The Tragedy of the Digital Commons

Article *in* Ethics and Information Technology · November 2003

DOI: 10.1007/s10676-004-2895-2 · Source: CiteSeer

2 authors:

Gian Maria Greco
Autonomous University of Barcelona
18 PUBLICATIONS **120** CITATIONS

SEE PROFILE

Luciano Floridi
University of Oxford
467 PUBLICATIONS **10,639** CITATIONS

SEE PROFILE

Ethics and Information Technology **6:** 73–81, 2004.

The tragedy of the digital commons

Gian Maria Greco[1,3] and Luciano Floridi[2,3]
[1]*Dipartimento di Filologia Classica e di Scienze Filosofiche, Università degli Studi di Lecce, Italy E-mail: gianmaria.greco@ateneo.unile.it*
[2]*Dipartimento di Scienze Filosofiche, Università degli Studi di Bari, Italy*
E-mail: luciano.floridi@philosophy.oxford.ac.uk
[3]*Information Ethics research Group, Computing Laboratory, Oxford University, Great Britain*

Abstract. In the paper it is argued that bridging the digital divide may cause a new ethical and social dilemma. Using Hardin's *Tragedy of the Commons*, we show that an improper opening and enlargement of the digital environment (*Infosphere*) is likely to produce a Tragedy of the Digital Commons (TDC). In the course of the analysis, we explain why Adar and Huberman's previous use of Hardin's *Tragedy* to interpret certain recent phenomena in the Infosphere (especially peer-to-peer communication) may not be entirely satisfactory. We then seek to provide an improved version of the TDC that avoids the possible shortcomings of their model. Next, we analyse some problems encountered by the application of classical ethics in the resolution of the TDC. In the conclusion, we outline the kind of work that will be required to develop an ethical approach that may bridge the digital divide but avoid the TDC.

Key words: artificial agents, digital divide, Garrett Hardin, Infosphere, tragedy of the commons, tragedy of the digital commons

Abbreviations: DD – digital divide; ICTs – information and communication technologies; TC – tragedy of the commons; TDC – tragedy of the digital commons

Introduction

"For the magnitude of the sea is such, as to be sufficient for the use of all nations, to allow them without inconvenience and prejudice to each other the right of fishing, sailing, or any other advantage which that element affords". Thus Hugo Grotius in 1625.[1] In those days, it seemed that the resources of the oceans were inexhaustible, and hence that regulating their use was unnecessary. Grotius' approach was correct, but his conclusions were not. He rightly considered the sea as a *common*, available to everybody, at anyone's discretion. But he failed to see that the exponential increment of the population, the indiscriminate use of natural resources and their comparatively slow process of regeneration could cause irretrievable exploitation. As Hardin (1998) remarks, the foreseeable exhaustion of worldwide fisheries is an instructive example (actually, fisheries management and other problems in the use of a common are sometimes interpreted as instances of a *Tragedy of the Commons* (Hardin 1968), but more on this presently). In this paper, we argue that, like Grotius when talking about the sea, we should consider the environment produced by digital resources (*Infosphere*) as a common, but that, unlike Grotius, we should be aware of, and possibly prevent, the potential problems implicit in unregulated exploitation of the Infosphere. Pursuing a global extension and enlargement of the Infosphere in order to bridge the digital divide may cause a *Tragedy of the Digital Commons*.

Here is a brief summary of the paper. In the next section, we review the original version of the *Tragedy of the Commons* (henceforth TC). In the third section, we analyse two of its features that are relevant to the application of the model to the Infosphere. In the fourth section, we briefly outline the well-known problem of the digital divide; we discuss some of the strategies currently pursued in order to solve it; and we show how bridging the digital divide may bring about a *Tragedy of the Digital Commons* (henceforth TDC). In the fifth section, we draw on Hardin's TC to model agents' behaviours in the Infosphere. In the sixth section, we criticise a previous use of the TC by Adar and Huberman, showing why it may not be applicable to the Infosphere. In the seventh section, we offer an improved

[1] Grotius 1625. Grotius is also cited in Hardin 1998.

version of the TDC that avoids the possible shortcomings of Adar and Huberman's model. In the eighth section, we summarise the debate between ecocentric and anthropocentric ethics, in view of the analysis of some difficulties encountered by these approaches when they are applied to the Infosphere. The paper is intentionally problematic and it focuses only on the TDC and its causes. Thus, in the conclusion, we only hint at the kind of work that will be required to develop an ethical approach that may avoid the TDC, a task left to a paper we hope to complete in the near future.

