Finding Optimal Strategies for Influencing Social Networks in Two Player Games

MAJ Nick Howard, USMA

Dr. Steve Kolitz, Draper Labs

Itai Ashlagi, MIT

Network Science Center
West Point

Problem Statement

- Given constrained resources for influencing a Social Network, how best should one spend it in an adversarial system?

Agents

- Each node in the network is an agent
- 4 classes of agents: Regular, Forceful, Very Forceful, Stubborn (immutable)
- Scalar belief for each agent: $X_i(t) \in [-0.5, 0.5]$

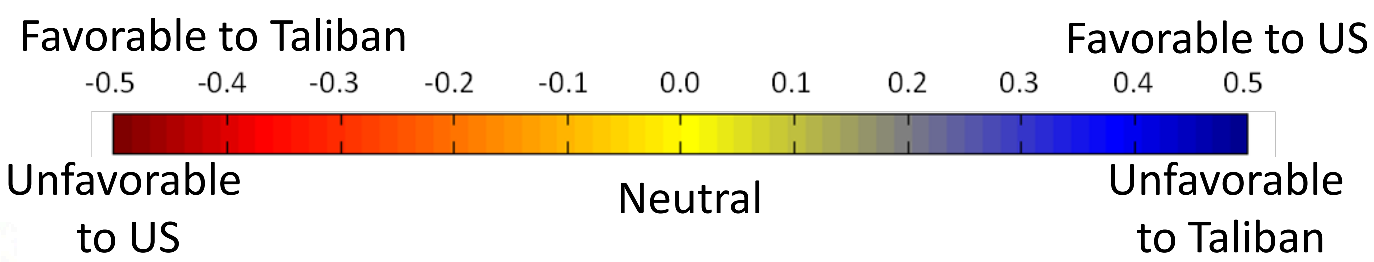

Network and Interactions

- Arcs represent communication
- Stochastic interactions occur with one of three types:
  - Forceful w.p. α_{ij}

$$X_i(t+1) = X_i(t)$$
$$X_j(t+1) = \varepsilon_{ij} \cdot X_j(t) + (1 - \varepsilon_{ij}) \cdot X_i(t) \qquad 0 \le \varepsilon_{ij} \le 0.5$$

  - Averaging w.p. β_{ij}

$$X_i(t+1) = X_j(t+1) = \frac{X_i(t) + X_j(t)}{2}$$

  - No Change w.p. γ_{ij}

$$X_i(t+1) = X_i(t)$$
$$X_j(t+1) = X_j(t)$$

Data Parameterization

- We choose a set of influence parameters that generally make influence flow ‘down’
- However we found through simulation that the solution to our game is highly robust to changes to these parameters.

