On the Manipulability of Voting Systems: Application to Multi-Operator Networks

François Durand
Inria
23 avenue d'Italie
75013 Paris, France
Email: francois.durand@inria.fr

Fabien Mathieu and Ludovic Noirie
Alcatel-Lucent Bell Labs France
Centre de Villarceaux
91620 Nozay, France
Email: first.last@alcatel-lucent.com

***Abstract*—Internet is a large-scale and highly competitive economic ecosystem. In order to make fair decisions, while preventing the economic actors from manipulating the natural outcome of the decision process, game theory is a natural framework, and voting systems represent an interesting alternative that, to our knowledge, has not yet being considered. They allow competing entities to decide among different options. In this paper, we investigate their use for end-to-end path selection in multi-operator networks, analyzing their manipulability by tactical voting and their economic efficiency. We show that *Instant Runoff Voting* is much more efficient and resistant to tactical voting than the natural system which tries to get the economic optimum.**

I. Introduction

The emergence of new services, along with the constant growth of existing ones, has led to the explosion of Internet traffic. The network of networks has become a huge economic ecosystem, on which lot of companies are making revenues. This includes network operators, that interconnect their infrastructures to make the Internet, but also service providers, that monetize their services on it. In this context, it is important to ensure fairness in decisions that involve many competing actors, with the goal to reach some kind of global economic optimum.

Indeed, participants of a decentralized network often have to make global decisions based on local interests. For instance, Internet routes span over multiple Autonomous Systems, while they result from local, arbitrary decisions. Many fields related to game theory have been proposed in order to have a better understanding of distributed decision-related issues ([1]). However, to the best of our knowledge, one of these fields remain mostly unexploited: voting systems, that allow competing entities to decide among different options.

The main goal of this paper is to investigate the usage of voting systems within the Internet economic ecosystem on a given use-case. In particular, we focus on the question of manipulability (i.e., vulnerability to tactical voting), which is crucial in a decision-making context: How hard is it for participants to change the decision by lying about their own interests? It is known that except for a few degenerated cases, all voting systems can be cheated ([2], [3]). However, the practical performance of existing algorithms is mostly limited to high level properties. This study applies the voting system framework to a networking use-case, and quantifies manipulability and its effects on such a scenario. This paper is organized as follows. Section II gives a generic, self-contained, framework of voting systems, with a focus on manipulability aspects. Section III defines and models the networking use-case of our paper: path establishment in multi-operator networks. Lastly, in Section IV, we analyze the results we obtain, illustrating how and why voting systems can be very interesting for practical use in an economic ecosystem such as the one of multi-operator networks.

II. Generalities on voting systems

This section presents the general framework of voting systems that we use in this study. It first describes the notion of elector's preferences, the general definition of voting systems and some examples that we use in our study. Manipulability criteria are then given, as well as some general results that are known about manipulability of voting systems.

A. Preferences

We note $\mathcal{E} = \{e_i\}_{i=1}^{n}$ the set of electors and $\mathcal{C} = \{c_j\}_{j=1}^{m}$ the set of candidates. Following Von Neumann-Morgenstern approach [4], the preferences of an elector e_i are represented by a utility vector $\mathbf{U}_i = (u_{i,j})_{j=1}^{m}$, with:

- $u_{i,j} > u_{i,k}$ means that e_i strictly prefers c_j to c_k,
- $u_{i,j} = u_{i,k}$ means that e_i equally values c_j and c_k,
- For $\lambda \in [0,1]$, $u_{i,k} > \lambda u_{i,j} + (1-\lambda)u_{i,l}$ means that e_i strictly prefers c_k to a lottery where c_j is chosen with probability λ and c_l with probability $1-\lambda$.

Thus, $\mathbf{U}_i$ does not only represent elector e_i's order of preferences but also the compared strengths of her preferences.