The tragedy of the commons

The TC is a thought experiment used by Hardin (1968) in order to show how the single-minded pursuit of the individual good may cause destruction of the common good. As it is well-known, we shall recall it briefly.

Imagine some herdsmen who have access to a common pasture. Each of them can increase his herd by at least one sheep. This would bring to each of them a positive advantage and cause only very marginal damage (e.g., less forage), which would be shared with the rest of the community. The herdsmen behave like fully rational and self-interested agents, so each of them decides to pursue his individual strategy by bringing one more sheep to the commons. The process is repeated and the herds increase progressively. The situation becomes a tragedy when the herdsmen's rational decision finally causes exploitation from which the pasture cannot recover. The excessive search for the individual good has resulted in the destruction of the common good.

Although the tragedy is typically caused by a process that decreases the amount of (renewable) resources available to an increasing number of exploiters, Hardin shows that it can also be the outcome of a process of *pollution*. In this case, the exploiters still gradually impoverish the available resources, ending up by destroying them, yet they do so not by taking something from them, but by adding something that spoils their nature or equilibrium. The reasoning is the same in both cases: "The rational man finds that his share of the cost of the wastes he discharges into the commons is less than the cost of purifying his wastes before releasing them" (Hardin 1968: 1245). According to Hardin, the principal source of the problem of the TC lies in the *population*. In his view, excessive exploitation or pollution of the commons are the effects of *indiscriminate access* to it and of the *lack of control on demographic development*.

The tragedy of the commons and the infosphere

The TC is an instance of the Prisoner's Dilemma, applied to a conflict among several agents sharing the same bounded environment and limited resources.[2] Given its analytic power and its intuitive character, TC has been used over the years as a theoretical framework to model and interpret analogous situations in various disciplines.[3] What makes the TC so interesting is that the use-benefits the agents derive from the environment is a common good. In this paper, we consider the Infosphere as such a common good and, in the fifth section, we shall see that the TC can be usefully applied to the Infosphere as well. For this purpose, in this section we wish to call the reader's attention to two key features of the model: the concept of "bounded environment" and the assumption that the agents' behaviour is entirely self-interested.[4]

TC occurs in a bounded environment, whose resources are exhaustible, that is, finite, not immediately renewable and not totally resilient. Consider, for example, the contrast stressed by Hardin between the limited resources of the Biosphere and their free use. When a woodcutter cuts down a tree, he consumes the common, because that tree cannot grow again. Likewise, when someone discharges his wastes into the common, he destroys it because the environment may not recover at all or only too late. Now, it may be argued that the Infosphere is a totally different kind of environment, for it lacks effectively fixed dimensions.[5] When the agents "put" something into the Infosphere, like an email, one may contend that they are actually "expanding"[6] the digital space, since the latter can be seen as being equivalent to the

[2] In general, the Prisoner's Dilemma is a strategic situation where the rational pursuit of each individual's preferences leads to sub-optimal results. The TC is a *n*-person Prisoner's Dilemma related to common goods or bounded environments with limited resources (see Ostrom 1986; Fletcher and Zwick 2000).

[3] See, for example, Turner (1991) for the study of conflicts in systems of distributed artificial intelligence; Pfeiffer et al. (2001) on the analysis of the mechanisms of ATP production; Robert and Sorci (2001) on the evolution of obligate interspecific brood parasitism in birds.

[4] See Ostrom 1977, 1986.

[5] Such a characterisation of the Infosphere has been under attack from a variety of quarters over the past few years. A justification of this model has been carried out in Floridi (2003). For its use related to the digital divide see also Floridi (2002). We are grateful to Charless Ess for having called our attention to the importance of this point.

[6] The quotes are there to remind us of the slightly different meaning of the word in this context.

totality of objects that constitute it. Similarly, when the agents "take" something from the Infosphere, as they do when they download a file, one may object that they are not necessarily destroying the file itself and, consequently, destroying the digital environment. For this reason, care must be exercised in applying the TC model to the Infosphere. We shall do so *not* isomorphically – as if the Infosphere were simply just another environment exactly like the Biosphere – but homomorphically, that is, by paying attention to the peculiarities of the digital environment as well.

