

Integration of Life Cycle Assessment Into Agent-Based Modeling

Toward Informed Decisions on Evolving Infrastructure Systems

Chris Davis, Igor Nikolić, and Gerard P. J. Dijkema

Keywords:

bioenergy
complexity
industrial ecology
matrix
network
ontology

Summary

A method is presented that allows for a life cycle assessment (LCA) to provide environmental information on an energy infrastructure system while it evolves. Energy conversion facilities are represented in an agent-based model (ABM) as distinct instances of technologies with owners capable of making decisions based on economic and environmental information. This simulation setup allows us to explore the dynamics of assembly, disassembly, and use of these systems, which typically span decades, and to analyze the effect of using LCA information in decision making.

We were able to integrate a simplified LCA into an ABM by aligning and connecting the data structures that represent the energy infrastructure and the supply chains from source to sink. By using an appropriate database containing life cycle inventory (LCI) information and by solving the scaling factors for the technology matrix, we computed the contribution to global warming in terms of carbon dioxide (CO₂) equivalents in the form of a single impact indicator for each instance of technology at each discrete simulation step. These LCAs may then serve to show each agent the impact of its activities at a global level, as indicated by its contribution to climate change. Similar to economic indicators, the LCA indicators may be fed back to the simulated decision making in the ABM to emulate the use of environmental information while the system evolves. A proof of concept was developed that is illustrated for a simplified LCA and ABM used to generate and simulate the evolution of a bioelectricity infrastructure system.

Address correspondence to:

Chris Davis
Jaffalaan 5
2628 BX Delft
the Netherlands
c.b.davis@tudelft.nl

© 2009 by Yale University
DOI: 10.1111/j.1530-9290.2009.00122.x

Volume 13, Number 2

Introduction

Understanding issues related to sustainability often involves recognizing that the world is composed of a mix of economic, environmental, and social systems. These systems are not generally mutually exclusive but are interconnected with each other. For example, the economy is inherently dependent on the environment for raw resources, and it is also the result of a multitude of social decisions.

Modeling has been useful as an exploratory and analytical technique for understanding these systems. One of the challenges in dealing with sustainability, however, relates to the modeling of interconnected systems. Modeling these can be very difficult, as it is not unusual for a change or decision in one part of the system to have rippling effects elsewhere throughout the system. For example, it can be difficult to understand how a management decision on investment in a certain type of technology could have environmental impacts due to economic conditions further up the supply chain for that technology.

To approach these problems of system interconnectedness and interdependency, in this we examine whether and how the combination of two types of tools can provide a step forward. One of these tools is life cycle assessment (LCA), which has become a standard tool for environmental analysis of our technological production systems. The other is agent-based modeling (ABM), a simulation tool that is useful for generating complex network systems that result from the decision making of individual entities. LCA is very good at analyzing complex network structures, although in its current form, it is a tool for linear modeling of static systems. Furthermore, traditional LCA does not examine economic and social concerns (Guinée 2002). Conversely, ABM provides a means to create nonlinear dynamic systems, which can be specified to include social and economic aspects.

We hypothesize that combining these tools can result in an interesting synergy. In this article, we present a proof of concept that this integration is possible and provides a means for modeling interdependent sociotechnical systems. Our intent is to show that this combination can provide additional information about sociotech-

nical systems to aid in decision making, rather than specifying the exact decision that should be made. It is still up to the decision maker to use his or her own judgment in interpreting the different aspects of this information.

The modeling framework developed has been used in a case study investigating bioelectricity production in the Netherlands. This is a topic of much interest, because using biomass as a fuel is assumed to offset fossil carbon dioxide (CO₂) emissions. Because CO₂ emissions are contributing to a global problem, however, one needs to examine the biomass supply chains to quantify the gains or even losses that may be occurring. The agents in the simulation were defined on the basis of technologies in these supply chains.

The sections below examine the foundations for LCA and ABM, the methods used for integrating these two tools, and present then a proof of principle demonstrating repeated accounting-type LCAs for a dynamic supply network. This is followed by a discussion of insights gained, along with the future outlook and conclusions.

Foundations

Both LCA and ABM are tools that employ systems approaches. That is, both of them are used to investigate systems with multiple interacting elements. Simply put, they both examine “things interacting with other things.” LCA looks at interacting technologies in the form of supply chains, whereas ABM is a more general tool that looks at the interactions of a group of agents. Although they are different tools, they are not fundamentally dissimilar in their approaches to representing systems. LCA stores the interactions of technologies in a matrix format. An agent-based model can be set up to represent interactions in a network format, where agents are represented by nodes and edges represent connections between the agents. From this insight, one of the keys to integrating these tools was the realization that the network data structure of the agent-based model is equivalent to the matrix data structure of the LCA, which allows the modeled system to be represented in both formats.

One large difference between LCA and ABM is that within an LCA, the connections between technologies are fixed, whereas an agent-based

model does not have to define fixed connections. Within an LCA, we know exactly which technologies use the output from another technology. Within an agent-based model, however, agents can be defined as having different options concerning those with whom they interact. Essentially, LCA has a static structure, whereas ABM can allow for changing and evolving structures.

On the basis of this, we suggest using ABM to generate a complex dynamic system and an LCA to analyze it at discrete simulation time steps. Instead of defining the world as a set of technologies with fixed connections, this approach breaks those connections. Technologies are now represented as agents whose owners must trade with other agents to buy inputs and sell their outputs. Supply chains within the simulation self-assemble due to these interactions between individuals. By appropriate definition of agents' decision rules and identity and by inclusion of a sufficient number of actors in the simulation, this concept can be extended to emulate socioeconomic decision making and yield dynamic models of supply networks. By extension, ABM could ultimately be used to model whole economies or societies.

The next two subsections give a brief overview of LCA and ABM to illustrate their approaches, limitations, and opportunities for integration.

Life Cycle Assessment

LCA is a methodology and tool used to analyze the environmental burden of products at all stages in their life cycle—from the extraction of resources, through the production of materials, product parts and the product itself, and the use of the product to the management after it is discarded, either by reuse, recycling or final disposal. (Guinée 2002, 5)

This comes from a recognition that environmental impacts have a systemic origin. In other words, by choosing one particular product or service, we are indirectly supporting environmental impacts that may occur several stages away from us in the supply chain that brings a product into being or provides a service.

Today's production networks have emerged through the accumulation of choices made by a multitude of actors driven by a variety of ratio-

nales. LCA provides an accounting of the total environmental impact of this network and relates it to a particular product or service. Although it thus may not be immediately apparent, LCA is at its heart a type of network metric or a means of calculating characteristics of a network. In particular, it is a way of understanding a network's structure and performance from the viewpoint of one node within a network that represents the functional flow or reference product.

LCA has been a valuable tool in helping us holistically approach the problem of emissions reductions. It is not enough to reduce emissions at a local level; we must be aware of emissions that occur beyond our horizon as a consequence of our actions. We cannot fully achieve systematic emissions reductions without recognizing these networks of interdependence between industries.

The idea behind LCA is very powerful, although in its current implementation, we believe its potential is not fully utilized. By connecting an LCA to an agent-based model, we anticipate that we can overcome many limitations.