We note $\mathcal{U} = \mathbb{R}^m$ the space of the possible preference profiles for an elector. Knowing elector e_i's utility vector $\mathbf{U}_i \in \mathcal{U}$, we can deduct e_i's order of preferences with the canonical surjection σ from the utility space $\mathcal{U}$ to the set of weak orders over the candidates.

B. Voting system definition

A voting system allows competing entities (the electors) to select one option among several ones (the candidates). An elector e_i can choose a deterministic *strategy*, also called *ballot*, from a strategy set S_i. This would be called a *pure strategy* in usual game theory. Once each elector has chosen

a strategy, the voting system picks out a candidate, applying a voting rule function $f : S_1 \times \ldots \times S_n \to \mathcal{C}$. In order to link utilities to voting systems, we also need to define sincerity. For each i, we assume that there is a function $g_i : \mathcal{U} \to S_i$ that describes the "spirit" of the voting system: If your preferences are $\mathbf{U}_i$, you are supposed to vote $g_i(\mathbf{U}_i)$ and doing so will be considered *sincere*. Actually, most voting systems admit a simple, canonical, sincerity function:

- When S_i is equal to $\mathcal{U}$, we use $g_i = Id$;
- When S_i is the set of weak orders over the candidates, we use $g_i = \sigma$, the canonical surjection.

Note that a great part of the literature limits *voting systems* to the case where S_i is limited to the set of strict orders over the candidates [3], [5], [6], [7], [8].

C. Examples of voting systems

Voting systems are a huge family, and studying all of them is far beyond the scope of this paper. For reasons that will be detailed in Section III, we focus on the following systems:

1) Range Voting (RV): Each elector e_i communicates a vector of marks, $\mathbf{S}_i = (s_{i,j})_{j=1}^m$. The candidates c_j who maximizes $\sum_{i=1}^n s_{i,j}$ wins. We consider that e_i is sincere when she communicates her utility vector $\mathbf{U}_i$ as vector of marks.

2) Exhaustive Ballot (EB): The protocol proceeds through a series of $m-1$ elimination rounds. At the beginning of each round, each elector communicates her preferred candidate among the remaining ones, the candidate with least votes is eliminated. The winner is the last remaining candidate.

3) Instant Runoff Voting (IRV): Each elector communicates her order of preferences once and for all. Then, the protocol emulates a series of $m-1$ elimination rounds. In each round, the candidate who is ranked first by the least number of electors is eliminated and each ballot in her favor is transferred *automatically* to the best ranked remaining candidate.

It is straightforward that IRV is indeed an emulation of EB. In particular, if all electors vote sincerely, both systems give the same result. The main difference is that in EB, an elector can contradict herself from one round to another, like changing the candidate she says she prefers between two rounds even if she has not been eliminated.

D. Manipulability criteria

We need properties that describe the manipulability of voting systems by tactical voting. We consider a given set of electors, candidates, utilities, and a voting system with an associated sincerity function. $v \in \mathcal{C}$ denotes the candidate who is elected when all electors vote sincerely. The following definitions hold.

1) Coalition manipulability (CM): A subset of electors, by casting insincere ballots, can make the result of the vote strictly better from their point of view. That is, there exists a challenging candidate $c \in \mathcal{C} \setminus \{v\}$ such that the electors strictly preferring c to v can cast their ballots so that c gets elected, assuming that other electors don't change their own ballots.

2) Trivial coalition manipulability (TM): There is a challenging candidate $c \in \mathcal{C} \setminus \{v\}$ that gets elected if all electors preferring c to v use their *trivial strategy*. An elector e uses her trivial strategy for candidate c against candidate v when she pretends that all candidates in $\mathcal{C} \setminus \{v, c\}$ have a low utility (while preserving their respective positions), that v has an even lower utility and that c has an infinite utility. For instance, the trivial strategy in RV consists of favoring the candidate c and to put the other at a disadvantage. In IRV or EB, the trivial strategy means putting c at the top and v at the bottom, while keeping one's sincere and personal order of preferences over the other candidates.

