

The banality of simulated evil: designing ethical gameplay

Miguel Sicart

Published online: 14 July 2009
© Springer Science+Business Media B.V. 2009

Abstract This paper offers an analytical description of the ethics of game design and its influence in the ethical challenges computer games present. The paper proposes a set of game design suggestions based on the Information Ethics concept of Levels of Abstraction which can be applied to formalise ethical challenges into gameplay mechanics; thus allowing game designers to incorporate ethics as part of the experience of their games. The goal of this paper is twofold: to address some of the reasons why computer games present ethical challenges, and to exploit the informational nature of games to suggest how to develop games with ethics at the core of their gameplay.

Keywords Information ethics · Computer game design · Level of abstraction · Game design · Methodologies · Ethics · Simulation · Gaming

Introduction

Computer games are steadily becoming the paradigm of 21st century entertainment. Their economic success fuels not only a creative industry, but also the technological

drive to push the boundaries of some fields of computation, especially those related to simulation and graphics.

At the same time, computer games are often maligned as causes of antisocial behaviours. Computer games have a reputation as unethical forms of entertainment, which corrupt the solid values of society by promoting both the practice of violence and the wrong personal and social virtues (Smith et al. 2003; Funk et al. 2004; Carnagey et al. 2007; Brey 1999; Coeckelberg 2007; Wonderly 2008). This reputation is due to the fact that contemporary computer games, which are no longer targeted exclusively to children, have found, in the simulation of violence, a perfect vehicle for the expression of conflict inherent to all games (Caillois 2001; Salen and Zimmerman 2004; Juul 2005). The tension between the perception of computer games as children's' entertainment and the demands of a mature audience for mature content is a key element of the alleged unethical nature of computer games.

In any discussion about moral concerns it is important to question its origins and assumptions. This paper introduces an Information Ethics (Floridi 1999, 2002, 2003, 2008a, b; Floridi and Sanders 2001, 2002) analysis of the ethical challenges that computer games pose. My goal is to argue for an interpretation of the ethics of games as designed software systems that are experienced by moral agents, and to provide a set of design inspirations oriented to the development of ethical gameplay within computer games. These two goals will be illustrated by the games *Defcon* (Introversion Software 2006), *Grand Theft Auto IV* (RockStar North 2008), and *Shadow of the Colossus* (Team Ico 2006).

The first section of this paper discusses Hannah Arendt's concept of the banality of evil as presented in *Eichmann in Jerusalem*. This is a key concept for understanding the ethical challenges that computer games pose. These ethical

M. Sicart (✉)
Center for Computer Game Research, IT University
of Copenhagen, Copenhagen, Denmark
e-mail: miguel@itu.dk
URL: www.miguelsicart.net; http://www.itu.dk;
http://web.comlab.ox.ac.uk/oucl/research/areas/ieg/
http://philosophyofinformation.net/centre/gpi/index.html

M. Sicart
IEG, Oxford University, Oxford, UK

M. Sicart
GPI, University of Hertfordshire, Hertfordshire, UK

challenges are then described in the second section of this paper, illustrated with a close analysis of *Defcon*. The third section presents computer games as informational systems. The Information Ethics (I.E henceforth) approach employed in this section makes the ontology of computer games as informational systems the starting point for research on their ethical properties and effects. The concept of Levels of Abstraction (Floridi 2008a, b) is presented here in its classic I.E meaning as part of the ethics of computer game design.

The fourth section presents some I.E concepts as tools for designing computer games. The paper concludes with a section suggesting possible game design approaches that could be applied to the design of computer games with ethical gameplay.

In the conclusions, some of the possible benefits of this approach to the understanding of the ethics of game design are outlined.

The banality of evil

In 1963, Hannah Arendt was commissioned by *The New Yorker* to write a report on the trial of Otto Adolf Eichmann, one of the main coordinators of the Holocaust. Arendt's work, *Eichmann in Jerusalem. A Report on the Banality of Evil*, presents Eichmann not as a monster or a psychopath, but as an efficient bureaucrat that followed the orders that the Party issued without questioning them, in a perverse interpretation of Kant's categorical imperative (Arendt 2006, pp. 135–137). While Arendt's portrait has been contested, her work has provided insight in to the inner workings of totalitarian states; specifically how their citizens survive and implicitly support atrocities and a constant lack of liberties. The most relevant concept developed by Arendt in this book is the banality of evil.

Eichmann is portrayed as a simple cog within a vast machinery—a tiny unit with great power and influence on how the state performed the so-called “Final Solution”. That position is at the centre of the concept of the banality of evil: in totalitarian states, bureaucrats are alienated within the whole machinery of the state, performing activities that lead to atrocities by blindly following the orders of the political apparatus that does not allow, nor provide, any feedback: “He did his *duty*, as he told the police and the court over and over again; he not only obeyed *orders*, he also obeyed *the law*” (Arendt 2006, p. 135).

The banality of evil describes the application of industrial procedures to mass extermination. Eichmann was not “evil” in the conventional sense of the term: he was determined to fulfil his orders, to complete his assignments as a moral duty, regardless of their ethical nature. The concept of banality of evil has inspired analysis of the role of the post-industrial state apparatus in war, genocide, and

even ecological disasters. All these acts are determined by the necessity to keep a system running without discussing the moral nature of the orders received, rather than guided by evil intention. In this essay I will suggest an alternative, expanded version of the concept of the banality of evil, using the theoretical framework of Information Ethics. This informational explanation of the banality of evil will be used to illustrate the ethical challenges that computer games pose, and how those challenges can be actually used as design inspiration to create interesting ethical ludic experiences.

As I have mentioned before, Eichmann is portrayed as an efficient bureaucrat who did not witness the consequences of his zeal for making the state machinery run efficiently (Arendt 2006, p. 89). Actually, this is a key element for understanding how the systems that encourage this banality of evil operate: by selecting the type and amount of feedback that a specific agent receives, these systems detach that agent from the informational environment of the actions. In other words, the agents in these systems do not perceive their actions as anything other than what is needed to keep the system running efficiently, and the system only provides feedback about how well it is functioning.

The banality of evil is a consequence of systems designed to obscure the causality of decisions. The purpose of this design is the limitation of (ethical) agency within that system. The banality of evil can be defined as a designed limitation of ethical agency in complex multi-agent, hierarchical systems. It is “designed”, because the system is created with the intention of limiting ethical agency; and “multi-agent” and “hierarchical” because these systems tend to operate with a number of agents that have operative power over both the system and the other agents in it. Agents within these kinds of systems can engage in unethical actions without receiving any feedback on the morality, or consequences of those actions. Unethical behaviour, due to the feedback structure of the system, is perceived as “necessary actions”, hence no ethical reflection is required. The question is, are computer games systems of this kind?