Interaction Simulation

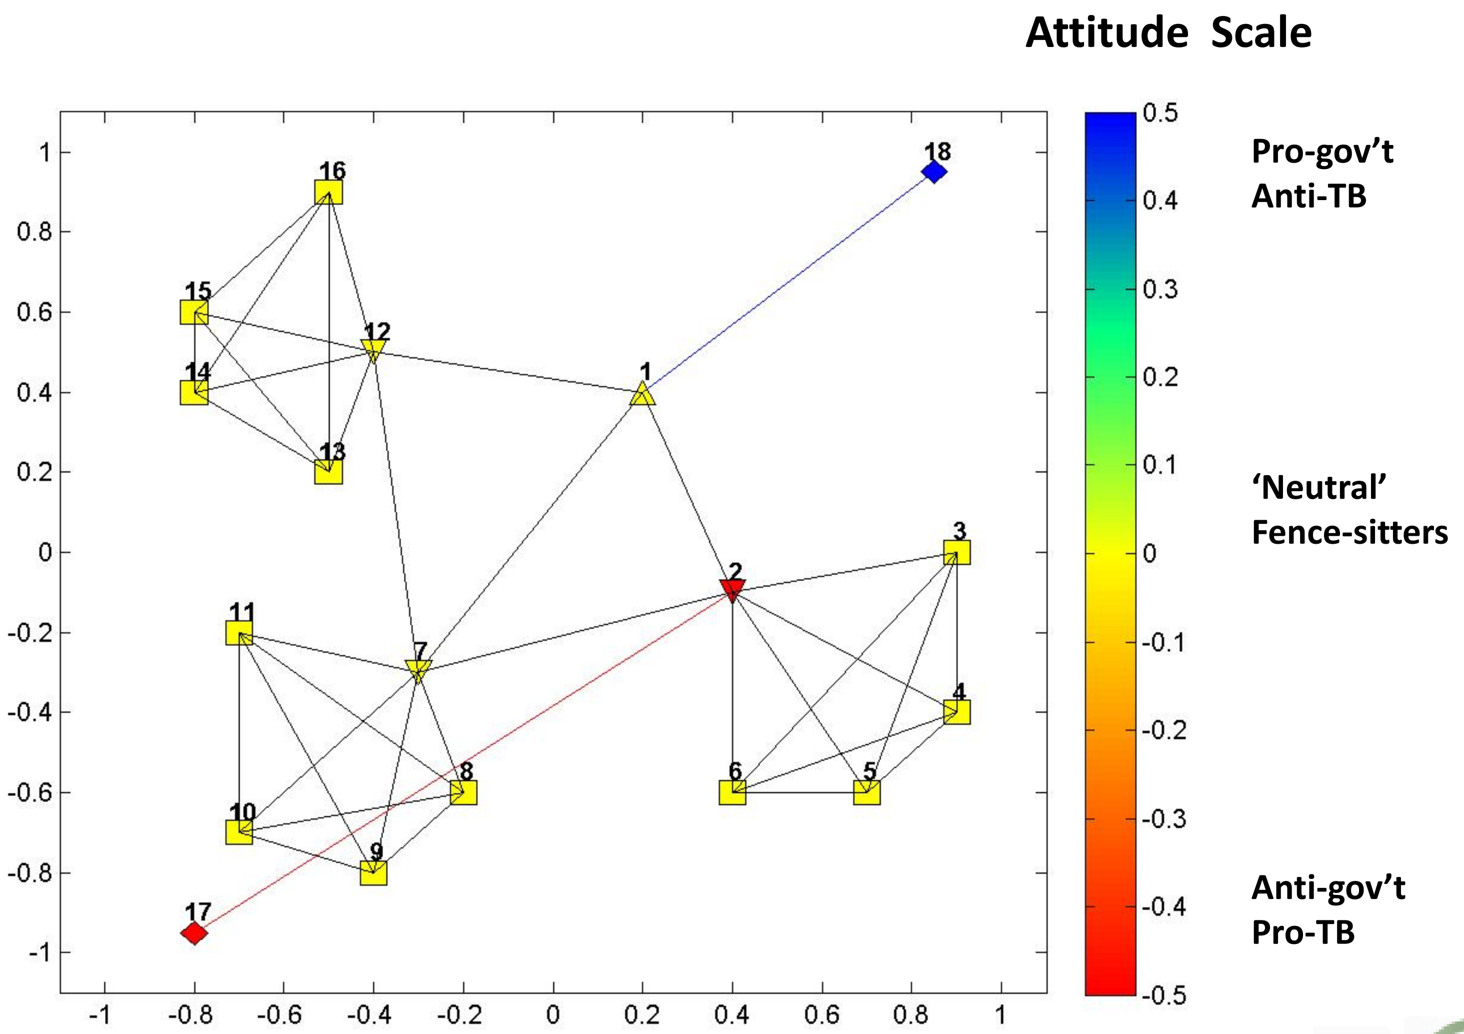

West Point

Interaction Simulation

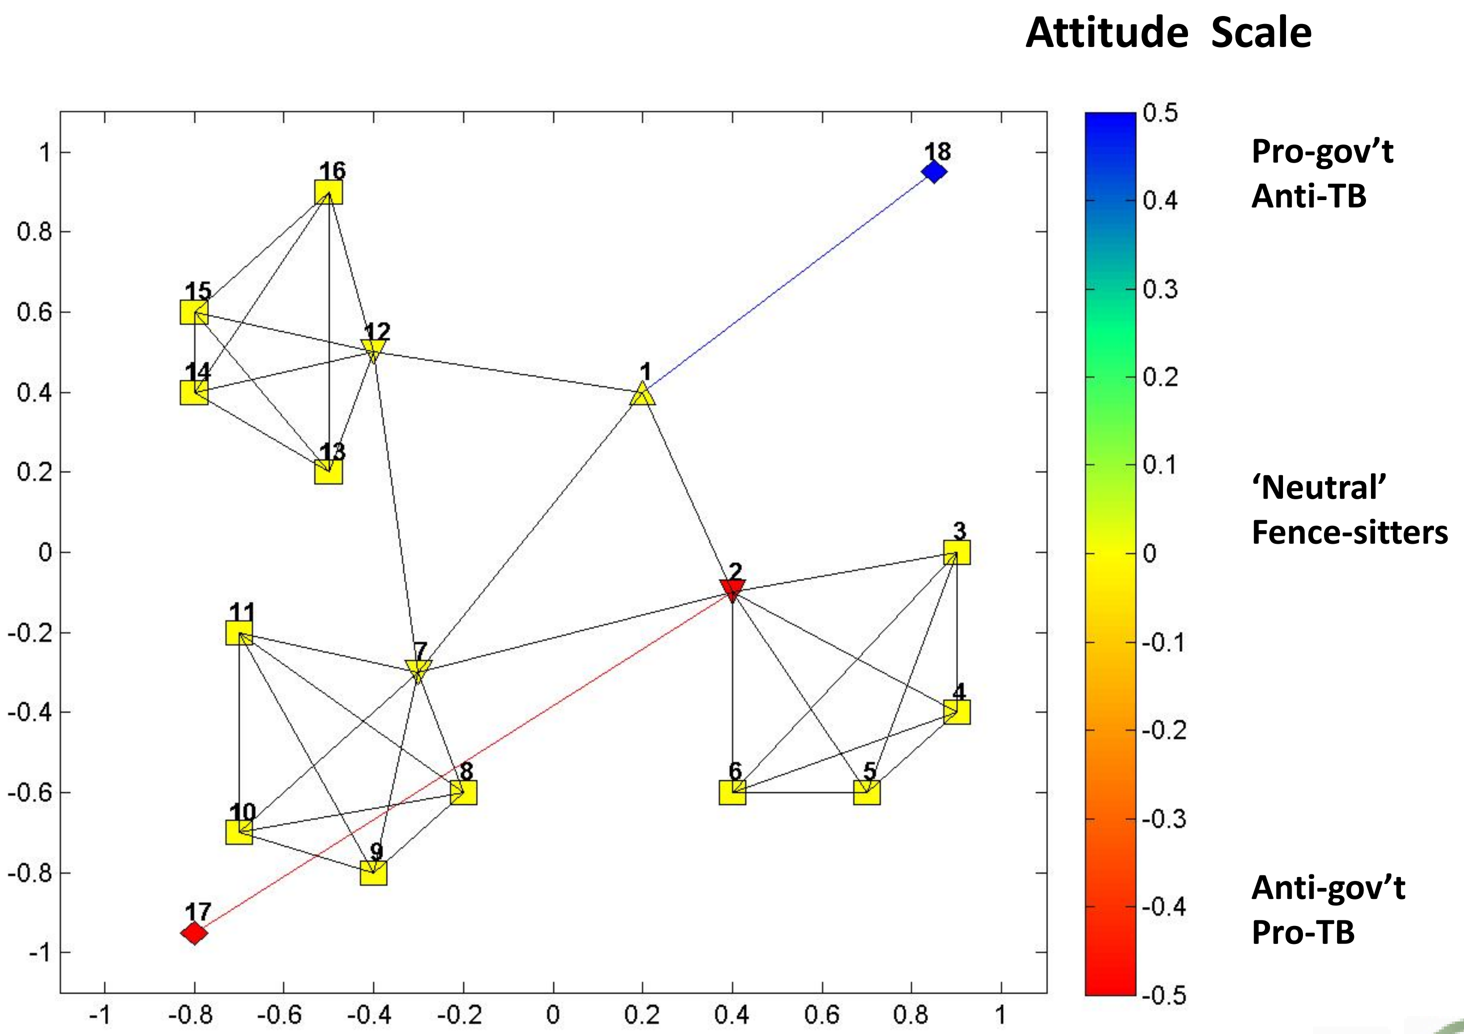

Interaction Simulation

