A Review of The Price of Malice and the Price of Anarchy

Chris Dabrowski

December 24, 2006

1 Introduction

I will review two concepts in this project. These two concept are: the price of anarchy and the price of malice. The price of anarchy is the ratio between, the worst Nash equilibrium possible in a game, and the optimal solution achieved when all agents collaborate. So it is comparing the Nash equilibrium which gives the worst social welfare to the solution that gives the best social welfare. The price of malice is the ratio between, the social welfare achieved when there are malicious agents in the system, and the social welfare when all the agents act selfishly. I will explain these two concepts in relation to a virus inoculation game, with respect to two different conditions. The first condition is the oblivious model where the selfish agents are unaware of the existence of the Byzantine or malicious agents, and the second condition is where they are aware of their existence. I will show the results that have been achieved with respect to these concepts and then I will do a critique of the work done and discuss my opinions.

2 The Game

The game is played in a grid. There are n agents each of which corresponds to a node in an undirected grid $G[r, c]$ of r rows and c columns. The upper left hand corner of the grid is $G[0, 0]$. There are selfish players defined by the set S and Byzantine players defined by the set B. $|S| = s$ and $|B| = b$, $n = s + b$.

2.1 Actions

There are two actions available to each selfish player, either inoculate or do nothing. If an agent inoculates then he protects himself from a possible virus infection, but at a cost associated with inoculation. If the agent decides to do nothing then it does not cost him anything but he leaves himself vulnerable to a possible virus infection. These actions can be summarized by $\vec{a} \in \{0, 1\}^n$, where $a_i = 1$ signifies that agent i installs the anti-virus software and $a_i = 0$ that he does not. Nodes i with $a_i = 1$ are secure. The set of secure nodes is denoted by $I_{\vec{a}}$.

After the nodes have made their selection the adversary picks some node uniformly at random to infect. The virus then propagates through all the non-secure agents that are attached to the agent who was the starting point of infection. The attack graph is the network graph associated with $I_{\vec{a}}$ and is denoted by $G_{\vec{a}} = G \backslash I_{\vec{a}}$. It is the network graph with all the secure nodes and their incident edges removed.

Byzantine players submit to the following strategy: they report themselves as inoculated but really they are not. We define $I_{\vec{a}}^{self}$ as the set of selfish agents who are really inoculated and secure. We have $G_{\vec{a}} = G - I_{\vec{a}}^{self}$.

2.2 Costs

Installing the anti-virus software has a cost of 1. If an agent decides to not install and she gets infected she suffers a cost of L. If she does not install and does not get infected then there is no cost to her. The cost can be summarized by the following equation $cost_i(\vec{a}) = a_i + (1 - a_i) * L * \frac{k_i}{n}$. k_i is the size of the connected component in $G_{\vec{a}}$ that contains i. The term $\frac{k_i}{n}$ is the probability that one of the nodes in the connected component that i is in is the originator of the virus. And therefore it is the probability that i gets infected conditioned on the fact that i is not secure. The social cost of the strategy profile $\vec{a}$ is

$$Cost(\vec{a}) = \sum_{j \in S} cost_j(\vec{a})$$

We do not consider the cost of Byzantine agents as they are not trying to maximize their utility. If we were to consider their costs in the social cost then they could just drive it up by inflicting harm on themselves.

The perceived individual cost $\widehat{cost}_i(\vec{a})$ is the cost expected by player i given his knowledge of the Byzantine players. This cost depends on the underlying model. There are two models, oblivious and non-oblivious.

Oblivious: In the oblivious model selfish players are not aware of the existence of Byzantine players. Selfish players assume that all other players are selfish too. The cost is the same as shown previously. In this model the perceived cost may be lower than the actual cost.

