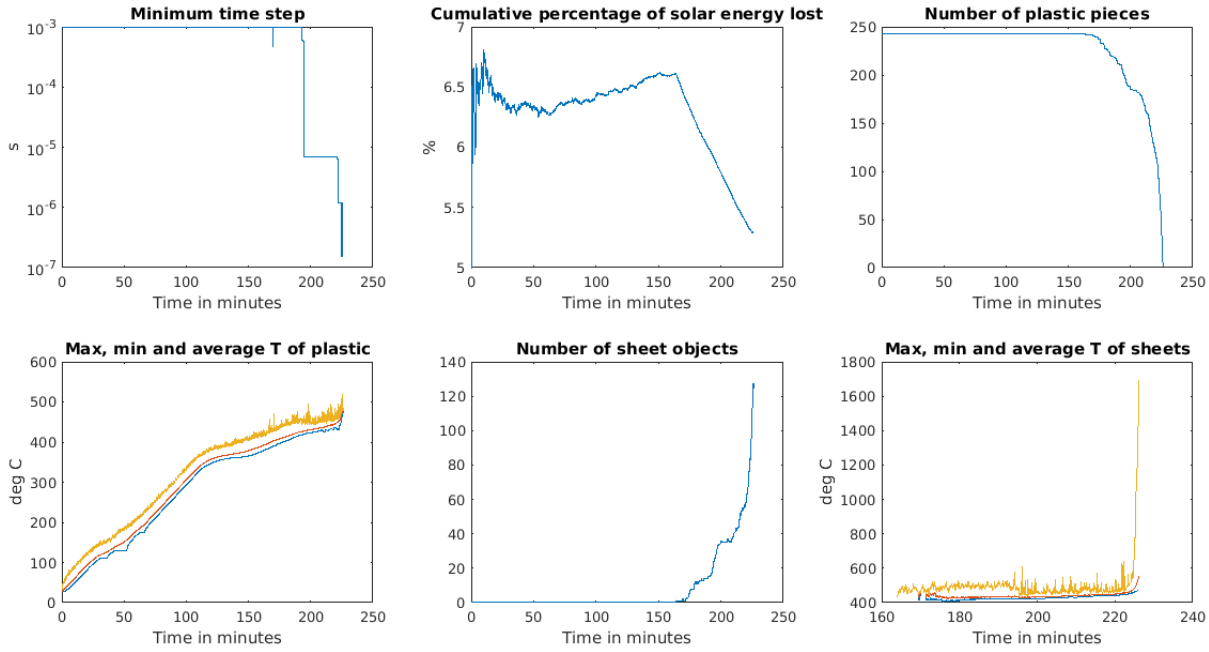

Figure 4.15: Overview of results for case with three times as much calcium carbonate filler

4.2.7 Sensitivity to granule size

While shredding the plastic to a size of around 3mm should be possible and seems to be sufficient, it is nevertheless interesting to check what the minimum size should be, in case a good enough shredder is not available or would cost much more. To this end, four additional simulations were run where the granule diameter was increased to 4mm, 5mm, 7mm and 10mm. The scale factor was increased accordingly, so that all of these simulations have exactly the same amount of participating granules as the base case. Figures 4.16 to 4.19 show the results obtained, respectively.

It can be seen from these results how the thermal lag increases significantly with increasing granule size. While in the base case, the maximum and minimum temperatures stay fairly close to each other throughout, the gap gets bigger and bigger as the size increases. Even just going from 3mm to 4mm almost doubles the difference. For 5mm, the lag is already quite big. It seems to not present too much

of a problem though, since the plastic still pyrolyses together and only one granule gets left over for a short while at the end in both those cases. This might not happen in real life though, and the solar input could be turned off towards the end to prevent evaporation of sheet objects.

However, at 7mm already six non-absorbing granules get left over at the end, and they remain isolated