Currently, LCA views the world as being composed of static connections between technologies. It is also a linear model in that if production of one good is increased, the flows of the upstream technologies are scaled proportionally as well. Additionally, it does not include any time or location aspects, which means that impact assessment for certain emissions may be too high, especially if those emissions are spread out widely over different time periods and regions (Guinée 2002).

Agent-based Modeling

The concept of representing elements of a system as individual agents is central to ABM. As ABM is used in very diverse scientific fields, however, it has no exact definition. One explanation is given by Shalizi (2006, 35), who states that

an agent is a persistent thing which has some state we find worth representing, and which interacts with other agents, mutually modifying each others' states. The components of an agent-based model are a collection of agents and their states, the rules governing the interactions of the agents, and the environment within which they live.

Kauffman (Ball 1999 as cited in Shalizi 2006, 35) provides a much more succinct definition by stating, “An agent is a thing that does things to other things.”

In this sense, agents serve as a flexible abstraction of the real world. The key idea is that system behavior is controlled not by a single entity but by the accumulated effects of decisions made by the individual agents. When these agents interact with each other, they form networks, which can then give rise to complex systems. For the agent-based model in this article, agents must trade to buy raw materials and sell their products. Agents are only limited in that they must find a suitable trading partner who makes a feedstock they need or is willing to buy what they produce. If an agent is not profitable, it will “die” and be removed from the simulation. In the ABM simulation, the technology options available and the decision-making rules specified for the agents are the ingredients used to generate the overall system structure and content.

ABM is especially suited for types of problems for which equations either cannot be solved or are impossible to formulate (Axtell 2000). Compared to other tools, such as system dynamics, ABM allows the modeling of more complex dynamics, because the system structure can also change during the simulation. In system dynamics, a fixed interaction structure is defined and maintained, which means that even before the simulation is started, one must define how the parts are connected. Using an ABM, one only defines an interaction space in the form of the types of interactions that the agents are allowed to have. This characteristic of ABM is important, given that many of the largest problems we face today are difficult to solve because they exhibit characteristics of complex adaptive systems (CASs). These have been defined by Holland as


a dynamic network of many agents (which may represent cells, species, individuals, firms, nations) acting in parallel, constantly acting and reacting to what the other agents are doing. The control of a CAS tends to be highly dispersed and decentralized. If there is to be any coherent behavior in the system, it has to arise from competition and cooperation among the agents themselves. The overall behavior of the system is the result of a huge number of decisions made every moment by many individual agents (Waldorp 1992, 145).

We use this concept of an agent as a basis for modeling of technological infrastructures. We expand the concept by incorporating insights from studies on large-scale sociotechnical systems (Bijker et al. 1987; Nikolić et al. 2007, 2008). Companies are defined as agents who own and manage technologies, as illustrated in figure 1. These agents must make decisions on how to operate their technologies and trade with other agents to sell their products and buy feedstocks. The agents encounter real-world constraints in that to survive, they must remain profitable. Technologies owned by agents are constrained in that they must balance their mass, energy, and monetary flows. As the technologies operate, economic and environmental performance data are gathered and then fed back to the agents to influence their management decisions.

Integration of Life Cycle Assessment into Agent-based Modeling

Both ABM and LCA deal with networked structures. Given a suitable system conceptualization, network metrics can be calculated. In an agent-based model, the network is represented as a graph consisting of nodes connected by edges.

Figure 1 Example of sociotechnical systems modeled. Nodes in the Technology Owners box are companies. CO₂ = carbon dioxide; CH₄ = methane; N₂O = nitrous oxide.






Figure 2 Equivalence of graph and matrix implementations.

In an LCA, one conceptualizes the network as a matrix where nonzero values represent the magnitude of flows between nodes, whose identities are indicated by the row and column locations. These two concepts are equivalent, as shown in figure 2, where a graph is mapped to a matrix. One of the keys to combining ABM and LCA has been to interchangeably represent the network structure in both forms. Graphs are very useful for keeping track of complex evolving structures, and algorithms can be written to navigate and retrieve relevant information from these structures. The matrix representation of LCA is the only computationally feasible means for calculations involving large systems, however (Heijungs and Suh 2002).

Combining ABM and LCA requires that we examine the data structures used by both. A typ-

ical data structure for an LCA is illustrated in figure 3. Most of the data are stored in two matrices. The technology matrix has information about material flows between technologies, and the emissions matrix has information about the amount of pollution produced by each technology. More detail on this is given in the *Calculations* section below.

Ontologies

Figure 4 shows a sample portion of the data structure used in the agent-based model. As mentioned previously and illustrated in figure 2, the agent-based model employs a graph-based data structure. This description is simplified in that the graph is called an ontology, which is a way of logically structuring and storing information so

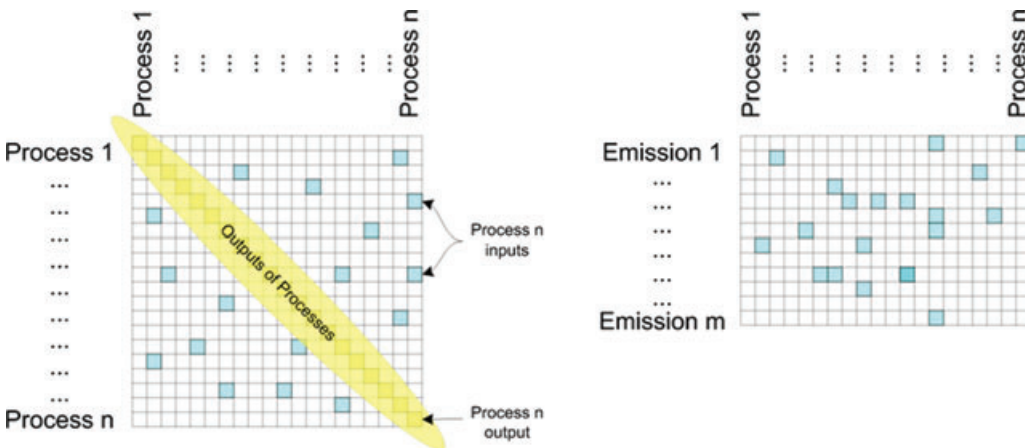


Figure 3 Illustration of life cycle assessment data structure. The left matrix is for material flows between technologies. The right matrix is for environmental interventions from technologies. Squares represent nonzero matrix elements.

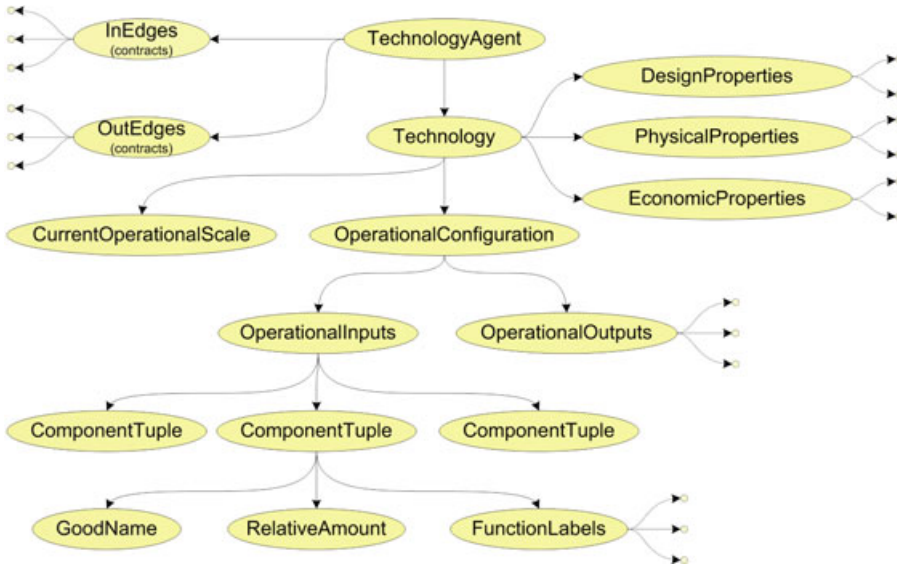


Figure 4 Partial illustration of the ontology used for agent-based modeling data structure.