CM is about being able to manipulate the system, no matter how hard it is. It answers the question: Can the system be manipulated by omnipotent manipulators, which have a complete knowledge of the system and can cast coordinated ballots? TM, on the other hand, answers the question of an almost-zero-knowledge, decentralized, manipulation: Once the challenger c is chosen, all electors of the coalition can cast their own ballot independently.

E. Gibbard-Satterthwaite theorem

We say that an elector e_i is a *dictator* iff, for any possible outcome $c \in f(S_1 \times \ldots S_n)$, elector e_i can cast a ballot so that the elected candidate will be c, whatever other electors vote. We say that a voting system is *dictatorial* if there is a dictator.

It is *manipulable* iff there exists a situation such that at least one single elector can benefit from casting an insincere ballot (which implies CM).

The following theorem was proved by Satterthwaite [3] for voting systems based on permutations and by Gibbard [2] in the general framework of strategies:

Every non-manipulable voting system with at least three possible outcomes is dictatorial.

III. Multi-operator networking case

In this section, we define the multi-operator end-to-end path establishment problem. We first explain how voting systems can be applied to solve it, and how the set of candidate paths is selected. Then we define the cost, gain and utility modeling, which quantifies the willingness of operators to carry or not a given path. Finally we explain how the voting systems are applied and how the manipulability is calculated.

A. Voting systems for multi-operator path establishment

The problem to solve is the following:

In multi-operator networks, for given ingress and egress, which end-to-end path should be selected if one takes operators' preferences into account?

Figure 1 represents an example[1] of interconnection of multi-operator networks, with one operator network per European country, interconnected by some geographical neighborhood.

[1]This example is not representative of a real multi-operator network interconnection. For instance, such a flat interconnection topology differs from the historical BGP hierarchical topology of Internet. Nevertheless, our goal here is to give a simple model where costs derive from some sort of underlying metric, so a geographical basis is a natural choice. Note that some recent studies show an evolution towards flatter topologies ([9], [10]).

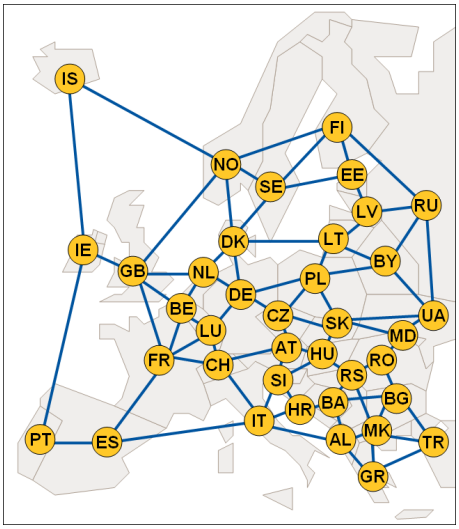

Fig. 1. Multi-operator network interconnection example

We apply the voting systems using the following modeling:

- The electors are the operators.
- The candidates are the feasible routing paths for a given demand or a subset of them.
- The preferences of each operator is represented by a utility vector, from which a ballot is defined (cf §II).
- The election result is the selected path for the client request. Any voting system can be used. Here we suppose there is a trusted and independent entity, called *supervisor*, in charge of the election process. The supervisor collects the ballots from all the operators and processes the voting algorithm to decide the winner path (but other options are possible: e.g., the operators cooperatively participating to the whole voting scheme).

Some operators may have some knowledge about the utilities of their competitors:

- by public knowledge about the utilities;
- by cooperation between some operators (coalition);
- by inference on previous votes (learning);
- by spying or information interception.

So the operators may use that knowledge to lie about their own preferences, in order to improve their benefits, maybe against the global interest. Studying the manipulability of the voting systems is thus important in that context.

B. Candidate paths

In our study case, we consider the interconnection of $n = 38$ operator networks of figure 1.

A *demand* is a request for end-to-end connections, with the only constraint that the connection must start at a given operator network (ingress) and end at another one (egress).