Attitude Scale

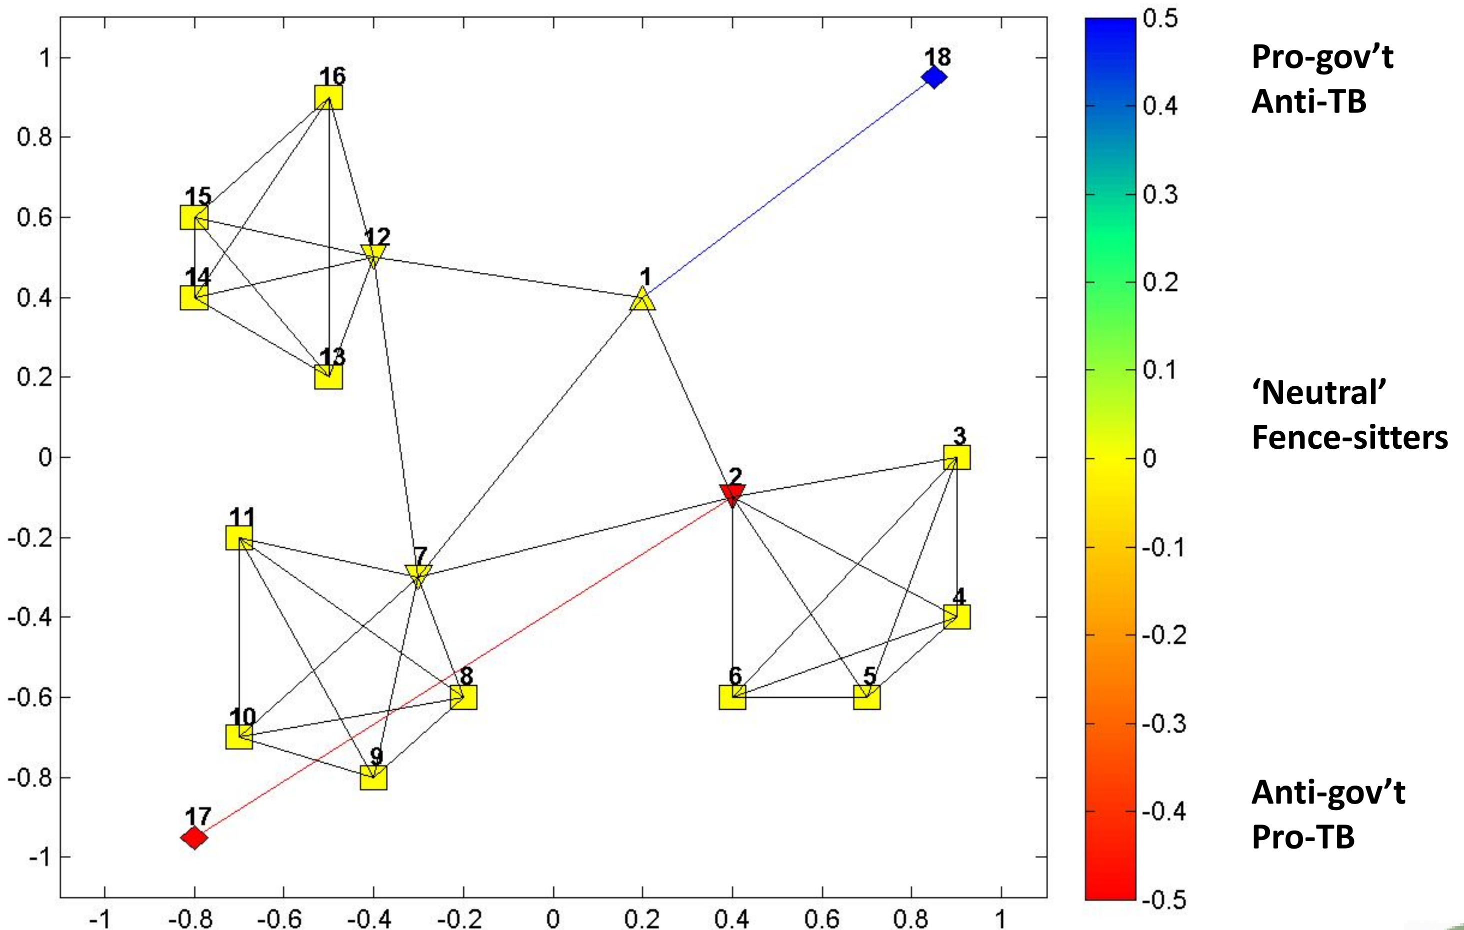

Interaction Simulation

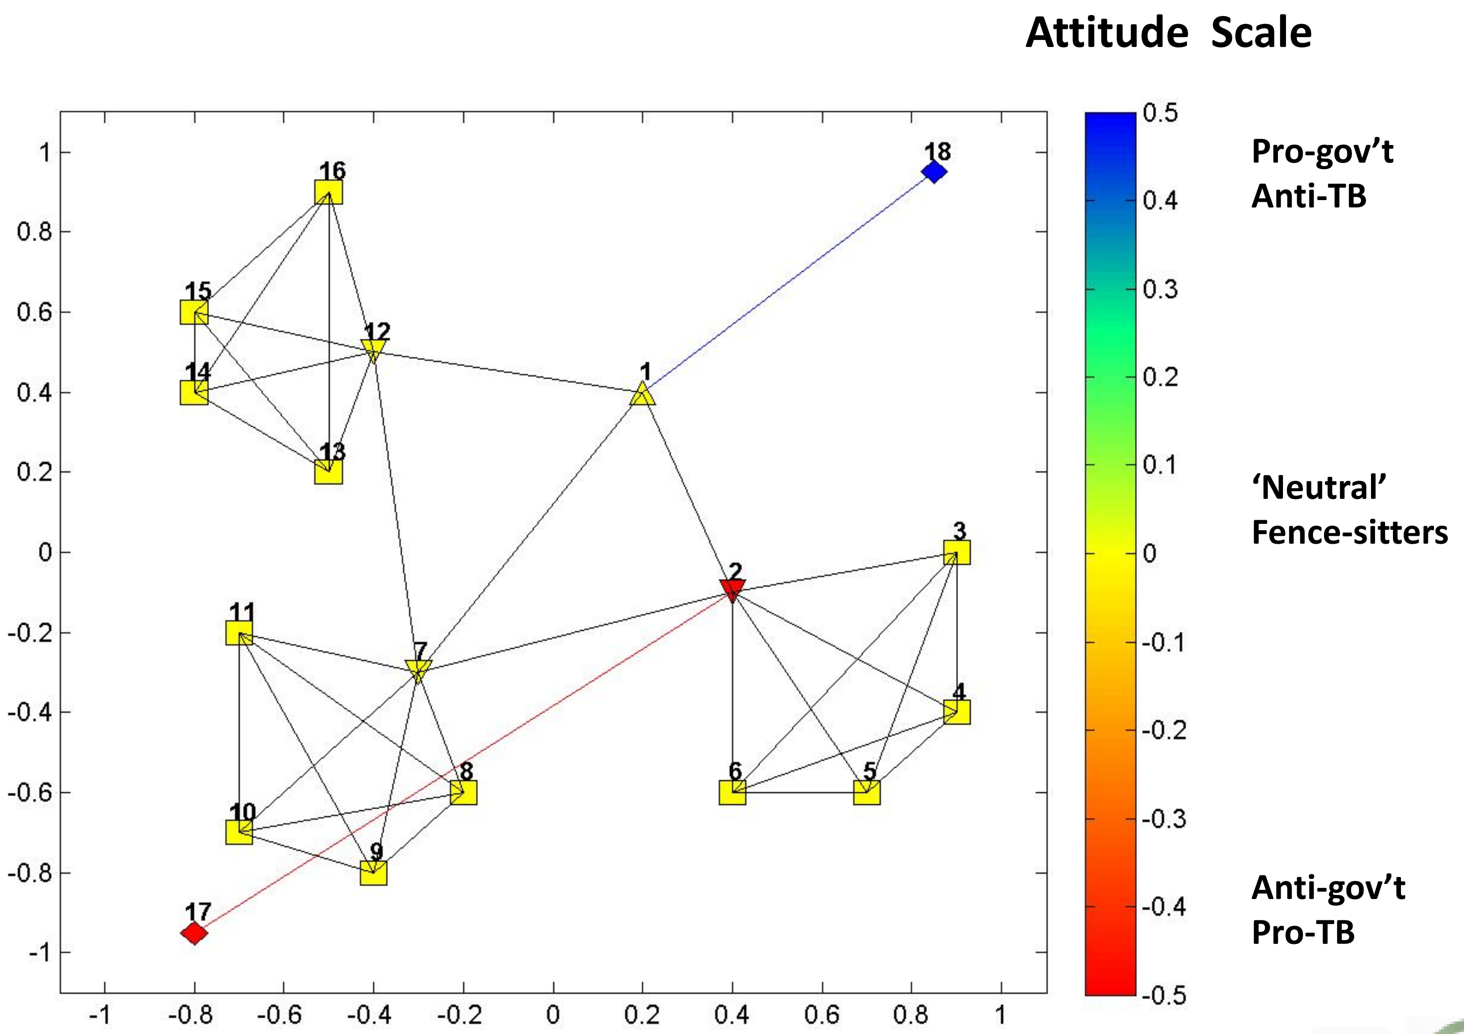

Interaction Simulation