that it can be understood by humans and implemented in and retrieved by executable computer code. An ontology often specifies “is a” and “has a” relationships for objects. For instance, an apple is a fruit and has a red color. In particular, an ontology is composed of classes, properties, and instances. A class can be used to describe a generic type or group that something belongs to, such as a *fruit* class. This fruit class can then have properties describing attributes, such as name, shape, and color. This defines the basic data structure through which we can describe fruit. We can then create multiple instances, such as for “this apple” and “this pear,” using the template specified by the class to indicate the specific name, shape, and color of this apple and this pear.

We use the ontology classes to describe the data structure of the agent, which defines the types of information it can know about itself. For instance, all agents have an *OperationalConfiguration* containing a list of inputs and outputs, represented as *ComponentTuples*, which hold information about the amount of goods needed for manufacturing and the amount of goods produced. The agents also have *InEdges* and *OutEdges*, which are used for contracts for items bought from and sold to other agents. This common data structure allows agents to communicate

with each other about their properties and aids in retrieval of information later on, whether by the agents themselves or for data-logging purposes.

The instances contained in the ontology serve as a database. When we create an agent, it gets its own data structure, which is partially filled in with information, such as its required inputs and outputs, construction costs, and fixed costs. As it operates during the simulation, it fills in other information, such as its profits and lists of contracts made.

Assembling the Network

Because the other agents have the same type of data structure defined by the ontology, we can see the contracts as facilitating the creation of the network of actors trading among one another. The information contained in the contracts is used in calculation of the LCA, as described further below in the *Technology and Emissions Matrices* section.

For *TechnologyAgents* to sell their goods and procure the necessary inputs, they must trade with other agents. When a *TechnologyAgent* needs an input, it asks all other *TechnologyAgents* for a possible contract. The agents offering goods are able to set the price in the contract on the basis of

their own decision making. Once an agent that needs an input has the complete list of offered contracts, it will decide to accept the contracts that it deems best.

To help keep track of all the mass flows that one will need in calculating an LCA, the concept of a *WorldMarket* and an *Environment* agent is used. Essentially, these agents ensure that everything in the simulation comes and goes from somewhere. The *WorldMarket* agent serves to connect the *TechnologyAgents* to the outside market. Although *TechnologyAgents* can buy and sell among themselves, they may find that no one is supplying the good they want or that no one is buying what they produce. In these cases, they can trade with the *WorldMarket*. The *Environment* is the agent that collects all emissions from the agents and provides the same inputs as the real environment does, such as air for processes involving combustion.

Constructing the LCA

To provide additional information for the LCA calculations, in this proof of concept we employ a hybrid approach whereby the *WorldMarket* agent's supply chains are defined on the basis of structures contained in a life cycle inventory (LCI) database. This is illustrated in figure 5: The evolving system is represented by the agent-based model, which is linked with the appropriate processes and supply chains in the database containing LCI information. Whenever an agent buys from the *WorldMarket*, due to the connection with the LCI database, we are thus able to retrieve the upstream environmental emissions resulting from that purchase. This linkage results in a coupling of systems, so that we now have a very large economic network composed of a core of dynamic actors (the *TechnologyAgents*) surrounded by static actors (based on the technologies described in the database). This allows a simulation in which we have a "lens" of dynamic behavior and evolving structures surrounded by a static definition of the world, as depicted in figure 6, which shows how conceptually the model structure changes over time. As indicated, the set of agents in the agent-based model can increase or decrease with time, whereas during the whole simulated period we are able to trace each



Figure 5 Systems representation in the hybrid model. Dynamic supply chains in the agent-based model connect to static, predefined supply chains.

agent's upstream flows through the structure defined in the LCI database and linked to the ABM data structure.

Technology and Emissions Matrices

Figure 7 illustrates the data structure used to combine LCA within ABM. The simulation is paused at every time step, and all the information about inputs and outputs (i.e., traded flows and emissions) is retrieved from the simulation's data structure. This information is then processed and placed into a technology and an emissions matrix, where calculations can begin. At this stage, economic allocation occurs for technologies with multiple outputs (multifunctional processes). The implementation of this particular allocation method is relatively easy, with the calculation performed automatically by the software, and the only user input is a set of initial prices for goods. With additional programming, other types of allocation could be used. For this proof-of-concept stage, however, economic allocation was used, with the recognition that programming for other forms of allocation could be worthwhile extensions to the methodology and models in the future.

In both of the matrices, one section is composed of information from the external database




Figure 6 Hybrid model over time. The agent-based model evolves, and predefined supply chains are used to provide additional information about environmental impacts that exist outside it.

containing LCI information. This portion stays constant throughout the simulation. Additional rows and columns are added to represent the agents themselves. Because part of the matrix is composed of a dynamic system, the matrix will grow and shrink over time in accordance with the number of active agents present at each time step.

Figure 7 illustrates how the links are made between different processes within the technology matrix. The diagonal represents the total output of each process. By reading down a column, we can see the different amounts of inputs used for a single process. By reading across a row, we can see how one process supplies several different processes. The top left quadrant is composed solely of the LCI database, whereas the top right

quadrant contains information about where an agent buys goods from the *WorldMarket*, which is connected to the LCI database. The bottom right quadrant shows trading between individual agents. The bottom left quadrant is not filled in, because the processes in the LCI database do not buy goods from the agents and only use predefined supply chains. An emissions matrix is set up that is similar, except that the columns represent individual processes, and rows represent types of emissions.

Calculations

An LCA system can be written in the form shown in equation 1 (Heijungs and Suh 2002). Here, \tilde{A} is the technology matrix, which is also shown in figure 7. The vector y is the demand on the LCA system and represents the functional unit. For instance, if one wanted to perform an LCA on 1 kilowatt hour (kWh) of electricity from a specific process, this vector would be all zeros, except for a 1 placed at the location corresponding to the process of interest. The vector s represents the scaling factors that one must apply to the technology matrix to provide the functional unit. When we perform an LCA, the values for \tilde{A} and y are already defined by the user, and we must solve for s , which results in equation (2).

$$\tilde{A}s = y \quad (1)$$

$$s = y\tilde{A}^{-1} \quad (2)$$

Solving for s is a nontrivial exercise. Although equation (2) shows a simple, elegant formula for solving a complicated system, there are quite a few implications behind this. Performing an LCA calculation means that one must balance supply and demand by appropriately scaling each technology in the matrix. Because an LCA is done over a supply chain, scaling one technology means that one must scale the connecting technologies as well to ensure that the mass balance is satisfied for the entire system. For complex interdependent supply chains that include feedback loops between technologies, this turns into a very complicated operation that truly necessitates the use of a matrix.

