



Distributed Systems Part II

Solution to Exercise Sheet 8

1 DHT

There are no “correct” and “wrong” answers for this task. Just answers that are better than others. The answers given here are just examples of potential solutions.

- a)
- The peer starting the search sends many messages in different directions, hoping they take a different path.
 - Each peer receiving the message sends an acknowledgement and the address of the next peer back to the one peer that started the search. The starting peer notices if the message gets lost or takes the wrong path. It then reacts, for example by restarting the search. The drawback of this solution are the additional connections that have to be made.
 - Assuming we have authentication. Every process receiving the message sends back an acknowledgement to the previous peer. The previous peer then forwards the message to its predecessor. If a peer does not receive two times an acknowledgement (from the correct peers), then its message might have been lost. In this case the peer sends the message again, but in another direction.
- b)
- When receiving an update for the routing table, the peer first contacts the new entries to ensure they really exist.
 - A peer only accepts an update if received from enough other peers, for example if received $f + 1$ times.
 - A peer keeps its old table for some time. Should the new table prove to be bad (e.g. messages are often delayed), then the peer switches back to the old table.
- c)
- Some trusted servers are used for joining. But this would mean that the DHT is no longer fully decentralized.
 - The joining peer contacts $3f + 1$ servers. These servers run a consensus-protocol. The result of the consensus is the place and the routing table of the joining peer.
 - The joining peer contacts $f + 1$ servers. If they do not agree, the peer randomly chooses another set of $f + 1$ servers. This works only if f is small compared to the total number of servers.

2 Selfish Caching

- a)
- i. The best response strategies are
 - A*: cache only if nobody else does. (B1)
 - B*: cache if neither *A* nor *D* cache. (B2)

C: cache unless *A* caches. (B3)

D: cache if neither *A* nor *B* cache. (B4)

Nash equilibrium. If we assume that *A* plays $x_A = 1$ (*A* caches) the system can only be in a NE if $x_B = x_C = x_D = 0$ due to (B1). Since for all *B*, *C*, and *D* it is the best response not to cache if *A* does, $x = (1000)$ is a Nash equilibrium. If $x_A = 0$ then (B3) implies $x_C = 1$. If furthermore, $x_B = 1$ it must hold that $x_D = 0$ due to (B2). This does not conflict with (B4), and (0110) constitutes another NE. Last, if $x_B = 0$ then (B2) implies $x_D = 1$, which is also okay with (B4). Hence (0011) is also a NE.

$$NE = \{(1000), (0110), (0011)\}$$

Price of anarchy. The social optimum is achieved in strategy profile (1000), namely $OPT = cost(1000) = 4 + 2 + \frac{3}{2} + 3 = 10.5$. Since (1000) is also a Nash equilibrium we immediately get that $OPoA = 1$. The worst-case price of anarchy is

$$PoA = \frac{cost(0110)}{OPT} = \frac{2 + 4 + 4 + 3.5}{10.5} = \frac{9}{7} \approx 1.286$$

ii. The best response strategies are

A: cache only if nobody else does. (B1)

B: cache unless *A* caches. (B2)

C: cache unless *D* caches. (B3)

D: cache if neither *A* nor *C* cache. (B4)

Nash equilibrium. If we assume that *A* plays $x_A = 1$ (*A* caches) the system can only be in a NE if $x_B = x_C = x_D = 0$ due to (B1). However, $x_D = 0$ implies that $x_C = 1$ due to (B3), and hence there can be no NE with $x_A = 1$. In any NE it must hold that $x_A = 0$. Consequently, it must hold that $x_B = 1$ from (B2). Now if $x_C = 1$ (B3) implies that *D* does not cache. This does not infringe rule (B4), and thus $x = (0110)$ is a Nash equilibrium. If $x_C = 0$ then (B4) implies that *D* caches. As thus, rule (B3) is not violated $x = (0101)$ is also a Nash equilibrium.

Price of anarchy. The social optimum is achieved in strategy profile (0110), namely $OPT = cost(0110) = \frac{1}{3} \cdot 2 + 10 + 10 + \frac{1}{2} \cdot 2 = 21.\overline{66}$. Since (0110) is also a Nash equilibrium we get that the optimistic price of anarchy is 1. The worst-case price of anarchy is

$$PoA = \frac{cost(0101)}{OPT} = \frac{2/3 + 10 + 2 + 10}{21 + 2/3} = \frac{68}{65} \approx 1.046$$

b) This game has two players, $N = \{A, B\}$. Each player can either cache ($x_i = 1$), or choose not to cache ($x_i = 0$). Thus $X_A = X_B = \{1, 0\}$. The bi-matrix of the game with peer *A* as the row player and *B* as the column player looks as follows.