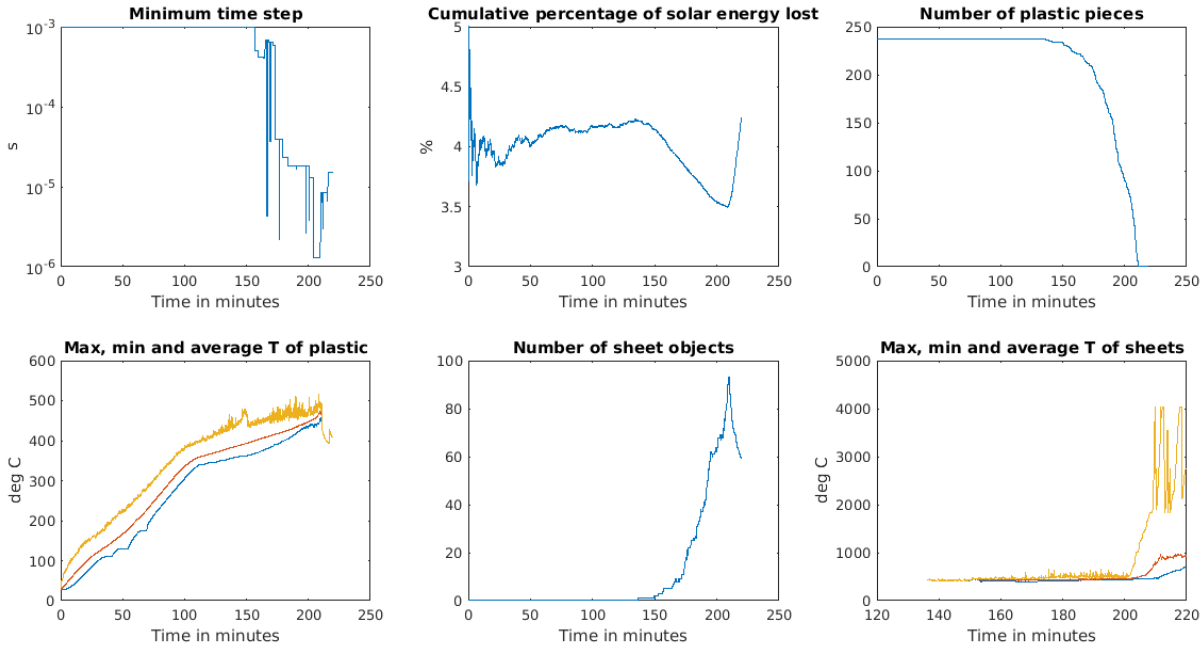

Figure 4.16: Overview of results for case with 4mm granule diameter

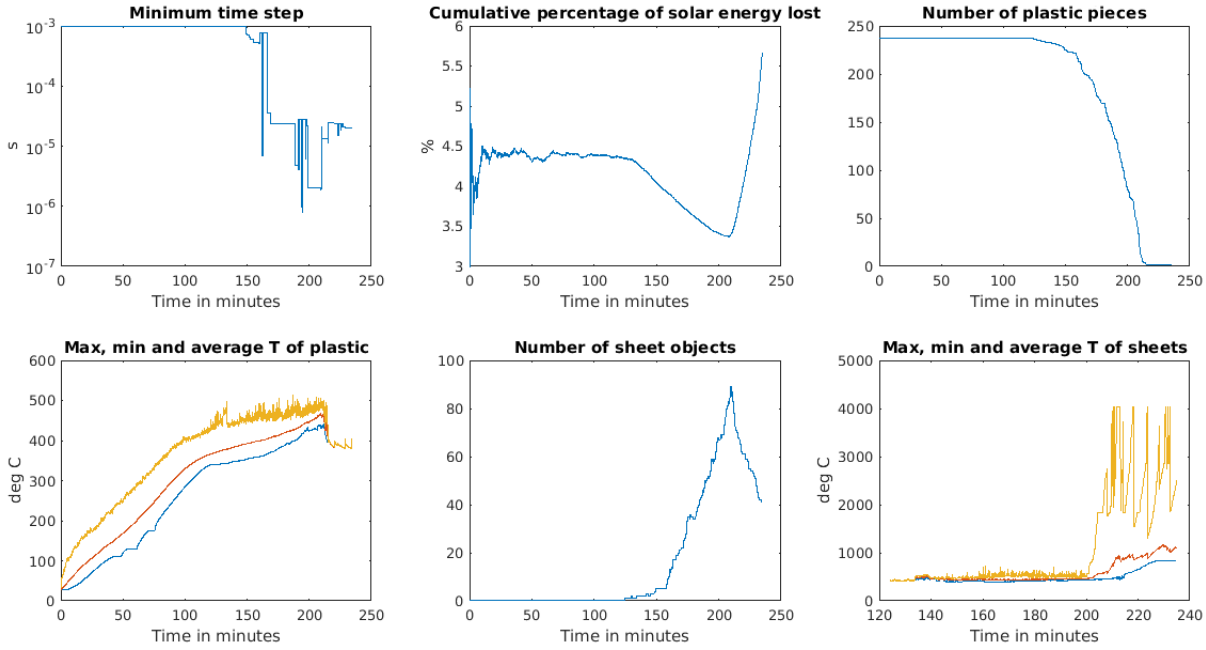

Figure 4.17: Overview of results for case with 5mm granule diameter

with slowly sinking temperatures for 20 minutes, at which point the simulation stops because it has a safeguard to not continue forever in such a case, and the maximum time with no change is set to 20 minutes. Quite probably they still had enough thermal energy to finish pyrolysing at that point, so overall

the whole process would not fail, but it is quite inefficient and wastes time, so this granule size is really not recommended.

Even less recommended is a size of 1cm, as the thermal lag is quite extreme here, reaching 400K.

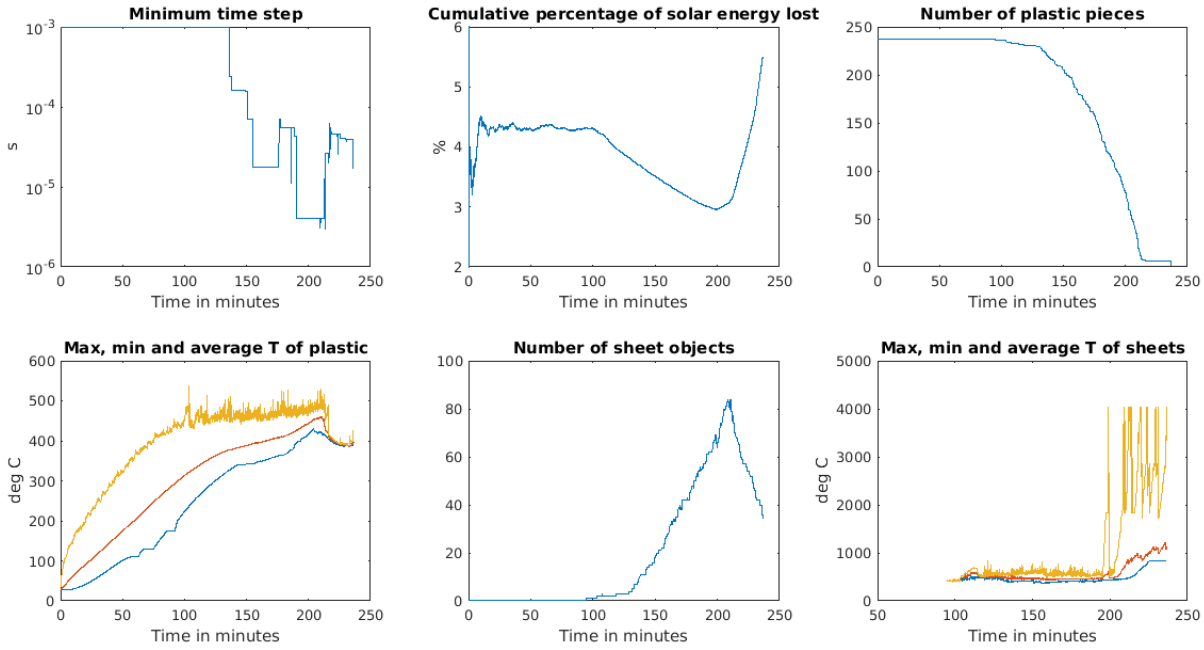

Figure 4.18: Overview of results for case with 7mm granule diameter

The losses go up to almost 8%, and a lot of energy is also wasted evaporating almost all of the sheet objects. The simulation takes 56 minutes longer because quite a few non-absorbing granules get left over at the end. The high temperatures that some of the sheet objects could reach might damage the reactor walls, and it may be that those remaining granules did not have enough energy left to finish pyrolysing by themselves if the input energy was turned off prematurely. The code would have to be extended to allow such a switching off to check. In a hypothetical setup with less dishes than reactors, the dish would then switch to a different reactor. Although, it seems unlikely that anyone would want to operate in such a way anyway. Shredders can do better, and if someone spends a lot of money on a solar reactor and parabolic dish, then you would assume that they would also invest in a good enough shredder so that they could utilise their reactor in an efficient way.