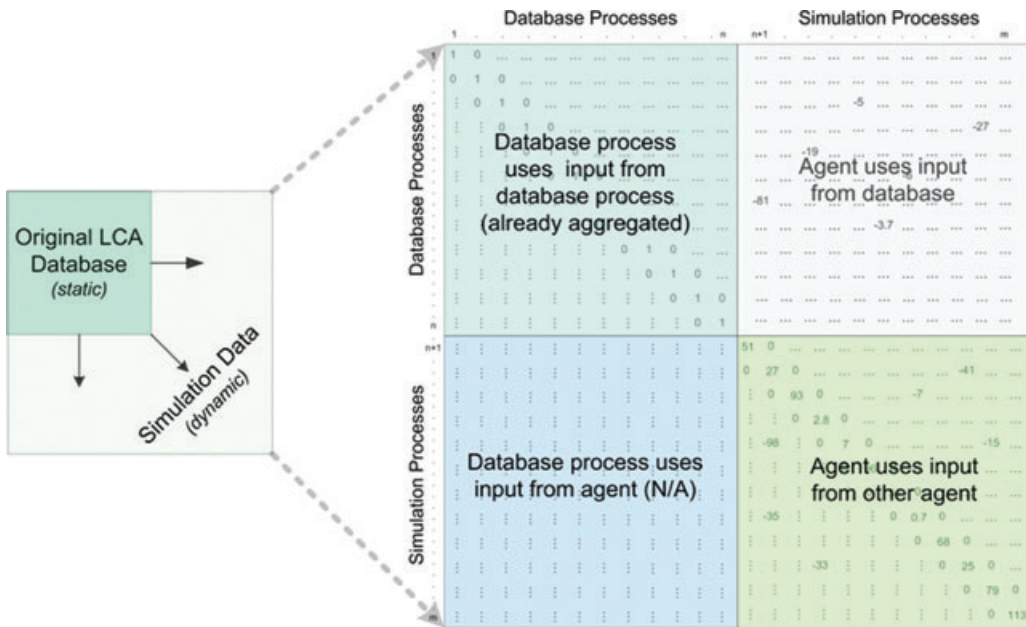


Figure 7 Creating the technology matrix for life cycle assessment (LCA) calculations. The figure corresponds to \tilde{A} shown in equation (1).

Although a matrix makes these types of calculations computationally feasible, there are still issues with the amount of time needed, especially if one desires to perform multiple LCAs on a dynamic system, as demonstrated in this article. A major bottleneck lies in calculating the inverse of the technology matrix, as required by equation (2). In the worst case, most matrix inversion algorithms are said to run in $O(n^3)$ time, which means that as the size of the matrix increases by n , the maximum calculation time needed to invert it increases at a rate of n^3 . In other words, one may need a fraction of a second to invert a technology matrix of 100 processes, whereas one may need 2 minutes for one involving 1,000 processes (Heijungs and Suh 2002).

The iterative matrix inversion algorithm proposed by Peters (2006) was designed to overcome this problem. This particular algorithm runs in Cn time, where C is generally a very large constant. This means that algorithms running in $O(n^3)$ time have an advantage when used for small matrices, but this iterative algorithm has a significant advantage in dealing with large matrices. In practice, this algorithm has been in-

valuable in the creation of the model, as it is not uncommon to conduct dozens of simplified LCAs during each simulation time step.

Within the simulation, a simplified LCA is generated for every agent during every time step. The functional unit is chosen on the basis of a single unit of the main flow from the agent and currently remains the same throughout the simulation. In other words, power plants will have a functional unit of 1 kWh of electricity, whereas another agent that makes palm oil will have a functional unit of 1 kg of palm oil. We performed a simplified LCA to evaluate climate change on the basis of emissions of CO₂, methane (CH₄), and nitrous oxide (N₂O). We used the characterization factors of the LCA methodology to estimate the global warming potential (GWP) in terms of CO₂-equivalents. The models could readily be extended to more extensive LCAs considering the full LCI and more impact categories, and these are considered areas of future extension of the methodology. At this proof-of-concept stage, however, the models were limited to this simplified LCA. Only contribution to climate change was evaluated, as it was considered

to be the most important impact category for this case study, which focuses on energy provision.

Overview of Steps

To summarize, the list below gives an overview of the steps by which LCA was integrated into the agent-based model, as described in the paragraphs above. A more complete description is given in the original work (Davis 2007).

1. Run the simulation for one time step, allowing the agents to trade with each other.
2. Construct the technology and emissions matrices for each agent.
 - (a) Collect information on all the flows in the system.
 - (b) If an agent trades with the *World-Market*, link it to the supply chain described in the LCI database.
 - (c) Perform economic allocation if necessary.
3. Calculate an LCA for each agent.
 - (a) Determine the functional unit to be used for each agent, on the basis of the main product produced. An agent's functional unit is presently constant, as agents currently do not produce different products over different time steps.
 - (b) Calculate an LCA for each agent on the basis of functional unit.
 - (c) Aggregate CO₂, CH₄, and N₂O into a single category on the basis of their GWP values.
4. Send results from the LCA to the agent for use in decision making.
5. Run the simulation for the next time step, and repeat.

As has been shown, it is possible to use a simplified LCA within an agent-based model. This opens up new realms when we consider that sustainability is an emergent property of a network (Allenby 1999), which LCA, to an extent, inherently recognizes. Sustainability cannot be measured at the level of individuals; it can only be measured at the level of the total system itself. The value of this setup is that it allows us to create a dynamic model that can explore the possible effects of independent as well

as interconnected decisions. This can allow for a dynamic representation of the system, including feedback loops and other dynamic system properties.

Proof-of-Principle Illustration

Bioelectricity

The types of agents used in this case study were based on inventory data gathered during a separate study investigating bioelectricity production in the Netherlands (Van der Voet et al. 2008). This study was performed to create a calculation tool to allow for comparison of different bioelectricity supply chains with regard to their greenhouse emissions. The greenhouse gas accounting was limited to CO₂, N₂O, and CH₄. Several different electricity production methods appropriate to the Dutch situation were chosen, in addition to the types of biomass that would likely be used to supply them. Inventory data were collected for all stages, from biomass production to transportation, processing, and final conversion to electricity. Several types of biomass were chosen from different parts of the world, and the electricity production methods ranged from small scale (10 MW) to large scale (500 MW). From this information, an LCA was performed of the different chains, which allowed for comparison of the greenhouse gas emissions resulting from 1 kWh of electricity production from each of these different possible supply chains. Other aspects, such as land use change, biodiversity impacts, and economic development, were not investigated.

This study by Van der Voet and colleagues (2008) was, in turn, part of a larger investigation commissioned by the Dutch government to explore issues related to the use of biomass for energy and material production. This was performed by the Sustainable Production of Biomass Project Group (Commission Cramer), which investigated criteria for the evaluation of sustainable production of biomass (Projectgroep Duurzame productie van biomassa 2006).

We give a brief explanation here to illustrate the system and the nature of the problem that was modeled.

Initial Settings

Each simulation starts with three large fossil-based power plants. These technologies are coal cofired with biomass, coal cofired with syngas, and natural gas and heavy oil cofired with bio-oil. Each of these large technologies is able to cofire biomass at a fixed ratio, in addition to being able to use pure fossil fuels.