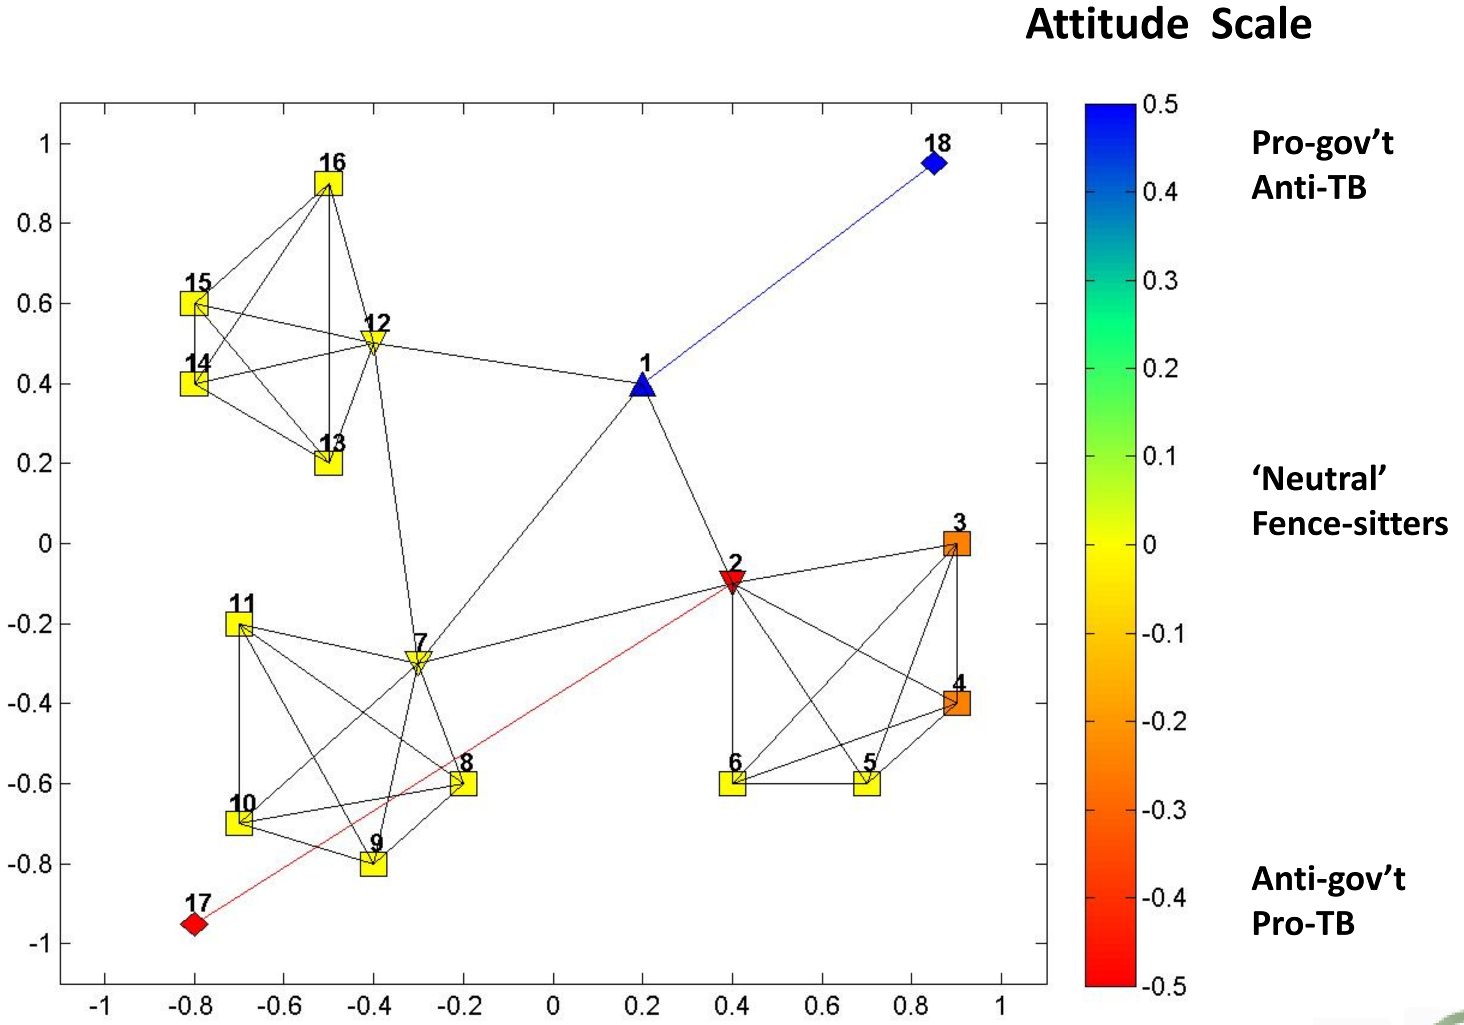

Interaction Simulation

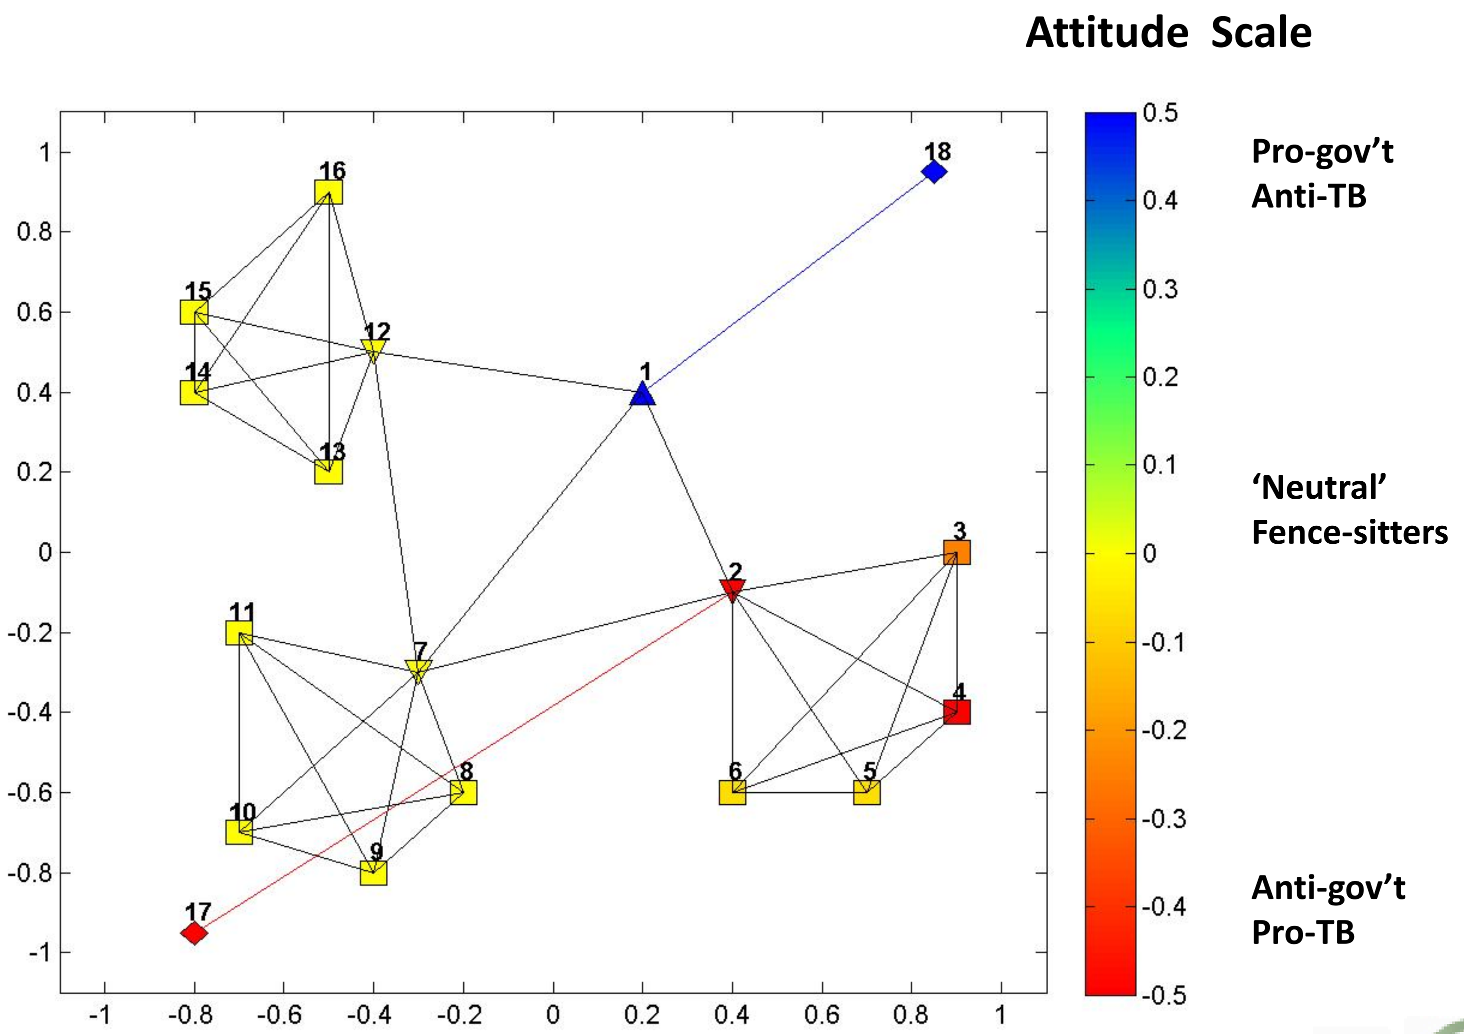

Interaction Simulation

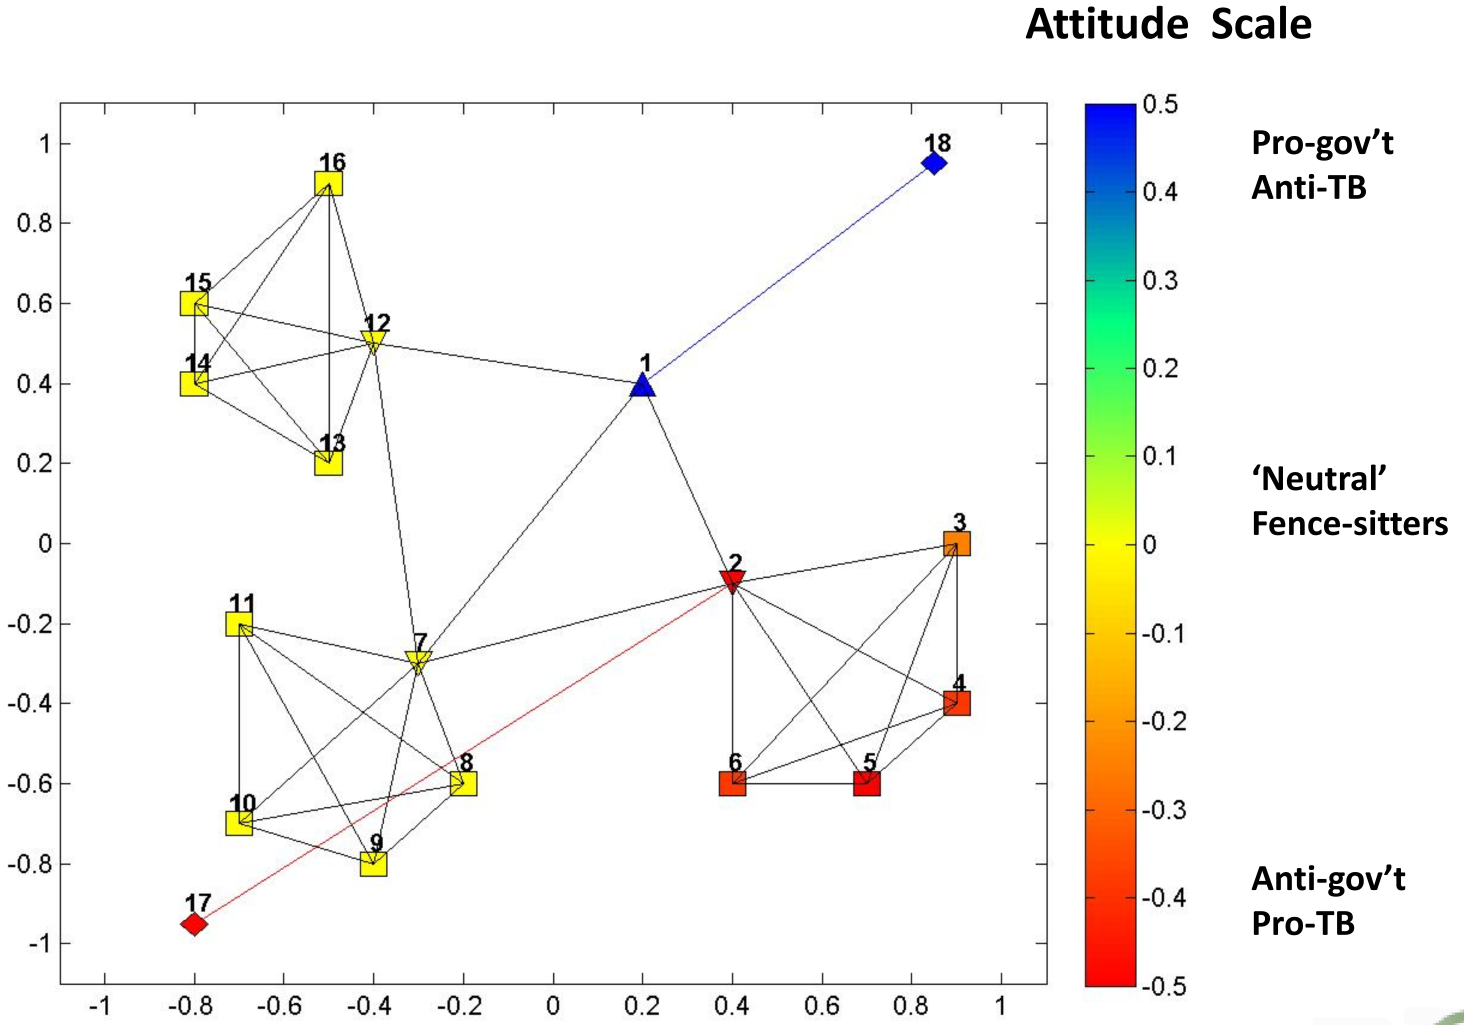