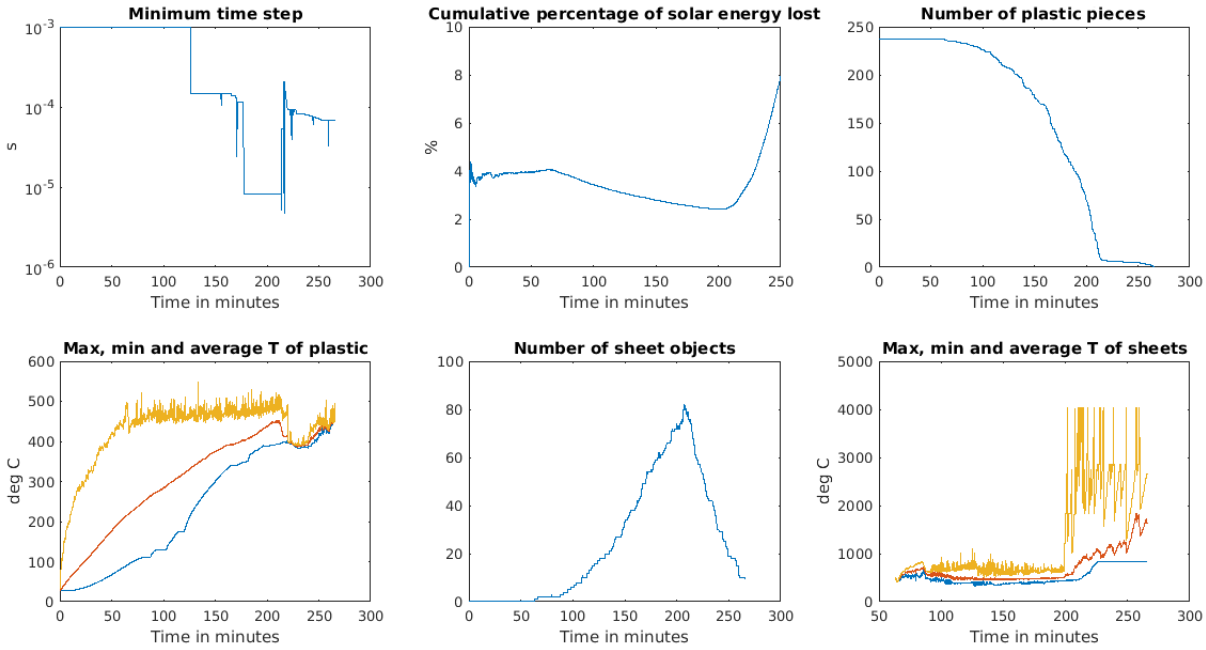

Figure 4.19: Overview of results for case with 10mm granule diameter

Chapter 5

Conclusions

Plastic pyrolysis via direct solar absorption seems promising, since the losses by escaping rays are quite low in most simulated scenarios, around 1% to 3.5%. Hence, it does not seem necessary to use indirectly irradiated reactors.
Also, at 3mm particle size, despite the low thermal conductivity of plastic, the heat transfer between the absorbing and non-absorbing granules and sheets is efficient enough. No plastic gets left over at the end and it seems unlikely that the majority of well absorbing pigments might evaporate and leave behind a layer of calcium carbonate at the surface, because they are continuously transferring their heat to the non-absorbing materials and never get the chance to get near their evaporation or decomposition points. Despite all the simplifications and assumptions that had to be made, these basic conclusions should hold.
4mm and 5mm should still work too, although 3mm is recommended if possible, and should be achievable with available granulators. And even if some batches of waste plastic contain a much lower pigmented fraction than normal and much more white, it should not be a big problem. The losses will be slightly higher, around 7% or so, but not very high, and the overall heat transfer is efficient enough that none of the white or clear granules will get left over for long.

5.1 Achievements

The conceptual framework for solar plastic pyrolysis was developed. The probably first ever code to simulate real waste plastic pyrolysis by direct absorption of sunlight was written. It was proven that despite the low thermal conductivity of plastic, the heat transfer is efficient enough at realistic granule sizes, and the overall losses are low and do not get very high even if the plastic contains little absorbing pigment or a lot of materials that scatter light.

5.2 Future Work

Since the results obtained so far look very promising, further research is recommended. Some suggestions are given below.

- Parallelize the code so that bigger simulations can be run in a reasonable time
- Make the code output more detailed progress information, like masses of all additives
- Check if implicit heat transfer makes the code faster than with explicit
- Include more additives and decomposition products, especially more organic pigments which decompose before the plastic does
- Take variation of heat capacity, thermal conductivity and density with temperature into account
- Measure and incorporate the crowding effect, i.e. that scattering by a pigment is less efficient when the particles are close together
- Measure and model the gradual, stepwise decomposition of some of the organic pigments like the phthalocyanins, and the changes in colour at each stage
- Take losses from reflection by reactor walls into account
- Obtain a sample of real waste plastic from a local waste collection company and determine prevalence of additives and scattering coefficients of pigments in plastic
- Heat up a pigmented plastic and observe if/how quickly the pigment particles sink down through it
- Purchase samples of the most common additives and measure all the missing properties
- Perform own experiments to get accurate DSC of all included plastics, and TGA experiments to determine the plastic pyrolysis kinetics under a given heat flow, first constant and then varying in time
- Perform experiments to find out the difference between pyrolysis occurring due to conductive heat input and solar radiation input, to see if bond breakage from the UV rays is significant
- Measure the scattering coefficients of pure polyolefins as function of their crystallinity, and include this aspect in the simulation
- Take into account that heat transfer will be less efficient when both particles are solid due to their irregular shapes
- Model the flow of argon gas, taking into account convection losses and that the presence of an inlet and outlet will cause higher radiation losses
- Take emission and absorption of thermal radiation into account

- Build a prototype reactor and compare experiments with simulation results, as well as determining the typical product spectrum obtained under such conditions
- Measure the cross-diffusion coefficients of the different plastic types and take this process into account, as this would also impact the kinetics
- Simulate passing of a cloud and its effect on the end products
- Perform some research on catalysts to determine which ones absorb the most sunlight, and could thus be used without the need to supply external heating
- Develop a kinetic model for conversion to TiC and compare with own experiments
- Irradiate white plastic with UV light in an inert atmosphere to check for possible photocatalysis from TiO_2
- Vary the shape of the top part of the reactor and see if this has a significant effect on the radiation losses
- Include additives that bind contaminants like chlorine, sulphur and bromine
- Compare windowed and non-windowed reactor to determine the most suitable type for this application
- Estimate the potential of different countries, based on their DNI and waste plastic composition
- Maybe also diversify and determine feasibility of solar pyrolysis of waste tyres and electronic plastic waste

Bibliography

[1] Plastics – the Facts 2012. An analysis of European plastics production, demand and waste data for 2011. Technical report, PlasticsEurope, 2012. URL http://www.plasticseurope.org/documents/document/20121120170458-{%}0Afinal{_}plasticsthefacts{_}nov2012{_}en{_}web{_}resolution.pdf.

[2] Michael Bolgar, Jack Hubball, Joe Groeger, and Susan Meronek. *Handbook for the chemical analysis of plastic and polymer additives*. CRC Press, 2008. ISBN 978-1-4200-4487-4.

[3] Laura Parker. Eight Million Tons of Plastic Dumped in Ocean Every Year, 2015. URL https://news.nationalgeographic.com/news/2015/02/150212-ocean-debris-plastic-garbage-patches-science/.

[4] Carolyn Barry. Plastic Breaks Down in Ocean, After All – And Fast, 2009. URL https://news.nationalgeographic.com/news/2009/08/090820-plastic-decomposes-oceans-seas.html.

[5] William Reusch. Polymers, 2003. URL https://www2.chemistry.msu.edu/faculty/reusch/VirtTxtJml/polymers.htm.

[6] Gaelle Gourmelon. Global Plastic Production Rises, Recycling Lags, 2015. URL http://www.worldwatch.org/global-plastic-production-rises-recycling-lags-0.

[7] Brickner Gershman and Bratton Inc. Gasification of Non-Recycled Plastics From Municipal Solid Waste in the United States. Technical Report September, The American Chemistry Council, 2013.

[8] World Nuclear Association. Heat Values of Various Fuels, 2016. URL http://www.world-nuclear.org/information-library/facts-and-figures/heat-values-of-various-fuels.aspx.