	1	0
1	10 , 10	10 , 6
0	3 , 10	100 , 100

Note that the numbers in this bi-matrix represent costs. Find the dual version where we use utility instead of cost below:

	1	0
1	90 , 90	90 , 94
0	97 , 90	0 , 0

The pure Nash equilibria are (10) and (01). If the players randomize their pure strategies, and draw their choice from a probability distribution there might be additional mixed NEs. Let peer A choose to cache with probability p . Thus she chooses not to cache with probability $1 - p$. Let the probability that B chooses to cache be q . Thus, $1 - q$ is the probability that she chooses not to cache. The expected cost of each peer now depends on the choices of p and q . We denote by $\Gamma_i(p, q)$ the expected cost function of player i . For the given game we have for peer A

$$\begin{aligned}\Gamma_A(p, q) &= p \cdot 10 + (1 - p)(q \cdot 3 + (1 - q) \cdot 100) \\ &= -90p + 97pq - 97q + 100 \\ &= (97q - 90) \cdot p - 97q + 100\end{aligned}$$

and for peer B

$$\begin{aligned}\Gamma_B(p, q) &= q \cdot 10 + (1 - q)(p \cdot 6 + (1 - p) \cdot 100) \\ &= -90p + 94pq - 94q + 100 \\ &= (94p - 90) \cdot q - 97p + 100.\end{aligned}$$

Peer A wants to minimize $\Gamma_A(p, q)$. If the term $(97q - 90)$ is positive A 's best response is to choose $p = 0$, if it is negative it should choose $p = 1$, and if the term equals 0 it can choose any p within $0 \leq p \leq 1$. The best response of peer B is to play $q = 0$ if the term $(94p - 90)$ is positive, $q = 1$ if it is negative, and $0 \leq q \leq 1$ if $(94p - 90) = 0$. In a Nash equilibrium, both players play a best response to each other's choice. If A plays $p = 0$ then B 's best response is $q = 1$. Since $p = 0$ is also a best response for A to B playing $q = 1$, $(p, q) = (0, 1)$ is a mixed NE. If A plays $0 < p < \frac{90}{94}$ then B 's best response is $q = 1$. As A 's best response to $q = 1$ is $p = 0$, no p between 0 and $\frac{90}{94}$ can constitute a Nash equilibrium. If A plays $p = \frac{90}{94}$ then B can play anything she wants. If B chooses $q = \frac{90}{97}$ then any choice of p is a best response for A . Thus if $p = \frac{90}{94}$ and $q = \frac{90}{97}$ both peers play a best response, and none of them has an incentive to deviate. Hence $(p, q) = (\frac{90}{94}, \frac{90}{97})$ is a mixed NE. $(p, q) = (1, 0)$ is the third mixed NE.

$$NE = \{(0, 1), (1, 0), (\frac{90}{94}, \frac{90}{97})\}$$

Note that $(0, 1)$ and $(1, 0)$ correspond to the two pure NEs observed earlier.

- c) Let us derive the social optimum in a first step given a line topology of n peers. Naturally, the best configuration will be one where the distance between two caching nodes is roughly the same. For simplicity let us assume that the number of non-caching peers between two caching peers is always x . Moreover, let the first caching node be $P_{\frac{x}{2}+1}$. The second one would be $P_{\frac{x}{2}+x+2}$, the third $P_{\frac{x}{2}+2x+3}$, and so on. For simplicity we assume that the last caching node is followed again by $\frac{x}{2}$ non-caching nodes. Let $k = n/x$ be the number of caching nodes. The cost for such a configuration is

$$\begin{aligned}cost(x) &= 2 \cdot \left(\frac{x}{2} + \left(\frac{x}{2} - 1 \right) + \dots + 1 \right) + (k - 1) \cdot \left(\alpha + 2 \left(\frac{x}{2} + \left(\frac{x}{2} - 1 \right) + \dots + 1 \right) \right) + \alpha \\ &= k \cdot \left(\alpha + 2 \sum_{i=1}^{\frac{x}{2}} i \right) = k \cdot \left(\alpha + \frac{x}{2} \left(\frac{x}{2} - 1 \right) \right) = \frac{n}{x} \left(\alpha + \frac{x}{2} \left(\frac{x}{2} - 1 \right) \right) \\ &= \frac{\alpha n}{x} + \frac{nx}{4} - \frac{n}{2}\end{aligned}$$

Computing the root of the first derivative of $cost(x)$ yields the optimal value of x .

$$\begin{aligned}\frac{d}{dx} cost(x) &= -\frac{\alpha n}{x^2} + \frac{n}{4} \stackrel{!}{=} 0 \\ \frac{n}{4} &= \frac{\alpha n}{x^2} \\ x &= 2\sqrt{\alpha}\end{aligned}$$

Choosing $x = 2\sqrt{\alpha}$ yields an optimal cost of $OPT = n\sqrt{\alpha} - n/2$.