Non-oblivious: In this model selfish players know b, the number of Byzantine players, but they do not know the location of these players. Selfish players are highly risk-averse and they aim to minimize their maximal individual cost. Let D be the set of all possible distributions of Byzantine players among all of the players. Then the perceived cost in the non-oblivious model is

$$\widehat{cost}_i(\vec{a}) = \max_{d \in D} \{cost_i(\vec{a}, d)\}$$

where $cost_i(\vec{a}, d)$ is the actual cost of i if the Byzantine agents are distributed according to $d \in D$. In the non-oblivious case the perceived cost might be higher than the actual cost.

So now we can write the social cost as

$$Cost(\vec{a}) = \sum_{j \in S} cost_j(\vec{a}) = |I_{\vec{a}}^{self}| + \frac{L}{n} \sum_{i=1}^{l} k_i l_i$$

Where the first term is the cost due to inoculation and the second term is the infection cost. k_1, k_2,...,k_l are the sizes of the non-secure components in

$G_{\vec{a}}$ and l_1, l_2,...,l_l are the sizes of the same components without counting the Byzantine nodes.

The optimal social cost $Cost_OPT(I)$ for a problem instance I is the sum of all the players individual costs when the players collaborate perfectly and therefore achieve the lowest possible social cost.

3 Definitions

3.1 Byzantine Nash Equilibrium

The Byzantine Nash Equilibrium (BNE) is defined for every possible strategy a'_i as,

$$\forall i \in S : \widehat{cost}_i(\vec{a}) \leq \widehat{cost}_i(\vec{a}[i|a'_i])$$

Where $\vec{a}[i|x]$ is the same strategy profile as $\vec{a}$ except the i-th component is replaced by x. The cost of the Byzantine Nash Equilibrium $Cost_{BNE}(I, b)$ is the social cost of the worst BNE of a problem instance I, with b Byzantine agents. Therefore it is the BNE with the highest social cost.

3.2 Price of Anarchy and Price of Byzantine Anarchy

The Price of Anarchy expresses the degradation of the socially optimal performance of a system due to the selfish behavior of the agents. This is defined ∀ problem instances I as,

$$PoA = \max_I \frac{Cost_{BNE}(I, 0)}{Cost_{OPT}(I)}$$

The Price of Byzantine Anarchy is defined ∀ problem instances I as,

$$PoB(b) = \max_I \frac{Cost_{BNE}(I, b)}{Cost_{OPT}(I)}$$

Therefore $PoA = PoB(0)$.

3.3 Price of Malice

The Price of Malice expresses how much the presence of malicious agents deteriorates the social welfare of a system consisting of selfish players. The Price of Malice is defined as,

$$PoM(b) = \frac{PoB(b)}{PoB(0)}$$

A high Price of Malice indicates that a system is particularly sensitive to Byzantine agents. A low Price of Malice indicates that a system is relatively stable and is not influenced greatly by the introduction of Byzantine agents.

4 Analysis

In the case where there are no Byzantine agents it holds that

$$\frac{1}{4} * \Big(\frac{s}{L}\Big)^{1/3} \leq PoA \leq \frac{6}{\sqrt{\pi}} * \Big(\frac{s}{L}\Big)^{1/3}$$

In the case with Byzantine agents with the oblivious model we have

$$PoB(b) \in \Theta\Big(\Big(\frac{s}{L}\Big)^{1/3}\Big(1 + \frac{b^2}{L} + \frac{b^3}{sL}\Big)\Big)$$

and

$$PoM(b) \in \Theta\Big(1 + \frac{b^2}{L} + \frac{b^3}{sL}\Big)$$

for $b < \frac{L}{2} - 1$, otherwise it holds that

$$PoB(b) \in \Theta(s^{1/3}L^{2/3})$$

and $PoM(b) \in \Theta(L)$

In the non-oblivious model we have

$$PoB(b) \geq \frac{1}{8}\Big(\Big(\frac{s}{L}\Big)^{1/3} + \frac{b}{2}\Big(\frac{L}{s}\Big)^{2/3}\Big)$$

and

$$PoM(b) \geq \frac{\sqrt{\pi}}{48}\Big(1 + \frac{bL}{2s}\Big)$$

for $b < \frac{n}{2L}$. For all b we have $PoB(b) \geq \frac{1}{8}(\frac{s}{L})^{1/3}$ and $PoM(b) \geq \frac{\sqrt{\pi}}{48}$.