The system is driven by the demand for electricity, which has an effect on the number of agents added. The actual initial demand for electricity is set slightly above the production capabilities for the three main fossil electricity plants, which allows for additional, smaller producers to join in to fill the gap.

Agent Addition and Removal

At the beginning of each simulation time step, new *TechnologyAgents* have the opportunity to join the simulation. During this period, candidate types of *TechnologyAgents* are picked at random and asked whether they want to invest. Once 10 in a row have declined, the simulation continues, as it is considered unlikely that any others may want to invest.

A candidate *TechnologyAgent's* investment decision is based on several factors. First, the demand-to-supply ratio for its reference product must be higher than a specified value to avoid oversaturating the market. If this condition is met, then the candidate will evaluate whether it can be profitable under current market conditions by calculating its expected revenue from the sale of its products and then subtracting its fixed operating costs and variable costs of inputs. If this profit is positive, then the *TechnologyAgent* will decide to invest and join the simulation at that time step.

Agents will be removed if they are consistently unprofitable for a specified length of time. They will also be removed if the lifetime of their technology has run out.

Carbon Dioxide Tax

The CO₂ tax is directly related to the fossil CO₂ output for each of a technology's *OperationalConfigurations*. This CO₂ output refers to

local emissions, not the ones indicated on the basis of its GWP, calculated from the LCA. For the simulations run, a parameter sweep was performed using different values of the CO₂ tax in order to find transition points where system changes could be observed.

Agent Operational Decision Making

Many of the agents have multiple *OperationalConfigurations*. An individual *OperationalConfiguration* is a list of the types and amounts of input and output flows that pass through the agent. For example, a coal-fired electricity plant will use a specified amount of coal and limestone as inputs and will generate a specified amount of electricity, ash, and several types of emissions. The sizes of the flows are fixed for the agents and are representative of the amount of the flows in a year. Each of these configurations represents the use of a different feedstock. For example, if an agent wants to cofire soybean oil instead of palm oil, it will have to switch to a different *OperationalConfiguration*. Furthermore, the agent can only choose a single *OperationalConfiguration* per simulation time step and cannot perform actions such as adjusting the ratio of cofiring of biomass to fossil fuels.

Each agent keeps a record of profit and LCA scores associated with each of its *OperationalConfigurations*. When the agent is initialized, it first collects these data by iterating through all of its *OperationalConfigurations* on consecutive time steps. Essentially, the agent wakes up and assesses the economic and environmental state of the world from its own point of view. Once the agent has completed this stage, it will then select future *OperationalConfigurations* on the basis of the decision behavior that has been specified for it. In other words, once it has looked through its options, it will now, at every consecutive time step, make a decision about what it thinks is best. Every time it makes a decision, it will then update the profit and calculate GWP for the *OperationalConfiguration* just chosen. For every *OperationalConfiguration*, there is only a single value stored for profit and a single value for the calculated GWP. Older values are overwritten, and the agents do not employ prediction or optimization techniques but just look at the last values encountered. For this simulation, *TechnologyAgents* are




Figure 8 Global warming potential (GWP) over time; the figure shows results for a single type of agent under different decision types.

able to pick feedstocks on the basis of those resulting in the most profit or the greatest reduction in CO₂-equivalent emissions. In this case, maximizing profit simply means picking the cheapest feedstock, and minimizing GWP means picking the feedstock with the lowest calculated LCA emissions.

In the current setup, all agents were specified to have the same type of decision making throughout a simulation. That is, we ran a simulation in which they would all try to maximize their profit or in which they would all try to minimize their greenhouse gas emissions, as measured through the GWP calculated from the LCA.

Simulation Results

The simulation setup allowed two primary means for the system to change. First, the overall portfolio of electricity-producing technologies could change on the basis of the CO₂ tax that was chosen. All agents would have to pay a tax based on their fossil CO₂ emissions. Second, agents could choose particular feedstocks to max-

imize profits or environmental benefit. As for data quality, the main focus in the case study was on gathering LCI data. The economic data gathered were more uncertain in comparison.

Effect of Decision Making

Figure 8 shows the effect that different styles of decision making can have given a fixed CO₂ tax. When the agents start out, they must evaluate the different feedstocks available to them and record information related to the profitability and LCA scores of each. The agent does not have a long-term memory of these values but rather stores the last values encountered. For example, the number stored for profitability is only based on the agent's profit calculated in the previous simulation time step. No prediction or optimization is done on the basis of previous values. As shown in the graph, once the evaluation stage is complete where the agent iterates through all feedstocks, the agents will consistently choose a feedstock on the basis of whether they are set to maximize profit or minimize their LCA score per each simulation time step.

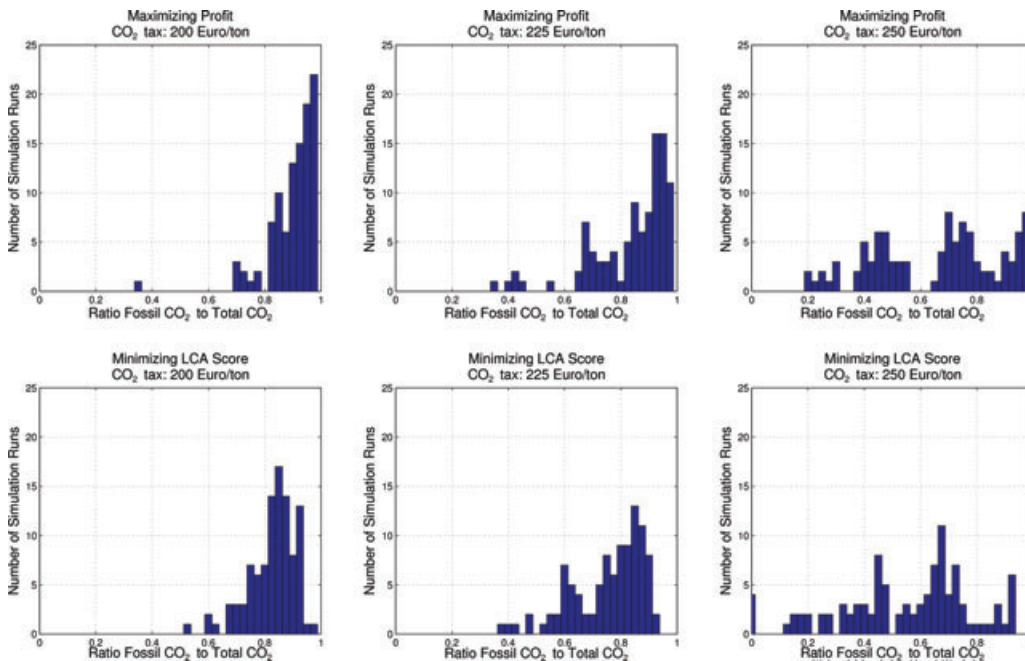


Figure 9 Exploring simulation outcomes under different combinations of carbon dioxide (CO₂) tax and decision-making types. Each histogram represents 100 model runs after 20 simulated years. LCA = life cycle assessment.