Our second caveat concerns the analysis of the agents' behaviour. Hardin assumes that each agent in the Biosphere follows his own strategy, while disregarding entirely all other agents, their interests and their actions. Social interrelations are mentioned only when the single herdsman reminds himself of the existence of the others in order to subdivide the damage. Hardin considers each herdsman as if he were isolated from the social context in which he works. One may object that this is one of the weakest features of the model, and that TC provides no more than an idealised model that, at best, is only partly applicable to real-life situations. However, the objection may be answered by pointing out that, first, this is true of all models, and second, that in the case of the Infosphere it is reasonable to argue that the "selfish" assumption appears largely justified. The Infosphere and its digital resources are such that they easily incline an ordinary user to be oblivious to the presence of others; this is the solipsistic nature of the web so often stressed in the literature. Consider, for example, the common situation of using an Internet connection. Typically, each user tends to use all the *bandwidth*[7] he has, without considering the presence or the needs of other users, who are consuming bandwidth at the same time. Each user considers the presence of other users only when there is saturation of the bandwidth, because he is then reminded that other agents are sharing the same limited resources. This is exactly what happens with Hardin's herdsmen, as we saw in the previous section.

We shall presently argue that there are good reasons to believe that a new version of the TC may arise in the Infosphere. However, applying the TC to the Infosphere without considering the previous two features means misusing Hardin's model and may easily lead to false problems and ineffective solutions.

The digital divide

A good way to approach the problem of the "digital tragedy" correctly, and hence provide a solid basis for its interpretation and for the development of a strategy for its resolution, is by looking first at its cause, the *digital divide* and the strategies for bridging it.

Information and Communication Technologies (ICTs) offer enormous possibilities for development and improvements but they also create new social disparities, usually analysed under the general label of *digital divide*.[8]

The digital divide (DD) is caused by the various difficulties encountered by people in their (lack of) interactions with ICTs and hence in accessing information contents, services and resources. The immediate effect of the DD is a discrimination between those who can be denizens of the Infosphere and those who cannot, between insiders and outsiders, between information rich and information poor.[9] In the long run, the DD redesigns the map of worldwide society, generating or widening generational, geographic, socio-economic and cultural divides. These are not merely reducible to the gap between industrialised and developing countries, since the DD is also a problem internal to many advanced societies. According to the Human Development Report

[7] "Bandwidth" is a technical term in ICT that refers to the difference between the highest and lowest frequencies of a transmission channel (the width of its allocated band of frequencies). However, following common practice and adopting the terminology used in the referred literature, in this paper we use "bandwidth" as synonymous for "connectivity", to refer to the data rate or capacity – that is, the volume of information or amount of data per unit time – that a system can process through a given communication channel, and hence as an informal measure of network capacity that is often wasted by users (waste of bandwidth).

[8] See eEurope 2000 and 2001; Human Development Report 2001; *The Information Society* 2003.

[9] The Human Development Report 2001 of the United Nations Development Programme (UNDP) shows a critical and difficult situation in the spread of the Internet. Internet users make up only 6.7% of the world population. Of those 54.3% are in the United States, 28.2% in the high-income OECD (excl. US), 3.2% in Latin America and the Caribbean, 2.3% in East Asia and the Pacific, 3.9% in Eastern Europe and CIS, 0.6% in the Arab States, 0.4% in Sub-Saharan Africa and 0.4% in South Asia. According to the Nielsen//NetRatings Global Internet Trends Q4 2002, Internet Population is circa 9.6% of World's Population, and it is so distributed: 29% United States, 23% Europe, 13% Asia and Pacific, 2% Latin America, 33% Rest of World (where: Europe covers France, Germany, Italy, Netherlands, Spain, Sweden, Switzerland and United Kingdom; Asia and Pacific covers Australia, Hong Kong, Japan; Latin America covers Brazil).

(2001), for example, Internet users are still predominantly *urban and located in certain regions, better educated and wealthier, young, male*.

Bridging the DD ought to be one of the goals of the international community. Thus, the UN, UNESCO and the European Union have organised action plans and task forces to study the problem and attempt to solve it. A case in point is the Digital Opportunity Task Force (DOT Force) of the US Government. The European Commission, on the other hand, has recognised that "closing the 'digital divide' between developed and developing countries is a key goal for the European Union" (eEurope 2000: 4) and it has instituted a special action plan, called *eEurope*, that aims at resolving the DD through the extension of Internet access to all fields of European society, in order "to bring European citizens on-line in all aspects of their lives, allowing them to participate in and benefit from the possibilities offered by digital technologies" (eEurope 2001: 3).