For each demand, if one considers the whole set of possible paths without loop, the size of that set increases exponentially with the number of nodes in the network. So we apply the following reasonable rule to limit the number of proposed paths (only the interconnection topology is known):

- The supervisor fixes a minimal number of candidate paths m_{min}, and an initial threshold δ_{min}.
- It knows the link topology (but not the costs!). So it computes the minimal number of hops h_{min} to satisfy the demand.
- It takes all paths without loop with a number of hops less than or equal to $h_{min} + \delta_h$, where δ_h is the smallest threshold that is greater than or equal to δ_{min} and that selects at least m_{min} candidates.

It is important to note that:

- The candidate paths are fully determined by the demand and the parameters (m_{min}, δ_{min}). Their exact number depends on the demand.
- For each candidate path c_j, only a subset $\mathcal{E}_j$ of the operators $e_i \in \mathcal{E}$ are concerned by the candidate path.
- Among the operators, some may be concerned by only a subset of candidates.

For numerical evaluations, we use two limitation options: $(m_{min}, \delta_{min}) \in \{(5, 0); (10, 1)\}$. With the above rules, the first option gives on average 9.94 candidate paths per request ($\min = 5$, $\max = 43$), and the second option gives on average 21.25 paths per request ($\min = 10$, $\max = 127$).

C. Multi-operator cost and gain modeling

We define the utilities of each operator as the difference between his gains and costs for each possible path. Various cost and gain models could be defined, but the main goal of this study is to evaluate the manipulability of voting system in that context, so we used the following simplified ones [2]:

1) Cost: The cost for an operator e_i in $\mathcal{E}_j$ to carry the path c_j is noted $\alpha_{i,j}$. We define $\alpha_{i,j}$ as the sum of half the cost of the incoming interconnection link (null for the ingress operator) and half the cost of the outgoing interconnection link (null for the egress operator). For the cost of the interconnection link between two adjacent operators a and b, we choose a linear function $C_0 + d_{a,b}/d_0$ of the distance $d_{a,b}$ between a and b (a and b being the capital cities in our multi-network example). In our numerical study, we considered three cost options: dominated by the constant cost C_0 ($C_0 = 1$ and $d_0 = 100 \times \max(d_{a,b})$), purely linear ($C_0 = 0$, $d_0 = \max(d_{a,b})$), and intermediate ($C_0 = 1$, $d_0 = \max(d_{a,b})/3$).

2) Gain: Concerning the gain for the operators, we consider that the client pays a fixed amount A for a given demand (flat fare). If path c_j is selected, this amount A is equally distributed between the concerned operators $e_i \in \mathcal{E}_j$. We fixed the value A in such a way that, when considering the least cost path for each request, the average global revenue represents 140% of the average global cost (i.e., benefit of 40%).

3) Utility: The sincere utility value for operator e_i to carry the candidate path c_j can be defined as the net income (positive or negative) for the operator if this candidate path is selected:

- $u_{i,j} = \dfrac{A}{card(\mathcal{E}_j)} - \alpha_{i,j}$ if $e_i \in \mathcal{E}_j$,
- $u_{i,j} = 0$ if $e_i \notin \mathcal{E}_j$ (the operator is then *indifferent*).

[2] If one would like to test more accurate cost or gain models, this is feasible. In our opinion, they should bring only minor changes on the major trends observed, but it may be interesting to test them in future works.

4) Global income: The global income of a given operator depends on the distribution of the demands. In our study, we consider that the demands are uniformly distributed among each pair of different ingress and egress operators.

D. Voting systems for the multi-operator study case

1) RV: With the utilities previously defined, the most natural choice for a voting system is the one that maximizes the global revenue, which is the sum of the utilities of all operators contributing to the selected path. This corresponds exactly to RV defined in Subsection II-C, so RV is the natural reference system. In details:

- The operators give their utilities to the supervisor;
- The supervisor sums the marks for each candidate path, and select the path with the maximal value.

