Eur. Phys. J. Special Topics (2016)

DOI : 10.1140/xxxxx (will be inserted by the publisher)

The Offer Network Protocol: mathematical foundations and a roadmap for the development of a global brain

Francis Heylighen[*]

Global Brain Institute, Vrije Universiteit Brussel, 1050 Brussels, Belgium

Abstract. The world is confronted with a variety of interdependent problems, including scarcity, unsustainability, inequality, pollution and poor governance. Tackling such complex challenges requires coordinated action. The present paper proposes the development of a self-organizing system for coordination, called an "offer network", that would use the distributed intelligence of the Internet to match the offers and needs of all human, technological and natural agents on the planet. This would maximize synergy and thus minimize waste and scarcity of resources. Implementing such coordination requires a protocol that formally defines agents, offers, needs, and the network of condition-action rules or reactions that interconnect them. Matching algorithms can then determine self-sustaining subnetworks in which each consumed resource (need) is also produced (offer). After sketching the elements of a mathematical foundation for offer networks, the paper proposes a roadmap for their practical implementation. This includes step-by-step integration with technologies such as the Semantic Web, ontologies, the Internet of Things, reputation and recommendation systems, reinforcement learning, governance through legal constraints and nudging, and ecosystem modeling. The resulting intelligent platform should be able to tackle nearly all practical and theoretical problems in a bottom-up, distributed manner, thus functioning like a Global Brain for humanity.

1. Introduction

Our society is confronted with a variety of deep and widespread problems. These include poverty, inequality, instability, waste, pollution, unsustainability, conflict over scarce resources, and ineffective governance. At the same time, technology—and information technology in particular—proposes a range of powerful tools for tackling such problems. It thus promises a world of abundance in which all human needs are satisfied (Diamandis & Kotler, 2012; Drexler, 2013; Rifkin, 2014). The difficulty, however, is how to get from the present situation to this envisioned utopia.

[*] email: fheyligh@vub.ac.be

This difficulty is perhaps best captured by the *VUCA* acronym: the situation is Volatile, Uncertain, Complex and Ambiguous (Beigi, 2014; Bennett & Lemoine, 2014). It is *volatile* in the sense that technology has accelerated all developments (Heylighen, 2008), positive as well as negative: both new threats and new opportunities appear and propagate at an unprecedented rate. It is *uncertain* because no one can say which of these new challenges will have a big impact, and which will merely dissipate as quickly as they appeared. It is *complex* because all these challenges interact with each other, so that they cannot be understood in isolation. Finally, it is *ambiguous* because no one really knows what the precise meaning or value of these various developments is.

Up to now, the promises of technology have tended to be piecemeal: some new solution (e.g. genetic modification of mosquitoes) is proposed to tackle a particular challenge (e.g. stop mosquitoes from transmitting malaria). But this is just a tiny part of the overall problem. Moreover, because of the complex interactions between challenges, an intervention in any one may create new ones elsewhere (e.g. the genetically modified mosquitoes may upset the ecosystem in which they are released). What we need is a method for tackling all challenges together, in a *coordinated* manner. However, because of the VUCA nature of the situation, it seems impossible to formulate an overall plan that would tackle all the world's problems simultaneously.

Nevertheless, it is possible to coordinate activities without a plan. *Stigmergy* is a method of coordination used by social insects, such as ants or termites (Heylighen, 2016a; Parunak, 2006; Theraulaz & Bonabeau, 1999). It allows them to efficiently collect food or construct complex nests without any centralized command or control. The coordination *self-organizes* because the partially finished work of some individuals signals to the others what to do next. This indirect information exchange is made possible by a shared medium. The traces left on the medium by the agents can be interpreted as either partial solutions (this is already available) or remaining problems (this still needs to be done), thus efficiently directing the further activity. The resulting coordination explains the impressive *collective intelligence* exhibited by swarms of insects whose individual intelligence is very limited (Bonabeau, 2009; Heylighen, 1999).

The Internet provides an amazingly powerful medium for the coordination of actions, as it allows the registration, storage and propagation of any information, to be shared between any human or technological agents (e.g. robots, sensors, or software). In principle, it could be used to locate and alert the agents most competent to solve a problem the moment it arises, and to ensure that a heterogeneous group of agents collaborate efficiently on a common project. Thousands of applications, websites and platforms have arisen over the past decades to support such coordinated activity, from email via forums, wikis, social media and online markets to community sharing and crowdsourcing. Its most impressive applications, such as the collaboratively authored Wikipedia encyclopedia, have successfully applied the principles of stigmergy (Heylighen, 2007).

Several authors have proposed that these local forms of Internet-supported collective intelligence will eventually coalesce into one planet-wide superintelligence: the *Global Brain* (Bernstein, Klein, & Malone, 2012; De Rosnay, 2000; Goertzel, 2002; Heylighen, 2008, 2011; Mayer-Kress & Barczys, 1995). Such a Global Brain would at last have the capabilities to deal with the immensely complicated and wicked problems of the VUCA world. Thus, it would be able to realize the world of abundance promised by technological visionaries (Heylighen, 2015).

Yet, the question remains how we get from here to there. The present amalgam of coordinating platforms is itself everything but coordinated, as they all use different protocols, paradigms, and interfaces, while addressing different communities and different types of problems (Heylighen, 2016b). The situation is similar to the one before the advent of the Internet in the 1980's, when there was a proliferation of local area networks, bulletin board systems, commercial providers and academic networks. The introduction of the very robust and flexible TCP/IP protocol made it possible to interconnect and eventually subsume all these local computer networks into the global Internet.

The present paper wishes to formulate the logic behind an envisaged new protocol, provisionally called "offer networks" (Goertzel, 2015; Heylighen, 2016b). This protocol would do for the Global Brain what TCP/IP did for the Internet, and HTML/HTTP for the World-Wide Web: define a universal standard that makes the hodgepodge of existing systems interoperable, so that they can merge into a system much larger and more effective than the sum of its parts. This protocol would enable efficient coordination between all agents and all activities on the planet, and thus support a distributed intelligence that should eventually be able to solve all major problems of our time.

Next to a formal foundation for such a protocol, the paper will sketch a first roadmap for implementing the corresponding coordination platform. This roadmap will survey the different steps that are to be performed in order to build such a system. As more steps are made, the effectiveness of the platform will increase, until it may eventually reach the level of intelligence of a true Global Brain. It is of course very difficult to predict which social and technological obstacles or innovations may or may not appear on the road to a Global Brain. Therefore, this roadmap should of course not be taken as a detailed scenario for what will happen when. Its intention is merely to bridge the gap between the present situation and the envisaged Global Brain organization, by listing a number of concrete and realistic steps that are likely to get us from here to there. Thus, it is intended to convince the reader that the Global Brain scenario is ready to be transformed from a utopian vision (Heylighen, 2002, 2015) into a practically realizable project.

2. Distributed Intelligence and the complexity of value

The concept underlying both the offer network protocol and the Global Brain is that of *distributed intelligence* (Heylighen, 2014, 2016c). Intelligence is traditionally conceived as the ability to solve specific problems. We will here define it more generally as *the ability to recognize and tackle challenges*. A challenge is defined as an aspect of a situation that carries potential value. A challenge can be a problem to be solved (negative value) or an opportunity to be exploited (positive value), or some combination. Recognizing challenges requires gathering *information* about the situation and processing that information in order to make *sense* of it: in which respects does it signal a meaningful opportunity (potential value) and/or a problem (potential loss of value)? Thus, our concept of information is *pragmatic* (Gernert, 2006) rather than syntactic or semantic.

The more diverse and the more complex the challenges a system can tackle, the more intelligent it is. Tackling challenges requires selecting and performing an appropriate combination of actions. Performing these actions provides the system with benefit relative to not performing them. Thus, the ultimate measure of intelligence is the amount of benefit or value that the system can obtain by effectively dealing with the challenges it encounters.

Intelligence is *distributed* when different parts of this process of challenge tackling are performed by different agents. For example, intelligence in the human brain is distributed, because the processing and interpretation of incoming information and the formulation of strategies for dealing with the perceived challenges are performed collaboratively by a myriad of neurons residing in different brain regions (Rumelhart & McClelland, 1986). No neuron or brain module is in control of the process; the process is fully decentralized and parallel.

Distributed problem-solving is even more obvious in society. Groups, organizations, and societies can perform incredible feats, such as putting a human on the moon, that no single individual would ever be able to achieve. This distributed intelligence is supported by the division of labor: different individuals have complementary skills, which they bring to bear on a common problem. But it is not just knowledge and ability that are distributed...

2.1. The complexity of value

Tackling challenges requires sensing the different aspects of a challenge. This is typically performed by different people, often supported by instruments. It requires not just gathering information about the situation, but interpreting that information in terms of *value*: what does it mean? In how far are its implications positive or negative? Is there a resource to be exploited or a danger to be evaded? This sense-making and valuing ability is distributed as well, since different individuals have different perspectives, preferences and needs, and what one may take for granted (e.g. having a variety of clothes, or being able to read) another may experience as a major deficiency.

In the brain too, valuing is distributed across a variety of neurophysiological systems. These include various receptors for pain, pleasure, pressure and heat, triggers for anger and fear in the amygdala, "feelgood" neurotransmitters such as endorphins, dopamine and serotonin, and the cognitive representation of moral values in the neocortex. The complex character of value is illustrated by the fact that we typically only become aware of how valuable something is (say, warm clothing) when we miss it, and by the fact that we tend to overestimate the value of something we desire (say, a new car), with the result that we often are disappointed when we get it.

The *complexity of value* (LessWrong, 2013; Muehlhauser & Helm, 2012; Yudkowsky, 2011) is a fundamental problem that is typically overlooked in the attempts to build an Artificial General Intelligence (AGI). An AGI is a software agent that would be able to decide and act autonomously in any situation, just like a human being. The standard assumption in AGI research, economics, and decision-making is that value can be represented as a *utility function*. Such function maps every situation onto a real number that measures how valuable it is to the agent. The agent would then decide what to do simply by calculating the expected utility of the different options, and selecting the one with the highest value. Its intelligence could then be defined as the degree to which it manages to maximize utility.

However, the number of possible situations and their potentially relevant aspects is infinite, and these aspects interact and evolve in an unpredictable manner. Therefore, a realistic goal system or utility function would be so complex as to be non-computable (Ashby, 1972; Heylighen, 2015). Moreover, as the examples above illustrate, in practice values are volatile, uncertain, complex and ambiguous, changing with previous experience (e.g. the value of food diminishes after a copious meal), the agent, and the situation. Therefore, a general utility function cannot even be defined, because the same thing will have one value in one context, and a different value in another context. This means that a true intelligence, which would make sense of its situation and dependably create value, must be *open-ended*, constantly adapting its preferences to a changing world, rather than having preprogrammed goals (Weinbaum & Veitas, 2016).

The problem of defining a universal value function (Ashby, 1972; Yudkowsky, 2011) does not arise in distributed intelligence, because no centralized decisions have to be made. Different local decisions are made depending on the subjective needs of the local agents, who know best what is most valuable in their specific situation. If their choice turns out to be unsatisfactory, they will simply express a new, updated need, which can again be tackled by the distributed intelligence. The only requirement for the system to work efficiently is that the different challenges and the actions performed to address them are *coordinated*, so that they result in cooperation or synergy rather than conflict (Heylighen, 2013). In other words, the sum of all value produced for all agents should be as much as possible positive (*synergy*), rather than zero (competition) or negative (conflict or *friction*) (Heylighen, 2008; Heylighen & Campbell, 1995; Wright, 2001).