The principal strategies pursued to bridge the DD are spreading the availability of ICTs and increasing their accessibility and usability. The goal is to open the Infosphere to the largest number of people. This is tantamount to enlarging the population of the Infosphere. However, following Hardin's reasoning, the excessive and undisciplined increase of population in an environment may bring about the destruction of the environment itself. So, improving the Infosphere by bridging the digital divide, without paying due attention to the necessary responsibility of the users, is likely to bring about an increase in bandwidth saturation and in information pollution. As Hardin (1968) says: "The pollution problem is a consequence of population" (p. 1245). The DD is certainly an issue that demands the maximum attention and engagement. However, its resolution risks being the cause of a Tragedy of the Digital Commons.

The tragedy of the digital commons

By considering the Infosphere as an environment and a public good, we can show the dilemma of the TDC and use the TC model to interpret it.

As in the case of Hardin's TC, the TDC arises in two ways. First, the average user of the Infosphere behaves in much the same way as Hardin's herdsman. He is inclined to take excessive advantage of the resources of the Infosphere, without paying attention to the consequences of his behaviour. Consider, for instance, the problem of the overloading of the net and the related slowing down of traffic. Huberman and Lukose (1997) have shown that Internet congestion is a direct effect of selfish reasoning: "because Internet surfers are not charged in proportion to their use, it often appears rational for them to consume bandwidth [sic] greedily while thinking that their actions have little effect on the overall performance of the network. Because every individual can reason this way, the whole Internet's performance can degrade considerably, which makes everyone worse off" (p. 535).

Second, the problem of pollution arises in the Infosphere too. It concerns (what might be considered, depending on the level of abstraction adopted) an indiscriminate and improper use of technologies and digital resources, with the consequent overproduction of data – therefore an excess of information, often redundant – and corruption of communications, which may be transformed into mere noise. A concrete example is *spam*. Apparently, more than 45% of email traffic consists of junk messages and an increasing number of ISPs are being forced to use software filters in order to try at least to reduce it. "America OnLine is now blocking an average of 780 m junk e-mails daily" (*The Economist* 2003: 58). The spam problem is more than a simple nuisance. It has been estimated that "spam will cost American organisations alone more than 10 billion dollars this year in lost productivity and extra spending to combat it. World-wide costs are much larger" (*The Economist* 2003: 58). The problem is becoming so huge that some governments, e.g. the EU, are passing strict legislation against spammers.

As a consequence of the gradual bridging of the DD, the increase in the population of the Infosphere may bring about a *bandwidth exploitation* and an *information pollution* of the digital environment. This is a tragedy of the digital commons no less worrying and urgent than the one occurring in the Biosphere. In the end, it is the classic case of jumping out of the frying pan into the fire: in order to avoid the DD we may jump into the TDC. Clearly due attention must be paid to preventing this. We shall return to this point in the conclusion.

The tragedy of the digital commons and peer-to-peer networks

Our attempt to identify a TDC in the Infosphere has been predated by Adar and Huberman (2000), who claim to have identified a new version of Hardin's TC, which they call the "Tragedy of Digital Commons". By analysing user traffic on the Gnutella network, Adar and Huberman have discovered the

presence of an excessive number of free riders.[10] Sending a message on the network and evaluating its effects over a period of 24 h, they have determined that approximately 70% of Gnutella users do not share files, and that approximately 50% of all the answers to the original demand come from a mere 1% of the sharing hosts, while 98% come from 25% of the network hosts. The consequence is an increase both in the degradation of system performance and in the "vulnerability" (their term) of the network. The performance degradation results from (a) the decrease in the number of peers who share most of the requested files and (b) the increase in the number of hosts within the search horizon.[11] On the one hand, since the bulk of the shared files is contributed by a small number of peers, steady demand for these files results in a dramatic overload of the relevant hosts, i.e., the files' owners. On the other hand, the increase in the number of hosts makes the identification and searching of the available files more difficult and slower, with the result that an increasing number of files become unreachable for a growing number of users. In the long run, this situation makes the network more "vulnerable". For example, in order to protect the copyright of the files it becomes sufficient to identify the few peers hosting them and to stop them sharing the files, without trying to identify every peer on the network. Adar and Huberman conclude that "if this trend continues copyright issues might become moot compared to the possible collapse of such systems". In the presence of numerous free riders, they find a new version of the TC: "since files on Gnutella are treated like a public good and the users are not charged in proportion to their use, *it appears rational for people to download music files without contributing by making their own files accessible to other users*. Because every individual can reason this way and free ride on the efforts of others, the whole system's performance can degrade considerably, which makes everyone worse off – *the tragedy of the digital commons*" (Adar and Huberman 2000, emphasis added).

