Report on Veronika Slivova's thesis by Nir Bitansky.

The thesis addresses two problems concerning the complexity class TFNP of total search problems with unique solutions. The results presented in this thesis are new and have been published in leading conferences in cryptography and algorithms. I can attest to the importance of most of the results in this thesis (at least those that fall in my area of expertise). The ideas and techniques introduced do prove creativity and ability to perform meaningful research.

I didn't have time to verify the correctness of all the results, but I found the high level approach sound. I did look deeper into some of the technical sections and I found that they can be improved (and I advise to do that).

Below I address the different parts of the thesis.

Sincerely,
Nir Bitansky
School of Computer Science, Tel Aviv University

The complexity of Arrival. The thesis shows that: (a) The problem is in UP intersect coUP, namely both instances and co-instances have a unique witness. (b) The corresponding search problem is in the class CLS. Previously, it was only shown that the problem is in NP intersect coNP and accordingly its search variant is in TFNP.

Context and background are missing, which made it hard for me to appreciate the importance of the result. I haven't encountered the problem in the past and would have wanted to know where it fits in the game-theoretic/combinatorial terrain and why people care about it. Indeed, it seems that people do care about it, and the fact the paper was accepted to ICALP attests to that. I strongly suggest adding context to the thesis or to a public version of the corresponding paper.

Aside from that, I do view the problem as natural/interesting and could appreciate its aesthetics. The ideas developed to achieve the result are not ground-breaking, but are nevertheless elegant and insightful, and certainly demonstrate creativity. The work identifies the right abstractions and gadgets needed to characterize the problem (in particular, the "last-visited induced graph").

Complexity of TFNP. The thesis shows two black-box impossibility results: (a) Hardness in TFNP (even worst-case) does not have a fully black-box reduction to average-case hardness in UP. (b) There are no "simple" fully black-box reductions to injective one-way functions.

This contribution falls under the bigger question of whether hardness of total search problems requires "structured hardness". The question is in fact quite nuanced. HNY showed that average-case NP hardness implies non-uniform total search problems and doing away w/

non-uniformity is only known under derandomization assumptions. A meaningful way to get a handle on why we cannot achieve hard total problems w/o sufficient structure is by ruling out natural types of reductions, which is what this thesis does.

Simple reductions seem quite restrictive, but nevertheless the step taken in this thesis is meaningful and far from trivial. In particular, it requires careful analysis, and a very good understanding of black-box separations.

Specific comments about section 3:

- You should elaborate on the high-level intuition behind the oracle solve. Your overview stops too early and there's more room for intuition. In particular, explain the intuition that we can imagine that inputs x of length n have no solution and consider how we would answer query i in that situation, and then that actually if you look at such an answer then it's likely to also be benign in the real experiment.
- A good mental exercise that can also help the reader is to explain why your separation approach wouldn't work if you say that a hard problem in NP intersect coNP (which indeed does imply hardness in TFNP)
- The quantification in Def 3.3.1-item3 seems off. The reduction only has to decide the relation W for which solve^W works. The issue in quantification seems to also be present the statement of Claim 3.4.1 where you say for every n, there is W.
- The statement and proof of Claim 3.4.1 don't compile. First the term |Q_{W,n}| isn't defined w/o specifying the oracle W, what you really mean here is a bound on this size (over all choices W). In the proof, Q_{w,n} (w/ lower case) appears a few times, and looks like a typo. In the proof you sum over elements in Q_{W,n} where it is not properly defined. It does seem to me that the approach is sound and that the proof is easily fixable (you basically want to take the maximum over W_{<n} and then look at the probability over the rest of the oracle W_{=> n}).
- You only claim to rule out reductions that run in quasi-poly time. Is this a limitation of your technique, or is there reason to expect that TFNP hardness can be based on, say, subexponential hardness of the corresponding primitives. Where does the proof fail if you allow the reduction to run is some mild exponential time? This should be addressed.
- The first time that S_{i,W} appears is before it is defined.