From these results we see that in the non-oblivious case the Price of Malice could actually be less than one. This shows that in some cases the introduction of Byzantine agents actually improves the social welfare. Intuitively this can be explained by the fact that as Selfish agents gain knowledge of the existence of Byzantine agents they are more likely to cooperate and inoculate. This leads to a concept called the $FearFactor$. The Fear Factor describes the overall gain in efficiency of a selfish system with the introduction of malicious players. It is defined as,

$$\Psi(b) := \frac{1}{PoM(b)}$$

The Fear Factor quantifies how much the threat of malicious agents can improve the social welfare of a system. In the virus inoculation game the minimum value the Price of Malice can take is $\frac{\sqrt{\pi}}{48}$ and therefore the maximum value that Ψ can take is $\frac{48}{\sqrt{\pi}}$. This shows that the improvement from the threat of malicious players is bounded and cannot exceed $\frac{48}{\sqrt{\pi}}$.

5 Discussion

There are a few points of discussion concerning [2]. The main problem exists with the analysis section when computing the upper and lower bounds on the cost in the oblivious case with Byzantine agents and the non-oblivious case with Byzantine agents. The problem lies with the choice of Nash Equilibria. In the two cases a different equilibrium is chosen each time. There is not much justification given for why these are the equilibria that are chosen. I think that a section is needed to explain these choices and to show that if we chose a different equilibria we would still get the same bounds.

Another interesting point is the specific game that was chosen for the analysis of this concept. While I think that this concept is somewhat limited as it has only been applied with regards to one game, the game that was used is fairly general. The game contains a few general concepts which are true of all games that this kind of concept would apply to. The agents have a choice of either cooperating(inoculating) or not. If the agents cooperate they suffer a cost for doing so. If they do not cooperate they do not have to pay anything. So it is advantageous to not cooperate. But if too many of the agents do not cooperate than they all suffer because the probability of infection goes up and the expected cost of infection exceeds the cost of inoculation. Therefore this game has that tragedy of the commons and prisoner's dilemma aspect to it where it is better to not cooperate but then if everyone does not cooperate everyone suffers more than if they had all cooperated.

The things that make this game specific are the way the agents are connected to each other and how the punishment for not cooperating works. We could imagine a game where the agents have power to affect all of the other agents and not just their direct neighbors. I think it is necessary to apply the concept of The Price of Malice to a few different games and maybe try to come up with some general conclusions about the Price of Malice that are not dependent on the specific game being analyzed.

6 Conclusion

In my opinion the I think the research presented here is valuable. It shows that in a specific game in a specific setting the introduction of malicious players is beneficial to the social welfare of the system. This gives valuable insight into mechanism design, as the threat of malicious players is a tool that could be used by the designer of the mechanism.

This research also spawns the idea of research into other types of agents that are not described by the extremes of just selfish or Byzantine. Agents that might have different qualities that could be analyzed, along with their impact on the social welfare of the system.

The Price of Malice and the Price of Anarchy have an interesting relationship. They are related in the sense that as the Price of Anarchy goes up and everything else stays the same the Price of Malice goes down. If the Price of Anarchy goes down and everything else stays the same then the Price of Malice will go up. So therefore in a system with a high Price of Anarchy the introduction of Byzantine agents may not be very harmful and might actually be beneficial.

References

[1] E. Koutsoupias and C. Papadimitrou. Worst-Case Equilibria. *Lecture Note in Computer Science*, 1563:404-413, 1999.

[2] T. Moscibroda, S. Schmid, and R. Wattenhoffer. When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game. In *Proc. 25th ACM Symposium on Principles of Distributed Computing(PODC)*, pages 35-44, 2006.

[3] C. Papadimitriou. Algorithms, games and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 749-753, 2001.

[4] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. In *STOC*, pages 428-437, 2002.