Interaction Simulation

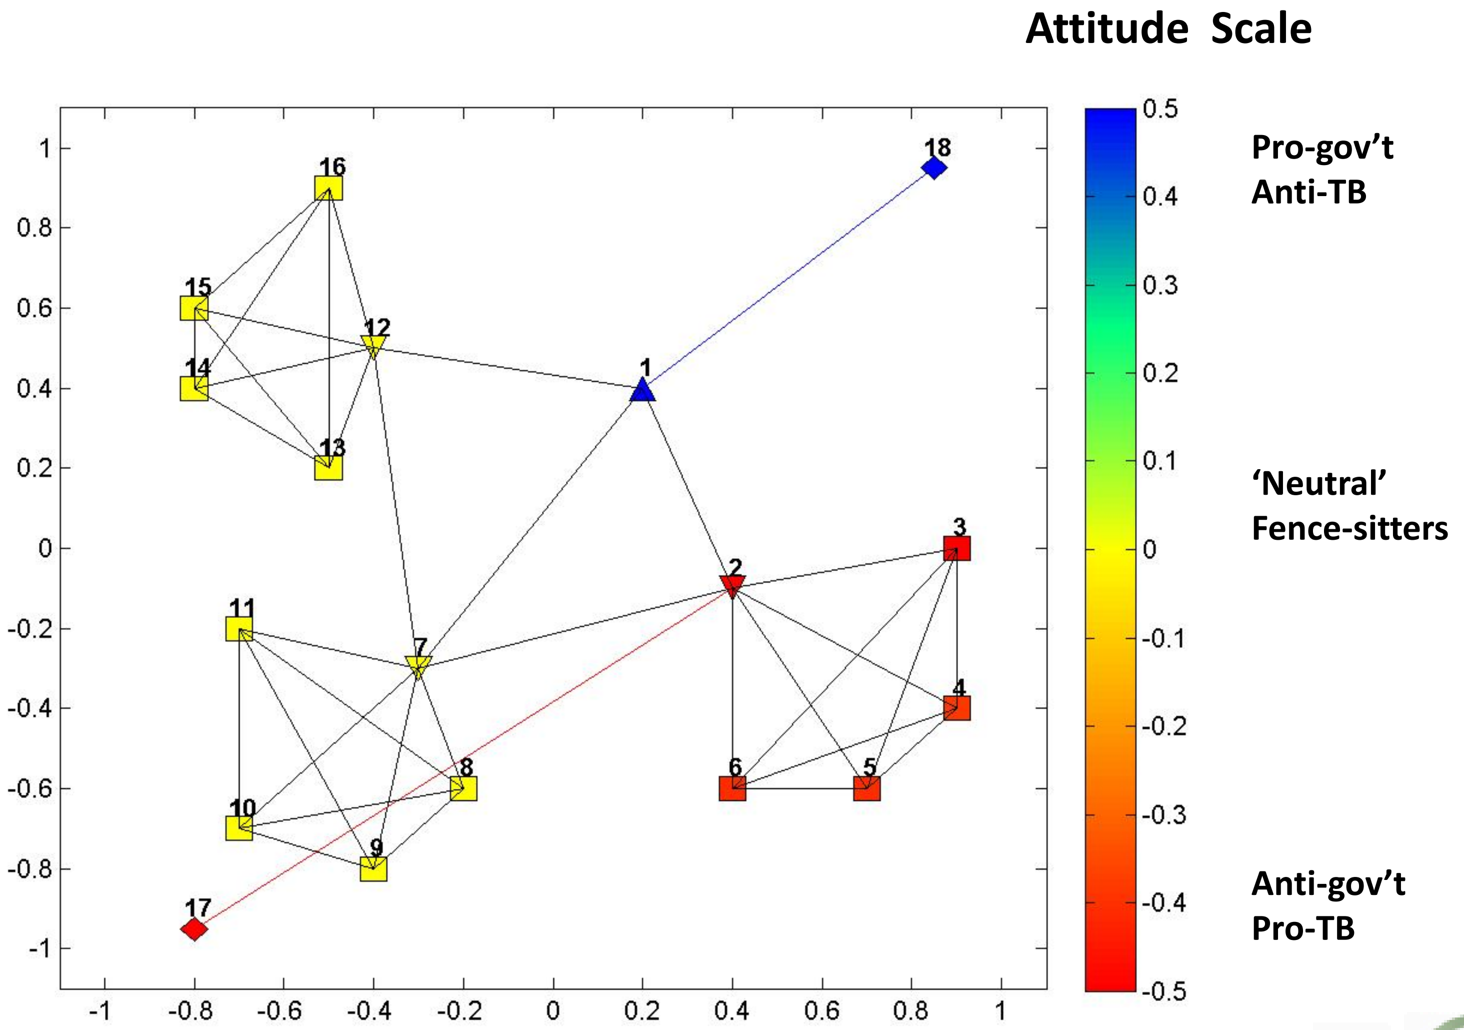

Interaction Simulation

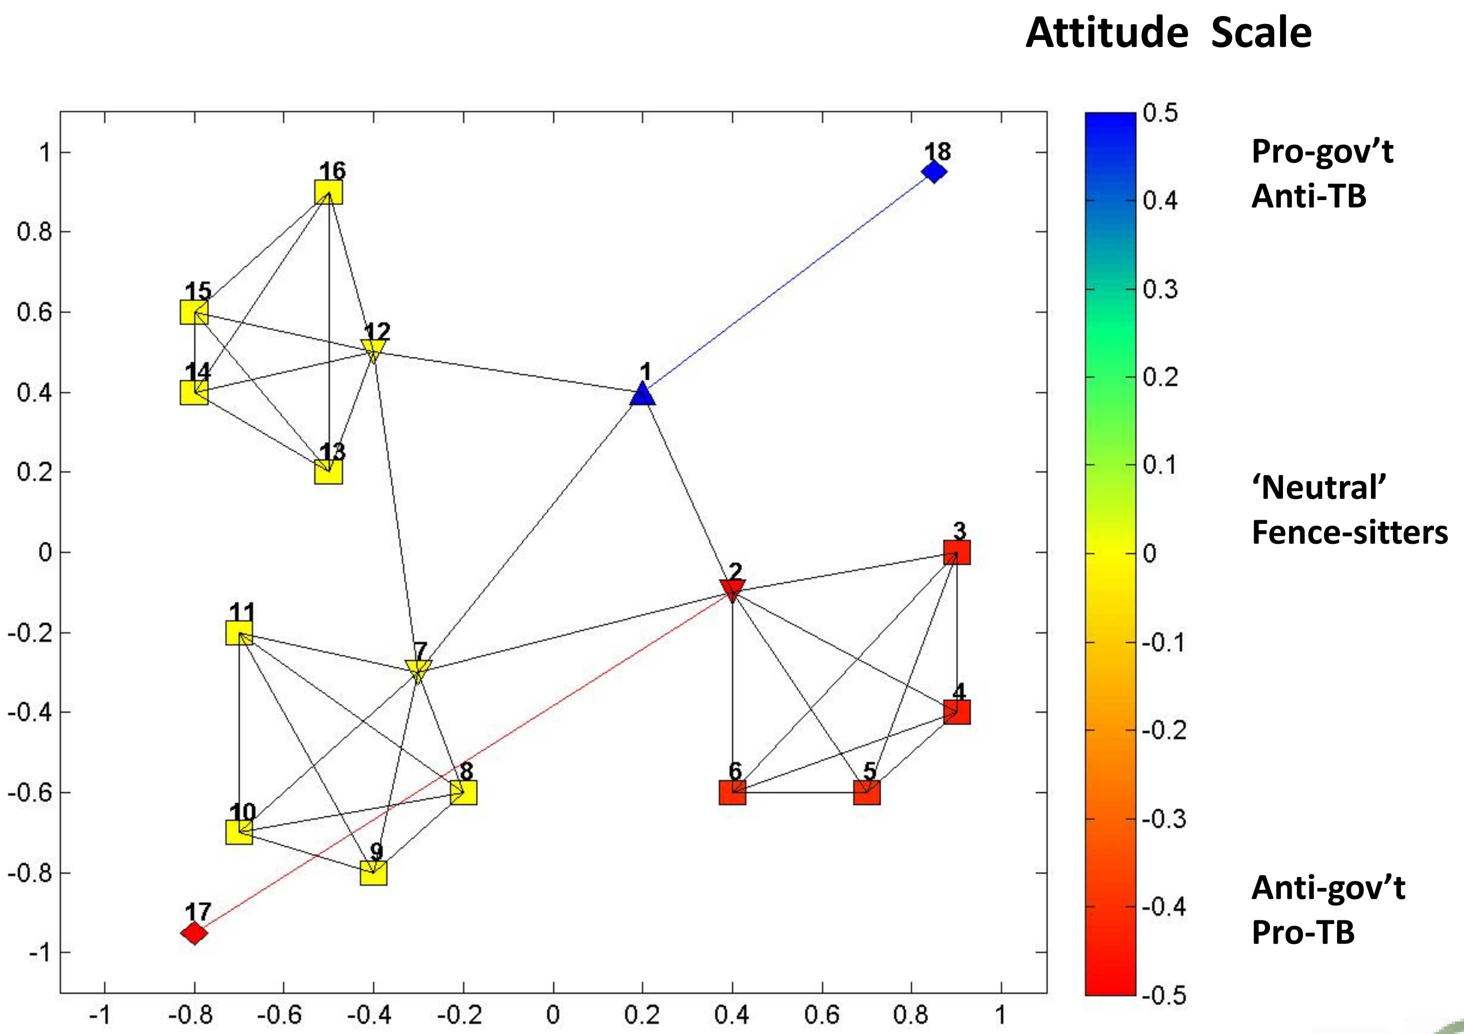

Interaction Simulation

Attitude Scale