Computer games as ethical challenges

Defcon can be a game of patience. Heavily inspired by the 1980s classic film *Wargames*, this game puts the player in command of a nuclear force with the goal of wiping out as many enemy units and civilians as possible, whilst minimising her loses. *Everybody dies*, reads the slogan for the game, and it is the truest summary of the gameplay: to win, players have to lose the least. *Defcon* is a political simulator of atomic warfare that makes clear what was implicit:

in nuclear war, the winning condition is still a losing condition.

Defcon's user interface presents the player with a relatively stylized map of Earth, one that clearly resembles that which popular media has always used for portraying doomsday bunkers where atomic war decisions are taken. The minimalist sound design uses ambient noises to give the impression that players are actually buried under tons of rocks taking decisions that will wreak havoc on the citizens of the Earth.

Defcon is a multiplayer game played over the Internet. In its default gameplay mode, the rules are very simple: players are given a fixed and limited number of units and resources. These have to be distributed on the map preparing for the nuclear showdown that inevitably will take place at the end of the round, when "Defcon One" is reached and nuclear warheads are launched. Players are rewarded with points according to the number of enemy units and civilians they have eliminated, while they lose points if their own cities are hit. The winner of the round is usually the player who loses the least whilst inflicting reasonable damage. Obviously, targeting large cities is one strategy which can plausibly yield good results.

Defcon is an example of the ethical challenges games pose. To an external observer, the obvious conclusion is that *Defcon* is a game about indiscriminate mass-murder with nuclear weapons; an end-of-the-world simulation where players compete to exterminate "civilians" and eliminate countries and cultures indiscriminately. These observations are, of course, correct to a certain extent, indicating the main arguments for the perception of videogames as unethical entertainment.

In the classic western research on games (Caillois 2001; Juul 2005), there is a certain understanding of games as something "separate", as an "unproductive" activity that takes place within boundaries that are set by rules agreed upon by players. Games are, in this outdated perspective, arbitrary systems of rules that establish a set of constraints that players have to accept in order to achieve goals, which will also determine the winning conditions of the game. Games are also entertainment, pastimes, vehicles for formal leisure with clear rules and unambiguous outcomes. Furthermore, games are also tools for education, like *Monopoly* once was and the *Lego* building blocks still are.

Games seem to share, in the western mind, a paradoxical double condition: on the one hand, they are unproductive entertainment; on the other hand, their rhetorical nature (Bogost 2007) makes them interesting tools for education, since they actually "force" their users to follow, and take as valid, a certain set of rules that will inevitably yield a specific set of outcomes. And it is in this dialectic that the roots for the understanding of games as ethically dangerous experiences are to be found.

Games are activities in which agents engage with a system designed to encourage the achievement of certain goals with predetermined means. But to play a game is to give supreme, albeit temporary, importance to these constraints. From a formalistic perspective, while playing, there is nothing more important than the rules we live by as players; rules that are embedded in a virtual world and in a system designed to enhance the experience of limited agency. In many cases, both the rules and the virtual world can be perceived by external observers as harmless: there is no moral risk in *New Super Mario Bros* because the "violence" is cartoonish, like there are no explicit challenges to our ethics in playing a game of *Buzz! The Big Quiz* (Relentless 2006), *Dance Dance Revolution* (Konami 1998) or similar social games. But games like *Defcon*, or more popularly *Doom* (id Software 1993), *Counter-Strike* (Valve Software 2000) or any other First Person Shooter seem to create immediate ethical concern, traceable to the concept of the banality of evil.

Some recent computer games like *Fable* (Lionhead Studios 2004) or *Knights of the Old Republic* (BioWare 2003) have even tried to implement ethics as a part of their gameplay experience. Some of the choices given to players are measured by the game system, which outputs an ethical evaluation of the player based on pre-determined moral parameters. In these games, evil acts of violence grant "evil" points, while more positive actions guarantee a higher "good" score. As I will argue later in this paper, even these games, which employ ethics as a part of their gameplay, are subject to a banality of evil critique: by alienating the player from reflecting about the ethics of their actions, and outsourcing moral evaluation to a closed, pre-designed system, the game effectively limits players' ethical agency. This limitation, by means of system design, recalls the concept of the banality of evil.

Killing millions (of units) is a matter of a mouse click in *Defcon*. As a matter of fact, killing millions is the only way of winning the game, and strategies for winning the game are optimizations of simulated massive nuclear destruction; an activity that in classic game theory terms presents no benefit beyond self-enjoyment. Killing in *Counter-Strike* is rewarded by money and weapons for the next round, or punished with an extended waiting time in the purgatory of spectator mode, waiting to resurrect so players are afforded another chance at killing each other. Yet they are only playing a game, obeying rules. Much like Eichmann did.

The banality of evil explains how humans can commit unspeakable acts of cruelty without remorse. In those situations where this concept is applicable, there is a system in place that detaches agents' actions from the perception of their consequences. Taken literally, many computer games could be understood as practicing the banality of simulated evil. This practice could potentially have an

unethical outcome: the desensitisation of players to violence, theft, and other unethical behaviours. The banality of simulated evil does not lead directly to criminality, but to the crisis of the reflective capacities of the agent when evaluating the ethics of those actions simulated by the game.

If games are considered unethical, it is partially because they can be identified with these totalitarian political and social structures that made policies just a detached game of matching numbers. In this line of thought, a game has to be unethical because the agent is presented with rewards for actions of (simulated) evil, while feedback on the consequences is not provided (nevertheless, it is necessary to say that the consequences of actions in games do have a system for feedback, embedded in the rule system, and usually tied to the winning conditions). The act of “killing” an opponent in a computer game becomes, from this perspective, an action that has no other consequence than that quantized by the game rules, and therefore the player is alienated from thinking what the actual consequences of her actions are. This is the potential unethical outcome of the banality of simulated evil.

Of course, in computer games players never actually “kill” anybody; players do interact with a system of rules that encourages a number of behaviours by rewarding them with tokens (such as points) meaningful only within the system. But computer games, understood as systemic environments, are often played in a virtual environment, a gameworld in which there is often simulated violence or unethical actions.