2.2. Creating synergy

We now come to the principle of an *offer network* (Goertzel, 2015; Heylighen, 2016b), which is an extension of what has been called a *web of needs* (Kleedorfer, Busch, Pichler, & Huemer, 2014), but which might also be called a *synergy web*. The idea is that all agents—whether human or technological—that are part of the system would express the challenges they experience. Negative challenges correspond to problems, desires, or *needs*, i.e. perceived deficiencies between what the agent has and what it ideally would like to have. Positive challenges correspond to resources, opportunities, or *offers*, i.e. perceived surpluses where the agent has something that could potentially satisfy a need (Heylighen, 2014).

Because of the distributed character of value, challenges are not objectively classifiable as positive or negative. For example, the cows kept by a dairy farmer produce a lot of manure. This is a problem because the runoff from manure pollutes rivers. It is also an opportunity because manure can be used to fertilize crops. The problem can be turned into a resource by arranging a transaction between the dairy farmer and a crop farmer who comes to collect the manure. Such coordination creates a synergy that solves two problems at once: the dairy farm gets rid of its manure, thus saving on costs for disposing of it, while the crop farm is provided with fertilizer, thus producing more value by growing more crops. This illustrates the Oxford Dictionary definition of **synergy**: "the interaction or cooperation of two or more organizations, substances, or other agents to produce a combined effect greater than the sum of their separate effects".

An offer network is an intelligent ICT platform intended to find such synergetic interactions. The crop farm advertises its need for fertilizer on the platform. The dairy farm publicizes its offer of manure in that same medium. The offer network incorporates a number of rules and algorithms that are smart enough to establish that the offer of manure can satisfy the need for fertilizer. It then brings the two farms into contact, and suggests the exchange.

Instead of struggling with the question whether manure has positive or negative value, the offer network finds a match between different challenges listed on the platform so that synergy, and therefore the total benefit for the different agents, is maximized. Thus, the complexity of value becomes an advantage rather than an impediment: because different agents have different needs and different offers at different moments, the one may complement the other so that everyone gains.

Of course, it is *a priori* not very likely that the needs of two agents would be precisely complementary. But when the pool of agents and their expressed needs is large enough, it becomes increasingly easy to find some (perhaps quite complicated) scheme of exchanges that satisfies most of them (Goertzel, 2015; Heylighen, 2016b). This can be demonstrated mathematically, as we will now try to show.

3. Mathematical foundations for offer networks

3.1. Agents and actions

The offer network protocol needs to represent the fundamental components of coordinated action: *agents*, *actions*, prerequisite *conditions*, problems to be solved (which we will call *needs*), and (partial) solutions proposed (which we will call *offers*). It moreover requires a medium in which these components can be expressed and a system of algorithms that would try to optimally match these components.

The first components of an offer networks are the agents: $A = \{ag^i \mid i = 1, \ldots, n\}$. Agents can be humans, organizations, machines, software agents, or even natural systems, such as plants or animals. Each agent has an unambiguous identity, and a corresponding address in the network. Agents are defined by their ability to produce actions, i.e. to change certain conditions without themselves being changed. Thus, each agent ag^i is characterized by a set of actions $\{a^i_j \mid \mathrm{j} = 1, \ldots m\}$ that it is able and willing to perform under the right conditions. Conditions $C = \{c_i \mid i = 1, \ldots, k\}$ are aspects of the agent's situation that can be perceived as challenges. This means that they can potentially incite an agent to act, and thus change the condition (Heylighen, 2014; Heylighen, Busseniers, Veitas, Vidal, & Weinbaum, 2012).

Agents act by transforming some initial conjunction of conditions (their input set C_I) into a subsequent conjunction of conditions (their output set C_O). In set-theoretic notation (where $P(C)$ stands for the power set of C), this becomes:

$$a^i_j : P(C) \rightarrow P(C): C_I \rightarrow a^i_j\,(C_I) = C_O, \text{ with } C_I = \{x, y, \ldots \in C\} \text{ and } \mathrm{C}_O = \{f, g, \ldots \in C\}.$$

Instead of set notation, we will henceforth adopt the notation used for chemical or physical reactions, and write a conjunction of conditions with the "+" sign:

$$a^i_j: \quad x + y + \ldots \rightarrow f + g + \ldots$$

We say that the conditions (or resources) $x, y\ldots \in C_I$ are *consumed* by the action a^i_j, while the conditions $f, g, .., \in C_O$ are *produced* by it. This representation is equivalent to what is called in AI a *condition-action rule* or a *production rule* (Anderson, 2014).

The simplest actions performed by an agent merely produce or consume a single condition, e.g. a^i_j: $x \rightarrow f$. This means that the agent transforms x into f. For example, a person may transform a sheet of paper (x) into a paper airplane (f), or a computer program may transform a list of numbers (x) into a bar diagram (f).

Note that the input and output sets can also be empty: $C_I = \varnothing = C_O$. In that case, the action of an agent is written with just a single condition, e.g.:

$$a_1: \quad \rightarrow f$$
$$a_2: \quad x \rightarrow$$

The first action a_1 can be read as the agent producing the new condition f without precondition. For example, a person can invent a story. The second action a_2 means that the agent consumes the condition x without producing anything in return. For example, a person can eat an ice cream.

Our general assumption is that agents will only act if this action is expected to be valuable to them, i.e. it provides some benefit relative to inaction (Heylighen, 2014; Heylighen et al., 2012). Thus, we can assume that agents are motivated to perform actions. However, in order to be able to perform an action, the input condition of that action must be fulfilled. Therefore, we may say that the input conditions C_I represent an agent's challenge, i.e. a set of conditions to which the agent would react, either to solve the implicit problem or to exploit the opportunity. The consumption of C_I leading to the production of C_O therefore produces value, in the sense that the new situation is more beneficial to the agent than the initial one.

But not all potential actions are equally valuable. When confronted with a situation that fulfills the conditions for different actions, the agent will need to choose which action to perform (first). This priority ordering of actions can be formalized by associating each action a_i with a weight or reaction constant $r(a_i)$ that expresses the relative preference, probability or intensity of performing this action: if $r(a_i) > r(a_j)$, then the agent will first perform a_i. If the agent can perform both actions simultaneously, then the reaction constant represents the relative intensity of each action, and therefore the amount of resources the agent produces through this action.

3.2. Self-sustaining networks of reactions

When an agent acts on a condition it thereby creates a new condition in the medium. This condition may be recognized by one or more agents as a trigger for subsequent actions, which produce subsequent conditions, and thus yet further actions. This process of actions eliciting more actions via the intermediate of the traces (changed conditions) they leave in their shared medium defines *stigmergy*, which is a self-organizing process of coordination between actions (Heylighen, 2016a). Thus, the provision of a medium via which agents share the conditions that determine their actions is in principle already sufficient for coordination to emerge. But let us analyze more profoundly how such coordination can evolve.

The triggering of one action by a previous one concatenates the condition-action rules of the different agents into a network of dependencies. Such networks can be analyzed mathematically with the help of *Chemical Organization Theory* (COT). COT is a very general formalism for describing the self-organization of interactions in any domain (physics, chemistry, ecology, society, …) (Dittrich & Fenizio, 2007; Dittrich & Winter, 2008; Heylighen, Beigi, & Veloz, 2015). Its basic elements are reactions R that transform combinations of resources $\{b, c, d, \ldots\}$ into new combinations, for example:

$$R_1: b \rightarrow d + 2\,e + f$$
$$R_2: h \rightarrow c$$
$$R_3: c \rightarrow h + 3\,f$$
$$R_4: d + f \rightarrow h + b$$
$$R_5: g + 2\,f + e \rightarrow h + d$$

The *action* of an agent a that transforms an input condition (say $e + f$) into an output condition (say d) can be rewritten as a *reaction* between the agent (functioning as the reaction's catalyst) and the conditions, here interpreted as *resources*. For example:

$$R_6: a + e + f \rightarrow a + d$$

The COT representation extends our previous production systems representation by also including agentless reactions (which could just be natural processes, such as oxidation or the growth of plants) and cooperative actions (in which more than one agent is involved). COT provides both qualitative (algebraic) modeling, by analyzing the topology of such reaction networks, and quantitative (dynamic) modeling, by additionally calculating the reaction rates and the variable amounts of resources present at any moment.

What it adds to traditional network formalisms is that reaction networks are directed *hypergraphs*, rather than ordinary graphs (Gallo, Longo, Pallottino, & Nguyen, 1993). This means that it can describe how *combinations* of challenges and agents together can produce new combinations, thus enabling complex coordination schemes.

The main strength of COT is that it provides an analytic method for identifying how reaction networks self-organize into self-sustaining subnetworks (called *organizations*). These are networks where all resources consumed by some reaction are also produced by some other reaction in the network, and this to a degree sufficient for the resource not to disappear. For example, all resources in the above reaction network, except *g,* that are consumed are eventually produced again. Therefore, the resource set $\{a, b, c, d, e, f, h\}$ as well as its subset $\{c, f, h\}$ are self-sustaining under the reaction set $\{R_1, R_2, R_3, R_4, R_5\}$. However, $\{c, g, h\}$ is not.

Organizations are *attractors* for the dynamics of a reaction network (Heylighen et al., 2015): if the process is started from any initial state (specific subset of available resources), it will end up in a state that is part of an organization (self-sustaining subset of resources). The mechanism is that the application of the reactions possible in the initial state will in general produce additional resources, while consuming some of the existing and newly produced resources, until the resource set stabilizes into one where no resource is consumed more than it is produced. Thus, organizations can be generated by starting with random combinations of resources, running the reactions applicable to these resource sets, until the system converges to a stable set, i.e. an attractor. This is equivalent to letting the reaction network self-organize.

The purpose of an offer network is to maximally match the offers and needs of all agents, so that everyone gets what s/he needs, and nothing is wasted. In COT language, this means that every resource that is consumed (thus satisfying a need) is also produced by some other reaction (thus exploiting an offer). For example, the case of the dairy farm can be written as follows:

R_1: dairy farm + cows → dairy farm + cows + manure + milk
R_2: crop farm + manure → crop farm + crops

Here, we see that the dairy farm (in collaboration with its cows) via R_1 produces an offer of manure, which is then exploited in R_2 to satisfy the need of the crop farm. Note that these two reactions still produce offers (milk and crops) that are as yet not exploited. However, these offers are likely to be consumed once we add reactions describing the needs of consumers for dairy and vegetables. Thus, by adding agents and the reactions in which they participate, we generally increase the total number of resources that are fully "recycled", i.e. produced and consumed in such a way that everything balances out.

3.3. Finding self-sustaining networks

Coordination is optimized when a maximum of needs and offers are matched. Formulated in COT language, this is equivalent to the problem of finding the largest organizations (self-sustaining subnetworks) for the network. Note that such largest organizations are in general not unique, and that they are likely to contain several smaller suborganizations. In principle the existence and reachability of such self-sustaining configurations can be proven analytically, but this will require a further elaboration of the mathematics of offer networks, so that it can deal simultaneously with discrete, one-time offers (e.g. a book, a bicycle), and quantitative, continuous offers (e.g. manure, electricity).

A global, "top-down" algorithm to find such an organization could be the following. Start with the set C of all resources in the network. Apply each reaction $r \in R$ to the resources $x \in C$. Now remove all resources in the set that are consumed by some reaction(s), but not produced with at least the same amount by some other reaction(s). This leaves a reduced set of resources, and therefore a reduced set of reactions that uses these resources to produce other resources. Again remove first all resources that are consumed but no longer produced, and then all reactions depending on those removed resources. Repeat until no further resources can be removed. The resource set has now become "self-sustaining": every resource that is consumed is also produced, and vice versa.