Now, in a Nash equilibrium the distance between two caching peers is at least α , since if a selfish peer is closer to a caching peer it is cheaper for it to access the object remotely than to cache it itself. The number x of non-caching nodes between two caching nodes can be up to 2α in the worst case. Hence the cost of the worst NE is roughly

$$\text{cost}(\text{worstNE}) \approx \text{cost}(2\alpha) = \frac{n}{2} + \frac{\alpha n}{2} - \frac{n}{2} = \frac{\alpha n}{2}.$$

Finally, we get for the price of anarchy

$$PoA = \frac{\text{cost}(\text{worstNE})}{OPT} = \frac{\alpha}{2\sqrt{\alpha} - 1}.$$

Remark: From this result, we can conclude, that in the worst case, namely if α is a constant fraction of n , the price of anarchy is in $\Omega(\sqrt{n})$.

3 Selfish Caching with variable placement cost

a) We define D_i to be the set of peers that cover peer i . A peer j covers peer i iff $w_i d_{ij} < \alpha_i$, i.e., peer i prefers accessing the object at peer j than caching it. Convince yourself that a strategy profile x is a Nash equilibrium iff for each peer i it holds that

- if $x_i = 1$ then $x_j = 0$ for all $j \in D_i$, and
- if $x_i = 0$ then $\exists j \in D_i$ with $x_j = 1$.

- i. $D_A = \emptyset$, $D_B = \{A, C\}$, $D_C = \{B\}$. D_A being empty implies $x_A = 1$. Hence $x_B = 0$, and $x_C = 1$. $NE = \{(101)\}$. $PoA = 1$ since (101) is also the social optimum strategy.
- ii. $D_A = \{B\}$, $D_B = \{A\}$, $D_C = \{A, B\}$. If $x_A = 1$ then $x_B = 0$ and $x_C = 0$. If $x_A = 0$ then $x_B = 1$. Hence $x_C = 0$. The equilibria are $NE = \{(100), (010)\}$.

$$PoA = \frac{\text{cost}(100)}{\text{cost}(010)} = \frac{3 + 1 + 8/3}{3/2 + 3/2 + 5/3} = \frac{40}{28} \approx 1.43$$

Dominant strategies. Every dominant strategy profile is also a Nash equilibrium. Hence we only have to check the computed NEs whether they consist of dominant strategies only. Furthermore, if there is more than one NE we already know that the game does not have a dominant strategy profile. This is because for at least one peer, the best strategy depends on the choices of the other peers. Thus game ii. has no dominant strategy profile. Profile (101) is no dominant strategy profile in game i. either, since, although $x_A = 1$ is the dominant strategy for A , $x_B = 0$, and $x_C = 1$ are not dominant strategies for B and C . If for instance, x_{-C} would be so that $x_B = 1$ it would be the best response of C to set $x_C = 0$.

b) Let I^n be an instance of $\mathcal{A}_{[a,b]}^n$ that maximizes the price of anarchy, i.e. $PoA(\mathcal{A}_{[a,b]}^n) = PoA(I^n)$. Let $x, y \in X$ be two strategy profiles in \hat{I}^n such that $PoA(I^n) = \text{cost}(y)/\text{cost}(x)$. We show the claim by constructing an instance $\hat{I}^n \in \mathcal{W}_{[\frac{1}{b}, \frac{1}{a}]}^n$ out of I^n for which it holds that $PoA(\hat{I}^n) \geq \frac{a}{b} PoA(I^n) = \frac{a}{b} PoA(\mathcal{A}_{[a,b]}^n)$. We construct \hat{I}^n by setting $w_i = 1/\alpha_i$, $\hat{\alpha}_i = 1$ where α_i are the placement costs of player i in I^n . All edges remain as in I^n . This game has the same Nash equilibria as I^n since the cover sets D_i for each peer stay the same. A peer j is in D_i iff $d_{ij} < \alpha_i$, or $d_{ij}/\alpha_i < 1$ respectively. We get the bound by comparing the performance of the two strategies x, y that produce the PoA in I^n in \hat{I}^n . Note that x is

not necessarily a social optimum, but y is a Nash equilibrium also in \hat{I}^n .