[10] *Gnutella* is a network where users can share files with everyone, independently on any server. More technically, Gnutella is a protocol designed for sharing peer-to-peer files in a distributed network, see http://www.gnutella.com/.

[11] "The search horizon is the farthest set of hosts reachable by a search request" (Adar and Huberman 2000). Technically speaking, a peer is a user on a network who has features equivalent to other users, and a host is a computer (or a computer system) that is connected to a network and that contains the data. For example, a peer is another computer running Gnutella. Following Adar and Huberman, in the present section we use the two terms as interchangeable.

Adar and Huberman's TDC has attracted a great deal of interest and discussion on the web. From our point of view, it may be subject to four main criticisms.

First, the model does not consider the importance of the dimensional feature of the environment (see third section) in order to apply the TC to this environment. The Infosphere and the Biosphere have different properties, and downloading a file from a peer-to-peer network, like Gnutella, does not involve the destruction or elimination or even the removal of the file from the network. The file sharing activity involves the use, through duplication, of the contents of the network. It does not imply a damaging or deleting of the downloaded data. So there is no analogy with the woodcutter's action mentioned above.

Second, the fundamental feature of Hardin's TC is the safeguarding of the common and not its development. The fact that most of the peers download public files without sharing any of their own does not bring about the destruction of the network, it is a mere "missed opportunity", for it does result in a lack of improvement of the ecosystem, but it is not an alteration of its equilibria. The network "vulnerability", stressed by Adar and Huberman, depends exclusively on external action against the hosts of the network (e.g., the shutdown of a host because of legal action by a record company). It does not depend on internal mechanisms.

Third, Adar and Huberman's analysis might be viewed as confusing levels in the manner, known in philosophy, since Aristotle, as "metabasis eis allo genos". In modern terminology, this consists in the unjustified transition from a certain level of abstraction, at which a system is analysed, to another. Adar and Huberman begin by adopting a level of abstraction at which the Infosphere is treated as a bounded environment, comparable to Hardin's Biosphere. They then proceed by analysing the relations among the informational objects in a way that presupposes an interpretation of the Infosphere as an unbounded space, lacking fixed dimensions. But when they draw their conclusions regarding the Tragedy of the Digital Commons, they implicitly go back to the original level of abstraction to interpret the Infosphere as a bounded environment. The confusion is obvious and the errors it causes invalidate their analysis.

Finally, and as a result of the previous confusion, Adar and Huberman analyse the behaviours of what are actually two kinds of agents, human and artificial, without distinguishing between them. This is a problem because the issues concerning the sharing of files involve the behaviour of peers, and so primarily

human agents, whereas the issues involving the improvement of the environment – namely the system's vulnerability and the degradation of the system's performance – concern digital features and informational objects, and so artificial agents as well (we shall return to this feature in the next section).

To conclude, the problem analysed by Adar and Huberman is not explicable in terms of a new TC, but in terms of a lack of improvement of the Infosphere. This does not mean that their analysis is mis-directed. As we have seen, the Infosphere may well be affected by a Tragedy of the Digital Commons, but it is important to delineate it correctly. Besides, having learnt from Adar and Huberman, we should pay attention to the behaviours of the artificial agents as well.

Extension of the tragedy of the digital commons

The TDC, analysed in the fifth section, describes exclusively human agents' actions. It does not thus far concern artificial agents, but Hardin's model can also be fruitfully applied to the possible actions and interactions of these agents in the Infosphere.[12]

Artificial agents may exploit or pollute the Infosphere. Consider a worm, that is, a computer program whose main actions are self-replication and the spreading of copies of itself within computer systems. The worm's multiplication may generate information pollution in the ecosystem. As in the case of spam, the worm's action may cause an over-production of redundant information and data (but if the worm is checking the network, the data produced does not pollute the environment). At the same time, by spreading copies of itself within the Infosphere, a worm consumes bandwidth. Since every copy is a self-sufficient program, it will create more copies and spread them through the environment. Consequently, the combined actions of all the worms may produce an irretrievable exploitation of the environment.[13]

Worms are not the only artificial agents that can exploit the environment "tragically". Perhaps an even more typical example is represented by a virus programmed to delete the files on a system. In this case, the damage is represented not by the exploitation of the environmental resources, like bandwidth, but by the destruction of parts of the digital environment, i.e., the informational contents of the Infosphere.[14]