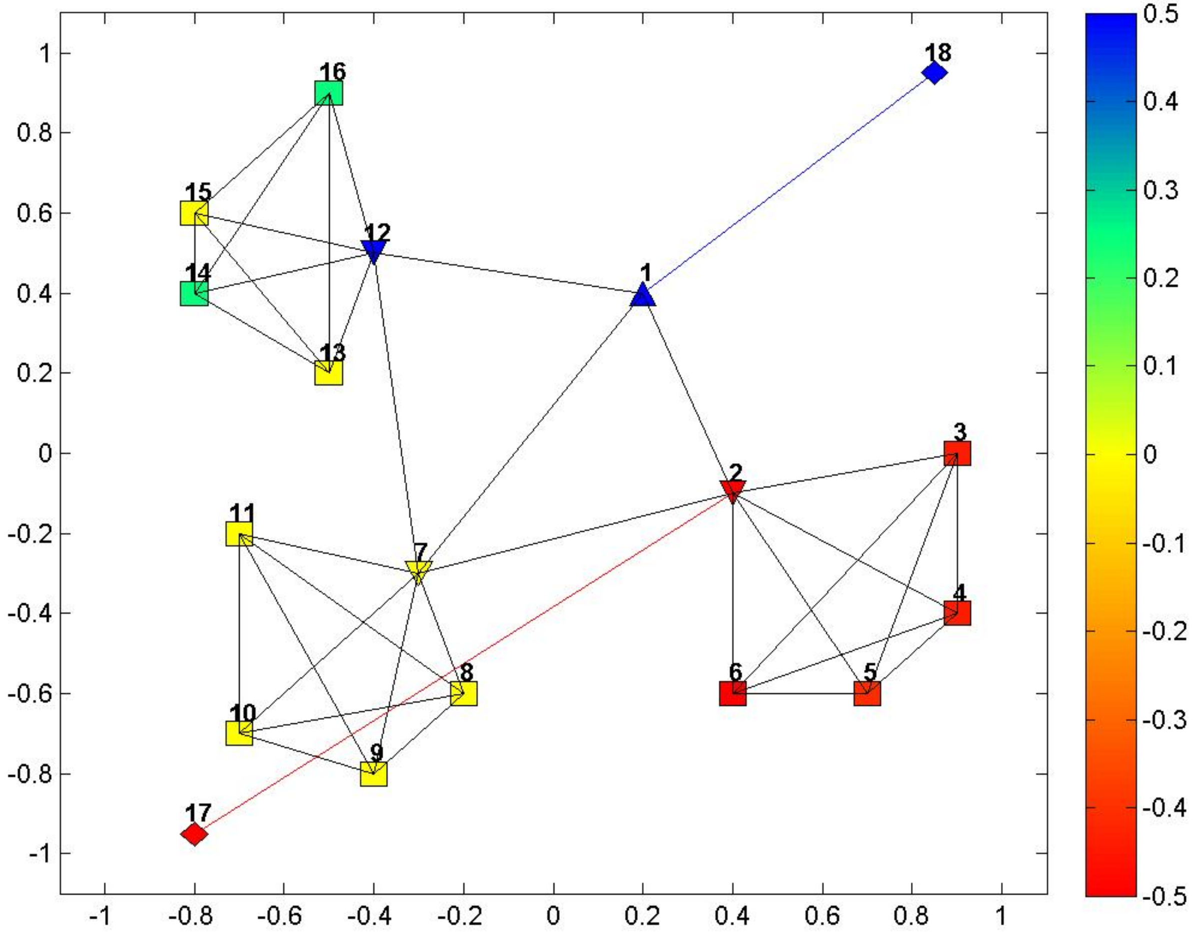

Pro-gov't Anti-TB

'Neutral' Fence-sitters

Anti-gov't Pro-TB

West Point

Interaction Simulation