Note that the decision as to whether a resource that is consumed is equally produced consists of two steps: 1) qualitative: do one or more reactions producing the consumed resource exist in the network?; 2) quantitative: can these reactions produce enough in order to compensate for the amount that is being consumed? When the first requirement is satisfied (which is trivially checked), the network is called "semi-self-maintaining". Checking the second requirement requires knowledge of the reaction constants determining the rate with which the reactions run, and the application of a basic algorithm equivalent to linear programming. When the network satisfies that requirement it is called "self-maintaining" (Dittrich & Fenizio, 2007). This will normally only be necessary for quantitative, continuous offers.

The top-down algorithm appears to be always applicable, but it would need to be proven mathematically whether the organization it produces is (Pareto) maximal, in the sense that no resources can be added that would satisfy additional needs without taking away at least as many other resources. For a very large, constantly changing set of resources and reactions, as can be expected in a global offer network, this algorithm may in practice be impossible to compute. Therefore, we will need to investigate various heuristics and other algorithms (Centler, Kaleta, di Fenizio, & Dittrich, 2008) that would produce satisfactory results even if they cannot be proven to find a maximal organization. Given that offers and needs are constantly added and removed from the network, it does not seem to make much sense to search for any "optimal" match between them. More important than optimization seems to be constant adaptation and rapid exploitation of new opportunities, so that needs can be satisfied within the shortest terms after they appear.

A more practical algorithm would be to start locally, or "bottom-up", with a few resources (preferably resources that are anyway produced but as yet not consumed). Applying the reactions that function under these conditions produces some additional resources, thus expanding the resource set. Repeat until no further resources can be added. Then perform the same iteration as above, removing step-by-step all resources that are consumed more than they are produced. This again leaves you with a self-sustaining set of resources that are produced as least as much as they are consumed. This set will be different depending on which initial resources you started out with, and is unlikely to be maximal. However, you can find additional self-sustaining sets by removing all resources in the already found set, and selecting some of the remaining ones to start the procedure anew to find another (small) self-sustaining set. This procedure can be repeated until no further self-sustaining sets can be found, or until new resources and/or reactions are added to the network. Thus, it can remain active continuously, locally searching for matching resource sets, which are then eliminated from the problem domain, without caring about the topology of the global network. It can also function in parallel, simultaneously exploring independent parts of the resource set, e.g. by subdividing resources according to their geographical region.

3.4. Some examples

Let us make these formal models more concrete by means of a series of examples. Our first example of synergy, in which a dairy farm tries to get rid of its manure, and a crop farm comes to fetch the manure to use it as fertilizer, can be represented most simply by the following two condition-action rules, with r standing for "removal of manure", and f for "fertilizer".

dairy farm: $r \rightarrow f$
crop farm: $f \rightarrow r$

This means that the dairy farm agent offers fertilizer on the condition that some other agent would satisfy its need for the removal of manure. The crop farm offers to remove manure on the condition that it can use it to satisfy its need for fertilizer. The offer network arranges the exchange between the two agents by simply concatenating the two condition-action rules into a cycle or loop: $r \rightarrow f, f \rightarrow r$. Both benefit and synergy is created. Thus, the offer of the one fulfills the need of the other. Everything is recycled and nothing is wasted.

This is a traditional one-to-one exchange, as happens during barter, market transactions, or an agreement for mutual aid. For such simple exchanges, the only advantage in using an offer network is that a much larger range of available offers and needs can be scanned by the algorithm so that it becomes easier to find two mutually matching ones. The true power of the formalism only comes to the fore when there are more than two parties involved. Goertzel's original proposal for offer networks (Goertzel, 2015) was intended to broaden the mechanism of barter to an unlimited number of agents whose offers and needs cannot be matched one-to-one, but can be matched many-to-many.

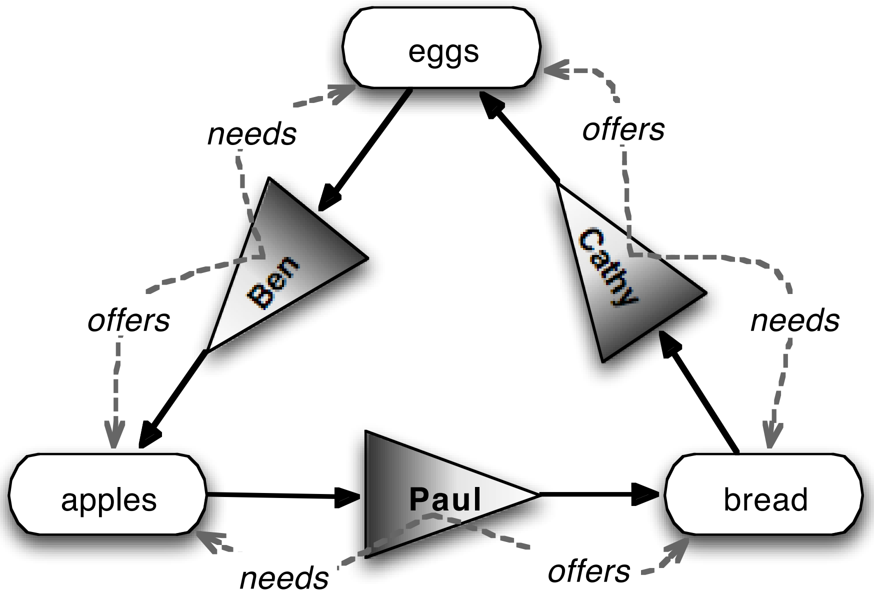

Fig. 1: a simple loop with three agents collectively satisfying each other's needs. Straight arrows depict the flow of resources (depicted as rounded rectangles) between the agents (depicted as triangles), with incoming arrows representing resources needed and eventually consumed by the agent, and outgoing arrows resources produced and offered to others. For example, Paul needs apples, which he receives from Ben, and offers bread, which is taken by Cathy. The broken arrows between agents and resources represent the initial offers and needs entered into the offer network. The loop represents the synergetic solution found by the network.

Let us consider the simplest possible example: a three-agent cycle. Consider the following condition-action rules. Each rule states that an agent (e.g. Ben) is willing to produce an offer (e.g. of apples) on the condition that its corresponding need (e.g. for eggs) is fulfilled:

Ben:	eggs → apples
Cathy:	bread → eggs
Paul:	apples → bread
Marta:	bread → cheese

A simple reordering of these rules shows that the first three can be arranged as a cycle (see Fig. 1). This means that the first three agents (Ben, Cathy and Paul) receive what they need, and in turn give what they offer. The fourth agent, Marta, cannot profit from this synergy. But searching for further agents and their rules in the network may well find an additional loop that also satisfies Marta's need.

Let us consider a more complicated example, where several conditions need to be satisfied simultaneously. Here, the maximally synergetic arrangement cannot be a simple loop, as the connections between offers and needs branch, thus requiring a more complex "organization".

John:	bread + apples → onions
Mary:	onions + eggs → bread
Bert:	beer → apples + milk
Ann:	milk → beer + eggs
Karen:	bread → cheese
Tom:	cheese + wine → eggs

Here, Tom's needs cannot be fulfilled as no one is offering wine. This means he also won't take the cheese offered by Karen, and therefore Karen won't find anybody to give her bread since she cannot give anything in return. Therefore, these agents and their reactions cannot be part of the organization. However, the first four agents can form a self-sustaining organization, because each resource needed by one of them is offered by one of the others. This closed organization is depicted in Fig. 2. (Note that as the number of reactions increases, such a graphical depiction quickly becomes unwieldy). Like in the previous case, searching for further agents or rules in the network is likely to create a match for the needs of Karen and Tom as well, by extending the closure of the network. For example, it would be sufficient to find an agent that is willing to exchange eggs for wine and bread.

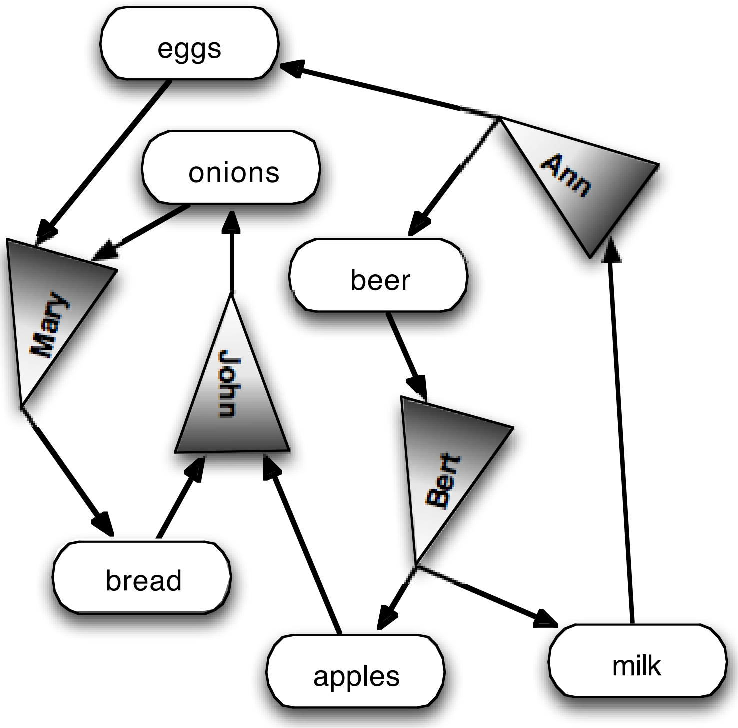

Fig. 2: a more complex closed network, where each of the resources needed or consumed by some agent is also produced or offered by another agent. Thus, each of the resources has both an incoming (production) and an outgoing (consumption) link.

As a final example, we will include unconditional offers ("gifts") and needs (Heylighen, 2016b). Suppose that Alice is an old lady who needs the big tree in her garden to be pruned. She cannot do that on her own, and she does not have the means to pay or otherwise reward someone for doing it. Therefore, her offer is empty. Mary, on the other hand, is ready to lend out a saw, on the condition of getting it back, but cannot prune trees either. Still, she has a surplus of eggs, which she is willing to donate to anyone who needs them. Martin is willing to prune the tree on the condition that he can use a saw, and get a beer. By letting the offer network search for additional, complementary condition-action rules, we may come to the following closed network in which everybody's needs are satisfied:

Alice: prune tree →
Mary: saw → saw
Mary: → eggs
Martin: beer + saw → prune tree + saw
John: egg → apple
Henry: apple + egg → beer

4. A roadmap for developing a global offer network

We have defined offer networks abstractly, and argued that they can potentially solve the problem of global coordination between the needs, offers and actions of all human and technological agents. To realize this promise, many concrete steps will have to be taken, so as to implement this general approach into an increasingly powerful technological platform. We will now briefly survey the most important of these steps, thus outlining a roadmap towards building such a platform.

4.1. Identifying users

The first step is to define a general protocol for representing agents, their needs and their offers. This protocol should be public, non-proprietary and standardized, so that it can be used by anyone using any hardware or software on any website implementing the protocol without need for entering website specific data such as user names or preferences. This would make the offer network system transparent and universal, so that it can interconnect and integrate thousands of as yet independent platforms (Heylighen, 2016b). A prototype of such a protocol and a platform that supports it has already been implemented as the *Web of Needs* (Kleedorfer et al., 2013, 2014).

Human agents are here represented as "user proxies", i.e. formal entities that list a particular individual's needs and offers. This demands a dependable connection between a user and its proxies. That in turn requires a universal, open protocol for defining personal identity (Evrard, Erdmann, Holmquist, Damon, & Dietrich, 2015; Recordon & Reed, 2006), such as WebID (http://webid.info/). It should allow individuals to be simply, securely and unambiguously identified across any platform, without need for separate logins. The ID protocol would ensure that the requests or offers are routed to the right people, so that they can realize their offer or receive what they need. The offer network would associate each resource with one or more identities/proxies, and thus build up a general picture of what this person would like to give or to receive.