The histograms in figure 9 show how accumulated individual agent behavior affects the overall system behavior. These six graphs show permutations of three levels of CO₂ taxation and the two types of decision making. Each graph represents multiple simulations where the agents are only allowed to employ a single specified type of decision making. That is, they will all, during the entire simulation, try to maximize their profit or minimize their LCA score, as specified by the modeler at the start of the simulation. The CO₂ tax increases as one moves from the left column to the right column of figure 9. The top row represents agents trying to maximize profit, whereas in the bottom row all the agents try to minimize their LCA scores. For each permutation of a CO₂ tax and decision-making type, a set of 100 simulations is completed wherein the order in which agents act during each simulation time step is randomly varied. This is our attempt to see whether trends are general and consistent, not an artifact of model construction. The simulation setup thus reflects that agents are not guaranteed to have the same or the best trading partners. For exam-

ple, someone could have bought up all the supply from a preferred partner. As described above, the agent addition process is also randomized, which adds to the inherent variability to be expected in the simulation results.

At the end of each simulation, we tabulated and processed the total emissions for all the agents to calculate the ratio of fossil CO₂ to total CO₂ (i.e., fossil + biogenic CO₂). This gives a rough indicator of the amount of bio-electricity being produced versus electricity being produced from fossil fuels. We then placed each of these measurements in a histogram to see the distribution of possible system outcomes over the set of simulations for a combination of a type of decision making and a specific CO₂ tax rate. These can be understood as showing system attractors that indicate the range of values that may be encountered, rather than showing specific outcomes.

Starting from the left column of figure 9, we see that the agents' decision making does not make a great deal of difference. There is some shift away from complete use of fossil fuels, although

it is not necessarily significant. This is due to the dominance of large electricity producers, which can only cofire a limited percentage of biomass. Although they try to minimize their LCA score, the current technology as defined is incapable of achieving significant emissions reductions. As we move from left to right and look at the scenarios with higher CO₂ tax rates, the types of emissions undergo a noticeable shift, indicating that more bioelectricity is being used, instead of electricity produced from fossil fuels. It also should be noted that the distribution of outcomes is flatter for higher tax rates, which indicates less certainty of one specific outcome. This is due to transitions occurring in the types of technologies in the simulation as fossil-based electricity producers can no longer afford to pay the CO₂ tax. Due to the nature of the simulation, there is no one single way this transition could happen, which leads to a larger variety of potential outcomes over the simulated time frame.

Essentially, we see that the ability of the system to change is defined by the dynamics of competition between the different technologies. In the simulation, new agents are added on the basis of the ratio of supply to demand and the possibility that they can operate profitably. This means that some of the larger fossil-based electricity producers can block out the entry for some of the smaller biobased electricity producers if the market is saturated. Some of these smaller biobased electricity producers can fill in the gaps as old technologies go offline, although they may not be able to make up a large percentage of total electricity production.

From the simulation runs, it was found that the most effective way to reduce CO₂ emissions was to impose a high CO₂ tax rate, rather than only having the agents pick the feedstock that would lead to the lowest emissions. This means that for drastic CO₂ reductions to occur, the cofiring limitation mentioned above is a barrier unless it can be overcome. Otherwise, given the assumptions and framework used for this case study, one must change the portfolio of technologies to achieve these reductions.

In viewing these results, readers should note that these outcomes were case study specific and thus tied to the definitions of technologies used. One characteristic to observe is that the tech-

nologies currently defined mostly formed linear supply chains. Electricity-producing agents had a greater choice of possible feedstocks than agents further up the supply chain. In other words, the number of possible supply chain network configurations was somewhat limited. This hints that other case studies involving different industries with more variety of possible connections could lead to much different results, as a greater variety of networks could emerge. In these cases, supplying information on environmental impacts to agents may have a much greater impact, as the agents would not be so constrained in their options.

Readers should also remember that an agent-based model can generate a large number of data. For instance, while the model is running, we use data logging to record all the monetary, environmental, and mass flows between the agents over time. Although a subset of these data is used to calculate a simplified LCA, many other ways of interpreting the data are possible. The modeler can decide to aggregate the data into system-level indicators or to filter them to examine properties of individual agents. This ability to process data in different ways can aid in the evaluation of different policies or decision-making types.

Discussion

Developers of ABM and LCA may find this merger of tools interesting from their own perspectives. From the point of view of LCA, ours is a modeling structure that allows for the simulation of dynamic, evolving supply chains. Although this is not a true dynamic LCA, it does provide a step forward. For those in the ABM community, this article shows a means to represent environmental impact information from the outside world through the inclusion of an LCI database. Additionally, LCA allows for examination of the supply chains that emerge from the simulation.

What Does Life Cycle Assessment Gain From Agent-based Modeling?

Platform for Analysis of Dynamic Systems

Traditional LCA includes several simplifications, such as linear definitions of technologies, analysis based on steady-state situations, and no

spatial differentiation of environmental interventions or impacts. Heijungs and Suh (2002) talk about challenges researchers encounter in overcoming these simplifications by creating further extensions to LCA in the form of nonlinear models, spatially differentiated models, and dynamic models. Although the method presented in this article does not solve all of the issues Heijungs and Suh mention, it does represent a first step in providing a platform that could conceivably be extended to address them.

Nonlinear models of technologies can be used within an agent-based model if appropriate rules are defined. Such rules could state why individual agents change the production levels of their technologies. Additional rules could specify why agents decide to enter and leave the market and invest in new technology. To define these rules, one must formalize them and program them within the model.

Researchers can deal with spatial differentiation of environmental interventions and impacts to an extent by assigning a location to each instance of a technology. For example, power plants can be given specific geographical coordinates, which allows for maps of emissions to be created. This solution does not deal with diffusion of emissions but rather allows the researcher to at least have data on the sources of emissions. Using ABM with geographic information systems is becoming more common (Gimblett 2002), and this combination of tools could facilitate a more dynamic look at the effects of emissions on a local scale.

The analysis described in this article is dynamic in the sense that the structure of the system generating electricity changes with time. Calculations are done that involve the network state at every simulation time step. This is not a true dynamic LCA, according to the definition of Heijungs and Suh (2002, 194), who stipulate that "in making LCA a dynamic model, processes must be specified according to the time at which they are active for the product under review."

The calculations performed do not incorporate time delays and thus do not track individual goods over time but rather examine the pathways through which goods flow. The agent-based model could conceivably be extended to include waiting times; however, one must also encode in

the model the different factors that create these delays. This can potentially lead to enormous data requirements.

As shown, we have created a data representation and modeling structure that allows for the simulation of dynamic, evolving supply chains. Although it is not a true dynamic LCA, it does provide a step forward, and further extensions to the modeling framework offer a platform for addressing several of the simplifications of traditional LCA discussed above.

Uncertainty Analysis

The use of a simplified LCA within ABM has interesting implications for a type of uncertainty analysis. Some LCA programs include the ability to perform an uncertainty analysis, for instance, by specifying a probability distribution for the emissions from a process. Such a feature can be useful, given the uncertainties due to measurement accuracy and system variability that are often encountered in the collection of inventory data. The uncertainty analysis can be quite important for balanced interpretation of results, as it provides the researcher a means to see how uncertainty from individual system components can propagate to give a range of LCA scores. For researchers comparing two different systems, this analysis will indicate whether their results are statistically distinct or whether there is so much overlap that they cannot confidently determine which one is better.