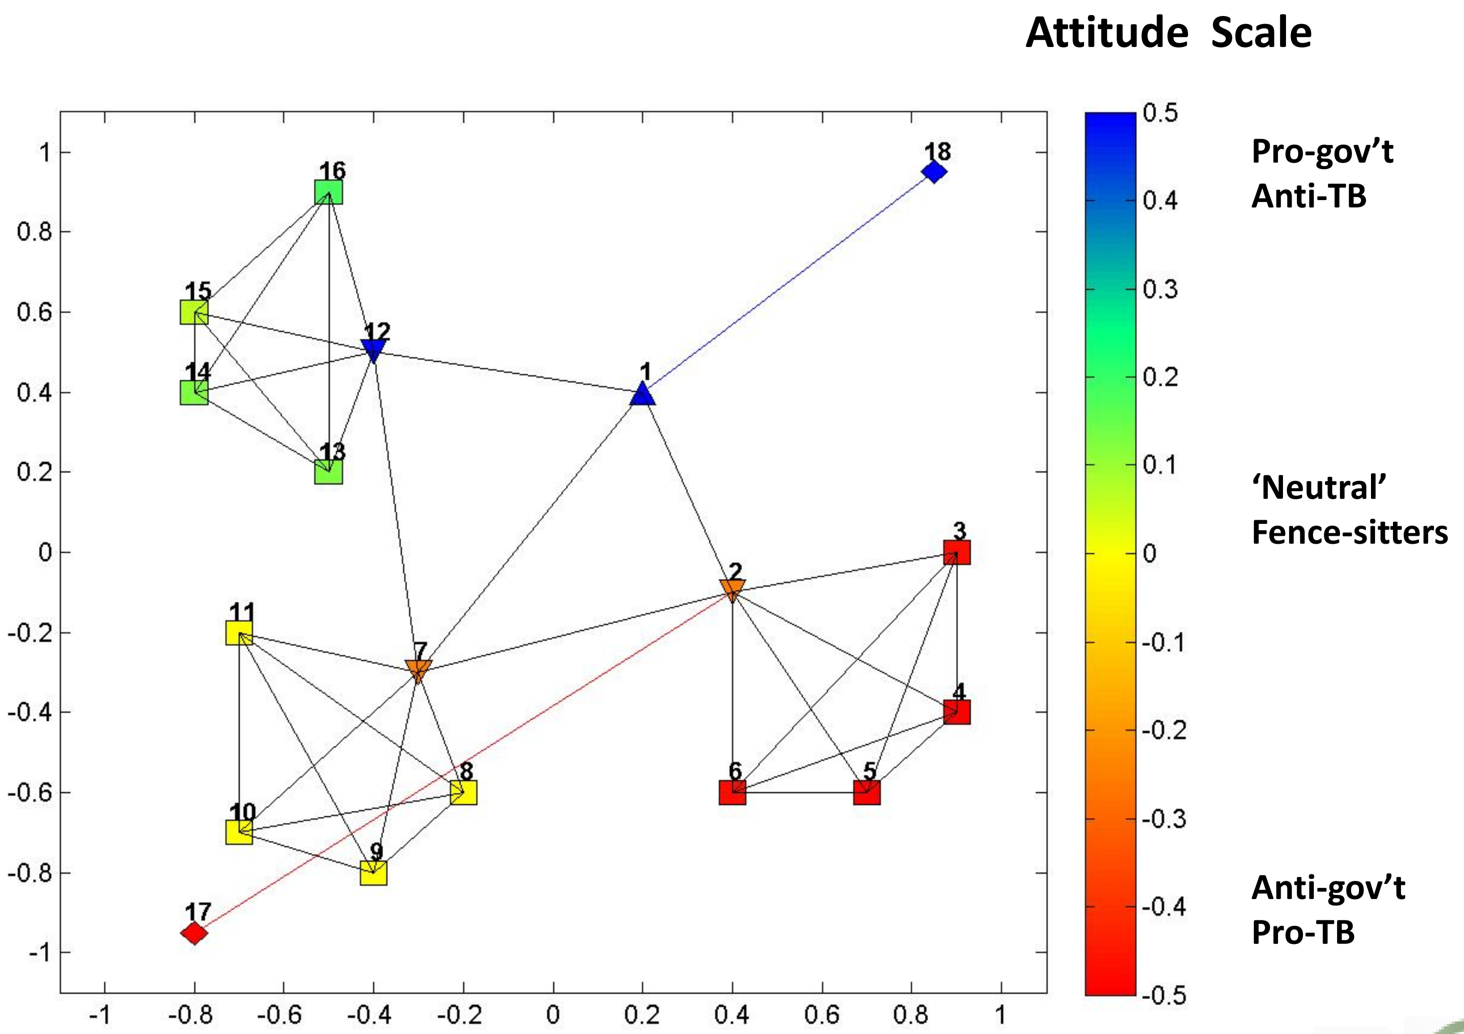

Interaction Simulation

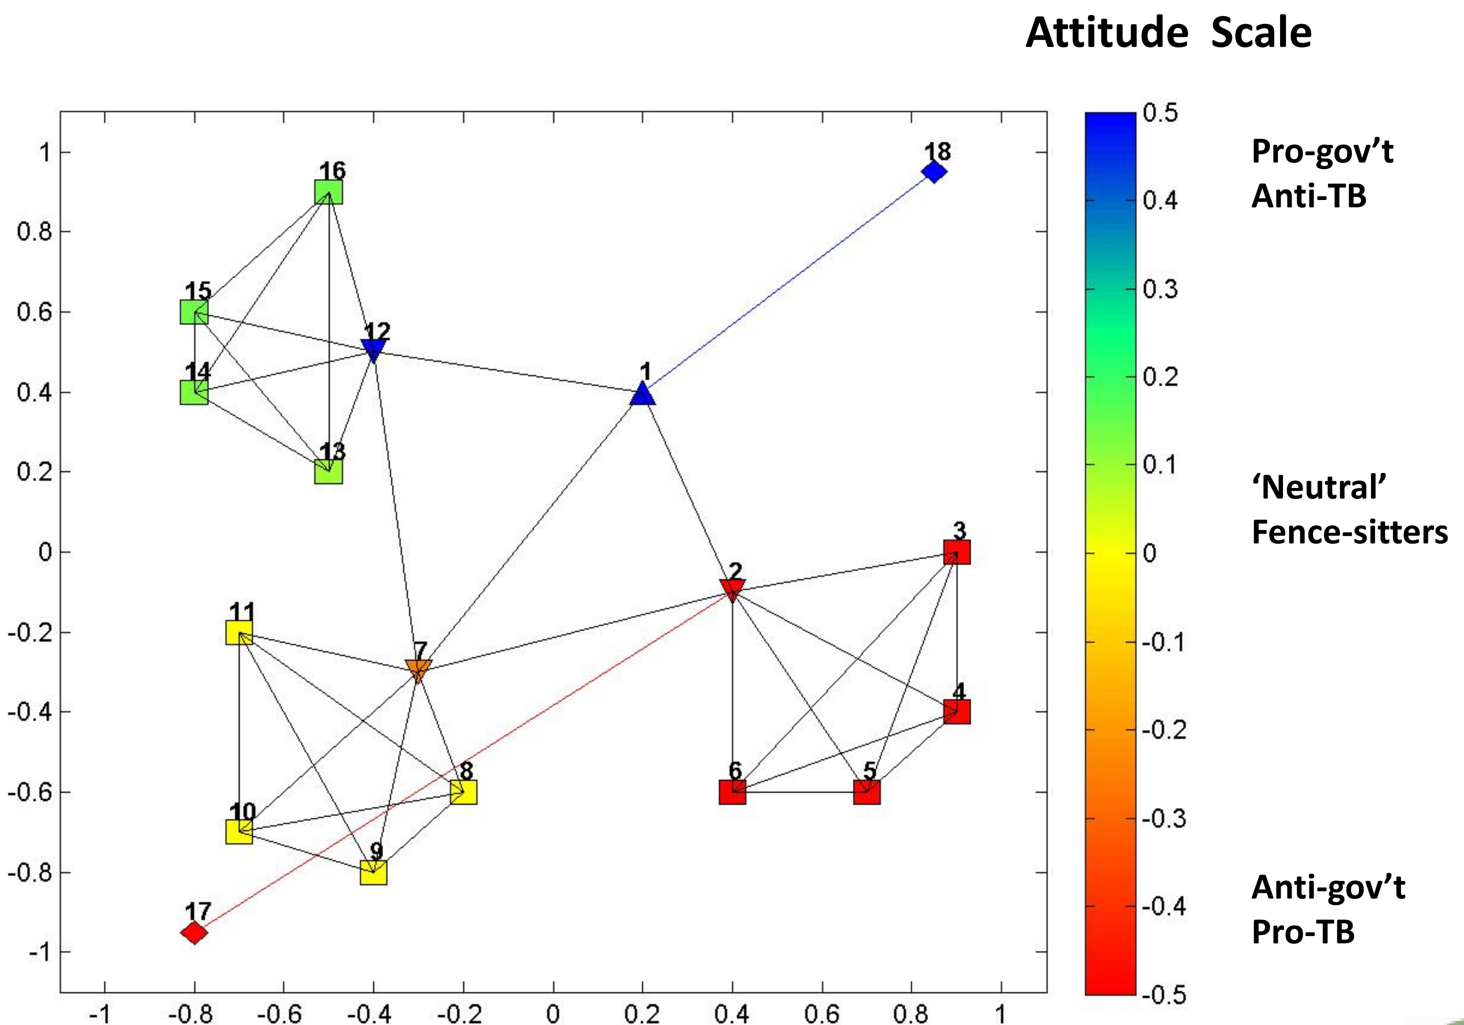

Interaction Simulation

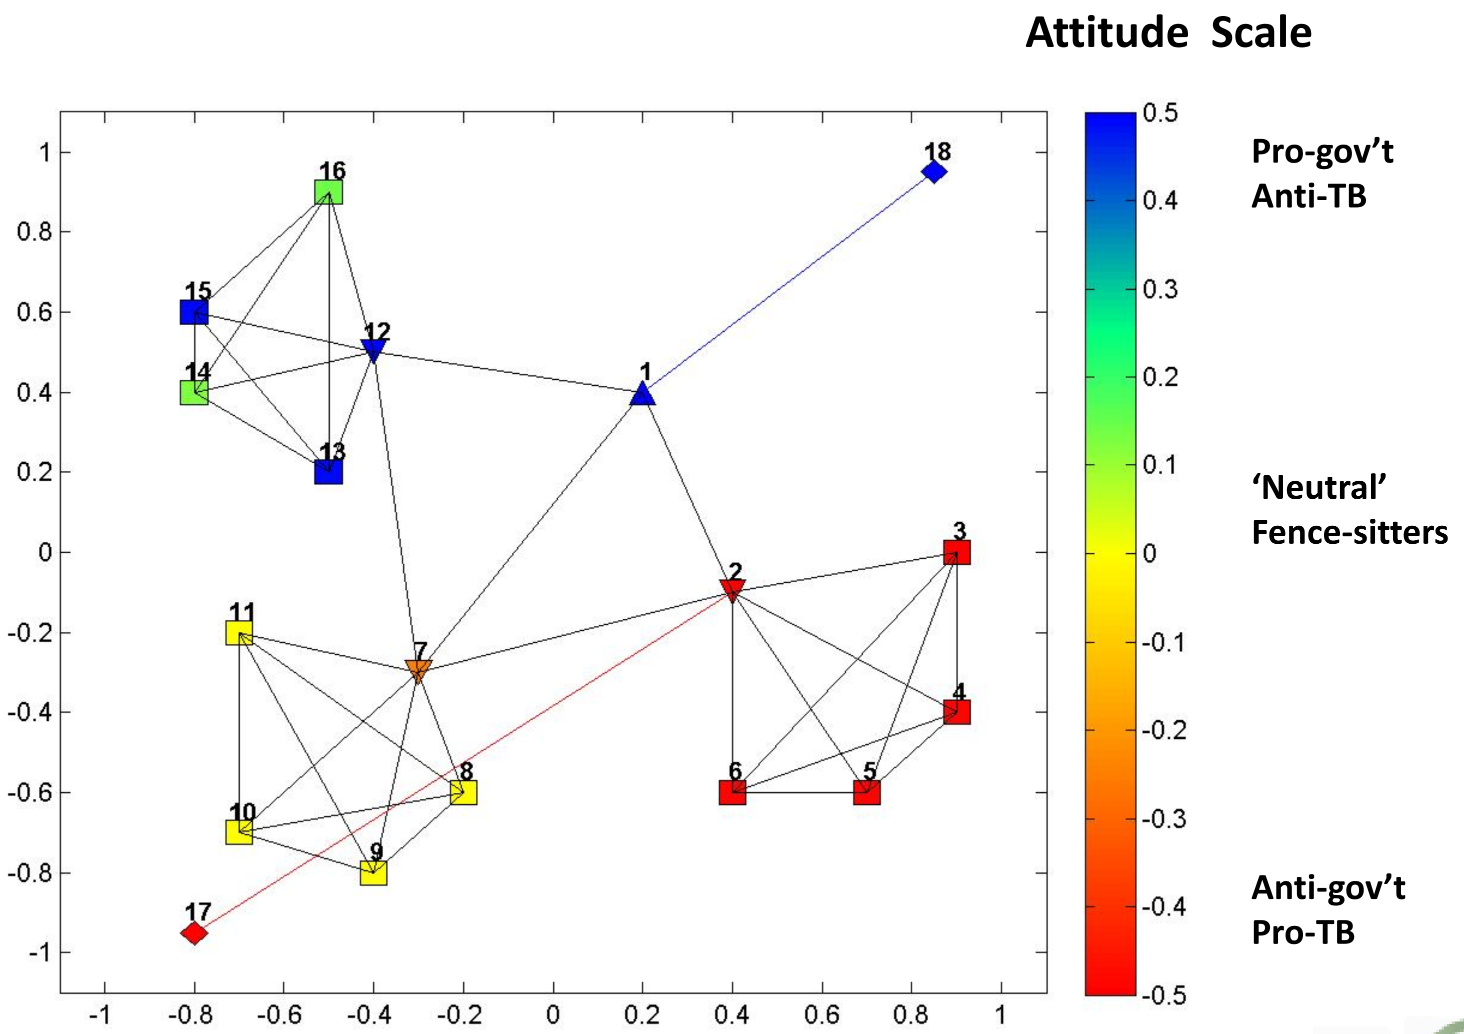