[9] Jeffrey Morris. Recycling versus incineration: an energy conservation analysis. *Hazardous Materials*, 47:277–293, 1996.

[10] Raymond H J M Gradus, Paul H L Nillesen, Elbert Dijkgraaf, and Rick J Van Koppen. A Cost-effectiveness Analysis for Incineration or Recycling of Dutch Household Plastic Waste. *Ecological Economics*, 135:22–28, 2017. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2016.12.021. URL http://dx.doi.org/10.1016/j.ecolecon.2016.12.021.

[11] Incineration of Municipal Solid Waste: Understanding the Costs and Financial Risks, Fact Sheet 4. Technical report, Pembina Institute, 2007. URL https://www.pembina.org/reports/Incineration{_}FS{_}Costs.pdf.

[12] S S Verma. Roads from plastic waste. *The Indian Concrete Journal*, (November):43–44, 2008.

[13] Nicolas Piatkowski, Christian Wieckert, Alan W. Weimer, and Aldo Steinfeld. Solar-driven gasification of carbonaceous feedstock—a review. *Energy & Environmental Science*, 4:73–82, 2011. doi: 10.1039/c0ee00312c.

[14] Ruby Ray and RB Thorpe. A Comparison of Gasification with Pyrolysis for the Recycling of Plastic Containing Wastes. *INTERNATIONAL JOURNAL OF CHEMICAL REACTOR ENGINEERING*, 5, 2007.

[15] S M Al-Salem, A Antelava, A Constantinou, G Manos, and A Dutta. A review on thermal and catalytic pyrolysis of plastic solid waste (PSW). *Journal of Environmental Management*, 197:177–198, 2017. ISSN 0301-4797. doi: 10.1016/j.jenvman.2017.03.084. URL http://dx.doi.org/10.1016/j.jenvman.2017.03.084.

[16] Shafferina Dayana, Anuar Sharuddin, Faisal Abnisa, Wan Mohd, Ashri Wan, and Mohamed Kheireddine Aroua. A review on pyrolysis of plastic wastes. *Energy Conversion and Management*, 115:308–326, 2016. ISSN 0196-8904. doi: 10.1016/j.enconman.2016.02.037. URL http://dx.doi.org/10.1016/j.enconman.2016.02.037.

[17] Cynar, 2014. URL http://www.cynarplc.com/.

[18] RES Polyflow, 2017. URL http://www.respolyflow.com/our-solutions/.

[19] Plastic2Oil. No Title, 2017. URL http://www.plastic2oil.com.

[20] E Koepf, I Alxneit, C Wieckert, and A Meier. A review of high temperature solar driven reactor technology: 25 years of experience in research and development at the Paul Scherrer Institute. *Applied Energy*, 188:620–651, 2017. ISSN 0306-2619. doi: 10.1016/j.apenergy.2016.11.088. URL http://dx.doi.org/10.1016/j.apenergy.2016.11.088.

[21] Joseph Zeaiter, Mohammad N Ahmad, David Rooney, Bechara Samneh, and Elie Shammas. Design of an automated solar concentrator for the pyrolysis of scrap rubber. *Energy Conversion and Management*, 101:118–125, 2015. ISSN 0196-8904. doi: 10.1016/j.enconman.2015.05.019. URL http://dx.doi.org/10.1016/j.enconman.2015.05.019.

[22] Nicolas Piatkowski and Aldo Steinfeld. Solar Gasification of Carbonaceous Waste Feedstocks in a Packed-Bed Reactor—Dynamic Modeling and Experimental Validation. *AIChE*, 57(12), 2011. doi: 10.1002/aic.

[23] J Matsunami, S Yoshida, O Yokota, M Nezuka, and M Tsuji. Gasification of waste tyre and plastic (PET) by solar thermochemical process for solar energy utilization. *Solar Energy*, 65(1):21–23, 1999.

[24] Morocco World News. Moroccan Engineers Convert Waste into Energy, 2015. URL https://www.moroccoworldnews.com/2015/12/174571/moroccan-engineers-{%}0Aconvert-waste-into-energy/.

[25] Bikram D Shakya. *Pyrolysis of waste plastics to generate useful fuel containing hydrogen using a solar thermochemical process*. Master thesis, University of Sydney, 2007.

[26] Gartzen Lopez, Maite Artetxe, Maider Amutio, Javier Bilbao, and Martin Olazar. Thermochemical routes for the valorization of waste polyolefinic plastics to produce fuels and chemicals. A review. *Renewable and Sustainable Energy Reviews*, 73(October 2016):346–368, 2017. ISSN 1364-0321. doi: 10.1016/j.rser.2017.01.142. URL http://dx.doi.org/10.1016/j.rser.2017.01.142.

[27] H Bockhorn, J Hentschel, A Hornung, and U Hornung. Environmental engineering: Stepwise pyrolysis of plastic waste. 54:3043–3051, 1999.

[28] Shafferina Dayana, Anuar Sharuddin, Faisal Abnisa, Wan Mohd, Ashri Wan, and Mohamed Kheireddine Aroua. Energy recovery from pyrolysis of plastic waste: Study on non-recycled plastics (NRP) data as the real measure of plastic waste. *Energy Conversion and Management*, 148:925–934, 2017. ISSN 0196-8904. doi: 10.1016/j.enconman.2017.06.046. URL http://dx.doi.org/10.1016/j.enconman.2017.06.046.

[29] Maria Laura Mastellone and Umberto Arena. Bed defluidisation during the fluidised bed pyrolysis of plastic waste mixtures. *Polymer Degradation and Stability*, 85(3 SPEC. ISS.):1051–1058, 2004. ISSN 01413910. doi: 10.1016/j.polymdegradstab.2003.04.002.

[30] Plastics – the Facts 2016. Technical report, PlasticsEurope, 2016. URL http://www.plasticseurope.org/documents/document/20161014113313-plastics{_}the{_}facts{_}2016{_}final{_}version.pdf.

[31] Bidhya Kunwar, H N Cheng, Sriram R Chandrashekaran, and Brajendra K Sharma. Plastics to Fuel: A Review. *Renewable and Sustainable Energy Reviews*, 54:421–428, 2016. doi: 10.1016/j.rser.2015.10.015.

[32] Guicai Liu, Yanfen Liao, and Xiaoqian Ma. Thermal behavior of vehicle plastic blends contained acrylonitrile-butadiene-styrene (ABS) in pyrolysis using TG-FTIR. *Waste Management*, 61:315–326, 2017. ISSN 0956-053X. doi: 10.1016/j.wasman.2017.01.034. URL http://dx.doi.org/10.1016/j.wasman.2017.01.034.

[33] Yafei Shen, Rong Zhao, Junfeng Wang, Xingming Chen, Xinlei Ge, and Mindong Chen. Waste-to-energy: Dehalogenation of plastic-containing wastes. *Waste Management*, 49:287–303, 2016. ISSN 0956-053X. doi: 10.1016/j.wasman.2015.12.024. URL http://dx.doi.org/10.1016/j.wasman.2015.12.024.