This type of uncertainty analysis only looks at one piece of the puzzle, however, as it assumes a static supply chain with predefined connections between technologies. If one wants to examine other possible supply chain configurations, one must manually construct them. For the implementation described in this article, the supply chains assemble themselves on the basis of economics and other system conditions. This flexibility can show us the range of possible supply chain configurations that can exist given the conditions in the simulation. In other words, we can test which supply chain configurations are likely to emerge given specified simulation parameters. These parameters may take the form of a decision-making type or a certain tax regime. One may find that there is only a single possible configuration or perhaps many. Additionally, one

could test whether a policy inadvertently encourages the emergence of undesirable supply chains.

Running the simulation multiple times with fixed parameters can lead to a type of uncertainty analysis in which we record the results of LCA calculations by the agents. These results will not necessarily be the same over time, and we can then use these data to create a distribution of possible values. This distribution will show the range of emissions possible given changes in the evolving supply chain.

Because ABM is inherently a dynamic tool, this development leads to interesting implications. Not only can we perform an uncertainty analysis based on structural dynamics, but we can do it under different system conditions, such as policy regimes or economic circumstances, that we choose to program into the model.

Limitations

The use of an LCA within an ABM only makes sense in certain circumstances, on the basis of the type of research question being investigated. This technique is especially suited for looking at environmental performance of evolving systems. For example, if one seeks to find out the environmental performance of current systems in a specific configuration, then a traditional LCA would be more useful, due to the ease of implementation and the nature of the investigation.

This methodology represents a level of complexity above a normal, traditional LCA due to additional data requirements. The integration of LCA into ABM can be very valuable when there are many possible permutations of supply chain elements that can lead to nonlinear effects. The same is true if decision making can lead to a wide variety of outcomes. If these conditions do not exist, then this type of analysis may not be very valuable. This is similar to LCA itself, for which studies may involve different levels of detail depending on the needs of the research.

The current methodology employed does not attempt to find the globally optimal solution but merely seeks a solution in which each agent is selecting a feedstock with the lowest cost or GWP. What we have shown does, however, provide a platform that can be expanded to test different

theories about how system-level emissions reductions may be realized.

The implementation described in this article has many simplifications that can be dealt with in future work. For instance, agents only make decisions on the basis of current information and do not consider historical information or make any type of prediction. There are no lead times for technologies, and the construction of new technologies is instantaneous. Additionally, innovation is currently not considered.

What Does Agent-based Modeling Gain From Life Cycle Assessment?

Developers of agent-based models can benefit from concepts of LCA in several ways. Two such benefits are explored below. First, the LCA methodology allows for a type of structural analysis to be performed. Second, the use of the LCI database in this article demonstrates how this information can be used to abstract the world outside the simulation.

Structural Analysis

The value of ABM lies in its ability to generate complex emergent behavior. Although this is a benefit compared to other types of modeling techniques, it can also present a challenge, as one needs a means to analyze the systems that emerge. LCA is one such method. It is important to realize that LCA is really a particular implementation of a class of algorithms meant to analyze network structure. LCA builds on the achievements of the fields of ecology and economics, which have long been concerned with studying money, material, and energy flows within systems (Suh 2004). To put this in perspective, we should not just think about combining ABM and LCA. Instead, we should think about combining ABM with various forms of material and energy flow analysis techniques.

Abstracting the World

As already mentioned, the *WorldMarket* of the agent-based model is actually a portal to an LCI database. This integration means that we can simulate a system in which we have a “lens” of complex dynamic behavior where agents interact with a static definition of the world. One

can further refine this concept to create a more sophisticated representation of the world. For instance, one could allow the *WorldMarket* to use macroeconomic input–output tables. Such a combination would allow the agents to operate within a more realistic macroeconomic environment while also allowing one to simulate microeconomic effects (Peters and Brassel 2000).

Outlook

The work shown is the result of trends in the development of information technologies, which are changing the ways we handle massive amounts of complex data. This is coupled with advances in the understanding of complexity science and a shift toward what some would call the generative sciences, in which the question “Can you explain it?” is replaced by the question “Can you grow it?” (Epstein 1999). Claims of understanding are backed by a model whose structure emerges as a result of rule sets based on those claims. The model can then act as a learning tool that allows us to test different theories about how the system operates. This can turn into an educational feedback loop, where the results of the model inform the designer, who then updates the model to match observed real-world behavior.

We can get closer to this vision through expanding the work mentioned in several ways. First, we can use more realistic decision making and scenarios to better model the dynamics of the systems we are interested in. Second, the integration of existing industrial ecology tools can help us understand the patterns that emerge and provide the use of real-world data to define the world that agents exist within.

More Realistic Decision Making and Scenarios

The decision making employed by the agents in this study was rather basic, although it could be made much more sophisticated. For instance, instead of only examining information from the previous simulation tick, the agents could have a memory of information encountered throughout the simulation for use in their decision making. One could also use different optimization techniques to try to better balance complex trade-offs.

As for more complex scenarios, one could consider analyzing effects from economies of scale or incorporating learning effects. The learning effects could take the form of reduced investment costs and greater efficiency with each additional instance of a particular technology. One could also test different types of policy regimes and economic circumstances to elucidate their effect on the performance of the system. For instance, one might investigate what could happen if the price of a certain feedstock suddenly went up or if financial incentives were given for certain types of technologies.

Integrating Tools

As more processing power becomes available, it is easier to overcome limitations and criticisms of the current analysis tools through more sophisticated calculations that are able to handle increasing amounts of data. This is seen with implementations such as hybrid LCA, in which economic input–output tables are included as a means to overcome the problem of limited system boundaries.

Although it is clear that the capabilities and sophistication of tools such as LCA are expanding, we should also consider the shape of this expansion. In essence, are we making this tool deeper or broader? Are we diving deeper and trying to calculate more accurate numbers, when perhaps a different, broader insight into the system is necessary? Are we closer to understanding the complex relationship among human equity, economy, and the environment? It is understandable that LCA was initially developed at a time when only static situations could be analyzed; however, the growth in computing power gives us the opportunity to do so much more.

Figure 10 illustrates an environmental analysis “toolbox” commonly used for exploring problems at different resolutions of flows and industrial networks. Material flow analysis (MFA) examines flows of resources of a region on an aggregated mass basis. An MFA examining flows at a substance level (e.g., chlorine, cadmium) is referred to as a *substance flow analysis* (SFA). A process-level MFA (PMFA) is an MFA done at the level of individual processes. Physical input–output tables (PIOTs) are similar

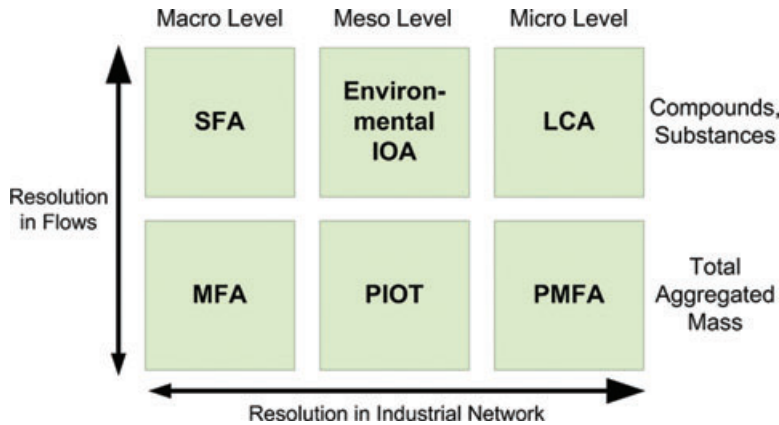


Figure 10 Approaches of quantitative materials and energy flow analysis in industrial network systems (Suh 2004, used with permission). SFA = substance flow analysis; IOA = input–output analysis; LCA = life cycle assessment; MFA = material flow analysis; PIOT = physical input–output analysis; PMFA = physical material flow analysis.

to the input–output tables (IOTs) that describe monetary flows between industries, except that they document physical flows. Environmental input–output analysis is also similar in that it links environmental statistics to the information contained in IOTs (Suh 2004).