However, the ID protocol should not compromise privacy, in the sense that senders or receivers of offers may want to remain anonymous. Therefore, users of the network do not *a priori* need to know the real name of the person with whom they are interacting. In principle, individuals should be able to create different identities (e.g. one for dating, one for professional communication, one for playing games) if they prefer not to mix up their different "personas" or preferences, or to create an extra layer of anonymity. A limitation of such a setup, though, is that, as we will discuss further, a smart offer network should be able to learn the tastes and preferences associated with a certain identity in order to provide that person with more appropriate offers and recommendations. The more identities a person has, the less information per identity the network will gather, and therefore the less personalized the recommendations it can make.

Initially, identification can be achieved by a simple login, where the user enters username and password to identify herself. In a later stage, identification of users may happen by automatic recognition of biometric data, such as fingerprints, iris scans, or voice. Then, it would suffice that a user sits in front of her computer or smartphone in order to be able to formulate offers and needs.

4.2. Integrating offer networks with the Semantic Web

Initially, offers and needs can be expressed in natural language (e.g. "I need a bike and I offer an aquarium"). But natural language tends to be vague and ambiguous, and may not capture all the relevant aspects of the requested or offered resource. Disambiguation can be performed with the help of a semantically structured list of words and their senses. For example, the very extensive WordNet database (Fellbaum, 2012) proposes "motorbike" and "bicycle" as distinct possible meanings for the word "bike". For each of these meanings, it holds a "synonym set", i.e. a list of words that can have this same meaning (e.g. "cycle", "bike" and "bicycle"). This ensures that an offer of a "cycle" can be matched with a properly disambiguated request for a "bike".

Such meanings correspond to clearly defined concepts or categories. A system of categories that is formally specified so that it can be used by a computer program without need for human interpretation is called an *ontology* (Livet, 2012). The concepts in an ontology are linked by specific types of relationships that determine their meaning, thus forming a *semantic network*. Semantic networks such as WordNet situate a concept such as "bicycle" within a neighborhood of associated concepts. These include *hypernyms* (more general, encompassing categories, such as "human-propelled wheeled vehicle"), *hyponyms* (more specific subcategories, such as "mountain bike" or "foldable bike"), *sister terms* (similar categories, such as "tricycle" or "scooter"), *meronyms* (terms for parts, such as "bicycle wheel" or "saddle") and *holonyms* (terms for larger wholes, such as "aquarium kit with accessories") (see Fig. 3).

All of these related concepts could be suggested to users as potentially more accurate formulations of their offer or need. For example, perhaps the user, rather than wanting a bicycle in general, may be more interested in the concrete problem of replacing his bicycle's broken saddle, or in getting a foldable bike, or perhaps even a different kind of wheeled vehicle without motor. The more precise the characterization of offers and needs, the easier it becomes to find a good match.

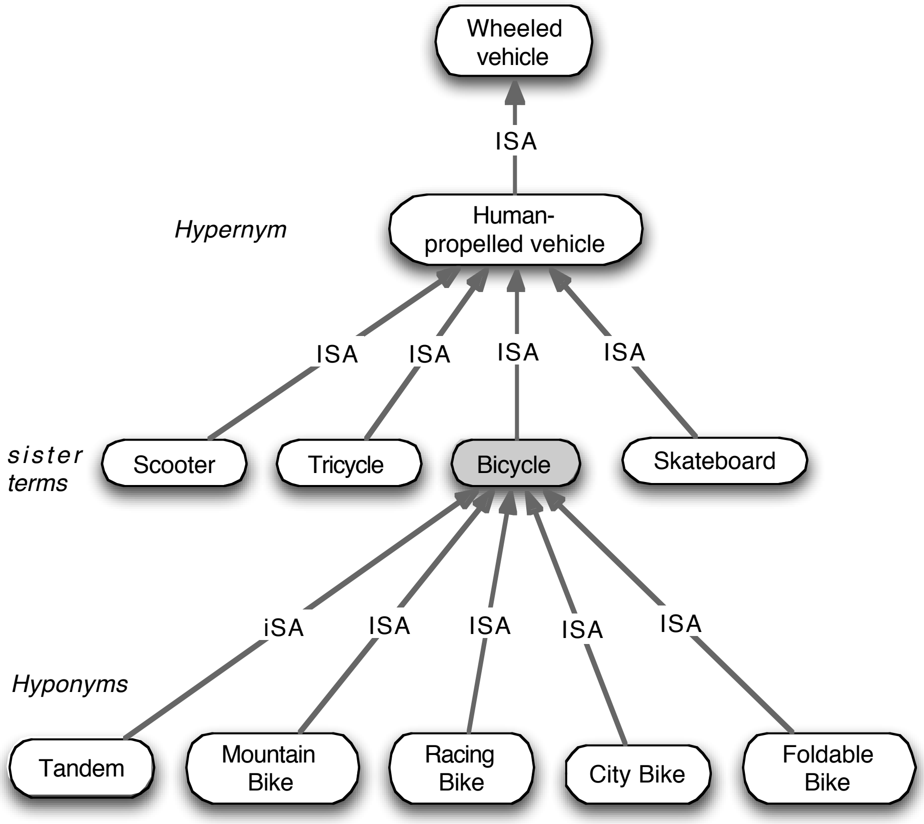

Fig. 3: the category of "bicycle" situated in part of its semantic neighborhood, including encompassing categories (hypernyms) and subcategories (hyponyms), connected to each other by "ISA" links.

In semantic network terminology, a concept and its hypernym are connected by an "ISA" type of link: a mountain bike IS A bicycle. A concept and its meronym are connected by a "HASPART" relationship: a bicycle HAS as PART a saddle. Both types of links can be interpreted as production rules (Heylighen, 2001): IF there is a mountain bike, THEN there is also a bicycle:

mountain bike → bicycle
bicycle → bicycle saddle

By incorporating such semantic links, the offer network increases its power of finding solutions to challenges. For example, using these production rules the network can infer that one person's offer of a "mountain bike" can satisfy someone else's need for a "bicycle" and that an offer of a "bicycle" may satisfy the need for a "bicycle saddle", if no saddle would be available elsewhere in the network.

For more precise matching between offers and needs, resources should be situated in the ontology with their specific attributes or features. Ideally, each concept recognized by the system should call up a template that lists its most relevant attributes. For example, for the category “fish tank/aquarium”, relevant attributes are volume, length, width, height, and material (with possible values “glass” or “plastic”). For the category “book”, the attributes include author, title, year, genre, publisher and edition. The system should by default propose the most common values for these attributes (e.g. glass for aquarium material), while taking into account what it already knows about the resource. For example, entering an author’s name would bring up a list of book titles by that author in the order of popularity, and selecting a title would add the year and publisher of first publication as default values for these attributes.

Concepts, attributes and values (including semantic links between concept) are best formulated in the protocols of the *Semantic Web* (Allemang & Hendler, 2011; Berners-Lee & Fischetti, 1999). This is an extension of the World-Wide Web standard intended for machine-readable, formal specifications of data and knowledge that is still in development. It has already engendered a number of more specific standards such as OWL (Web Ontology Language) and RDF (Resource Description Format) that may be used for the representation of offer network components without requiring many additions. The use of such emerging standards will ensure the easy interchange of data between the offer network and other platforms and databases, thus fulfilling the web’s initial aim of supporting a universal knowledge system.

Demands for ontological specification should not put any unnecessary burden on the user. They should in particular make it easy to ignore irrelevant features. For example, if there are differing prints of a book that all possess the same content, then most users would not care which year the book was printed. Or perhaps a user may not care which brand a bike is, as long as it is the right type and size. Depending on how precisely the offer or request is described, the matching algorithm will have a smaller or a larger choice of potential matches to propose. For the user it seems advisable to be less specific for requests, as this will find more matching offers, and be more specific for offers, as the matching algorithm will normally prefer offers that match a specific request in detail, while ignoring these details for more general requests.

When no precise match can be found, the offer network can use the ontology structure to suggest available resources that are semantically similar. For example, if no city bike is available, it may suggest a mountain bike. For concepts with attributes, the semantic neighborhood would include instances where some of the desired attributes are different or absent. For example, the network could propose a plastic aquarium if no glass one is available, or one of 100 liter volume when no 80 liter one is available.

For such a smart matching program to function well, an increasingly elaborate ontology with good defaults and suggestions will have to be developed for nearly all resources that someone might request or offer. This is likely to evolve out of the merging of hundreds of both general (e.g. WordNet) and specialized ontologies (e.g. for books, chemical products, plants, songs, machine parts, cars, hotel rooms, etc.). In order to guarantee an open-ended system, this should be complemented by the ability of users to create templates for new categories that do not exist in the system yet. Moreover, machine learning algorithms may find recurrent patterns in common offers and requests, and propose these patterns as new categories and templates. Eventually, this will produce a global ontology that categorizes about all commonly used resources and needs present on Earth, with their most relevant attributes and relationships. This would allow users to describe any challenge with just its most characteristic aspects, and the network to infer all other aspects not explicitly entered.

4.3. Integrating the Internet of Things

Another step in our roadmap for developing a planetary-scale offer network is the incorporation of the *Internet of Things*. This is an emerging infrastructure for the exchange of information between people and physical objects, and even between the objects themselves (L. Atzori, Iera, & Morabito, 2010; Rifkin, 2014). Soon, most fabricated objects will have an in-built tag (e.g. RFID) that can be wirelessly consulted to ascertain the identity, features and state of the object. Thus, it may be sufficient to check the Internet address of the bicycle that you would like to offer or to request in order to get all its relevant attributes, so that they can be published on the offer network. (Note that at present you can already point a scanner at a bar code—e.g. on a book—and get a description of the product it is attached to, but this still requires local, physical interaction.)

Objects themselves will eventually be able to publish their needs. For example, a bicycle with a flat tire could request the right size and type of tube for a replacement. An object that is overheating might request fire-protection measures, such as the activation of sprinklers. Objects should also be able to advertise their offers. For example, a publicly shared bicycle could announce its availability to all people needing transport within a certain neighborhood of its location.

Future intelligent artifacts will be able to offer a variety of products and services. For example, 3-D printers can produce a wide range of objects that they assemble out of plastic materials on the basis of a specification of the precise three-dimensional shape and composition (Lipson & Kurman, 2013). That shape specification can be coupled to the ontological specification of the corresponding category of objects, such as bicycle saddles. The attributes could then be used to specify variable properties, such as width, length, material or color. In this case, users would merely need to select their preferred values for the attributes from within the template proposed by the offer network, and the requested object would be immediately produced by the printer in their place of residence (Heylighen, 2015; Rifkin, 2014).

More complex requests (e.g. for a tailor-made car or suit of clothing) could be directed by the offer network to industry-level assemblers in factories, and delivered by autonomous vehicles. Requests for specific services, such as transportation or catering, could also be satisfied by smart objects, such as robots, self-driving vehicles, drones, or remotely controllable devices such as coffee machines or ovens.

4.4. Incorporating knowledge and intelligence

As yet, we have only discussed applications of offer networks where the match between offer and need is straightforward, because they both belong to the right category with the correct attributes. This is similar to how a user enters a query in the catalog of a web shop, and the search engine returns the matching entries. But the condition-action rules that connect needs with the conditions necessary to satisfy them enable a more intelligent form of problem solving or need satisfaction. Here, the role of the search engine is extended to what in AI is called an *inference engine* (Singh & Karwayun, 2010). Such an engine can chain together a branching sequence of rules so as to find a fit encompassing several conditions. Thus, it may assemble the solution to a complex need by combining several independent offers that are distributed across different agents and locations.