[34] K Lovegrove and J Pye. Fundamental Principles of concentrating solar power (CSP) systems. *Concentrating Solar Power Technology*, pages 16–67, 2012.

[35] Franz Trieb, Christoph Schillings, Marlene O'Sullivan, Thomas Pregger, and Carsten Hoyer-Klick. Global Potential of Concentrating Solar Power. Technical Report September, German Aerospace Center, Institute of Technical Thermodynamics, 2009.

[36] D.W. van Krevelen and Klaas te Nijenhuis. *Properties of Polymers*. Elsevier, 4 edition, 2009. ISBN 9780080548197.

[37] Heraeus Quartzglas. transmission_calculator @ www.heraeus.com, 2017. URL https://www.heraeus.com/en/hqs/fused{_}silica{_}quartz{_}knowledge{_}base/t{_}calc/transmission{_}calculator.aspx?selection=infrasil{_}301{_}302{&}chart=2{&}rangeX=120,5000{&}rangeY=-5,2.

[38] Marc Röger, Christoph Rickers, Ralf Uhlig, Frank Neumann, and Christina Polenzky. Infrared-Reflective Coating on Fused Silica for a Solar High-Temperature Receiver. *Solar Energy Engineering*, 131(May 2009):2–8, 2009. doi: 10.1115/1.3097270.

[39] Keith Emery. Reference Solar Spectral Irradiance: Air Mass 1.5. URL http://rredc.nrel.gov/solar/spectra/am1.5/.

[40] Mark Fedkin and John A. Dutton. Parabolic Dish CSP Technology. URL https://www.e-education.psu.edu/eme812/node/648.

[41] Michael Tolinski. *Additives for Polyolefins, Getting the most out of Polypropylene, Polyethylene and TPO*. Elsevier, 2 edition, 2015. ISBN 9780323266987.

[42] Chuanbin Zhou, Wenjun Fang, Wanying Xu, Aixin Cao, and Rusong Wang. Characteristics and the recovery potential of plastic wastes obtained from landfill mining. *Journal of Cleaner Production*, 80:80–86, 2014. ISSN 0959-6526. doi: 10.1016/j.jclepro.2014.05.083. URL http://dx.doi.org/10.1016/j.jclepro.2014.05.083.

[43] Stuart Foster. Domestic Mixed Plastics Packaging Waste Management Options. Technical report, WRAP, 2008.

[44] Ahmet Gürses, Metin Acikyildiz, Kübra Günes, and Sadi Gürses. *Dyes and Pigments*. Springer, 2016. ISBN 9783319338903.

[45] R & R Trading Limited. Sell HDPE blow moding grade dirty regrinds-natural, 2013. URL http://www.scrapmonster.com/selloffer/sell-hdpe-blow-moding-grade-dirty-regrinds-natural/10862.

[46] Tanya Levdikova. World Titanium Dioxide Production to Reach 6.7 Mln Tonnes in 2017, According to New Report by Merchant Research & Consulting, 2014. URL http://www.prweb.com/releases/2014/03/prweb11667192.htm.

[47] Ltd CNNC HUAYUAN Titanium Dioxide Co. Introduction to the Application of Titanium Dioxide in Plastic, 2013. URL http://www.sinotio2.com/node/60.

[48] Statista. Production of plastics worldwide from 1950 to 2015, 2017. URL https://www.statista.com/statistics/282732/global-production-of-plastics-since-1950/.

[49] Tanya Levdikova. World PET Supply to Exceed 26.29 Mln Tonnes in 2016, According to New Report by Merchant Research & Consulting, 2014. URL http://www.prweb.com/releases/2014/02/prweb11590546.htm.

[50] Dominique Einhorn. The Honda-Fujishima effect: Using titanium dioxide as a photocatalyst, 2015. URL https://born2invest.com/articles/the-honda-fujishima-effect-using-titanium-dioxide-as-a-photocatalyst/.

[51] Shengying Li, Shihong Xu, Lijun He, Fei Xu, Yonghong Wang, and Li Zhang. Photocatalytic Degradation of Polyethylene Plastic with Polypyrrole/TiO2 Nanocomposite as Photocatalyst. *Polymer-Plastics Technology and Engineering*, (49):400–406, 2010. doi: 10.1080/03602550903532166.

[52] Robert M Christie. *Colour Chemistry*. The Royal Society of Chemistry, 2nd edition, 2015. ISBN 9781849733281.

[53] R W J Westerhout, J Waanders, J A M Kuipers, and W P M van Swaaij. Kinetics of the Low-Temperature Pyrolysis of Polyethene, Polypropene, and Polystyrene Modeling, Experimental Determination, and Comparison with Literature Models and Data. *Ind. Eng. Chem. Res.*, (1):1955–1964, 1997.

[54] Pallab Das and Pankaj Tiwari. Thermal degradation kinetics of plastics and model selection. *Thermochimica Acta*, 654(March):191–202, 2017. ISSN 0040-6031. doi: 10.1016/j.tca.2017.06.001. URL http://dx.doi.org/10.1016/j.tca.2017.06.001.

[55] Shai Vaingast. im2graph, 2017. URL www.im2graph.co.il.

[56] Anabela Coelho, Luis Costa, Maria das Merces Marques, Isabel Fonseca, Maria Amelia Lemos, and Francisco Lemos. Using simultaneous DSC/TG to analyze the kinetics of polyethylene degradation—catalytic cracking using HY and HZSM-5 zeolites. *Reac Kinet Mech Cat*, (99):5–15, 2010. doi: 10.1007/s11144-009-0114-1.

[57] Jing Li and Stanislav I Stoliarov. Measurement of kinetics and thermodynamics of the thermal degradation for non-charring polymers. *Combustion and Flame*, 160(7):1287–1297, 2013. ISSN 0010-2180. doi: 10.1016/j.combustflame.2013.02.012. URL http://dx.doi.org/10.1016/j.combustflame.2013.02.012.

[58] W J JR Frederick and C C Mentzer. Polymers by Differential Scanning Calorimetry. *Applied Polymer Science*, 19:1799–1804, 1975.

[59] Tiziano Faravelli, Giulia Bozzano, Mauro Colombo, Eliseo Ranzi, and Mario Dente. Kinetic modeling of the thermal degradation of polyethylene and polystyrene mixtures. *Anal. Appl. Pyrolysis*, 70:761–777, 2003. doi: 10.1016/S0165-2370(03)00058-5.

[60] I Zupancic, G Lahajnar, and R Blinc. NMR Self-Diffusion Study of Polyethylene and Paraffin Melts. *Polymer Science*, 23:387–404, 1985.

[61] Gerald Fleischer. Self Diffusion in Melts of Polystyrene and Polyethylene Measured by Pulsed Field Gradient NMR. *Polymer Bulletin*, 158:152–158, 1983.

[62] Suhas Patankar. *Numerical Heat Transfer and Fluid Flow*. McGraw-Hill, 1980.

[63] Michael F Modest. *Radiative Heat Transfer*. Academic Press (Elsevier), 2nd edition, 2003.

[64] Ch Sasse and G Ingel. The role of the optical properties of solids in solar direct absorption process. *Solar Energy Materials and Solar Cells*, 31:61–73, 1993.