One of the drawback with RV is that the operators are required to give all the information about their cost to the election supervisor. Even if it is a trusted and independent entity, they may wish to avoid giving this kind of information.

2) IRV: Numerous other voting systems may be applied. Previous works (see for instance [11] or [12]) suggest that IRV belongs to the least manipulable voting systems known, so we used IRV as the second voting system for our use case[3]. Path selection with *IRV* works as follows:

- Each elector (operator) gives the order of preferences on the candidate paths that pass through her network. and if she likes it (financial gains) or dislikes it (financial losses) in order to be able to place them relatively to the candidate paths for which she is not concerned (forced indifference for them).
- The supervisor processes the *IRV* mechanisms as defined in Subsection II-C, with the following rule with equal preferences: In each round of *IRV*, for a given operator (elector), the vote is equally divided between the candidates with the highest preference. In each round, the election supervisor eliminates the candidate path with the least votes (the number of votes may be not an integer value in this case). In case of ties, we choose to eliminate the candidate path with the lowest index [4].

E. Manipulability algorithms

First, for both voting systems, the manipulation of the vote by a operator is limited to the candidate paths concerning this operator: The operator cannot pretend to like or dislike a candidate path to which he must be indifferent.

1) RV: When electors preferring c to v try to make c win, their best strategy is obviously the trivial one: Give the maximum mark to c and the minimum mark to the other candidates, except the marks of the paths for which they are indifferent, which must remain unchanged. Because of that, TM and CM are equivalent. The maximum mark is set to $+A$ and the minimal mark to $-A$ in our study.

[3]Actually, we also tried other voting systems in experiments not presented here, and verified that IRV was the most promising candidate with respect to manipulability for our concerns.

[4]It could be done randomly, but the impact is marginal in our problem.

2) IRV: Finding out whether there is a way to manipulate with IRV is much more difficult. In fact, the problem is known to be NP-hard [13]. So we use simple methods that test the manipulability (or non-manipulability) of a given demand. When the tests are inconclusive for a given demand, we can only answer *maybe*, but when considering all demands, this allows to give lower and upper bounds for manipulability.

In order to prove that IRV is manipulable, we just try trivial manipulation. This gives a lower bound of the manipulability.

In order to prove that IRV is not manipulable, our algorithm is an adaptation of [14]. The idea is to use a variation of the voting system that gives more power to manipulators and for which CM can be exactly computed. If the altered system cannot be manipulated, IRV cannot either. In details:

- At each round, we authorize the manipulators to change their vote. This shifts the voting paradigm from IRV to EB. The interest is that eliminating candidate a then b or b then a lead to the same situation, whatever the manipulators have done to get there. This permits an iterative approach instead of a recursive one.
- At each round, each manipulator can share her vote between several candidates, even non equally, for instance $\frac{1}{3}$ vote for one candidate and $\frac{2}{3}$ vote for another one. This allows a water-filling approach and avoid Knapsack-like issues.
- We authorize electors to lie even about the paths they don't belong to. So all manipulators are symmetric in right and we don't have to manage them individually.

With these modifications, we can manage the group of manipulators globally: At each round, we can divide their votes as we like between the candidates, in order to eliminate the candidate we want to.

When manipulation is impossible with these adapted rules, it is also impossible with the rules given in III-D: This provides an upper bound for IRV manipulability.

As we will see in next section, lower and upper bounds tend to be close to each other, so we get a good estimate of manipulability.

Remark: Actually, the previous algorithm is still very costly (virtually in 2^m, where m is the number of candidates). When there are more than 25 candidates and that trivial manipulation does not work, we don't try to prove the impossibility and we directly consider the test inconclusive.

IV. Multi-operator networking results

In this section, we analyze the results we obtained on the multi-operator network case described in Section III.

For both *RV* and *IRV*, we first zoom on one specific scenario that will serve as reference, and then we extend the results for several parameters. We observe manipulability and economic efficiency, with sincere and insincere ballots.