Ethics in the infosphere

The *Tragedy of the Digital Commons* comes in two versions: one involves human agents, the other involves artificial agents. In both versions, TDC causes ethical problems that concern the behaviours of agents online and their effects on the digital environment. The second version seems less problematic. For, although its solution involves the engineering of "ethical" artificial agents – and this is not a trivial matter – the problem itself does not constitute a dilemma, that is, its solution is not coupled with the emergence of another problem. Unfortunately, this is the case with the first version, which raises an ethical dilemma about bridging the DD without, at the same time, creating the conditions for the irreversible exploitation and pollution of the Infosphere. Both problems can be approached from two ethical perspectives, one ecocentric, the other anthropocentric.

The ecocentric perspective concentrates its attention on the environment, its properties, preservation and (possibly) intrinsic values. The fundamental tenet is that an ethical theory should be based on the principle that an action is morally good or evil depending on how it affects the environment, its contents and inhabitants. So, an ecocentric ethics argues that it is possible to find an exact and objective definition of "good" and "evil" for the environment and then evaluate actions in terms of their environmental effects. Moral evaluation is the result of human activity but its specific nature does not depend on it.

The anthropocentric approach is based on the principle that "good" and "evil" are not just identified by human beings, but depend strictly on human judgement, interests and perspectives. The idea that there may be a mind-independent measure of evil or good is considered to be meaningless. Ethical facts are the result of human judgement about the nature and importance of the elements involved in specific situations. Values are not simply grasped more or less successfully through reflection and experience, but

[12] For a definition of "artificial agents", as used in Information Ethics, and an analysis of their (moral) actions, see Floridi and Sanders' forthcoming paper.

[13] Since bandwidth is not only a characteristic of the Infosphere (the way to access to the Infosphere itself), but also a resource that is owned, bought, sold, and leased, the issue of worms "chewing up" bandwidth in the Infosphere is not simply a tragedy of the commons: it is also a major issue of business and economic concern. We owe this remark to Charles Ess.

[14] Note that this can also be accomplished by an human agent. At this level of abstraction, the two actions would be fully assimilable because the possible differences in the nature of the two agents would be either unperceivable or irrelevant.

they are constituted by the ethical discourse, and the latter is a human product. So there would be no values if there were no humans.

Defined classically as being antinomic, both perspectives disagree more about the genesis and mind-independent nature of values than on their normative conclusions, on which they can easily converge. The anthropocentric idea that "the personal wants of the individuals … should guide the use of resources" (Kneese and Bower 1979: 4–5) is the basic principle of liberal politics. In Hardin's model, for example, the moral value of actions is grounded on their consequences for the environment. Following this consequentialist assumption, Hardin's solution is based on the fundamental principle of Authoritarianism, that is, on "mutual coercion mutually agreed upon" (Hardin 1968: 1247). On the other hand, the ideas that the ecosystem is "a diversity of life forms existing in a dynamic and complex but stable interdependency" (Marietta 1979: 197), and that "all living things are alike in having value in their own right, independent of their usefulness to human purposes" (Brennan and Lo 2002) ground the politics inspired by the "deep ecology movement". Both these general approaches can be interpreted as philosophies of "normative constraints" (Andrews 1982).

Choosing between the anthropocentric and the ecocentric approach in the Biosphere is difficult and problematic. This is so even in the case of the Infosphere. Consider the following example. Suppose there are two agents in the Infosphere. One is a worm. It produces copies of itself and spreads them through the network. The other is an antivirus program. It destroys the worm and its "products". Now, suppose that the worm has been created in order to check the security of the network and that the antivirus program has been installed in order to prevent an overflow on the same network. Two problems arise, both for the ecocentric and for the anthropocentric perspective. According to an ecocentric approach, the worm increases the digital space (and this may be considered "good"), but at the same time it exploits the environment by using too much bandwidth (and this may be considered "evil"). On the other hand, the antivirus prevents over-exploitation (good), but at the same time it prevents an increase in the space (evil). According to the anthropocentric approach, the worm is checking the network (good), but at the same time it decreases the bandwidth available for users (evil). On the other hand, by deleting every copy of the worm, the antivirus program avoids the over-use of users' connections (good), but at the same time prevents checking of the network security and the expansion of the digital space (evil).