[65] Ronnen Levinson, Paul Berdahl, and Hashem Akbari. Solar spectral optical properties of pigments—Part II: survey of common colorants. *Solar Energy Materials & Solar Cells*, 89:351–389, 2005. doi: 10.1016/j.solmat.2004.11.013.

[66] N A Hassine, J G P Binner, and T Cross. Synthesis of Refractory Metal Carbide Powders via Microwave Carbothermal Reduction. *Refractory Metals & Hard Materials*, 13:353–358, 1995.

[67] Aldo Steinfeld and M Schubnell. Optimum aperture size and operating temperature of a solar cavity-receiver. *Solar Energy*, 50(1):19–25, 1993.

[68] Erik Hansen, Nils H Nilsson, Delilah Lithner, and Carsten Lassen. Hazardous substances in plastic materials. Technical report, COWI and Danish Technological Institute, 2013.

[69] Scientific Polymer Products Inc. Refractive Index of Polymers by Index, 2013. URL `http://scientificpolymer.com/technical-library/refractive-index-of-polymers-by-index/`.

[70] M D Roussel, A R Guy, L G Shaw, and J E Cara. The Use of Calcium Carbonate in Polyolefins Offers Significant Improvement in Productivity. *PLACE Conference*, 2005.

[71] Thermal Conductivity of common Materials and Gases. URL `https://www.engineeringtoolbox.com/thermal-conductivity-d{_}429.html`.

[72] Ferdinand Rodriguez. Plastic, 2017. URL `https://www.britannica.com/science/plastic`.

[73] Shichao Wang, Shuangjun Chen, and Jun Zhang. Characterization, Solar Reflectance, and Crystal Properties of Polyethylene and Ethylene Copolymer after Thermal Treatment. *International Journal of Polymer Analysis and Characterization*, 18(4):257–268, 2013. doi: 10.1080/1023666X.2013.766782.

[74] NANJING LEPUZ CHEMICAL CO LTD. Cyasorb UV 531, 2017. URL `https://www.alibaba.com/product-detail/Cyasorb-UV-531{_}60514121292.html`.

[75] National Center for Biotechnology Information, US National Library of Medicine, National Institutes of Health, and U.S. Department of Health & Human Services Human. Pubchem Database, 2017. URL `https://pubchem.ncbi.nlm.nih.gov`.

[76] ChemBK Online Chemical Database. UV-531, 2015. URL http://www.chembk.com/en/chem/UV-531.

[77] Feng Peng and Steven W Effler. Characterizations of calcite particles and evaluations of their light scattering effects in lacustrine systems. *LIMNOLOGY and OCEANOGRAPHY*, 62:645–664, 2017. doi: 10.1002/lno.10451.

[78] Carlos Rodriguez-Navarro, Encarnacion Ruiz-Agudo, Ana Luque, Alejandro B Rodriguez-Navarro, and Miguel Ortega-Huertas. Thermal decomposition of calcite: Mechanisms of formation and textural evolution of CaO nanocrystals. *American Mineralogist*, 94:578–593, 2009. doi: 10.2138/am.2009.3021.

[79] Marcel Babin, Andre Morel, Vincent Fournier-Sicre, Frank Fell, and Dariusz Stramski. Light scattering properties of marine particles in coastal and open ocean waters as related to the particle mass concentration. *LIMNOLOGY and OCEANOGRAPHY*, 48(2):843–859, 2003.

[80] Chul Seoung Baek, Kye Hong Cho, and Ji-Whan Ahn. Effect of Grain Size and Replacement Ratio on the Plastic Properties of Precipitated Calcium Carbonate Using Limestone as Raw Material. *Journal of the Korean Ceramic Society*, 51(2):127–133, 2014. doi: 10.4191/kcers.2014.51.2.127.

[81] William J Tropf. Calcium Carbonate, Calcite (CaCO3). In *HANDBOOK OF OPTICAL CONSTANTS OF SOLIDS III*, pages 701–715. Elsevier, 1998. ISBN 0125444230.

[82] M. N. Polyanskiy. Refractive index database. URL https://refractiveindex.info.

[83] Theodore Gray. Technical data for Copper, . URL http://periodictable.com/Elements/029/data.html.

[84] Kaupo Kukli, Mikko Ritala, Timo Sajavaara, Timo Hanninen, and Markku Leskela. Atomic layer deposition of calcium oxide and calcium hafnium oxide films using calcium cyclopentadienyl precursor. *Thin Solid Films*, 500:322–329, 2006. doi: 10.1016/j.tsf.2005.10.082.

[85] N A Surplice and R P Jones. The thermal conductivity of alkaline earth oxides. *BRIT. J. APPL. PHYS*, 14:720–722, 1963.

[86] ChemicalBook. Calcium oxide. URL http://www.chemicalbook.com/ChemicalProductProperty{_}EN{_}CB2853017.htm.

[87] calcium oxide, 2017. URL http://webbook.nist.gov/cgi/cbook.cgi?ID=C1305788{&}Type=JANAFS{&}Plot=on{#}JANAFS.

[88] Jolyon Ralph and Katya Ralph. Goethite, 2017. URL https://www.gemdat.org/gem-1719.html?id=1719.

[89] Ki-iti Horai and G Simmons. Thermal conductivity of rock-forming minerals. *EARTH AND PLANETARY SCIENCE LETTERS*, 6:359–368, 1969.

[90] Igor Diakonov, Igor Khodakovsky, Jacques Schott, and Elvira Sergeeva. Thermodynamic properties of iron oxides and hydroxides. I. Surface and bulk thermodynamic properties of goethite (α-FeOOH) up to 500 K. *European Journal of Minerals*, 6:967–983, 1994.

[91] Sulata K Sahu, Manjulata Sahu, Raman S Srinivasa, and Thiagarajan Gnanasekaran. Determination of heat capacities of PbCrO4(s), Pb2CrO5(s), and Pb5CrO8(s). *Chemical Monthly*, 143(9): 1207–1214, 2012. doi: 10.1007/s00706-012-0756-y.

[92] Phthalocyanine green. URL `http://cameo.mfa.org/wiki/Phthalocyanine{_}green`.

[93] Mikako Takeda, Takashi Onishi, Shouhei Nakakubo, and Shinji Fujimoto. Physical Properties of Iron-Oxide Scales on Si-Containing Steels at High Temperature. *Materials Transactions*, 50(9): 2242–2246, 2009. doi: 10.2320/matertrans.M2009097.

[94] Dominika Trefon-Radziejewska, Justyna Juszczyk, Austin Fleming, Nicolas Horny, Jean Stéphane Antoniow, Mihai Chirtoc, Anna Kaźmierczak-Bałata, and Jerzy Bodzenta. Thermal characterization of metal phthalocyanine layers using photothermal radiometry and scanning thermal microscopy methods. *Synthetic Metals*, 232(July):72–78, 2017. doi: 10.1016/j.synthmet.2017.07.012.

[95] Reade International Corporation. Hematite Powder / Red Iron Oxide (Fe2O3), 2016. URL `http://www.reade.com/products/hematite-powder-red-iron-oxide-fe2o3`.