The argument would go as follows: if players are faced with simulated evil and rewarded when they engage in unethical actions in that gameworld, then there are risks that players will believe that these unethical actions have no consequences. This would be the first step to desensitisation, the loss of empathy that some researchers (Funk et al. 2003; Funk et al. 2004) and some media believe to be the most likely outcome of playing violent computer games. In other words, because the interaction with a system of rules that constitutes the core of any computer game experience is often covered with a semantic layer that we recognize as media violence, playing computer games is the experience of the banality of evil.

This is a very limited understanding of computer games, not to mention a very poor consideration of the ethical capacities of players as moral agents. In the following section I will present a description of videogames as informational systems that will lead us to reconsider this interpretation, and transform the alleged ethical shortcomings of computer games into interesting design tools for creating ethical experiences. I will also argue that (computer) games are not the detached, encapsulated systems of meaning that classic game research has argued for,

and that they actually have a strong presence in the configuration of our ethical and cultural being.

Computer games as informational systems

When describing a game to someone who has never played it, the first thing is to describe the rules and the game mechanics. We need to know, as players, what is possible and impossible, what is allowed and what is not allowed, in order to start playing. Once we have a grasp of the rules, we start playing: the initial states of the game are, in general, approximations to the rules and the environment where we play, trying to find the strategies to win while staying true to the rules or, in the case of some player types (Bateman and Boon 2006), the initial states of play are focused on finding ways to break the rules to see how far the game can be stretched as an experience without falling out of the rule system.

Mastering a game means understanding the rules and mechanics of the game, how they interact with each other and form behavioural patterns by which we play. In other words: mastering which information is relevant within the game experience, and how to manipulate it within the given boundaries. Even in games such as contact sports where mastery implies physical prowess and the amount of information is rather low, players have to understand the interrelation between physical skills and game information, and how to act upon it. Good tennis players seldom run, when in control of the game.

A game is then an *informational system*: a construction of rules that determine which actions are meaningful within a certain experience, and how those actions can be performed. By informational system, I am here advocating an ontology based on Information Ethics and its object oriented approach to understanding the morality of beings and information.

For Information Ethics, “the moral action itself can now be modeled as an information process, i.e. a series of messages (M), invoked by *a*, that brings about a transformation of states *directly* (...) affecting *p*, which may variously respond to M with changes and/or other messages, depending on how M is interpreted by *p*’s methods” (Floridi. 2002, p. 289). In this article, I will adapt this approach to encompass all kinds of *ludic* actions: a ludic action being any action taken by an agent within a game system that is evaluated by a game rule. A *ludic action* will be *any interaction within the gameworld via a game mechanic that produces an output from the game*. For example, placing the units in *Defcon*, firing a missile, or using the chat interface to communicate with other players.

In order to understand the ethics of computer games from an informational perspective, it is relevant to define

computer games within the terminology of Information Ethics. This will allow the use of the concept of Levels of Abstraction as both a descriptive and a normative pattern, which in turn can be translated to design methods.

A computer game is an *infosphere*, a “context constituted by the whole system of information objects, including all agents and patients, messages, their attributes and mutual relations” (Floridi 2002, p. 289). In more classic computer game terms, a game infosphere is constituted by all game elements: players and AI agents, environments and gameworld, rules and game mechanics, and the interaction modes in the space of possibility (Salen and Zimmerman 2004, pp. 66–67).

Within this framework, it is important to focus on the notion of agent more than on the concept of player. In computer games there are agents, some human, some controlled by an artificial intelligence, but they all operate in the same way: interacting with the game environment/state via game mechanics constrained by game rules. The act of playing a game is an act of agency within an infosphere, understanding agency as the interaction by means of exchange of information with a system and other agents. This exchange of information is conducted via the afforded methods (game mechanics) of the different game objects, as constrained by the game rules.

In *Defcon* players construct a strategy by carefully placing units in the places they are afforded by the rules, relative to their initial state in the game, and by doing so they are manipulating the gameworld and their relations to other players in ways sanctioned and determined by the game. Agency in *Defcon* is limited by the game rules and the mechanics afforded to players, and thus the relation with the game environment will take place within the boundaries of those mechanics. In *Defcon*, as in all games, the whole infosphere, and the agency level of players, is determined by the design of the game.

From an ethical perspective, it is crucial to understand that at some stage the infosphere was designed to afford agency and allow for the flow of information in particular ways. A game is designed with the intention of creating a ludic experience that allows players to interact with the environment, and with other agents, in interesting, yet pre-calculated ways. Any system that is designed to modify or enhance agency in particular, limited ways, any system that is *scripted* (Latour 2002; Akrich and Latour 1992), has to raise ethical awareness. Modifying or directing agency can have strong implications on the moral status of the agents: what we have to do is ethically relevant, and in games what we have to do is actually designed with a set of intentions in mind.

For instance, *Defcon* is designed to encourage conflict among players, and all the mechanics and rules present in the infosphere as a space of possibility for the player are

geared towards that conflict. Players of *Defcon* cannot find negotiated outcomes—furthermore, they are not interested in negotiated outcomes: the goal of the game is to simulate nuclear war and to reward those that annihilate more units of the opponent, and agents in the game can only do that. The design of the game, as with the design of any infosphere, is loaded with ethical values.

Agents in games, however, are not determined in their ethical configuration by the ethical values of the design. Any agent in an infosphere, and particularly any human agent, has to be considered a moral agent, capable of ethically relating to the whole system, reflecting on her own values and ethics, and capable of acting upon them and thus modifying the actual ethics of the infosphere as experienced. This means that while an infosphere can have ethical values imprinted in its design, it is the actual interaction of a moral agent with those ethical affordances and constraints that constitutes the ethics of a given infosphere. In other words: *we can only determine the informational ethics of such a system when it is experienced by agents, and not by its design*. Thus, we need to understand the active role of players as ethical agents in the configuration of the ethical experience of gameplay.

Agents exert their ethical capacities by what Floridi has defined as the creative stewardship of the *homo poieticus*, who “concentrates not merely on the final result, but on the dynamic, on-going process through which the result is achieved” (Floridi and Sanders 2005). Agents within an infosphere are not only in charge of exchanging information, but also of preserving the nature of the system, nurturing, protecting, and producing meaningful interactions in and with it. In games, this is translated as the ethical responsibility for playing without cheating, not allowing other players to harass the community, or developing interesting strategies for beating the game. Players as ethical agents are not mere providers of input: they understand the nature of their actions within the semantics of the infosphere, and they act upon that moral understanding.