The previous example is elementary but already suffices to illustrate how subtle and tricky choosing between the two approaches may be. An ethics grounded on one of the two approaches, without any further consideration for the specific features characterising the Infosphere, is not useful for interpreting and solving the new ethical problems caused by the ICTs, some of which we have explained in the present paper. What is needed is an ethics capable of taking into account the peculiarities of the new digital environment.

Conclusions and further work

We have shown how Hardin's TC can be applied to the Infosphere in order to uncover and model a new ethical dilemma, the Tragedy of the Digital Commons, and its main related problems. We have explained that, depending on how the digital divide is bridged, its solution may bring about a TDC as a side effect. We have also shown that Hardin's TC can be applied to the Infosphere, only if due attention is paid to the specific properties of the new environment, its differences with respect to the Biosphere, and the nature of the agents that inhabit it. As a consequence, we have criticised the previous use of TC by Adar and Huberman and explained why their model may not be applicable to the Infosphere. After that, we have provided a revised version of the TDC, which takes into account artificial agents. Finally, we have outlined some difficulties in finding a solution to the TDC through the use of two classical perspectives in ethics.

It is not the task of this paper to propose a solution for the TDC. However, by way of conclusion, we wish to sketch here a possible direction in which the research may be developed.[15]

ICTs have been changing human societies on a global scale and for some time now. They have modified the old social order and have been introducing new social factors. As Hardin said: "the laws of our society follow the pattern of ancient ethics, and therefore are poorly suited to governing a complex, crowded, changeable world" (Hardin 1968: 1245). This holds true for the information society as well, and may be extended to classic ethical approaches in computer ethics. Most of the efforts made to solve the new problems caused by ICTs seem to be informed only by legal needs and requirements. Ethical theorising seems largely subordinate to applications, codes and laws. However, a legal

[15] We are currently working on a possible solution of the TDC centered on the ontological properties of the Infosphere considered as common good.

approach, or even an approach too applications-oriented, does not seem to be of primary importance in the discussion of the issues analysed in this essay. There have been some attempts to elaborate an innovative ethical approach based on the idea that the new problems are evidence of a new moral context. It seems important that further research will be developed in this direction. To give an answer to the new problems on the Infosphere, such as the TDC, an ethical theory has to consider the peculiar features of the system under analysis. In ethical and social systems, the players are the environment, its contents and its inhabitants, the agents. But we know that when we say "agents" in the Infosphere we are not talking exclusively about human beings but also about artificial agents. And we also know that the digital environment is very different from the Biosphere. Classic ethical approaches do not seem well-equipped to deal with these novelties (Floridi 1999). The same holds true for standard legal systems. As Barlow said: "Digital technology is detaching information from the physical plane, where property law of all sorts has always found definition" (Barlow 1994: 1). The identification of a potential Tragedy of the Digital Commons allows us to understand, with clarity and precision, that ICTs cause new ethical and social problems, and that the search for their solution leads towards the development of a specific ethics for the Infosphere. Intellectual research in this direction will need to be provided in the future.

Acknowledgements

We are much in debt to Jeff Sanders for his many valuable suggestions, sharp comments and fundamental feedback on several versions of this paper. We like to consider him a co-author, although he bears no responsibilities for any remaining mistakes. We also wish to thank Carlo Dalla Pozza, Charles Ess and Guglielmo Forges Davanzati for useful comments; and Simona Capogna, Gianluca Paronitti and Matteo Turilli for helpful conversations. The anonymous referee of the journal provided a decisive feedback, without which the paper would be much worse. An early draft of this paper was given by G. M. Greco at the yearly national conference of the Società Italiana di Filosofia (Italian Philosophical Association), "Filosofia e Società della Conoscenza (Philosophy and Knowledge Society)", Foggia, 17–19 October 2002 and we are grateful for the feedback received from the participants. A previous draft has been delivered as research report for the Information Ethics research Group, Computing Laboratory, Oxford University. The final version was carefully copyedited by Paul Oldfield. G. M. Greco's work was supported by a grant from the Università degli Studi della Calabria (Italy), area Scienze Storiche, Filosofiche, Pedagogiche e Psicologiche, prot. N. 651/082 – 13/11/2002. We remain responsible for all scientific content and any possible mistake. For legal requirements Gian Maria Greco must be considered the only author of this paper.

References

E. Adar and B.A. Huberman. Free Riding on Gnutella. *First Monday*, 5(10), 2000. http://firstmonday.org/issues/issue5_10/adar/index.html.