$$PoA(\hat{I}^n) \geq \frac{\hat{cost}(y)}{\hat{cost}(x)} = \frac{\sum_{i=1}^n \left(y_i + (1 - y_i) \frac{d_i(y)}{\alpha_i} \right)}{\sum_{i=1}^n \left(x_i + (1 - x_i) \frac{d_i(x)}{\alpha_i} \right)} \quad (1)$$

$$= \frac{b \cdot a \sum_{i=1}^n \left(y_i + (1 - y_i) \frac{d_i(y)}{\alpha_i} \right)}{b \cdot a \sum_{i=1}^n \left(x_i + (1 - x_i) \frac{d_i(x)}{\alpha_i} \right)} \quad (2)$$

$$\geq \frac{a \sum_{i=1}^n (y_i \alpha_i + (1 - y_i) d_i(y))}{b \sum_{i=1}^n (x_i \alpha_i + (1 - x_i) d_i(x))} \quad (3)$$

$$= \frac{a \cdot cost(y)}{b \cdot cost(x)} = \frac{a}{b} PoA(I^n) \quad (4)$$

$\hat{cost}(x)$ denotes the cost function in \hat{I}^n . x_i , and y_i are either 1 or 0. x_i equals 1 if player i caches in strategy profile x , and 0 if she does not. For step (3) we exploit the fact that $b \geq \alpha_i$ and $a \leq \alpha_i$ for all i .

4 Matching Pennies

The bi-matrix of the game with Tobias as row player, and Stephan as column player looks as follows:

	H	T
H	1 , -1	-1 , 1
T	-1 , 1	1 , -1

This zero-sum game has no pure Nash equilibrium. For the mixed NEs, Tobias plays heads (H) with probability p , tails (T) with probability $1 - p$. Stephan plays H with probability q , and T with probability $1 - q$. We get the expected utility functions Γ :

$$\begin{aligned} \Gamma_T(p, q) &= p(q - (1 - q)) + (1 - p)(-q + (1 - q)) = (4q - 2) \cdot p + 1 - 2q \\ \Gamma_S(p, q) &= q(-p + (1 - p)) + (1 - q)(p - (1 - p)) = (2 - 4p) \cdot q + 2p - 1 \end{aligned}$$

If Stephan plays $q = 1/2$ the term $4q - 2$ equals 0, and any choice of p will yield the same payoff for Tobias. If Tobias plays $p = 1/2$ then any choice of q is a best response for Stephan. Thus $(p, q) = (1/2, 1/2)$ is a mixed NE. Note that for any choice of $p > 1/2$, Stephan's best response is to choose $q = 0$. For a $p < 1/2$ Stephan would choose $q = 1$. However, Tobias' best response to $q > 1/2$ is $p = 1$, and $p = 0$ if $q < 1/2$. Hence $(p, q) = (1/2, 1/2)$ is the only pair of mutual best responses.

5 P2P File Sharing

- a) **Temporal imbalance.** If two peers, each of which has data the other is interested in, want to trade then at least one of them has to start sending data before it has received anything in return. It takes the risk that the trading partner does not stick to the deal, and cuts the connection. If the two peers have the prospect of continuing trading in the future, this problem can be relaxed. However, if a peer lacks only one file block to complete its download, it has no incentive at all to return the favor after a trading partner has provided it with the missing block.

Bootstrapping. In swarm-based P2P file sharing systems like Bittorrent, a peer that starts a download has no trading capital yet when it enters the swarm. In order to use a

T4T strategy a new peer has to be provided with some file blocks for free before it can start proper trading. This can be exploited by a freeloader in that it always claims to be a new peer.

Moreover, statistics of today's file-sharing systems reveal that there are typically only a few peers that publish most of the content. Thus there is an inherent, and probably desired imbalance among the users that could not be supported by a strict T4T regime.

- b)
 - i. The freeloader can receive Δ different file blocks from each peer in the swarm. Hence the swarm size has to be at least $\lceil m/\Delta \rceil$.
 - ii. The swarm size n does not matter. The freeloader can only get Δ file blocks for free in total. Thus, it can only download the entire file for free if $\Delta \geq m$.
 - iii. The P2P paradigm means that all peers are equally privileged, equipotent participants in the application. To have a central authority contradicts this understanding of a P2P system. Thus, the central authority has to be emulated by the peers themselves. Unfortunately, this is not an easy task, and most reputation system that try to implement it are susceptible to cheating. One problem is that the actions that affect the reputation—or in the case of a Bittorrent-like system, the "global" balance—are typically observed only by very few peers. For instance, if a peer A claims that it did not receive a certain file block from peer B although B claims so it is hard for a third peer C to tell which peer is lying. One cumbersome possibility is to make majority votes among a large fraction of all the peers. However this is inefficient, brings a large overhead, and is still susceptible to colluding freeloaders, or Sybil attacks. There are promising attempts that use cryptographic primitives to install a distributed trusted authority through multiparty computation.

In the realm of P2P file-sharing, an alternative to a global balance that achieves the same effect is to choose the file blocks that a peer is allowed to get for free depending on a requesting peer's Id, or IP-address. If all peers use the same rule when computing this set of free blocks a freeloader will only get these blocks. One example would be to provide a peer with IP x with the blocks in the range $[\mathcal{H}(x) \bmod m, \mathcal{H}(x) \bmod m + \Delta]$, where \mathcal{H} is a hash function, e.g. SHA-1.