A. Reference scenario

Our reference scenario is the case with $(m_{min}, \delta_{min}) = (5, 0)$ (on average 9.94 candidate paths per demand) and the intermediate link cost model ($C_0 = 1$, $d_0 = \max(d_{a,b})/3$). For this scenario, we measure:

Voting system	RV	IRV
Sincere efficiency	100%	95 %
Manipulability	96%	$<$ 20% (TM: 18%)
Insincere efficiency (average)	37%	90%
Insincere efficiency (worst case)	-75%	89%

TABLE I
MAIN RESULTS FOR THE REFERENCE SCENARIO
$((m_{min}, \delta_{min}) = (5, 0), (C_0, d_0) = (1, \max(d_{a,b})/3)$)

- *Sincere efficiency*: global net income, in percentage of the optimal global net income that assumes sincere voting for all operators.
- *Manipulability*: proportion of the demands that are CM. For range voting, this can be exactly computed because CM is TM. For IRV, we use TM as a lower bound and the EB variant to give an upper bound (cf §III-E).
- *Insincere efficiency*: efficiency of the system when manipulations are allowed. For a given demand, several manipulations can occur. We can measure the average insincere efficiency, if one manipulation is selected at random, or the worst case situation if for each manipulable demand, one chooses the manipulation that minimizes the global net income (remember that by definition, the net income of the manipulators will be increased, though).

The results are presented in Table I. IRV presents many advantages over the more natural RV selection:

- For sincere voting, i.e., with no manipulation, RV gives the economic optimum as expected. While not having global optimum as a target, IRV still manages to achieve about 95% of the optimum. This slight efficiency decrease can be seen as the price to pay for robustness (see below).
- Almost all demands (96%) can be manipulated if ranged voting is used, against less than 20% for IRV.
- Considering the impact on efficiency, the degradation is very high for RV: It becomes only 37% of the optimum on average, down to -75% in the worst case scenario. On the other hand, IRV maintains a robust 90% on average (89% in the worst case scenario).

So one should retain that IRV is much less subject to manipulation than range voting, and that even when a manipulation exists, its impact on the global welfare is bearable.

B. Impact of parameters

We proposed in Section III two path-limitation parameters and three cost models, so we have six possible configurations. Results for these six configurations follow.

1) Manipulability: Figure 2 indicates the manipulability of the considered scenarios.

For both voting systems, the most numerous the path candidates are, the most manipulable they are. This is somehow expected, as more candidates mean more possible challengers for manipulation. So the supervisor should limit the number of proposed candidates to decrease the manipulability, while keeping enough candidates to allow a fair path selection. Concerning the link cost model, we observe that the flatter it is, the lower the manipulability (for both voting systems). But the most interesting result is that the manipulability of IRV stays much lower than the one of RV in all cases. While RV