For example, assume that your request is to organize a party. The inference engine may use some of the following relevant production rules:

guests + location + live music + drinks + food → party
list of friends + invitation to party + distribution of invitation to list → guests
piano on location + pianist → live music
piano + truck + mover → piano on location
caterer → food
soft drinks + water + juices + beer + wine → drinks
loft → location
bar → location
party hall → location

Note that some of these rules can be derived from standard semantic links, such as "a loft ISA location", or "a party HAS_PART guests", others from more specific causal relations, such as "a pianist having access to a piano CAN_PRODUCE live music".

The complex request for a party can now be recursively decomposed into a number of simpler requests, each of which can be matched against a range of available offers, such as pianos, trucks, beer, etc. The algorithm that moves from the desired outcome to the conditions necessary to achieve it, and from these conditions to even earlier preconditions, is called "backward chaining". The process will continue moving backward while exploring a branching network of possibilities until an offer is found for each of the required conditions. If more than one possible way is found to fulfill a condition, the first or ideally the best one will be selected. For example, if no one has a loft available and all party halls are fully booked, the nearby bar may come out as the most suitable location for the party.

Inference engines can also perform forward chaining, starting from an available offer and finding the requests to which this offer can best contribute. There are a wide variety of other intelligent algorithms for making such chains of inferences across a complex network of rules with overlapping conditions. To realize this intelligence, we will first need to define a protocol for formulating production rules in a standardized manner that different inference engines can understand. However, the inference algorithms themselves should not be included into the protocol yet, given that we do not know which algorithms would be most effective. It is better to let a variety of inference engines, algorithms and heuristics compete for their ability to quickly and efficiently match offers and requests. For example, users may have access to different AI agents implementing different heuristic methods, and find out through trial-and-error which ones work best for them. If different programs are better at different types of problems, their features may eventually be integrated into a more broad and versatile program.

Another necessary step in our roadmap towards global intelligence is the collection of an as broad as possible set of relevant reactions or rules. A good ontology will already incorporate semantic rules such as "an apple ISA fruit", "fruit HAS_PART sugar", "fruit HAS_PART vitamins". However, we will also need causal and procedural rules to describe the processes that may be needed to make predictions or to solve problems, such as "excess acid production in the stomach CAUSES ulcers" (causal), and "to suppress the production of acid in the stomach one can take omeprazole" (procedural).

There already exist extensive databases of common-sense rules, such as CYC (Lenat, 1995). These can be enriched by more specialized knowledge bases, e.g. with medical expertise. Further rules can be extracted by on-going knowledge acquisition from experts (Wagner, 2006), and sophisticated machine learning techniques, such as deep learning (Hinton, Osindero, & Teh, 2006), applied to "Big Data" collections. Users too should be able to propose as yet missing rules. For this, they may rely on an intelligent interface that helps them to structure the knowledge they enter by checking their proposals for ambiguities and inconsistencies with existing rules, and by suggesting more precise formulations (Heylighen, 2001). Like the offers and needs, all these rules should be made publicly accessible on a huge, shared information network, so that everyone can use them, but also examine their accuracy, and if necessary suggest corrections.

Once a variety of both specific and general rules are available, the offer network will not only be able to answer concrete requests, but also more abstract questions, such as "which warm-blooded animals can get ulcers?" or "what contributes to global warming?" Answering such questions will required sophisticated chains of inference across many rules entered by different experts across different fields. This turns the offer network into a universal knowledge system, integrating the whole of human and machine-derived knowledge into a coherent intelligence that can answer nearly any question for which an answer can be inferred. This is the Global Brain version of the old dream of the encyclopedists (Heylighen, 2011, 2015; Rayward, 1994; Wells, 1937).

4.5. Implementing theory into practice

An offer network is much more than a universal knowledge system, though, because it combines abstract knowledge with the ability to deliver concrete resources. For example, if you want to set up a healthcare center in your town, you would need more than the ability to ask health-related questions to a system incorporating medical knowledge. You would also need access to the medicines, machines, and professional service providers that can implement the solutions suggested by the knowledge-based system.

When tackling real needs, you cannot separate the theoretical solution to a problem (e.g. omeprazole may cure an ulcer) from its concrete implementation (e.g. this patient with an ulcer will need to take omeprazole capsules under the supervision of a qualified professional). The offer network will first need to decompose the initial need (a patient wishing to get rid of recurrent stomach pain) into a number of abstract preconditions (the symptoms can be explained by an ulcer, an ulcer can be cured by omeprazole, therefore the patient should better take omeprazole capsules). It then must locate an actual offer of omeprazole capsules, and a professional able to supervise the treatment.

In the simplest kind of offer network, the patient would formulate a direct request for omeprazole capsules. However, in general we cannot expect people to know exactly what they need in order to solve their problems. It is here that the distributed intelligence of the offer network shows it power, by filling in the knowledge gaps via its network of rules and by integrating widely distributed questions, answers, offer and needs into an efficient, coordinated plan of action that tackles as many challenges as possible.

The decomposition of basic needs into complex plans of action is not limited to individual requests. The desire to create a collective provision, such as a hospital in the poor African town of N., will require the coordination of many *human resources* (such as builders, nurses, doctors, administrators, technicians), *physical resources* (building materials, medical devices, power generators, medicine supplies, etc.), and *cognitive resources* (medical expertise, building blueprints, legal framework, management policies, etc.). All of these can be mobilized by the offer network, after it has selected resources matched for their complementarity, so that they are ready to work together synergetically.

The different people collaborating on such an enterprise do not need to know each other in advance: the offer network will have selected them for their willingness and ability to contribute to the overall goal of building a hospital and thus helping the town inhabitants. It is unlikely that any of these individuals would have formulated a specific offer involving hospital building in the city of N. However, they may have formulated more general offers of doing volunteer work in an African country, while specifying their concrete areas of competence (e.g. building, administration, nursing, and others). The offer network, using its ontology and its general knowledge rules, should be able to infer which of these general offers best match the specific needs that it has derived by decomposing the general need of hospital building in N.

While the work is on-going, the different collaborators will form an *ad hoc* organization. Since many people (e.g. architects working on a blueprint, volunteers donating medicines, lawyers formulating the legal framework) can collaborate on the project without actually meeting, they may form what has been called a "decentralized autonomous organization" (M. Atzori, 2015; Garrod, 2016). This is a social system that is held together merely by a system of coordinated rules, e.g. as implemented in the offer network. As soon as the project is completed, the organization is likely to disband as easily as it was assembled, freeing up its members for participation in different projects.

To start the process, it is sufficient that an individual or organization (e.g. an NGO) would formulate the general need for a hospital, and that a sufficient number of other parties confirm its importance in order to give it a suitably high priority. However, in a later stage the offer network itself may formulate such a general goal, by inducing the need for a larger, collective solution from the many unfulfilled individual requests for medical care in the city of N. In this scenario, the offer network can analyze the immense number of as yet unfulfilled needs it will have gathered by grouping together related needs (e.g. all individual needs involving medical care in the geographical region of N.). It can then use composition rules to reformulate this need-cluster as a single problem (lack of health care in N.) that demands a single, high-level solution (building a hospital).

4.6. Learning and reputation systems

Up to now, we have assumed that the people, rules, and description of offers and needs are dependable and truthful, so that you can trust that the offer network will deliver what it promised. This will obviously not always be the case: while most people tend to be honest, they can be forgetful, badly organized, inaccurate in their descriptions, and ambiguous or erroneous in the rules they formulate. Therefore, the system will need an on-going process of *quality control*, to bring forward the best and discourage or discard the not-so-good. Given the complexity of value, there are no objective, absolute criteria for assessing how valuable or important an offer or need is. The ultimate criterion is in how far the people who formulated a need are subjectively satisfied with the solution they receive. Therefore, all users will be invited to express their degree of satisfaction after they received an offer. This produces a feedback signal that allows the offer network to learn more effective ways of tackling the challenges, by reinforcing the use of the more productive elements, while reducing the contribution of the less productive ones.

In the simplest case, a need (e.g. for a means of transport) expressed by a single person (e.g. Mary) is satisfied by a single offer (e.g. a no longer used bicycle) coming from a single person (e.g. John). Mary evaluates the degree to which she is satisfied with the bicycle, e.g. by giving it a score between 0 (absolutely useless) and 5 (ideal). These points are added to John's reputation within the semantic category "offering physical objects" and more specifically "bicycles". Mary in general does not need to know John: the evaluation is for this specific offer, not for John as a person. After many different offers, though, John accumulates many reputation points spreading across different categories (e.g. "repairing cars", "answering questions on geometry", "offering fruit", …). When choosing between otherwise equivalent offers from different individuals, this allows the network to select the people that are most likely to deliver a satisfactory solution within this category.

There is in general no need to add these points together in order to determine an overall evaluation of John. However, when estimating how far John will be able to deliver a satisfactory offer for some need for which John has not been evaluated yet, the network may try to infer John's reliability based on his record in semantically neighboring domains (e.g. John may have a poor record in repairing cars, but be good in answering questions on geometry, and therefore he is likely to also be good in answering questions on algebra). John's overall score may only provide useful information when he makes a completely new offer, unrelated to anything he did before. This global score should remain invisible to the public, so that John cannot be judged by society as being superior or inferior to others in some general sense. This pre-empts the emergence of hierarchies and the resulting inequalities, rivalries and jockeying for position. Thus, reputation should not become the offer network equivalent of the one-dimensional value of monetary wealth (Heylighen, 2016b).

On the other hand, the fact that John knows that each good action he performs increases his reputation will motivate him to provide the best possible service, and to avoid not fulfilling his commitments (De Alfaro, Kulshreshtha, Pye, & Adler, 2011; Jøsang, Ismail, & Boyd, 2007). This motivating power of reputation can be increased by allowing people with a sufficiently high reputation to take more out of the system than they put in, thus profiting from unconditional offers. However, exploiting such gifts without adding anything in return should gradually reduce their reputation, so that they are not inclined to take out more than their fair share. This ensures that they do not turn into "free riders" that consistently abuse the system (Heylighen, 2016b). Instead, it motivates them to behave altruistically by unconditionally helping others just as they have been helped. Providing anyone who joins the network by default with a positive reputation makes the offer network flexible enough to supply resources to people who cannot offer anything in return yet because they are too poor. For people who are intrinsically needy, e.g. because of old age or poor health, this positive reputation may be upheld indefinitely.

The equivalent of reputation points can also be given to objects and rules, so as to assess how helpful or dependable they are in satisfying needs. For example, experts can directly evaluate the accuracy of rules entered into the system. The total score of the rule, preferably weighted by the reputation in the domain of the person who gave the points, can then be used as a first indication of the accuracy of the knowledge it represents. When different rules compete for execution, rules with a higher reliability score or "weight", or that produce more specific results, will normally be preferred to less reliable or more general rules (Holland, Holyoak, Nisbett, & Thagard, 1989).

Finally, such weights should be able to adapt depending on the actual results. This will normally require some analog of the backpropagation method for neural networks or similar "credit-assignment" algorithms that try to fairly allocate rewards or penalties to the different rules that cooperatively produced a solution (Holland et al., 1989). The same rule may have participated in several chains of inference. If the solution coming out of these chains was on average better than the solution coming out of chains that used an alternative rule, then the first rule will be reinforced relative to the alternative, and thus be considered as more reliable by the network.

In this way, the offer network will be constantly learning from its experience, just like the neural networks in our brain, and thus become ever more effective.