This moral understanding in games is determined by two elements: the player as agent, and the cultural being that experiences play (Gadamer 2004). As players, we construct our agent values with those ethical affordances and the constraints provided by the system. In *Defcon*, players build their moral agency based on how the system encourages betrayal and competition, and how it rewards the strategies that annihilate the opponents’ cities.

But players are not only input providers within the game system—the ethical configuration of players is also dependent on the ethics of the agent that becomes a player. In other words, the ethical values and agency outside of the gameworld are also crucial. Agents within a gameworld are not configured solely by the input/output structure of the game system—a crucial part of the infosphere of a game is

the world to which it relates by means of simulation and representation.

Defcon is not only a state machine that simulates the conditions of conflict: it simulates nuclear war and its outcomes, representing them by means of an aural output system; both the aural system and the simulation of nuclear war, the semantic layer of the game, are interpreted by the human agents using their cultural and ethical knowledge. The game infosphere comprises both the mechanical agents that interact with predetermined methods with the state machine, and the moral agents that evaluate the mechanical actions of input and output, and the cultural and ethical relevance of in-game agency.

Within this perspective, it is necessary to have tools that determine when and how we can define the action of an agent as moral, and what the extent and effect of that ethical agency is. These tools are the concepts of *Level of Abstraction* (LoA henceforth) and *Gradient of Abstraction* (GoA henceforth), as defined for Information Ethics¹ (Floridi 2008a, b). For computer games, there are two dominant GoAs: the first one is limited to the direct interaction between agents and the state machine by means of game mechanics. This gradient is concerned with all the input/output operations performed by, and for, the modification of the game state within the limitations of the rule system. For example, the action of selecting a unit in *Defcon* invokes a LoA in which the user interface, the class object of the unit, and the response from the state machine via the user interface are relevant. The set of LoAs that apply to understanding the input/output processes between agents and the state machine can be defined as the syntactic or procedural GoA: it is concerned with the inner mechanisms of the game, its procedural nature (Murray 1997; Bogost 2007) as a system.

The second GoA that is applicable to computer games comprises the game system as simulation and agents as ethical agents. This expands the previous GoA and adds a semantic layer. The game system is, in this layer, more than a simple state machine: it comprises all the aural/aesthetic levels, giving cultural meaning to the procedural elements of the game state machine. In *Defcon*, the state machine simulates nuclear war—all its behaviours plus

semantic levels, its procedural and aesthetic content, are meant to be understood as an interpretation of war commands in a nuclear bunker. The game as simulation is then the GoA that comprises its procedural systemic nature with the cultural/aesthetic layers that wrap it.

On the other hand, the agent within this GoA is capable, and has the duty, to reflect morally not only about the appropriateness of her actions providing input, but also about the meaning of those actions both within the perspective of the cultural meaning of the simulation, and outside of it, with her own cultural values. This statement seems to discard any agent that is not human, but that is not the case: there are some LoAs in that GoA that can apply to agents that are not human: every LoA that does not require attention to those values external to the simulation can apply to artificial agents (Floridi and Sanders 2001).

In *Defcon*, the ethical agent understands that the game is a simulation of nuclear war, and thus the game is not only loaded with ethical values, but also interpreted as such by an ethical agent. An ethical agent is then the actor that interacts with a simulation in the context of a ludic infosphere.

The semantic GoA comprises the procedural GoA, and all the different LoAs present in a computer game can then be visualized within those relations Fig. 1. This conceptualization can be modeled as follows:

At the heart of the ethical concerns with computer games lies the incapacity to understand how players navigate this

¹ A Level of Abstraction is defined as “a finite but non-empty set of observables. No order is assigned to the observables, which are expected to be the building blocks in a theory characterised by their very definition. A LoA is called discrete (respectively analogue) if and only if all its observables are discrete (respectively analogue); otherwise it is called hybrid” (Floridi 2008a, b, p. 309). A Gradient of Abstraction is defined as a “a formalism defined to facilitate discussion of discrete systems over a range of LoAs. Whilst a LoA formalises the scope or granularity of a single model, a GoA provides a way of varying the LoA in order to make observations at differing levels of abstraction” (Floridi 2008a, b, p. 311).

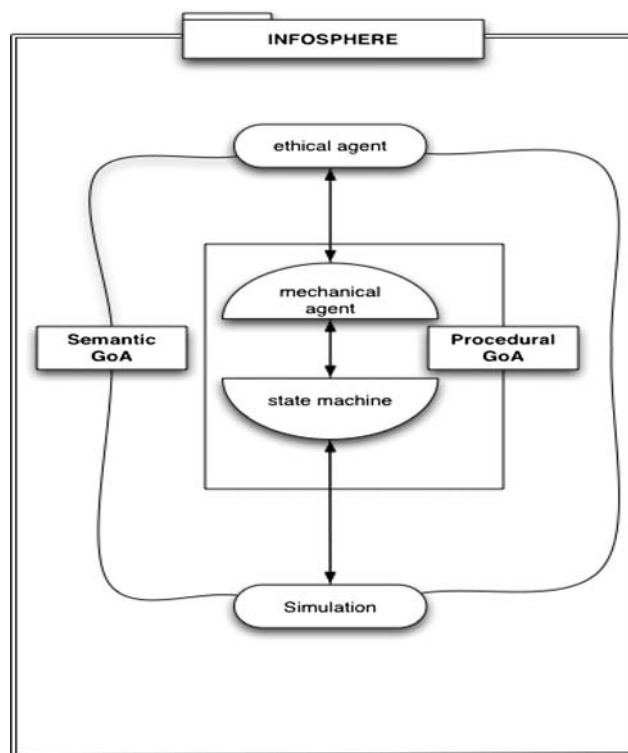


Fig. 1 An information ethics model for computer games

infosphere. Common accusations against computer games understand them only within the procedural GoA, and thus with an agent that is not concerned with anything other than providing the right input to modify the state machine (Funk et al. 2004; Wonderly 2008). Furthermore, there seems to be a common issue with the mapping of the simulation within the procedural GoA, which is incorrect: the mechanical agent is not interacting with the simulation; in other words, when agents interact with the simulation, they are also using their ethical capacities, since those are a part of their cultural resources needed to understand how and why to play a game.