[96] E. Takegoshi, Y. Hirasawa, S. Imura, and T. Shimazaki. Measurement of Thermal Properties of Iron Oxide Pellets 1. *International Journal of Thermophysics*, 5(2):219–228, 1984.

[97] Xue Dong Gong, He Ming Xiao, and He Tian. Comparative Studies on the Structures, Infrared Spectrum, and Thermodynamic Properties of Phthalocyanine Using Ab Initio Hartree–Fock and Density Functional Theory Methods. *International Journal of Quantum Chemistry*, 86:531–540, 2002. doi: 10.1002/qua.1112.

[98] Gerd Lobbert. Phthalocyanins. *Ullmann's Encyclopedia of Industrial Chemistry*, pages 181–213, 2012. doi: 10.1002/14356007.a20.

[99] Yasunori Taru and Kyo Takaoka. Thermal Stability of Organic Pigment. 55, 1982.

[100] A Stendal, U Beckers, S Wilbrandt, O Stenzel, and C von Borczyskowski. The linear optical constants of thin phthalocyanine and fullerite films from the near infrared up to the UV spectral regions: Estimation of electronic oscillator strength values. *J. Phys. B: At. Mol. Opt. Phys., IOPScience*, 29:2589–2595, 1996.

[101] Theodore Gray. Technical data for Carbon, . URL `http://periodictable.com/Elements/006/data.html`.

[102] Guidechem Chemical Trading Guide. CAS No. 6358-30-1 (Pigment Violet 23), 2017. URL `http://www.guidechem.com/cas-635/6358-30-1.html`.

[103] Janusz Nowotny. *Oxide Semiconductors for Solar Energy Conversion: Titanium Dioxide*. CRC Press, 2012. ISBN 9781439848395.

[104] ChemicalBook Inc. N, 2017. URL http://www.chemicalbook.com/ProductMSDSDetailCB3109508{_}EN.htm{#}2.

[105] Azo Materials. Titanium Dioxide - Titania (TiO2). URL https://www.azom.com/properties.aspx?ArticleID=1179.

Appendix A

Material Properties

The material properties read in by the code are presented here. If there is an asterisk after a number, it means that number is a random one since no values could be found, but something needs to be there else a division by zero will occur in the code, resulting in a NaN. In the case of pigment concentrations, they are educated guesses based mostly on the general information in [68].
A "-1" in the refractive index means that there is a separate file which has values for different wavelengths, and the code knows to look for this file if it encounters a "-1" in such a position.
A "-1" in the melting, evaporation or decomposition point means that this particular material does not have such a point, and the code will put a NaN in the object if it finds a "-1" there, and when checking for phase transitions it checks that those values are not NaN.
Forward scattering ratios of all pigments are taken from [65]. Materials that do not scatter have a FSR of 0.5, which is just to avoid potential problems in the code. The value in such a case does not matter because it will anyway be multiplied by zero.

	HDPE	LDPE	PP
Prevalence [fraction]	0.26 [28]	0.51 [28]	0.23 [28]
Molar Mass [$\frac{g}{mol}$]	100,000*	100,000*	100,000*
Refractive Index (real part)	1.49 [36]	1.49 [36]	1.49 [69]
Thermal Conductivity [$\frac{W}{m \cdot K}$]	0.51 [70] [71]	0.33 [71]	0.2 [71] [70]
Density [$\frac{g}{cm^3}$]	0.96 [72]	0.92 [72]	0.90 [72] [36]
Heat Capacity C_p [$\frac{J}{kg \cdot K}$]	2700 [70]	1650 [36]	1680 [70]
Melting Point [°C]	130 [5]	110 [5]	176 [5]
Latent Heat of Fusion [$\frac{J}{kg}$]	221,680	55,420	89,010
Decomposition Point [°C]	340 [56]	340	340 [57]

Table A.1: Properties of Plastics

The melting energies are calculated as the product of the melting energy of a theoretical 100% crystalline version and the estimated degree of crystallinity. However, crystallinity of plastics varies widely from one batch to another, as do the estimates that can be found for the amount. So these numbers are more like guesses. Crystalline melting energy of PE is 277.1 J/g [73], and for PP 207 J/g.

	UV531	Calcium Carbonate
Minimum Content [%]	0.2 [74]	10 [41]
Maximum Content [%]	0.5 [74]	60 [41]
Prevalence [%]	30*	10*
Molar Mass [$\frac{g}{mol}$]	326.43 [2]	100.086 [75]
Refractive Index (real part)	1.547 [76]	-1
Thermal Conductivity [$\frac{W}{m \cdot K}$]	1*	2.7 [70]
Density [$\frac{g}{cm^3}$]	1.16 [74]	2.711 [41]
Heat Capacity C_p [$\frac{J}{kg \cdot K}$]	1000*	860 [70]
Forward Scattering Ratio	0.5*	0.8 [77] [1]
Melting Point [°C]	48 [2] [74]	-1
Latent Heat of Fusion [$\frac{J}{kg}$]	10000*	0
Boiling Point [°C]	-1	-1
Latent Heat of Vaporization [$\frac{J}{kg}$]	0	0
Decomposition Point [°C]	457.9 [76]	825 [75]
Decomposition Energy [$\frac{J}{mol}$]	100000*	177800 [78]
Decomposition Product 1	Carbon Black	Calcium Oxide
Moles of DP1 per mol of material	5	1
Decomposition Product 2	none	none
Moles of DP2 per mol of material	0	0

Table A.2: Properties of Additives

The scattering spectrum of calcium carbonate is estimated from one value at 550nm and the general formula $b_p(\lambda) = B(\frac{\lambda}{\lambda_0})^\gamma$ given in [79]. Obviously, particles in plastics will differ somewhat from those in lakes, although the general size range is not too different (up to 7μm in the lake studied [77] versus 0.05μm to 5μm for precipitated calcium carbonate and bigger sizes for ground calcium carbonate [80]). But finding data for the scattering in plastics is difficult.

No values could be found for the absorption coefficient of calcium carbonate in the solar spectrum, because these values are so low that they are hard to measure over all the high scattering. [81] has plots using all the then known data for complex refractive indices (for ordinary and extraordinary rays), and the values of the imaginary part k disappear between 200nm and 6μm. The first major absorption band only appears at around 7μm, and then more in the far-infrared follow. So, a value of 100m^{-1} was used at all wavelengths that are part of the simulation, which is a very low value compared to those of the pigments.

The refractive index as function of wavelength of calcium carbonate was downloaded from [82].