4.7. Soft constraints and recommendations

For common needs, it is likely that the offer network will find many offers that satisfy the stated requirement. In this case, the network must do more than satisfy the "hard constraint" that offer and need should match at the ontological level: it must order the different offers according to some preference criterion so that the "best" offers are proposed first. Assuming that all the required hard conditions are fulfilled, so that there is no *a priori* difference in value between the offers, several soft criteria seem worth considering.

The most obvious one is the estimated quality or reliability of the offer. This may be computed as a duly weighted combination of the reputation of the person(s) offering, of the reputation of the ontological subcategory or resource being offered (e.g. a particular brand of mountain bike, or a particular answer to a commonly asked question), and of the weight of the rules that were used to determine in how far this offer would fit the need.

Another general criterion is the "nearness" of the offer, which will depend on the topology used to structure the offer space. The simplest topology is geographical distance. For physical offers that must be transported to the person that needs them, it is obviously better to select those that will require the shortest travel, as this will minimize time, effort, and energy spent in the transfer. However, this criterion is irrelevant for offers of informational goods or services, such as a music recording or a translation, which can be transmitted immediately over the Internet.

If no perfect match of ontological categories can be determined, e.g. because the need is unique, or is expressed vaguely or ambiguously, then semantic nearness seems to be an important criterion: selecting the offer that appears most similar to what was asked. This is typically how search engines answer questions, by returning the documents that seem closest to the meaning of the search string that was entered.

For transactions that involve some degree of social interaction, e.g. a request for a tennis partner, a date, or a collaborator, social nearness may be a useful criterion. This can be calculated on the basis of the requester's social network, in which close friends can be defined as a first neighborhood, acquaintance as a second level neighborhood, friends of friends as a third level neighborhood, etc. There exist reliable algorithms for calculating the "distance" between two nodes in such a network (e.g. Fouss, Pirotte, Renders, & Saerens, 2007), and the preference would be for minimizing this distance. More generally, most people might prefer giving or receiving offers to or from others that are within their extended community, as the exchange to some degree creates a bond, and you may prefer this to happen with someone that you are likely to interact with on some further occasion as well, or that you are more inclined to trust because of having mutual friends.

For very common needs, such "nearness" criteria can be used to reduce the search space in which the algorithm will be searching for matches. For example, at any moment millions of bicycles may be offered all around the globe. There seems to be little sense in examining all of these before selecting the best match. A more efficient method may be to start looking at all the bicycles offered within the same city or geographic area, and to focus on those offered by people within your extended community. Thus, depending on the relevant topology, a series of concentric neighborhoods around the present need can be defined, which the algorithm would explore in the order of increasing radius, until it finds a sufficient number of good matches.

Priority criteria should normally depend on the user, given that different individuals have different preferences. Adjusting the criteria can be done manually, by letting the user set the relative weights of various standard criteria. It can also be done automatically, by inducing implicit preferences from the choices the user made and the degree of satisfaction with these choices (Heylighen & Bollen, 2002). Thus, the offer network should be able to develop an increasingly accurate picture of a particular individual's needs, values and preferences. That would allow the network to propose offers even when no explicit need has as yet been formulated, thus attracting the user's attention to potentially valuable opportunities.

This is a function that is already well developed in so-called *recommender systems* (Ricci, Rokach, & Shapira, 2011), which e.g. propose books, videos or songs to a person depending on the kind of books etc. that they previously enjoyed. To foster diversity and exploration, such recommendations should not be limited to offers immediately similar to previously consumed offers: ideally there should be an element of *serendipity*, so that the user is occasionally challenged to try something new that only has an indirect connection with what that user typically likes or needs (Heylighen & Bollen, 2002).

4.8. Integrating social governance

In an even further, as yet more vaguely envisaged stage, the offer network protocol should be able to integrate the functions of governance, law and justice. Thus, it would ensure that transactions are to the benefit of society, and not just of the individual parties involved. This can be achieved by expressing laws, contracts, policies, and ethical or deontological norms as condition-action rules: *in the following conditions, the following actions should, or should not, be performed*. Thus, legal constraints, such as the obligation to fulfill a contract or the prohibition to deliver weapons or drugs, can be combined with practical procedures about how to tackle a particular problem in order to determine the best course of action. When different rules compete, the legally binding rules should have an absolute priority rather than just a higher weight. Commonly accepted policies or norms, on the other hand, may merely impose a high priority for certain actions, without a priori excluding actions that do not obey these norms.

As a general principle, the number of legal constraints should be minimized, so as to keep the system as flexible and transparent as possible. On the other hand, given the potentially immense power of adaptation and self-organization of offer networks, the legal conditions entered into such a system will probably engender much less rigidity and friction than present regulatory and legal systems. Part of the reason is that there will not be a need for a painstaking examination of all possible regulations that may apply whenever performing some transaction: the same algorithms that search through all ontologically relevant condition-action rules to match offers and needs will automatically take into account any regulatory rules that apply, making sure e.g. that applicable taxes are automatically withheld and deposited with the appropriate authorities.

The extra level of trust and dependability required for contracts and laws may be secured by emerging Internet encryption and security technologies, such as the Blockchain underlying Bitcoin. It has been proposed that these may eventually be used to create "Decentralized Autonomous Organizations" (Atzori, 2015; Garrod, 2016) that are held together by automated contracts that cannot be tampered with because the codes identifying them are fully distributed.

Offer networks can also implement a more flexible form of social governance that has been called "libertarian paternalism" (Thaler & Sunstein, 2003), "nudging" (Thaler & Sunstein, 2008), or "mobilization" (Heylighen, Kostov, & Kiemen, 2013). The principle here is that there are no absolute obligations or prohibitions: everyone has the freedom to choose from the full range of available options. However, the options are arranged in such a way that the ones that are most beneficial to society (e.g. consuming vegetables, or taking courses) become easiest, most apparent or most attractive to choose, while those that are least beneficial (e.g. consuming tobacco, junk food or fossil fuels) require some special effort or motivation to go beyond the default choice. Socially beneficial "nudging" can be achieved by adding soft constraints that represent *externalities*, i.e. the costs and benefits to others of a transaction between private parties (Cornes & Sandler, 1996; Heylighen, 2016b). These additional constraints would affect the order or priority of the solutions suggested by the offer network. For example, a less polluting option would get a higher priority even though it may fit the request somewhat less accurately.

Like all other rules in the network, both legally binding and "nudging" rules should be made public and explicit, so that anyone can scrutinize them, evaluate them, criticize them, and if necessary propose alternatives. This pre-empts the danger that people's choices would be invisibly manipulated by private concerns or totalitarian governments (Helbing, 2015), e.g. for commercial profit or political power.

4.9. Integrating the ecosystem

Perhaps the most fundamental externalities represent the costs of an action to the larger *ecosystem*, e.g. in terms of its contributions to global warming, exhaustion of scare resources, or destruction of natural habitats. Therefore, the governance realized by an offer network should eventually extend to the biosphere. However, deciding what these costs and benefits are should ideally not be left to a committee of lawmakers. A better option is to include the ecosystem itself into the network, on the same footing as the people who formulate their human offers and need.

This can be achieved by modeling the physical, chemical and biological processes that govern the planetary ecosystem in the form of reactions that consume and produce resources (Heylighen et al., 2015). For example, the large-scale contribution of organisms to the planetary system can be described abstractly by the following self-sustaining cycle of reactions:

plants + minerals + sunlight + CO_2 → plants + oxygen + waste
animals + plants + oxygen → CO_2 + animals + waste
waste + decomposers + oxygen → decomposers + CO_2 + minerals

Such "natural" reactions are then added to the "social" and "technological" reactions already present in the network, which include, e.g.:

fossil fuels + power plants → power plants + CO_2 + electricity

The network is then analyzed as a whole by the algorithms searching for self-sustaining organizations. In this case, the program may come to the conclusion that the overall network is not sustainable because of an overproduction of CO_2. Making it sustainable requires adding reactions that consume CO_2, or that produce energy without CO_2 as waste product, e.g.:

sunlight + solar panels → solar panels + electricity

The matching and recommendation algorithms may therefore increase the priority of such reactions, while downgrading the priority of CO_2 producing reactions. This ensures that the ecosystem's "services" (e.g. waste decomposition, CO_2 absorption), "offers" (e.g. of oxygen) and "needs" (e.g. for forested areas) are duly taken into account (Costanza et al., 1997), so that the global system including nature, society and technology is truly sustainable.

For this to work, we need an as accurate and comprehensive as possible representation of the reactions that govern the planetary ecosystem. The formulation of these production rules and their coupling into a global network may be provided in part by the emerging "Global Systems Science" with its interest in developing a "Living Earth Simulator" (Helbing, Bishop, Conte, Lukowicz, & McCarthy, 2012). In practice, an immense variety of more specific, local reactions (e.g. between a particular species of flower and the insect populations that pollinate it) will be entered into the system by a variety of experts, while being evaluated for their relative importance by other experts, and tested as to their adequacy by checking in how far the predictions inferred from them are accurate. Thus, they will be subjected to the same kind of learning mechanisms that should gradually broaden and deepen the overall offer network, ensuring that it becomes increasingly effective in satisfying all needs, social as well as natural.

The offer network protocol will need to ensure that ecological reactions are entered as easily into the system as human offers and needs. This may require the ontological specification of ecological "agents" that perform specific reactions, such as certain populations (e.g. of bees) or geographical areas (e.g. the Amazon rain forest).

5. Conclusion: tackling global challenges

A previous paper (Heylighen, 2016b) has argued that offer networks are potentially able to tackle the most fundamental problems of our present society, including inequality, fragility, financial instability, poverty and unsustainability. This requires changing the very foundations of our present socio-economic system: away from the accumulation of money and towards an intelligent system of sharing, collaboration and exchange based on directly matching offers and needs (Goertzel, 2015). The principle is to maximally exploit potential synergies between what different human, natural and technological agents have to offer, by efficiently coordinating their actions at the local and global level. The present paper has started to develop a concrete strategy for reaching this lofty goal.

The first step was to formulate the abstract elements of an *offer network protocol*, i.e. a universal system of formal representations and methods for expressing and matching offers and needs. The mathematical foundations for such a protocol may be derived from a synthesis of production systems (Anderson, 2014; Holland et al., 1989) with their condition-action rules and inference engines, and Chemical Organization Theory (Dittrich & Fenizio, 2007; Heylighen et al., 2015), with its procedures for finding self-sustaining "organizations" within networks of reactions. We performed a first step in that synthesis by showing the direct correspondence between *condition-action rules* and *reactions*: both transform a conjunction of necessary conditions or resources (needs) into a conjunction of newly produced conditions or resources (offers). Assuming that actions are only performed when they are beneficial to the agent performing them, the input conditions can be interpreted as part of that agent's needs, and the output conditions as its offer.

The next step will be to explore and integrate the different algorithms that can match offers and needs by chaining together such rules or reactions so that they reach *closure*, i.e. generate a subnetwork in which all resources that are consumed (needed) by some reaction(s) are also produced (offered) by some other reaction(s). We proposed some examples to illustrate how this can be done, and argued on theoretical grounds that this can in principle be achieved in the most general case. However, further research will need to experiment with a variety of algorithms and heuristics in order to find those that can do this most efficiently for huge networks of rules.

Given such an abstract model of what an offer network can do, the next stage is to develop concrete strategies for implementing these abilities, supported by existing and emerging technologies. These strategies define a *roadmap* towards the creation of a planet-wide offer network. As this network becomes more intelligent, encompassing, and powerful, it will increasingly start to play the role of a Global Brain, i.e. a nervous system for the planet that would be able to tackle about all problems we are confronted with (Heylighen, 2015). The process is likely to require several decades before it can fully realize these promises. However, most of the steps we sketched can be initiated right now, using available technologies.