Let us analyze two very different computer games from this informational perspective. *Grand Theft Auto IV* is the latest instalment in a series of highly popular computer games that deal with epic stories of urban North American crime in different historical periods. It is a game for adults, riddled with violent situations and criminal characters. However, *GTA IV* has arguably evolved into an ambitious tale about destiny and redemption—a tale structured around elements of ethical gameplay design. In *GTA IV*, players control the life of Niko Bellic, a criminal of Serbian descent who arrives in the United States of America lured by his cousins' stories of life, liberty, and the pursuit of happiness. But the stories are all false, and Niko is progressively trapped in the underworld of Liberty City, with no exit in sight.

From a purely mechanical perspective, the game consists of understanding a limited set of actions in order to succeed in different challenges and progress in the game narrative and space. These actions are “semantized” as violent behaviour—in the semantic level, players kill, hijack cars, and commit other crimes. In all previous *GTA* games, the ethical player was more or less left alone with her judgment to evaluate her ethical relation to these actions. Some players found previous instalments distasteful, while many enjoyed its escapist, heavily-clichéd take on urban criminals and their contemporary mock-epics. *GTA IV* took a turn towards tragedy—in the cutscenes, Niko Bellic is presented as a man who wants to stay away from violence, who has travelled to America to run away from who he was, even though he brought too a secret seed of vengeance. Yet, when players get to control him, the only possible goal is to further that spiral of violence from which Niko wanted to escape. The semantic levels of *GTA IV* are designed around this stark tension, forcing players to reflect upon the meaning of their actions, and the inevitability of the violence. *GTA IV* is also targeted to an ethical player; one who understands the design oscillating between the active player and the spectator, between empathising with Niko and playing the game, between the mechanical and the semantic levels of abstraction.

Shadow of the Colossus is designed with an entirely different approach. The story is less direct: players are presented with a young man carrying a dead girl to a far-away shrine. There, a godly voice commands the man to slay 16 colossi that peacefully roam the world in exchange of resurrecting the girl. Players are then in control of this man, and will have to fight these 16 majestic creatures in order to complete the game. *Shadow of the Colossus* cues certain elements of tragedy by its aesthetic design: every time a colossus is killed, sad music is played, and an animation is triggered showing how killing colossi almost kills the player avatar.

However, *Shadow of the Colossus* communicates ethical gameplay in a deeper way: in order to kill these colossi, players have to climb them. The climbing mechanic is based on a stamina meter, which is depleted as the avatar exerts more physical effort while holding the colossi. This stamina meter increases the more colossi the player kills. Hence, the mechanical player will understand that as a beneficial action, a desirable action, since it gets her closer to the successful ending of the game. However, the more colossi are slain, the sicker the avatar looks. This is an explicit dissonance between the mechanical level of abstraction and the semantic level of abstraction, one arguably designed to increase the ethical experience of *Shadow of the Colossus*. In this evocative, poetic title, ethics is not at the core, but on the outside, “enveloping the tale which brought it out only as a glow brings out a haze, in the likeness of one of these misty halos that sometimes are made visible by the spectral illumination of moonshine” (Conrad 2004, p. 10).

Ethical gameplay is, then, the outcome of designing the relations between the mechanical and semantic levels of abstraction in a game, considering that the ludic experience will be evaluated by an ethical epistemic agent.

The most common ethical concern raised by games has to do with a misinterpretation of the concept of the banality of evil applied to the informational being of games and the presence of players as agents within that infosphere. If the banality of evil is a concept that explains why, inserted in a system that obscured the outcome of their decisions, bureaucrats could take unethical decisions without understanding them at all as ethical choices, then the concept could apply to games: players would allegedly be providing input to the system without thinking about the consequences of their actions in the simulation, thus becoming desensitized from the connection between cause and consequence. In this context, *desensitisation could be defined as the crisis of the ethical tools that agents have to evaluate their conduct.*

However, given an Information Ethics understanding of the process of playing computer games, this concern can be counter-argued. Any agent in a game operates within two

distinct GoAs—one that cares about the procedural elements of the game, from user interface to score system, goals, rules and mechanics, and another one that encapsulates that procedural part of the game into a larger GoA in which the semantics of the game, its meaning, is extremely relevant for agents, as it guides its interactions with the system.

This process can be understood with an analogy to piano players: sometimes, players need to “warm” their fingers before playing, and they do so by playing on the piano a number of exercises. They are not concerned with the semantic elements of the music piece, only with the mechanical interaction with the piano. Once they start playing music, though, the semantic layer is also present, so they have to pay attention to tone, emotion, phrasing, and the aesthetic qualities of playing the piano. Both are reasonably different Gradients of Abstraction. Similarly, a computer game player performs a series of actions to interact with the game in order to achieve goals, but the actual gameplay experience requires players to understand the semantics of the game, the meaning of the simulation. Hence, to play is to interpret those mechanical actions in the light of both the game simulation, and the player as an epistemic moral agent.

This informational perspective on the ontology of games provides on one hand an interesting set of tools for the analysis of the ethics of computer games, and on the other hand explains why some computer games are not relevant for ethical theory: those games in which the procedural dominates over the semantic are not of interest for ethical agents. For instance, *Tetris* (Pajitnov 1985) is a game in which the importance of the semantics is really not crucial for the experience of the game: players don’t need to understand the simulation fully in order to interact with the system successfully. An understanding of the rules and mechanics is sufficient. Abstract games (Juul 2005, pp. 130–132), then, are those games that privilege the importance of the procedural over the semantic GoA.

This brings forth an interesting outcome for game design: it is possible, as designers, to manipulate the ways agents exist and act in the different GoAs—furthermore, it is possible to manipulate the importance of the different gradients in order to produce intended outcomes. *Tetris* does not require a strong semantic layer that conditions the agent, but adventure games like *Fahrenheit* (Quantic Dream 2005) often require a certain understanding of the semantic layer in order to proceed in the story or, in some cases, to make a choice that can affect the outcome of the game.

Within this perspective, then, games that include ethics as a game mechanic embedded in the game system, like the aforementioned *Fahrenheit*, *Fable* or *Knights of the Old Republic*, are fundamentally flawed: as I have explained

earlier, these games place in the procedural gradient what should be relevant for the semantic gradient. These games taunt players with ethical decision-making; understood as choosing between two or three options of varied ethical alignments, from good to neutral to evil. These types of ethical game designs are fundamentally flawed because their alleged ethical simulation is placed dominantly in the procedural gradient: “evil” is not understood as a dominant semantic condition but a procedural one—it is a state in the machine. Thus the ethical agents are not required to use their ethical values as agents within the semantic layer in order to take a choice: it is enough to understand the arbitrary ethics assigned to a particular game state, and let the game system evaluate your behaviour. It is a process of desensitising the agent to their ethical thinking about the simulation, and focusing it on the procedural layer. It is, then, a process similar to those described by the banality of evil concept: agents are deprived of their ethical capacities in favour of a procedural external system that will evaluate their choices.