[1] estimate from Fig. 3 for 0.2μm

	Calcium Oxide	Copper
Molar Mass [$\frac{g}{mol}$]	56.077 [75]	63.546 [83]
Refractive Index (real part)	1.75 [84]	-1
Thermal Conductivity [$\frac{W}{m \cdot K}$]	83.68 [85]	400 [83]
Density [$\frac{g}{cm^3}$]	3.34 [86]	8.92 [83]
Heat Capacity C_p [$\frac{J}{kg \cdot K}$]	900 [87]	384.4 [83]
Forward Scattering Ratio	0.5*	0.5*
Melting Point [°C]	2572 [86]	1084.62 [83]
Latent Heat of Fusion [$\frac{J}{kg}$]	1000	13100 [83]
Boiling Point [°C]	2850 [86]	2927 [83]
Latent Heat of Vaporization [$\frac{J}{kg}$]	100000*	300000 [83]
Decomposition Point [°C]	-1	-1
Decomposition Energy [$\frac{J}{mol}$]	0	0
Decomposition Product 1	none	none
Moles of DP1 per mol of material	0	0

Table A.3: Properties of Decomposition Products

Absorption and scattering spectra of calcium oxide were simply copied from calcium carbonate. Both tend to have a bright white colour, so they are clearly not totally different. However, measurements should be performed to determine the exact values.

Scattering from copper particles will depend on the size of the particles formed after the phthalocyanine pigment particles decompose, which will depend on the local conditions. These will vary a lot since the plastic in which they were contained will already have finished decomposing. Since it was thus not possible to estimate anything, this coefficient was simply set to zero. It should anyway not matter too much since the amount of copper will be very low, as organic pigments have a high tinting strength and thus are used in very small amounts.

Absorption coefficient and refractive index as function of wavelength for copper were taken from [82].

	Iron Oxide Yellow	Chrome Yellow	Iron Oxide Orange
Pigment Number	42	34	NA
Minimum Content [%]	2*	2*	2*
Maximum Content [%]	4*	4*	4*
Prevalence [%]	3*	2*	4*
Molar Mass [$\frac{g}{mol}$]	88.851 [75]	323.192 [75]	124.269
Refractive Index (real part)	2.4 [88]	2.66 [75]	2.6
Thermal Conductivity [$\frac{W}{m \cdot K}$]	2.91 [89]	10*	7.09
Density [$\frac{g}{cm^3}$]	3.8 [88]	6.3 [2]	4.4
Heat Capacity C_p [$\frac{J}{kg \cdot K}$]	900 [90]	470 [91]	800
Forward Scattering Ratio	0.79	0.72	0.805
Melting Point [°C]	-1	844 [75] [2]	-1
Latent Heat of Fusion [$\frac{J}{kg}$]	0	100000*	0
Boiling Point [°C]	-1	-1	-1
Latent Heat of Vaporization [$\frac{J}{kg}$]	0	0	0
Decomposition Point [°C]	180 [90]	230 [75]	180
Decomposition Energy [$\frac{J}{mol}$]	55200 [90]	50000*	35466
Decomposition Product 1	Iron Oxide Red	none	Iron Oxide Red
Moles of DP1 per mol of material	0.5	0	0.67875
Decomposition Product 2	none	none	none
Moles of DP2 per mol of material	0	0	0

Table A.4: Properties of Pigments Part 1

Iron Oxide Orange is assumed to be made of 50% each of Iron Oxide Yellow and Iron Oxide Red, by weight, so its properties are calculated from that.

	Iron Oxide Red	Phthalo Blue	Phthalo Green
Pigment Number	101	15	7
Minimum Content [%]	2*	0.1*	0.2*
Maximum Content [%]	4*	0.2*	0.3*
Prevalence [%]	5*	5*	5*
Molar Mass [$\frac{g}{mol}$]	159.687 [75]	576.082 [75]	1127.154 [75]
Refractive Index (real part)	-1	-1	1.4 [92]
Thermal Conductivity [$\frac{W}{m \cdot K}$]	11.27 [89] [93]	0.35 [94]	0.35*
Density [$\frac{g}{cm^3}$]	5 [95] [93]	1.62 [75]	2.1 [92]
Heat Capacity C_p [$\frac{J}{kg \cdot K}$]	700 [96]	850 [97]	850 [97]
Forward Scattering Ratio	0.82	0.82	0.86
Melting Point [°C]	1550 [95]	-1	-1
Latent Heat of Fusion [$\frac{J}{kg}$]	100000*	0	0
Boiling Point [°C]	2000*	550 [98]	-1
Latent Heat of Vaporization [$\frac{J}{kg}$]	100000*	100000*	0
Decomposition Point [°C]	-1	-1	600 [98]
Decomposition Energy [$\frac{J}{mol}$]	0	0	100000*
Decomposition Product 1	none	Carbon Black	Carbon Black
Moles of DP1 per mol of material	0	24	24
Decomposition Product 2	none	Copper	Copper
Moles of DP2 per mol of material	0	1	1

Table A.5: Properties of Pigments Part 2

The phthalocyanins actually decompose in two stages and partially sublimate over the temperature range 460 – 630°C [98] [99], but the code is not set up to model such a complex process, so one was made to decompose and the other to evaporate. Measurements of the change in optical properties during this process would need to be made as well.

The heat capacities of the two phthalocyanins are estimates from a paper that investigated pure, metal-free phthalocyanin. Nothing better seems to be available.

The refractive index of Phthalocyanine Blue as function of wavelength was read from the graph of [100] with the software im2graph [55]. The refractive index of Iron Oxide Red as function of wavelength is taken from [82]. The heat capacity of Iron Oxide Red varies considerably with temperature, as can be seen from Figure 5 of [96]. However since the code currently does not allow for any variation, an average value was used, that of around 250°C.

	Dioxazine Violet	Carbon Black	Titanium Dioxide (rutile)
Pigment Number	23	7	6
Minimum Content [%]	0.1*	2*	3*
Maximum Content [%]	0.2*	3*	5*
Prevalence [%]	4*	10	14
Molar Mass [$\frac{g}{mol}$]	589.476 [75]	12 [101]	79.865 [75]
Refractive Index (real part)	1.787 [102]	1.9 [82]	-1
Thermal Conductivity [$\frac{W}{m \cdot K}$]	0.5*	140 [101]	11 [103]
Density [$\frac{g}{cm^3}$]	1.53 [102]	2 [104]	4.26 [103]
Heat Capacity C_p [$\frac{J}{kg \cdot K}$]	1000*	710 [101]	690 [105]
Forward Scattering Ratio	0.37	0.99	0.8
Melting Point [°C]	-1	3550 [101]	2100 [105]
Latent Heat of Fusion [$\frac{J}{kg}$]	0	105000 [101]	900000 [105]
Boiling Point [°C]	-1	4027 [101]	2972 [103]
Latent Heat of Vaporization [$\frac{J}{kg}$]	0	715000 [101]	100000*
Decomposition Point [°C]	385 [102]	-1	-1
Decomposition Energy [$\frac{J}{mol}$]	100000*	0	0
Decomposition Product 1	Carbon Black	none	none
Moles of DP1 per mol of material	20	0	0
Decomposition Product 2	none	none	none
Moles of DP2 per mol of material	0	0	0

Table A.6: Properties of Pigments Part 3

Dioxazine violet actually decomposes through two stages (320 - 440°C, 460 - 740°C) in an inert atmosphere [99], but again this cannot be modelled by the code, so only one value is used. In practice what happens in the code though is that it anyway has to wait until its plastic has finished decomposing before doing so itself.

The refractive index of Titanium Dioxide as function of wavelength is taken from [82].