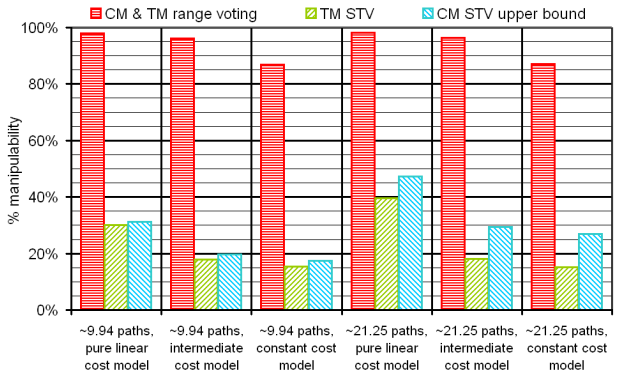

Fig. 2. Manipulability of IRV vs. Range Voting

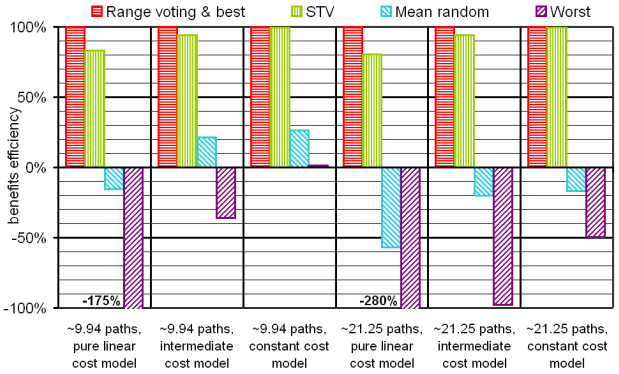

Fig. 3. Efficiency of IRV vs. Range Voting for sincere preferences

manipulation is always higher than 85%, IRV manipulation is about 30% or less except for one scenario (between 40% and 48% for the highest number of candidate paths with purely linear link cost model).

Note that for IRV, the lower and upper bounds for manipulability are relatively close. The differences are higher when the number of candidate paths increases, but this is due to the way we calculate the upper-bound (see Section III), skipping the evaluation for demands with too many candidate paths.

2) Sincere economic efficiency: Figure 3 gives the sincere efficiency of each voting system for the considered scenarios. By definition, RV gives 100% economic efficiency. But IRV gives an economic efficiency close to this optimum. For both low or high number of candidate paths, this efficiency is about 80% for purely linear link cost model, about 95% for intermediate link cost model, and more than 99% for constant link cost model. This confirms that, even if IRV may give a path which is not the optimal one for global economic benefits, the selected path is quite close to the optimal choice.

For completeness, Figure 3 also indicates the efficiency obtained when the path is chosen at random among the candidates, and when the worst candidate is chosen.

3) Insincere efficiency: To complete our study, Figure 4 displays the insincere efficiencies. For all scenarios, one observes a huge gain by choosing IRV instead of RV. For IRV, the economic efficiency is only slightly degraded compared to sincere preferences, and the worst case is never far from the average case. For RV, the economic efficiency is largely

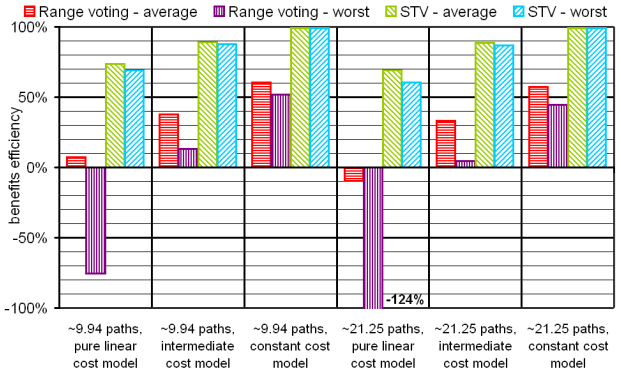

Fig. 4. Efficiency of IRV vs. Range Voting for manipulated preferences

degraded compared to sincere utilities: While the economic efficiency with sincere utilities is 100%, when considering manipulations, it falls to: about 0% on average and large negative values on worst case for the purely linear link cost model; about 35% on average and about 8% on worst case for the intermediate link cost model; about 60% on average and about 50% on worst case for the constant link cost model.

All these results confirm that IRV is really safer to preserve the economic benefits than RV.

V. Conclusion

Existing theoretical results show that, except for a few degenerated cases, any non-dictatorial system is susceptible, in some scenarios, to be manipulated, even by a single voter. In this paper, we proposed to quantify manipulability in practical scenarios and to measure its effects with respect to global welfare. We focused on the use case of multi-operator end-to-end path establishment. We compared two voting systems:

- *Range Voting*: RV maximizes the global net income of the system, provided that the carriers (the voters) give their sincere utilities on the proposed candidate paths (their own net income).
- *IRV*: Based on weak order preferences with elimination rounds, IRV has a "reputation" to be less manipulable ([11], [12]), which we were able to validate in our study.

In the end, our study highlights the interest of voting systems in the context of Internet economic ecosystems where many competing players are involved, with end-to-end path selection as a practical use case.