The procedural gradient comprises the design and implementation of the game as a state machine, with the basic mechanics and rules that determine the interaction of input agents with the system. The ethics of these agents is limited to the well-functioning of the game system, the informational balance of the system interaction.

The semantic gradient comprises the layers of meaning that we understand as the gameworld—the reasons why players are emotionally attached to the game, understand how to play it, and take choices informed by the consistency and informational feedback of the gameworld. Agents in this gradient are concerned with the community of agents, the cultural and ethical values of the game, and the connections of the game with the larger cultural infosphere outside of the game.

Is it possible, then, to create computer games that have ethics at the core of their gameplay? Yes, as long as the design takes into consideration the different gradients that configure computer games as infospheres, and how players ethically relate to these gradients. It is possible to translate this theoretical framework into a set of tools and practices that can contribute to the development of ethics as a gameplay mechanic.

Designing ethical games

Ethics and moral choice are progressively becoming of interest to game developers (Hocking 2007). While the capacities for creating ravishingly beautiful worlds with extreme graphic detail have increased with every computer hardware evolution, the experiences presented to players do not always match the ambitions of game designers.

Moral choice has been seen as one way of engaging players more deeply in deeply complex gameworlds. Unfortunately, ethics has been misunderstood and misused, and as such ethical gameplay has been confused with a variation of branching storytelling where the player takes choices based on alleged moral parameters evaluated by the game system.

As I have previously argued, those games in which agents just need to understand the procedural rules that determine the game state, without thinking about the actual moral implications of their actions, are deeply flawed in their ethical design. Creating games with ethics as a gameplay choice requires a more nuanced approach to the role of game designers via game mechanics, and agents within the systems.

Ethical decision making is the property of ethical agents, which in the case of computer games, means that this type of gameplay has to take place in the GoA correspondent to the simulation.² The simulation here has to be understood as both the aural/aesthetic layer and the state machine underlying it. The ethical agent, on the other hand, is not only the reflective agent that can perceive the simulation in context of the game, her own culture and ethical standing, but also the mechanic agent. In fact, designing ethical games requires focusing simultaneously on two different Levels of Abstraction: one comprising agents, and another one comprising the system. Designing ethics as gameplay means understanding the inner workings of both, and how they translate to the actual game infosphere as experienced by agents.

Game designers should be concerned with the ethical configuration of the agent in relation with the game infosphere and the game system. One of the key rules of game design is to understand the core audience of a game, and some work has been dedicated to game design based on psychological profiles of players (Bateman and Boon 2006). Similarly, designing ethical gameplay has to take into consideration players as ethical agents. In other words, the design of ethical games requires a model of players that can actively reflect upon the meaning of the ethical challenges posed by games. This means players that understand the encapsulation of the game infosphere within the larger infosphere of their being. These players also understand how to play ethically, how to interpret the interrelation

between their actions and the simulation, and how they can contribute to creating the values of the game by playing it.

What is, then, the model of the ethical player? As previously mentioned in this paper, the ethical player model is based on the Information Ethics concept of the *homo poieticus* (Floridi and Sanders 2005). This is a constructionist approach to the anthropology of agents, as opposed to the *homo ludens* (Huizinga 1950) approach that does not have an ethical dimension. By constructionist, I refer to the capacity and duty of agents within a ludic infosphere to constitute themselves as ethical agents. It is not merely pushing buttons: playing is actively configuring both the game state, by means of interaction, and the agent's ethical capacities and relations to other agents, the infosphere, and the impact of playing the game as perceived from outside the game.

The ethical player has to be considered as much more than a mere input provider. To have an ethical player, designers have to think about agents with constructionist capacities; agents that will determine who they are in the game, and how that being is related to the being outside the game, without being evaluated morally by the game. An agent has to be able to construct their ethics within a game infosphere.

This is why games like *Knights of the Old Republic*, or *Fable*, ultimately fail at promoting players' ethical agency in gameplay: not only do they not succeed in conveying the complexity of the ethical capacities of players, they also disallow the constructionist necessities of players. The ethical experience in these games is limited to a mere calculation of possibilities, numbers and choices that do not affect the ethical constitution of the player as an agent. They don't experience ethical gameplay, they play with ethics.

Given the model of the ethical player as a constructivist agent who wants, and has to, build her ethical capacities and discourses within the game experience, how can we design ethical gameplay? Designers will need to take into consideration the state machine and simulation gradient in order to create interesting play.

The key element to design this kind of gameplay is the ethical player model. Players have to be able to relate to the simulation with their own ethics, as constructed and relevant for the game. Thus, first the game has to provide a space for ethical agency. In other words, the world has to reflect moral choices. The world of *Tetris* is largely morally irrelevant; whilst the world of *Manhunt* (RockStar North 2004) is based on ethical challenges (Sicart 2006). Designers have to make clear to players that ethics is important in that gameworld, and that choices based on moral reasoning will actually have an impact. However, the impact of those choices has to be transmitted through the simulation, and not through values applicable to the state machine. The *Fable* type of games outputs to the ethical

² Since any agent in an infosphere is informationally ethical, the mechanical agent has some moral choice as well, mostly related to maintaining the informational balance of the game. While there could be interesting applications of this approach, I believe that it is more interesting to focus the ethics of games in the agent that has a cultural connection both to the simulation and to the outside of the game infosphere, since this places computer games as a relevant rhetoric tool on its own.

agent the consequences of her actions through a numeric system that is mostly relevant for the state machine. The world should change and react as it does, but players must never have an in-built function in the game system that quantizes their ethical values. It is their task, as moral agents, to understand and build their ethical presence in the game. The choices of the player affect the state machine, but the output through the simulation must respect and encourage the constructivist capacities of players. Decisions only matter when the player puts her own ethics at play.

Ethical gameplay can be built by exploiting the ontological tension between the player as agent within the gameworld, and the player as input provider for a state machine. Players construct their ethics by combining the embedded ethics in the design with their own values, both in their history as players and as beings outside the game infosphere. This process creates an ontological tension that can be exploited creatively: the actions that the player has to take in order to play the game, and win it, can be ethically challenging, provoking a constant balance of the decision making process between the optimization of the resources and the values attached to them.