A first step, already partly achieved (Kleedorfer et al., 2014), is to develop a basic protocol and platform that allows the identification of users, their offers and their needs. A next step is to help these users formulate their offers and needs more precisely, so that they can be matched more effectively. This can be realized with the help of *ontologies* of concepts and their semantic relations. These allow disambiguating natural language expressions and suggesting more accurate categorizations and relevant attributes. Practically, this is likely to happen by integrating the offer network protocol with the protocols of the emerging *Semantic Web*—such as RDF or OWL. Another step that co-opts presently emerging technologies is an integration with the *Internet of Things*. This would allow humans to exchange offers and needs with physical objects.

The gradual incorporation of ontologies and related procedural and causal rules would provide the offer network with an extensive knowledge of the world. This would enable it to answer theoretical as well as practical questions. Most importantly, it would boost its ability to tackle complex problems that require abstract reasoning in combination with the provision of concrete resources. For example, the offer network should be able to aggregate the individual needs for health care in a particular town, induce the collective need for building a hospital, and decompose that need in the form of a detailed project listing the different human, physical and cognitive resources that in coordination would be able to satisfy the need. Thus, it could create an *ad hoc* organization just to tackle this particular problem, without need for the establishment of permanent institutions.

More generally, an intelligent offer network should be able to implement a form of adaptive and distributed *governance* that would mobilize people and resources to work for the benefit of humanity, while preventing harmful activities such as fraud, waste or pollution. This can be achieved in part by including laws, policies and ethical codes into the network in the form of condition-action rules, in part by adding criteria that prioritize the offers that are most beneficial to society, i.e. that have positive externalities.

Perhaps the most important externalities for long-term sustainability are the costs and benefits imposed on the planetary *ecosystem*. These are best addressed by including natural agents, such as forests, lakes or animal populations, and the reactions they take part in, into the offer network. Thus, the network would take into account the offers and needs of the full biosphere—and not just those of its human population—when trying to find a maximally self-sustaining, synergetic organization.

In order too implement such adaptive governance, the offer network would need more than “hard”, formal categories. To function efficiently, it would also need a variety of “soft” criteria that can be used to rank or prioritize options that match the hard criteria equally well. Some of these criteria will depend on the individual tastes and preferences of the user, as expressed explicitly, or implicitly through previous choices. Others, representing the preferences of people collectively, can be deduced from generalized reputation scores, where individuals, categories, rules and actions collect rewards proportionally to the degree to which they have satisfied the needs of society.

Since these scores are constantly updated, while being open to scrutiny, criticism and improvement, they cannot congeal into a rigid “utility function” that is vulnerable to manipulation and to the *complexity of value problem* (LessWrong, 2013; Yudkowsky, 2011). Moreover, the priority ordering derived from such scores will be highly context-dependent, as for example the same person will have a different reputation depending on what kind of offer she makes, or the same offers made to different people will be ordered differently depending on their personal preferences at this moment.

Implementing all these steps should eventually transform the offer network into a true Global Brain. Such a system would be able to answer any theoretical questions for which the knowledge is available, whilst simultaneously implementing and monitoring practical solutions for all the problems that confront individuals, society and the global ecosystem (Heylighen, 2015).

The creation of such a global problem-solving system may seem too ambitious to be realistic. Yet, the envisaged offer network protocol makes it achievable in principle. The reason is that offer networks are founded on a theoretical framework that does not *a priori* separate cognitive, physical, social, economic, legal, ecological and other problems (Heylighen, 2011; Heylighen et al., 2015)—unlike practically any existing framework. Instead, the offer network protocol considers all these issues as components of one huge distributed network of interdependent conditions and actions. Because all conditions and actions are linked through a single medium, this network can be analyzed by intelligent computer programs so as to find the most synergetic solutions. Given the tremendous on-going progress in both hardware and software, it seems likely that such an intelligent system can be implemented on a planetary scale over the coming decades.

What makes tackling global challenges achievable in practice is the democratic and open-ended nature of the network. A protocol is a formally specified language for expressing and manipulating information. It does not prescribe which data, knowledge or values are to be expressed. Therefore, anyone can use it to express any need, any offer or any rule, or to assess the value of any proposed solution. The more people participate in this way, the richer the store of knowledge, values and resources the offer network can rely on, and the greater its distributed intelligence. The greater its intelligence, the more helpful it will be to people, and the more these people will be inclined to use it, and thus add to its store of shared resources and collective wisdom.

This virtuous cycle is similar to the positive feedback that led to the explosive growth of the World-Wide Web (Berners-Lee & Fischetti, 1999). This growth was made possible by the formulation of a public protocol (HTML/HTTP/URL) that could be used to freely create content in a public medium (the Internet) in such a way that everybody could profit from it—and this in a simple and flexible way, without needless technicalities or restrictions. The development of a universal and easy to use offer network protocol is likely to have a similarly revolutionary impact, but now going beyond the publication of information to the coordination of action. That should be sufficient to transform the World-Wide Web into a Global Brain, i.e. from a passive repository of knowledge to a system actively using that knowledge to solve any kind of problem.

A familiar criticism of this utopian vision is that people are too irrational, selfish, undisciplined or stupid to use something like an offer network protocol in an effective manner. This is the same kind of skeptical attitude that led many observers in the 1990's to assume that the Internet and Web may be useful tools for scientists and technologists, but that no one could expect ordinary citizens to adopt such complex technologies. The same skeptics were convinced that an open system like Wikipedia, which invites everybody to write or edit texts on any subject, could never become a dependable encyclopedia. These skeptics overlooked a number of factors that are worth reviewing here, because these same factors can facilitate the realization of an offer network.

First, technologies that are initially formulated in a highly abstract, technical manner eventually get implemented in an intuitive, user-friendly manner, which requires minimal effort from their users. Second, when a system becomes truly useful, people become increasingly motivated to learn to use it in the most effective way, putting in the necessary time and effort to assimilate all the tips and tricks. Third, such usage eventually becomes part of general education—just like reading and writing evolved from highly technical activities performed by a small elite of experts to general skills expected from everybody. Finally, we already alluded to the virtuous cycle of stigmergy (Heylighen, 2007, 2016a): the partial products (such as an incomplete Wikipedia article) of some, when clearly expressed in a common medium, stimulate and guide others to add their own contributions. Thus, contributions build on top of contributions, generating an increasingly rich and sophisticated system in which everyone may eventually find what s/he needs, and where poor material is eventually replaced by better material.

An additional answer to the skeptics is that while offer networks will require people to change their habits and assimilate new and potentially complex forms of interactions, the kind of transactions demanded from them are not intrinsically more complicated than those of our present market economy, with its intricate systems of marketing and financing. This economic system functions remarkably well, in spite of its complexity and the irrationality, selfishness, and "stupidity" of the people that participate in it. There is no reason why an offer network economy, which shares some characteristics with older forms of barter and gift economies and which uses self-organizing mechanisms at least as powerful as those of the market's "invisible hand" (Goertzel, 2015; Heylighen, 2016b), should not be able to cope equally well with human limitations.

In conclusion, the roadmap towards a global system for coordination and problem solving sketched here will require years, if not decades, of hard work by an increasingly large group of contributors to be realized. The complexity of the envisaged result is such that at present we have many more question marks than answers about how the different stages are best implemented. However, this is a characteristic of all far-reaching socio-technological transitions. The global spread of systems such as the Internet, the World-Wide Web, Wikipedia, cellular networks, and the Google search engine shows that in spite of the complexity of the eventual implementation, truly good ideas can become a reality remarkably quickly. I believe that the offer network protocol—or some future version of it—is such an idea with world-changing potential. The coming years will tell us in how far the present roadmap presents a realistic scenario for its deployment…

This research was supported by the Yuri Milner Foundation through its grant "Global Brain Institute". Several of its ideas were inspired by contributions from Ben Goertzel, Florian Kleedorfer, Tomas Veloz, Marta Lenartowicz, and other members of the Global Brain Institute.

References

Allemang, D., & Hendler, J. (2011). *Semantic Web for the Working Ontologist: Effective Modeling in RDFS and OWL*. Elsevier.

Anderson, J. R. (2014). *Rules of the Mind*. Psychology Press.

Ashby, W. R. (1972). Setting goals in cybernetic systems. In H. W. Robinson & D. E. Knight (Eds.), *Cybernetics, Artificial Intelligence and Ecology* (pp. 33–44). New York: Spartan Books.

Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, *54*(15), 2787–2805. Retrieved from http://www.sciencedirect.com/science/article/pii/S1389128610001568

Atzori, M. (2015). Blockchain Technology and Decentralized Governance: Is the State Still Necessary? *Available at SSRN*. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2709713

Beigi, S. (2014). *Mindfulness Engineering: A Unifying Theory of Resilience for Volatile, Uncertain, Complex and Ambiguous (VUCA) World* (PhD thesis). University of Bristol, Bristol, UK.

Bennett, N., & Lemoine, G. J. (2014). What a difference a word makes: Understanding threats to performance in a VUCA world. *Business Horizons*, *57*(3), 311–317. doi:10.1016/j.bushor.2014.01.001

Berners-Lee, T., & Fischetti, M. (1999). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper San Francisco. Retrieved from http://portal.acm.org/citation.cfm?id=554813

Bernstein, A., Klein, M., & Malone, T. W. (2012). Programming the Global Brain. *Communications of the ACM*, *55*(5), 1. Retrieved from http://cci.mit.edu/publications/CCIwp2011-04.pdf

Bonabeau, E. (2009). Decisions 2.0: The power of collective intelligence. *MIT Sloan Management Review*, *50*(2), 45–52. Retrieved from http://www.worldwideopen.org/uploads/groups_bull/files/30/decision%202.0.pdf

Centler, F., Kaleta, C., di Fenizio, P. S., & Dittrich, P. (2008). Computing chemical organizations in biological networks. *Bioinformatics*, *24*(14), 1611–1618. Retrieved from http://bioinformatics.oxfordjournals.org/content/24/14/1611.short

Cornes, R., & Sandler, T. (1996). *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press.

Costanza, R., D'Arge, R., de Groot, R., Farber, S., Grasso, M., Hannon, B., Limburg, K., et al. (1997). The value of the world's ecosystem services and natural capital. *Nature*, *387*, 253–260. doi:10.1038/387253a0

De Alfaro, L., Kulshreshtha, A., Pye, I., & Adler, B. T. (2011). Reputation systems for open collaboration. *Communications of the ACM*, *54*(8), 81–87. Retrieved from http://dl.acm.org/citation.cfm?id=1978560

De Rosnay, J. (2000). *The Symbiotic Man: A new understanding of the organization of life and a vision of the future*. Mcgraw-Hill. Retrieved from http://pespmc1.vub.ac.be/books/DeRosnay.TheSymbioticMan.pdf

Diamandis, P. H., & Kotler, S. (2012). *Abundance: The Future Is Better Than You Think*. Free Press.

Dittrich, P., & Fenizio, P. S. di. (2007). Chemical Organisation Theory. *Bulletin of Mathematical Biology*, *69*(4), 1199–1231. doi:10.1007/s11538-006-9130-8

Dittrich, P., & Winter, L. (2008). Chemical organizations in a toy model of the political system. *Advances in Complex Systems*, *11*(04), 609. doi:10.1142/S0219525908001878

Drexler, E. K. (2013). *Radical Abundance: How a Revolution in Nanotechnology Will Change Civilization*. PublicAffairs.