We also observed that all voting systems are not equivalent. When carriers cannot be trusted, IRV can largely outperform (in our framework) the economic gain maximization that RV is supposed to achieve:

- IRV manipulability probability can be as low as 20% while, in the same configuration, RV is close to 100%.
- With IRV, the operators do not need to give all the information to process the vote, the preference order with the indifference limit is enough, while for RV, they must give the whole information about their cost.
- Although IRV does not target the economic optimal choice, it remains very close to it. The price of non-manipulability is low (in the range of 5% in our study).
- With manipulations, the degradation on the economic efficiency is limited for IRV, while it is huge for RV.

Future work

The work presented in this paper opens many opportunities for future works:

- Analysis of the sensibility of the various parameters, other cost and revenue models.
- Analysis of other voting systems or proposition of new efficient voting algorithms to apply.
- Deeper theoretical analysis of manipulability in voting systems, using a more generic framework.
- Identification of other use cases in the context of Internet, with other kind of ecosystems.

Indeed, the framework of voting systems can be useful in any situations in which one can identify voters that need to decide among different options. We proved in our study case that, by correctly choosing the voting system, one can limit the manipulability by tactical voting of some coalition of voters and preserve the revenue for the global economic ecosystem.

Acknowledgment

All of the authors are members of the Laboratory of Information, Network and Communication Sciences (LINCS), Paris, France.

References

[1] N. Nisam, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, Sep. 2007.
[2] A. Gibbard, "Manipulation of voting schemes: A general result," *Econometrica*, vol. 41, no. 4, pp. 587–601, July 1973.
[3] M. A. Satterthwaite, "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of Economic Theory*, vol. 10, no. 2, pp. 187–217, April 1975.
[4] J. Von Neumann and O. Morgenstern, *Theory of games and economic behavior*. Princeton University Press, 1944.
[5] S. Barber, "Strategy-proofness and pivotal voters: A direct proof of the Gibbard-Satterthwaite theorem," *International Economic Review*, vol. 24, no. 2, pp. pp. 413–417, 1983.
[6] J.-P. Benot, "The Gibbard-Satterthwaite theorem: a simple proof," *Economics Letters*, vol. 69, no. 3, pp. 319 – 322, 2000.
[7] P. J. Reny, "Arrow's theorem and the Gibbard-Satterthwaite theorem: a unified approach," *Economics Letters*, vol. 70, pp. 99–105, Jan. 2001.
[8] T. Weber, "Alternatives vs. outcomes: A note on the Gibbard-Satterthwaite theorem," University Library of Munich, Germany, MPRA Paper, 2009.
[9] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proceedings of the ACM SIGCOMM 2010 conference on SIGCOMM*, ser. SIGCOMM '10. New York, NY, USA: ACM, 2010, pp. 75–86.
[10] A. Dhamdhere and C. Dovrolis, "The internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proceedings of the 6th International COnference*, ser. Co-NEXT '10. New York, NY, USA: ACM, 2010, pp. 21:1–21:12. [Online]. Available: http://doi.acm.org/10.1145/1921168.1921196
[11] J. R. Chamberlin, J. L. Cohen, and C. H. Coombs, "Social choice observed: Five presidential elections of the american psychological association," *The Journal of Politics*, vol. 46, pp. 479–502, 1984.
[12] T. Walsh, "An empirical study of the manipulability of single transferable voting," in *Proc. of the 19th ECAI, IOS Press*, 2010.
[13] J. J. Bartholdi and J. B. Orlin, "Single transferable vote resists strategic voting," *Social Choice and Welfare*, vol. 8, pp. 341–354, 1991.
[14] T. Coleman and V. Teague, "On the complexity of manipulating elections," in *Proceedings of the thirteenth Australasian symposium on Theory of computing - Volume 65*, ser. CATS '07. Darlinghurst, Australia, Australia: Australian Computer Society, Inc., 2007, pp. 25–33.