Games like *September 12th* (Newsgaming.com 2003), *Deus Ex* (Ion Storm 2000), or *Shadow of the Colossus* satisfy this principle: what the player has to do can be in collision with either her values external to the game (*September 12th* puts the ethical/political player in the dilemma of playing vs. not doing anything, and the consequences of the “war on terror”), her values within the game (some choices in *Deus Ex* force players to take temporary allies who may be opposite to the ethical approach to the gameworld enforced during gameplay), or her own actions as a player who wants to finish the game. Exploiting the difficult balance between what is good in a game, what players consider as good in that game, and what cultural beings consider as good in general is a question of careful balance but it is a relevant way of achieving interesting ethical gameplay.

Of course, including multiple ethical agents in the game enhances the moral relevance of the gameplay experience. As a matter of fact, when there are two ethical agents inserted in a game infosphere, in a certain LoA we will have an ethical game, even in abstract games. A multiplicity of agents is an obvious resource for creating ethical gameplay, but is one that has to be managed with care, since the temptation of disempowering agents of their constructivist capacities in favour of a better balance and stability of the game system is always present.

For example, the highly successful online game *World of Warcraft* (Blizzard Entertainment 2005) does not allow players to change the gameworld in any significant way. Furthermore, players are limited by an overtly general and

unclear end user license agreement that has in fact been used to ban accounts of players that protested, or that publicly expressed their sexuality.³ *World of Warcraft's* gameplay is centred on a specific progression ladder, and all forms of expression that diverge from the predetermined route the designers have allowed, are either impossible or prosecuted. There is limited ethical agency in *World of Warcraft*.

On the other hand, a game like *Eve Online* (CCP 2003) provides players with a world in which they can enact their ethical values with almost no limits. The game creates a world, a universe where players interact guided by their own constructivist capacities. Whatever happens in that world, it's the player's choice. As such, it is perhaps the most ethical virtual world available today.

A more radical, more dangerous approach for creating ethical gameplay is to understand the moral agent, then constrain it. *Manhunt* is, by all accounts, a disgusting game where the player has to commit gruesome assassinations to survive in a snuff-movie inspired environment. Yet it is a fascinating ethical experience since it is focused on creating disgust for the player, who cannot avoid feeling certain empathy for the character she is controlling. Limiting in interesting ways the constructivist capacities of players is a fundamental approach for designing ethical gameplay. However, this limitation has to be used very carefully, since any constraint of the poietic capacities of players, that is, their capacity to construct their own moral values for play, and act upon them, is unethical. Those games that explore this type of ethical gameplay should make clear to players that they are playing a game—they should operate almost as Brechtian experiences, forcing the player to permanent awareness of the act of playing a game and, as such, breaking engagement as a rhetorical trope for ethics-based gameplay.

Ethical gameplay design is a task that requires specific details for each game developed. However, in this article I have suggested a number of general principles that can be used to focus the design of morally relevant gameplay. These general principles are:

1. Create an ethically relevant game world.
2. Do not quantize your player's actions: let them live in a world that reacts to their values.
3. Exploit the tension of being an ethical player.
4. Insert other agents with constructivist capacities and possibilities.
5. Challenge the poietic capacities of players, by expanding or constraining them.

³ See for example http://news.cnet.com/2100-1043_3-6033112.html (retrieved 27/10/2008).

How can these five approaches actually be implemented in the game design? I would argue by means of a creative use of the levels of abstraction. Technically, game design is the craft of creating enjoyable abstractions: both the state machine and the simulation layers extensively use levels of abstraction to determine what is informationally relevant or not within a game experience (Juul 2007). A game is a designed gradient of abstractions intended to create a ludic experience of a kind in one or more agents.

Designing a game implies limiting the scope of the simulation: is it relevant that the weather is dynamic? What is the function of game time, and its speed? How many different paths will the player have? Which are the winning conditions, as opposed to any other choices or outcomes? All these questions are posed when designing a game, and they are only answered by applying an intuitive method of abstracting until the desired informational ecology is created.

If, at the core of computer game design, we have levels of abstraction as a design method, we should also think about them as the tools for creating ethical gameplay, both in terms of how to create the infosphere, and what to present to the player as relevant and desirable to achieve. These designed levels of abstraction can be mapped to the five approaches for designing ethical gameplay:

1. Designing a gameworld means delimiting the extent of what is informationally relevant for the game to be played. So, in order to create ethical gameplay, ethics has to be introduced as an important part of the world within any level of abstraction that is observable and relevant for the player. That is, players have to know that ethics is part of the gameworld.
2. The ethical gradient of abstraction must affect the gameworld, but the output of information to the player has to be exclusively focused through the world. In other words: it is the world that reacts to the player's ethical choices, and the player has to deduce from that what her moral stance is within that world, the ethical payoff matrix where actions are important not only in terms of winning strategies, but also in terms of morals.
3. The player's ethical agency has to be in tension, making the choices collide, or reflect, those that are problematic in the real world. Players are mature ethical beings that can be challenged, since they do understand, by means of playing, that they are embedded in a somewhat encapsulated informational environment. Designers should push the boundaries of ethical conventions while letting players exert full ethical agency within the gameworld: taunting player's ethical agency also means allowing them to construct their own ethical levels of abstraction relevant in the gameworld.
4. The game should be open to players creating and implementing their own values. A good ethical game allows the creation of ethical communities that import to the game experience their values.
5. A creative use of levels of abstraction for designing ethical games consists of closing them for player agency. Allow the player to understand that there is a moral world that has moral implications, but do not allow her to do in the gameworld more than what is, in that world, wrong. A good closed ethical design limits the constructivist capacities of the level of abstraction of the player, allowing her to play an on-rails experience of values that exploit the ontological tension between in-game and external ethics.

Playing with the level of abstraction as the experiential domain of a player's agency is the crucial element for designing ethical games. What a player experiences is a gradient of different abstractions she understands as relevant for experience as a player. Ethical gameplay is a ludic encouragement of the constructionist capacities of the player as a crucial part of their experience of the game. Designing ethical gameplay is a challenge for game developers: a challenge beyond conventional approaches to creating games, but that can create game experiences of seldom reached complexity and depth.