R.N.L. Andrews. Cost-Benefit Analysis and Regulatory Reform. In D. Swartzman, R. Liroff and K. Croke, editors. *Cost–Benefit Analysis and Environmental Regulations: Politics, Ethics, and Methods*, pp. 107–135. The Conservation Foundation, Washington, 1982.

J.P. Barlow. The Economy of Ideas. *Wired Magazine*, 2(3), 1994. http://www.wired.com/wired/archive/2.03/economy.ideas.html.

A. Brennan and Y. Lo. Environmental Ethics. *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/ethics-environmental, 2002.

eEurope. An Information Society For All. Action Plan prepared by the Council and the European Commission for the Feira European Council 19–20 June 2000, http://europa.eu.int/information_society/eeurope/action_plan/pdf/actionplan_en.pdf, 2000.

eEurope 2002. Accessibility of Public Web Sites and their Content, September 25, http://europa.eu.int/information_society/eeurope/news_library/pdf_files/communication_accessibility_en.doc, 2001.

J.A. Fletcher and M. Zwick. N-Player Prisoner's Dilemma in Multiple Groups: A Model of Multilevel Selection, In E. Boudreau and C. Maley, editors. *Proceedings of the Artificial Life VII Workshops*. Portland, 2000.

L. Floridi. Information Ethics: On the Theoretical Foundations of Computer Ethics. *Ethics and Information Technology*, 1(1): 37–56, 1999.

L. Floridi. Information Ethics: An Environmental Approach to the Digital Divide. *Philosophy in the Contemporary World*, 9(1): 39–45, 2002.

L. Floridi. On the Intrinsic Value of Information Objects and the Infosphere. *Ethics and Information Technology*, 4(4): 287–304, 2003.

L. Floridi and J.W. Sanders. On the Morality of Artificial Agents. *Minds and Machines*, forthcoming.

G.M. Greco. *La Tragedia dei Beni Digitali Collettivi*. Talk given at the yearly national conference of the Società Italiana di Filosofia (Italian Philosophical Association), "Filosofia e Società della Conoscenza (Philosophy and Knowledge Society)", Foggia, 17–19 October 2002. Forthcoming in the Proceedings.

G.M. Greco and L. Floridi. *The Tragedy of the Digital Commons*. IEG Research Report, IEG-RR1, Information Ethics research Group, Computing Laboratory, Oxford University, 14 October 2003.

H. Grotius. *On the Law of War and Peace*, John W. Parker, Cambridge, 1853 [1625].

G. Hardin. The Tragedy of the Commons. *Science*, 162(3859): 1243–1248, 1968.

G. Hardin. Extensions of "The Tragedy of the Commons". *Science*, 280(5364): 682–683, 1998.

B.A. Huberman and R.M. Lukose. Social Dilemmas and Internet Congestion. *Science*, 277(5325): 535–537, 1997.

A. Kneese and B. Bower. *Environmental Quality and Residuals Management*. Johns Hopkins University Press, Baltimore, 1979.

D. Marietta, Jr. The Interrelationship of Ecological Science and Environmental Ethics. *Environmental Ethics*, 1(3): 195–207, 1979.

Nielsen//NetRatings Global Internet Trends Q4 2002, 20 February 2003.

E. Ostrom. Collective Action and the Tragedy of the Commons. In G. Hardin and J. Baden, editors. *Managing the Commons*. Freeman, New York, 1977.

E. Ostrom. *How Inexorable is the "Tragedy of the Commons"? Institutional Arrangements for Changing the Structure of Social Dilemmas*. Workshop in Political Theory and Policy Analysis, Distinguished Faculty Research Lecture, Indiana University, Bloomington, 3 April 1986.

T. Pfeiffer, S. Schuster and S. Bonhoeffer. Cooperation and Competition in the Evolution of ATP-Producing Pathways. *Science*, 292(5516): 504–507, 2001.

M. Robert and G. Sorci. The evolution of obligate interspecific brood parasitism in birds. *Behavioral Ecology*, 12(2): 128–133, 2001.

Special issue on "Remapping the Digital Divide". *The Information Society*, 19(4), 2003.

Stopping spam. *The Economist*, 367(8321): 58, 2003.

R.M. Turner. *The tragedy of the commons and distributed AI systems*. Presented at the XII International Workshop on Distributed Artificial Intelligence, University of New Hampshire, Durham, 1991.

United Nations Development Programme. *Human Development Report 2001. Making new technologies work for human development*. Oxford University Press, New York-Oxford, 2001.