Evrard, A. E., Erdmann, C., Holmquist, J., Damon, J., & Dietrich, D. (2015). Persistent, Global Identity for Scientists via ORCID. *arXiv:1502.06274 [astro-ph, physics:physics]*. Retrieved from http://arxiv.org/abs/1502.06274

Fellbaum, C. (2012). WordNet. *The Encyclopedia of Applied Linguistics*. John Wiley & Sons, Inc. Retrieved from http://onlinelibrary.wiley.com.ezproxy.vub.ac.be:2048/doi/10.1002/9781405198431.wbeal1285/abstract

Fouss, F., Pirotte, A., Renders, J. M., & Saerens, M. (2007). Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation. *Knowledge and Data Engineering, IEEE Transactions on*, *19*(3), 355–369.

Gallo, G., Longo, G., Pallottino, S., & Nguyen, S. (1993). Directed hypergraphs and applications. *Discrete Applied Mathematics*, *42*(2-3), 177–201.

Garrod, J. Z. (2016). The Real World of the Decentralized Autonomous Society. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society*, *14*(1), 62–77. Retrieved from http://triple-c.at.dd29412.kasserver.com/index.php/tripleC/article/view/692

Gernert, D. (2006). Pragmatic Information: Historical Exposition and General Overview. *Mind and Matter*, *4*(2), 141–167.

Goertzel, B. (2002). *Creating internet intelligence: Wild computing, distributed digital consciousness, and the emerging global brain*. Kluwer Academic/Plenum Publishers.

Goertzel, B. (2015). Beyond Money: Offer networks, a potential infrastructure for a post-money economy. In B. Goertzel & T. Goertzel (Eds.), *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity* (pp. 522–549). Humanity+ Press. Retrieved from http://globalbraininstitute.github.io/onet/files/offer_networks.pdf

Helbing, D. (2015). Societal, Economic, Ethical and Legal Challenges of the Digital Revolution: From Big Data to Deep Learning, Artificial Intelligence, and Manipulative Technologies. *arXiv:1504.03751 [physics]*. Retrieved from http://arxiv.org/abs/1504.03751

Heylighen, F. (1999). Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map. *Computational & Mathematical Organization Theory*, *5*(3), 253–280. doi:10.1023/A:1009690407292

Heylighen, F. (2001). Bootstrapping knowledge representations: from entailment meshes via semantic nets to learning webs. *Kybernetes*, *30*(5/6), 691–725. doi:10.1108/EUM0000000005695

Heylighen, F. (2002). The global brain as a new utopia. *Zukunftsfiguren*. Suhrkamp, Frankurt. Retrieved from http://pespmc1.vub.ac.be/papers/GB-Utopia.pdf

Heylighen, F. (2007). Why is Open Access Development so Successful? Stigmergic organization and the economics of information. In B. Lutterbeck, M. Baerwolff & R. A. Gehring (Ed.), *Open Source Jahrbuch 2007* (pp. 165–180). Lehmanns Media. Retrieved from http://pespmc1.vub.ac.be/Papers/OpenSourceStigmergy.pdf

Heylighen, F. (2008). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain. *Globalization as evolutionary process: modeling global change*, Rethinking globalizations (p. 284). Routledge. Retrieved from http://pcp.vub.ac.be/papers/AcceleratingEvolution.pdf

Heylighen, F. (2011). Conceptions of a Global Brain: an historical review. *Evolution: Cosmic, Biological, and Social, eds. Grinin, L. E., Carneiro, R. L., Korotayev A. V., Spier F.* (pp. 274–289). Uchitel Publishing. Retrieved from http://pcp.vub.ac.be/papers/GB-conceptions-Rodrigue.pdf

Heylighen, F. (2013). Self-organization in Communicating Groups: the emergence of coordination, shared references and collective intelligence. In À. Massip-Bonet & A. Bastardas-Boada (Eds.), *Complexity Perspectives on Language, Communication and Society*, Understanding Complex Systems (pp. 117–149). Berlin, Germany: Springer. Retrieved from http://pcp.vub.ac.be/Papers/Barcelona-LanguageSO.pdf

Heylighen, F. (2014). Challenge Propagation: Towards a theory of distributed intelligence and the global brain. *Spanda Journal*, *V*(2), 51–63. Retrieved from http://cleamc11.vub.ac.be/papers/ChallengePropagation-Spanda.pdf

Heylighen, F. (2015). Return to Eden? Promises and Perils on the Road to a Global Superintelligence. In B. Goertzel & T. Goertzel (Eds.), *The End of the Beginning: Life, Society and Economy on the Brink of the Singularity*. Humanity+ Press. Retrieved from http://pespmc1.vub.ac.be/Papers/BrinkofSingularity.pdf

Heylighen, F. (2016a). Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research*. doi:10.1016/j.cogsys.2015.12.002

Heylighen, F. (2016b). Towards an Intelligent Network for Matching Offer and Demand: from the sharing economy to the Global Brain. *Technological Forecasting & Social Change*, *in press*. Retrieved from http://pespmc1.vub.ac.be/Papers/GB-OfferNetwork.pdf

Heylighen, F. (2016c). Distributed Intelligence Technologies: present and future applications. In M. Burgin & W. Hofkirchner (Eds.), *The Future Information Society: Social and Technological Problems*. World Scientific. Retrieved from http://pespmc1.vub.ac.be/Papers/GB-applications-survey-WS.pdf

Heylighen, F., Beigi, S., & Veloz, T. (2015). Chemical Organization Theory as a modeling framework for self-organization, autopoiesis and resilience. *International Journal Of General Systems*, *submitted*. Retrieved from http://pespmc1.vub.ac.be/papers/COT-ApplicationSurvey-submit.pdf

Heylighen, F., & Bollen, J. (2002). Hebbian Algorithms for a Digital Library Recommendation System. *International Conference on Parallel Processing* (p. 439). Los Alamitos, CA, USA: IEEE Computer Society. doi:10.1109/ICPPW.2002.1039763

Heylighen, F., Busseniers, E., Veitas, V., Vidal, C., & Weinbaum, D. R. (2012). *Foundations for a Mathematical Model of the Global Brain: architecture, components, and specifications* (GBI Working Papers No. 2012-05). GBI Working Papers. Retrieved from http://pespmc1.vub.ac.be/papers/TowardsGB-model.pdf

Heylighen, F., & Campbell, D. T. (1995). Selection of organization at the social level: Obstacles and facilitators of metasystem transitions. *World Futures*, *45*(1-4), 181–212. doi:10.1080/02604027.1995.9972560

Heylighen, F., Kostov, I., & Kiemen, M. (2013). Mobilization Systems: technologies for motivating and coordinating human action. In Peters M. A., Besley T. and Araya D. (Ed.), *The New Development Paradigm: Education, Knowledge Economy and Digital Futures*. Routledge. Retrieved from http://pcp.vub.ac.be/Papers/MobilizationSystems.pdf

Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, *18*(7), 1527–1554. doi:10.1162/neco.2006.18.7.1527

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1989). *Induction: processes of inference, learning, and discovery*. MIT Press.

Jøsang, A., Ismail, R., & Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision Support Systems*, Emerging Issues in Collaborative Commerce, *43*(2), 618–644. doi:10.1016/j.dss.2005.05.019

Kleedorfer, F., Busch, C. M., Grill, G., Khosravipour, S., Salcher, F., Tus, A., & Gstrein, E. (2013). Web of Needs–A New Paradigm for E-Commerce. *Business Informatics*

(CBI), 2013 IEEE 15th Conference on (pp. 316–322). IEEE. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6642894

Kleedorfer, F., Busch, C. M., Pichler, C., & Huemer, C. (2014). The Case for the Web of Needs. *2014 IEEE 16th Conference on Business Informatics (CBI)* (Vol. 1, pp. 94–101). Presented at the 2014 IEEE 16th Conference on Business Informatics (CBI). doi:10.1109/CBI.2014.55

Lenat, D. B. (1995). CYC: A large-scale investment in knowledge infrastructure.

LessWrong. (2013). Complexity of value. *LessWrong*. Wiki. Retrieved from http://wiki.lesswrong.com/wiki/Complexity_of_value

Lipson, H., & Kurman, M. (2013). *Fabricated: The New World of 3D Printing*. John Wiley & Sons.

Livet, P. (2012). Web Ontologies as Renewal of Classical Philosophical Ontology. *Metaphilosophy*, *43*(4), 396–404. doi:10.1111/j.1467-9973.2012.01757.x

Mayer-Kress, G., & Barczys, C. (1995). The global brain as an emergent structure from the Worldwide Computing Network, and its implications for modeling. *The information society*, *11*(1), 1–27.

Muehlhauser, L., & Helm, L. (2012). The Singularity and Machine Ethics. In A. H. Eden, J. H. Moor, J. H. Søraker, & E. Steinhart (Eds.), *Singularity Hypotheses*, The Frontiers Collection (pp. 101–126). Springer Berlin Heidelberg. Retrieved from http://link.springer.com/chapter/10.1007/978-3-642-32560-1_6

Parunak, H. V. D. (2006). A survey of environments and mechanisms for human-human stigmergy. In D. Weyns, H. V. D. Parunak, & F. Michel (Eds.), *Environments for Multi-Agent Systems II* (pp. 163–186). Springer.

Rayward, W. B. (1994). Visions of Xanadu: Paul Otlet (1868-1944) and Hypertext. *Journal of the American Society for Information Science*, *45*(4), 235–250. Retrieved from http://people.lis.uiuc.edu/~wrayward/Visions of Xanadu_JASIS.pdf

Recordon, D., & Reed, D. (2006). OpenID 2.0: A Platform for User-centric Identity Management. *Proceedings of the Second ACM Workshop on Digital Identity Management*, DIM '06 (pp. 11–16). New York, NY, USA: ACM. doi:10.1145/1179529.1179532

Ricci, F., Rokach, L., & Shapira, B. (2011). Introduction to Recommender Systems Handbook. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender Systems Handbook* (pp. 1–35). Springer US. Retrieved from http://link.springer.com/chapter/10.1007/978-0-387-85820-3_1

Rifkin, J. (2014). *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. Palgrave Macmillan. Retrieved from http://www.bookdepository.com/Zero-Marginal-Cost-Society-Jeremy-Rifkin/9781137278463

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. San Diego: University of California Press.

Singh, S., & Karwayun, R. (2010). A comparative study of inference engines. *Information Technology: New Generations (ITNG), 2010 Seventh International Conference on* (pp. 53–57). IEEE. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5501457

Thaler, R. H., & Sunstein, C. R. (2003). Libertarian Paternalism. *American Economic Review*, *93*(2), 175–179.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale Univ Pr.

Theraulaz, G., & Bonabeau, E. (1999). A Brief History of Stigmergy. *Artificial Life*, *5*(2), 97–116.

Wagner, C. (2006). Breaking the knowledge acquisition bottleneck through conversational knowledge management. *Information Resources Management Journal (IRMJ)*, *19*(1), 70–83. Retrieved from http://www.igi-global.com/article/information-resources-management-journal-irmj/1286

Weinbaum, D. (Weaver), & Veitas, V. (2016). Open-Ended Intelligence: On the Role of Individuation in AGI. In B. Steunebrink, P. Wang, & B. Goertzel (Eds.), *Artificial General Intelligence*, Lecture Notes in Computer Science (pp. 43–52). Springer International Publishing. Retrieved from http://link.springer.com/chapter/10.1007/978-3-319-41649-6_5

Wells, H. (1937). *World Brain*. Ayer Co Pub. Retrieved from http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0836920333

Wright, R. (2001). *Nonzero: The logic of human destiny*. Vintage.

Yudkowsky, E. (2011). Complex Value Systems in Friendly AI. In J. Schmidhuber, K. R. Thórisson, & M. Looks (Eds.), *Artificial General Intelligence* (Vol. 6830, pp. 388–393). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from http://link.springer.com/10.1007/978-3-642-22887-2_48