Conclusions

Computer games have the potential to create ethical experiences thanks to their unique informational nature. Understanding games as infospheres, and the task of designing them as the work of configuring a number of levels of abstraction available to ethical agents not only overrides the challenge of computer games as an exercise on the banality of evil, but also opens a number of possibilities for game designers interested in challenging players with new types of deep moral choices. This paper has argued for an analytical framework for games as informational environments, from which a set of very basic game design inspirations for creating ethical gameplay can be deduced. These design reflections are only a basic introduction to the key challenge of creating interesting ethical games, and they will probably benefit from a design research approach: these inspirations need to be developed through a game prototype that can be then be tested to validate the ethical design. This, however, is beyond the scope of this paper.

Computer games offer experiences of odd and seemingly impossible worlds. But the moral dimension of games, the possibilities of challenging players as ethical agents, is still in its infancy. With this paper I have

introduced a framework for understanding games as ethical systems, and how these systems can be designed. The possibility of creating engaging ethical gameplay should be the true promise of any next generation computer game.

References

- Akrich, M., & Latour, B. (1992). A summary of a convenient vocabulary for the semiotics of human and nonhuman assemblies. In W. Bijker & J. Law (Eds.), *Shaping technology, building society: Studies in sociotechnical change* (pp. 259–264). Massachusetts: The MIT Press.
- Arendt, H. (2006). *Eichmann in Jerusalem. A report on the banality of evil*. London: Penguin.
- Bogost, I. (2007). *Persuasive games. The expressive power of videogames*. Massachusetts: The MIT Press.
- Brey, P. (1999). The ethics of representation and action in virtual reality. *Ethics and Information Technology*, 1, 4–14.
- Caillois, R. (2001). *Man, play and games*. Urbana and Chicago: University of Illinois Press.
- Carnagey, N. L., Anderson, C. A., & Bushman, B. J. (2007). The effect of video game violence on physiological desensitization to real-life violence. *Journal of Experimental Social Psychology*, 43(3), 489–496.
- Coeckelberg, M. (2007). Violent computer games, empathy, and cosmopolitanism. *Ethics and Information Technology*, 9, 219–231.
- Conrad, J. (2004). *Heart of darkness*. Clayton: Prestwick House.
- Floridi, L. (1999). Information ethics: On the theoretical foundations of computer ethics. *Ethics and Information Technology*, 1(1), 37–56.
- Floridi, L. (2002). On the intrinsic value of information objects and the infosphere. *Ethics and Information Technology*, 4(4), 287–304.
- Floridi, L. (2003). On the intrinsic value of information objects and the infosphere. *Ethics and Information Technology*, 4(4), 287–304.
- Floridi, L. (2005). Information ethics: An environmental approach to the digital divide. *Philosophy in the Contemporary World*, 9(1), 39–45.
- Floridi, L. (2008a). Information ethics: Its nature and scope. In J. van der Hoven & J. Weckert (Eds.), *Moral philosophy and information technology* (pp. 40–65). Cambridge: Cambridge University Press.
- Floridi, L. (2008b). The method of levels of abstraction. *Minds and Machines*, 18(3), 303–329.
- Floridi, L., & Sanders, J. W. (2001). Artificial evil and the foundation of computer ethics. *Ethics and Information Technology*, 3(1), 55–66.
- Floridi, L., & Sanders, J. W. (2002). Computer ethics: Mapping the foundationalist debate. *Ethics and Information Technology*, 4(1), 1–9.
- Floridi, L., & Sanders, J. W. (2005). Internet ethics. The constructionist values of *Homo Poieticus*. In R. J. Cavalier (Ed.), *The impact of the internet in our moral lives* (pp. 195–215). New York: SUNY.
- Funk, J. B., Baldacci, H. B., Pasold, T., & Baumgardner, J. (2004). Violence exposure in real-life, video games, television, movies, and the internet: is there desensitization? *Journal of Adolescence*, 27, 23–29.
- Funk, J. B., Buchman, D. D., Jenks, J., & Bechtold, H. (2003). Playing violent video games, desensitization, and moral evaluation in children. *Journal of Applied Developmental Psychology*, 24(4), 413–436.
- Gadamer, H. G. (2004). *Truth and method*. New York: Continuum.
- Hocking, C. (2007). *Exploration—from systems to spaces to self*. Talk at the Game Developers Conference, San Francisco, March 2007.
- Huizinga, J. (1950). *Homo ludens*. Boston: Beacon Press.
- Juul, J. (2005). *Half-real. Video games between real rules and fictional worlds*. Massachusetts: The MIT Press.
- Juul, J. (2007). A certain level of abstraction, paper presented at the 3rd digital games researchers association conference, Tokyo, September 2007. Available online at <http://www.jesperjuul.net/text/acertainlevel/>. Accessed 27 Oct 2008.
- Latour, B. (2002). Where are the missing masses? Sociology of a few mundane artefacts. In W. Bijker & J. Law (Eds.), *Shaping technology, building society: Studies in sociotechnical change* (pp. 225–258). Massachusetts: The MIT Press.
- Murray, J. (1997). *Hamlet on the holodeck*. Massachusetts: The MIT Press.
- Salen, K., & Zimmerman, E. (2004). *Rules of play. Game design fundamentals*. Massachusetts: The MIT Press.
- Sicart, M. (2006). *Game, player, ethics*. Ph.D. Thesis, IT University of Copenhagen.
- Smith, S. L., Lachlan, K., & Tamborini, R. (2003). Popular video games: Quantifying the presentation of violence and its context. *Journal of Broadcasting and Electronic Media*, 47(1), 58–76.
- Wonderly, M. (2008). A humane approach to assessing the moral significance of ultra-violent video games. *Ethics and Information Technology*, 10, 1–10.

References: Games

- Bioware (2003). *Knights of the Old Republic*.
- Blizzard Entertainment (2005). *World of Warcraft*.
- CCP (2003). *Eve Online*.
- id Software (1993). *Doom*.
- Introversion Software (2006). *Defcon*.
- Ion Storm (2000). *Deus Ex*.
- Konami (1998). *Dance dance revolution*.
- Lionhead Studios (2004). *Fable*.
- Newsgaming.com (2003). *September 12th*.
- Pajitnov, A. (1985). *Tetris*.
- Quantic Dream (2005). *Fahrenheit*.
- Relentless (2006). *Buzz! The Big Quiz*.
- RockStar North (2008). *Grand Theft Auto IV*.
- RockStar North (2004). *Manhunt*.
- Team Ico (2006). *Shadow of the Colossus*.
- Valve (2000). *Counter-Strike*.