Although this article focuses on LCA, we believe that the integration of these other industrial ecology tools with the ABM platform should be explored as well. These tools can be used to analyze what has emerged or to add information about the outside world that the agents operate within. For example, just as an LCI database was connected to the *WorldMarket* agent, the IOT-based tools may be connected to it as well. Although this has not been attempted yet, the implementation may have parallels to the hybrid analysis techniques that seek to combine LCA and input–output analysis, as discussed by Suh and Huppes (2005).

It is interesting that input–output analysis allows the creation of potentially large, evolving networks containing connections based on monetary, mass, and environmental flows. We can grow these networks, and, because an ontology is used for their representation, we can understand what each of the edges in such a network represents. We can then create algorithms that navigate these networks to extract information, such as an MFA. In other words, we have the

potential to generate large amounts of data about evolving systems that can be analyzed in multiple ways while coupled to real-world data. This vision is not a reality yet, but the work shown indicates that it may be possible.

Conclusions

This article has demonstrated a method for the creation of a data representation and modeling structure that allows for the simulation of dynamic, evolving supply chains. An LCA has been used to evaluate this network in the form of a repeated accounting-type LCA. The networks that emerge are subject to and caused by different types of decision making, technological constraints, and economic factors. Some of the challenges addressed involve connecting the data structures of these two tools and reducing the computational time needed by the LCA. Some of the simplifications of traditional LCA can be overcome either by the demonstrated work or by extensions that have been proposed.

Future development can proceed in several directions. First, the current limited implementation of LCA could be expanded toward a full-featured LCA involving improvements such as inclusion of more environmental impacts and impact categories. Integration of the ABM platform with other industrial ecology

tools should be investigated as well. Work is also needed on incorporating more realism into the systems modeled. A large opportunity exists to incorporate more domain-specific knowledge in areas such as decision making and economics. For instance, one may program the agents to employ multicriteria decision making or to evaluate future investments based on standard economic calculations. Many further improvements can be made, and it is hoped that the presentation of the work shown can lead to further discussions of possibilities that can be realized by new types of tool development for industrial ecology.

Acknowledgements

We acknowledge Dr. Ester Van der Voet for her supervision during the thesis on which this article is based. Additionally, we thank Dr. Reinout Heijungs for discussions on life cycle assessment that led us to overcome technical barriers and enabled this work.

References

- Allenby, B.R. 1999. *Industrial ecology: Policy framework and implementation*. Englewood Cliffs, NJ: Prentice Hall.
- Axtell, R. 2000. Why agents? On the varied motivations for agent computing in the social sciences. Tech. rep. working paper no. 17. Washington, DC: Brookings Institution.
- Ball, P. 1999. *The self-made tapestry: Pattern formation in nature*. Oxford, UK: Oxford University Press.
- Bijker, W. E., T. P. Hughes, and T. J. Pinch. 1987. *The social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, MA: MIT Press.
- Davis, C. 2007. Integration of life cycle analysis within agent based modeling using a case study on bio-electricity. Master's thesis, Delft University of Technology, Delft, the Netherlands. <http://www.industrialecology.nl/GraduationReports/2007/Davis%20MScThesis.pdf>.
- Ehrenfeld, J. R. 2007. Would industrial ecology exist without sustainability in the background? *Journal of Industrial Ecology* 11(1): 73–84.
- Epstein, J. M. 1999. Agent-based computational models and generative social science. *Complexity* 4(5): 41–60.
- Gimblett, H. R., ed. 2002. *Integrating geographic information systems and agent-based modeling techniques for simulation social and ecological processes*. New York: Oxford University Press.
- Guinée, J., ed. 2002. *Handbook on life cycle assessment. Eco-efficiency in industry and science*, vol. 7. New York: Springer.
- Heijungs, R. and S. Suh. 2002. *The computational structure of life cycle assessment. Eco-efficiency in industry and science*, vol. 11. Dordrecht, the Netherlands: Kluwer Academic.
- Kauffman, S. A. 2002. *Investigations*. New York: Oxford University Press.
- Nikolić, I., P. J. Beers, and G. P. J. Dijkema. 2007. Facilitating interdisciplinary modelling of complex problems. Paper presented at HICCS-40 Hawaii International Conference on Systems Science, 3–6 January 2007, Manoa, HI.
- Nikolić, I., E. J. L. Chappin, C. Davis, and G. P. J. Dijkema. 2008. On the development of agent-based models for infrastructure evolution. Paper presented at the International Conference on Infrastructure Systems—Building Networks for a Brighter Future, 10–12 November, Rotterdam, the Netherlands.
- Peters, G. P. 2006. Efficient algorithms for life cycle assessment, input-output analysis, and Monte-Carlo analysis. *International Journal of Life Cycle Assessment* 12(6): 373–380.
- Peters, I. and K. H. Brassel. 2000. Integrating computable general equilibrium models and multi-agent systems—why and how. Paper presented at the 2000 AI Simulation and Planning in High Autonomy Systems conference, 6–8 March, 27–35. www.acims.arizona.edu/CONFERENCES/ais2000/Review/peters_i.pdf. Accessed February 2009.
- Projectgroep Duurzame productie van biomassa. 2006. Criteria voor duurzame biomassa productie [Final report of the project group Duurzame productie van biomassa]. www.senternovem.nl/mmfiles/w690_tcm24-280290.pdf. Accessed February 2009.
- Shalizi, C. R. 2006. Methods and techniques of complex systems science: An overview. In *Complex systems science in biomedicine*, edited by T. S. Deisboeck and J. Y. Kresh. New York: Springer.
- Suh, S. 2004. Materials and energy flows in industry and ecosystem networks. Ph.D. thesis, Leiden University, Leiden, the Netherlands.
- Suh, S. and G. Huppes. 2005. Techniques for life cycle inventory of a product. *Journal of Cleaner Production* 13(7): 687–697.

- Van der Voet, E., L. van Oers, C. Davis, R. Nelis, B. Cok, R. Heijungs, E. Chappin, and J. B. Guinée. 2008. Greenhouse gas calculator for electricity and heat from biomass. CML report 179. Leiden, the Netherlands: Institute of Environmental Sciences (CML), Leiden University. http://www.leidenuniv.nl/cml/ssp/publications/co2_tool.pdf. Accessed February 2009.
- Waldorp, M. 1992. *Complexity: The emerging science at the edge of order and chaos*. NY: Simon and Schuster, USA.

About the Authors

Chris Davis, Igor Nikolić, and Gerard P. J. Dijkema are affiliated with the Faculty of Technology, Policy and Management within Delft University of Technology in Delft, the Netherlands. Chris Davis and Igor Nikolić are Ph.D. researchers, and Gerard P. J. Dijkema is an associate professor.