Interaction Simulation

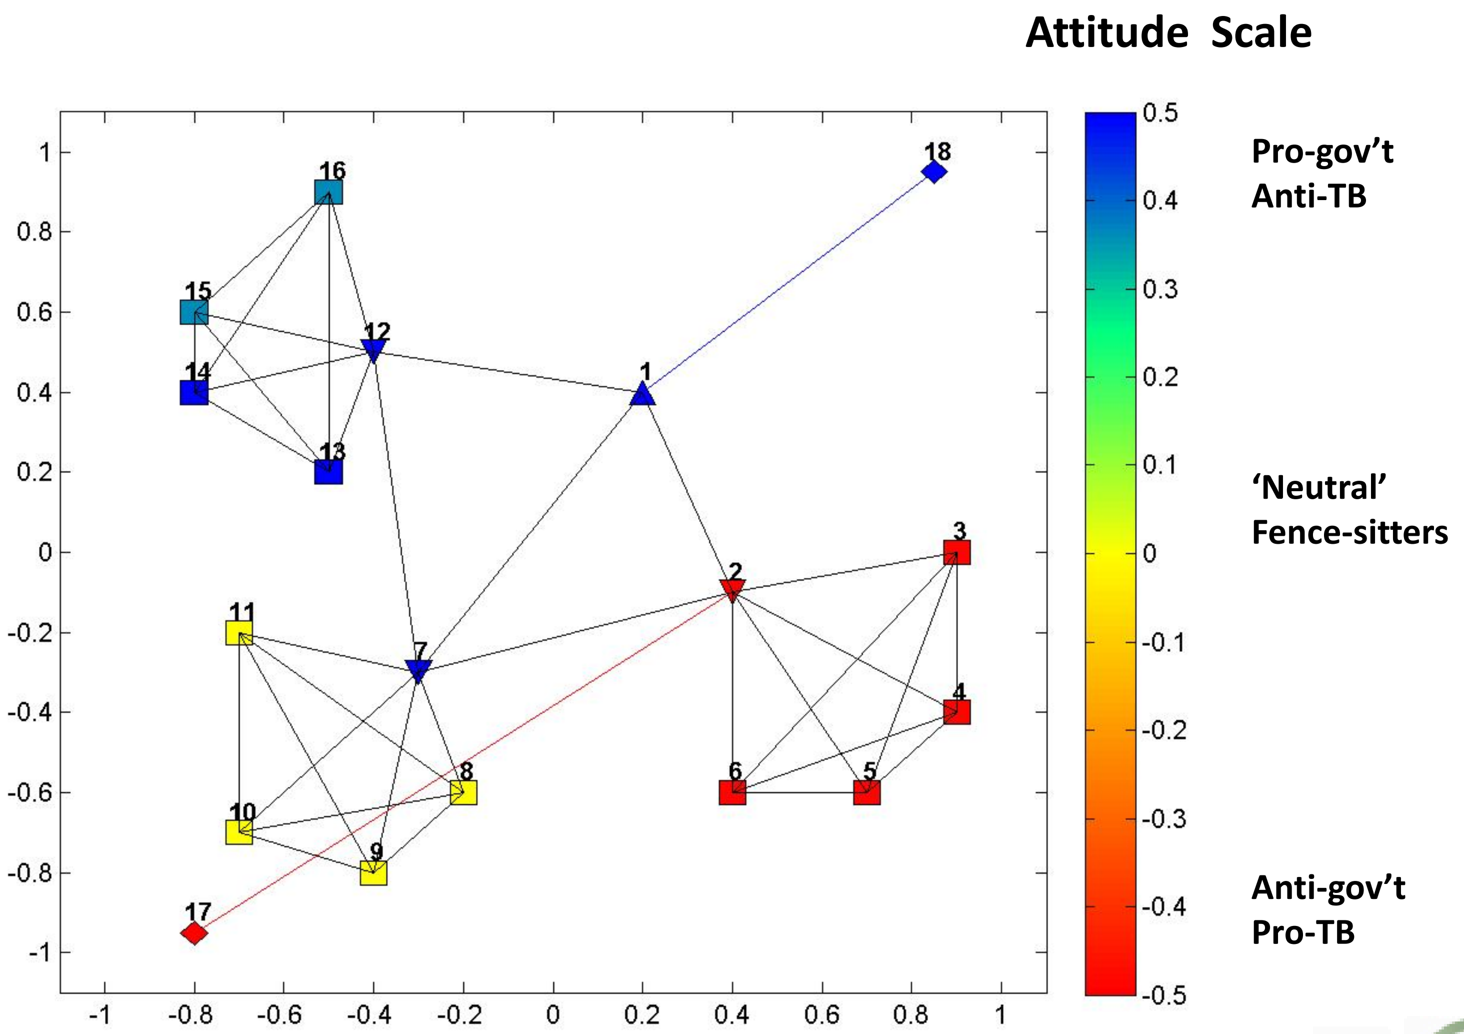

Interaction Simulation

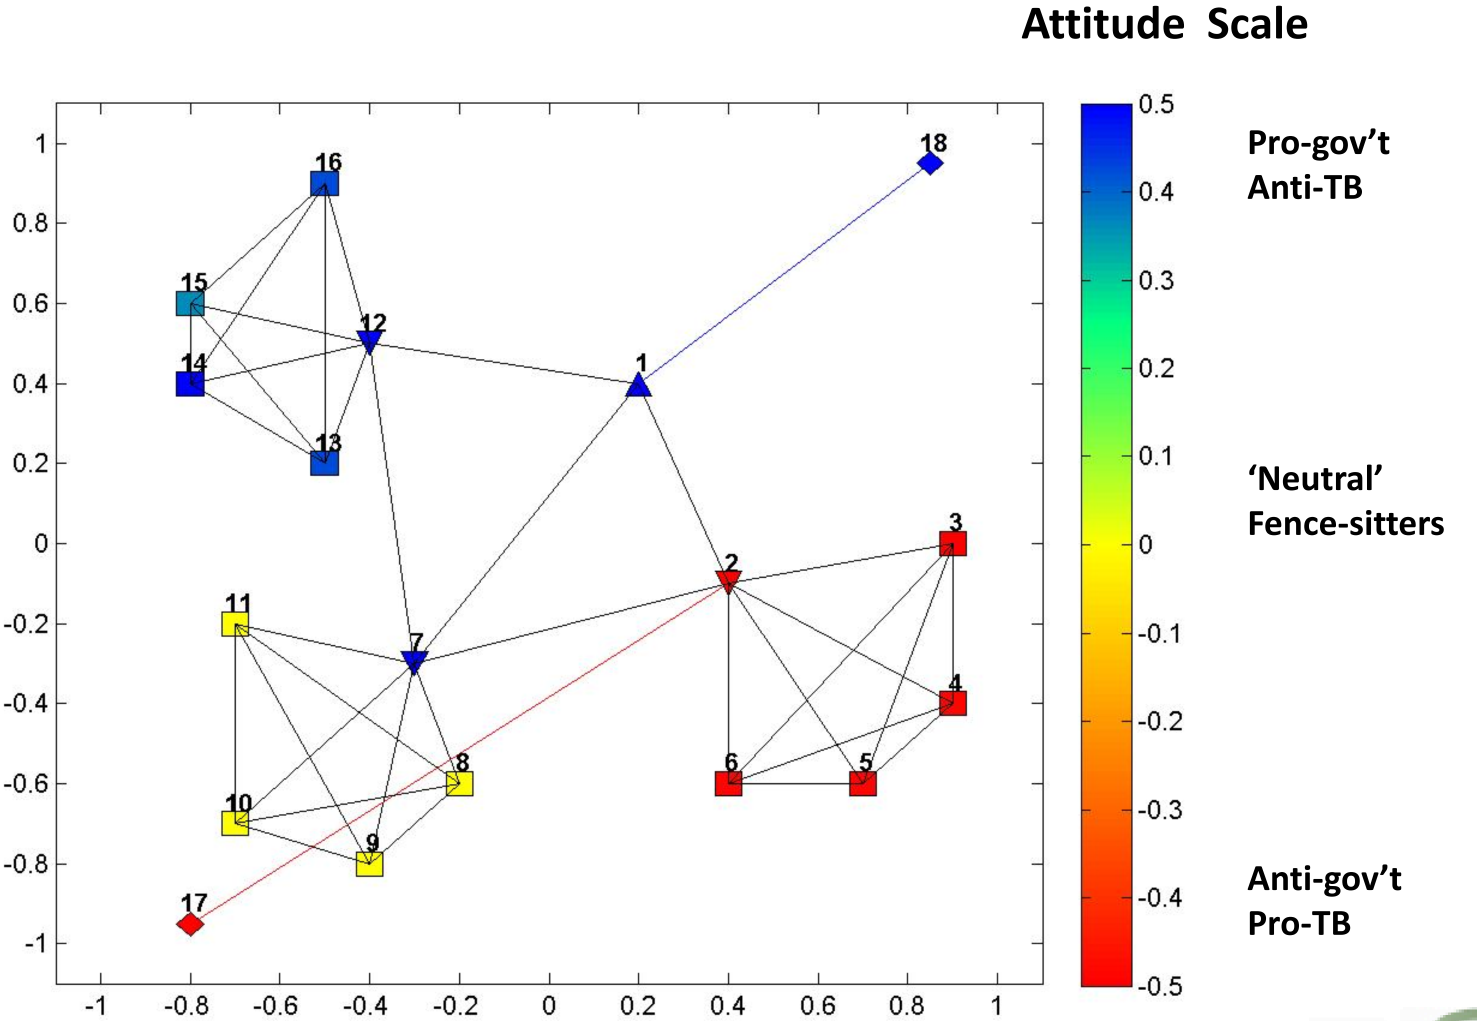

West Point

Interaction Simulation

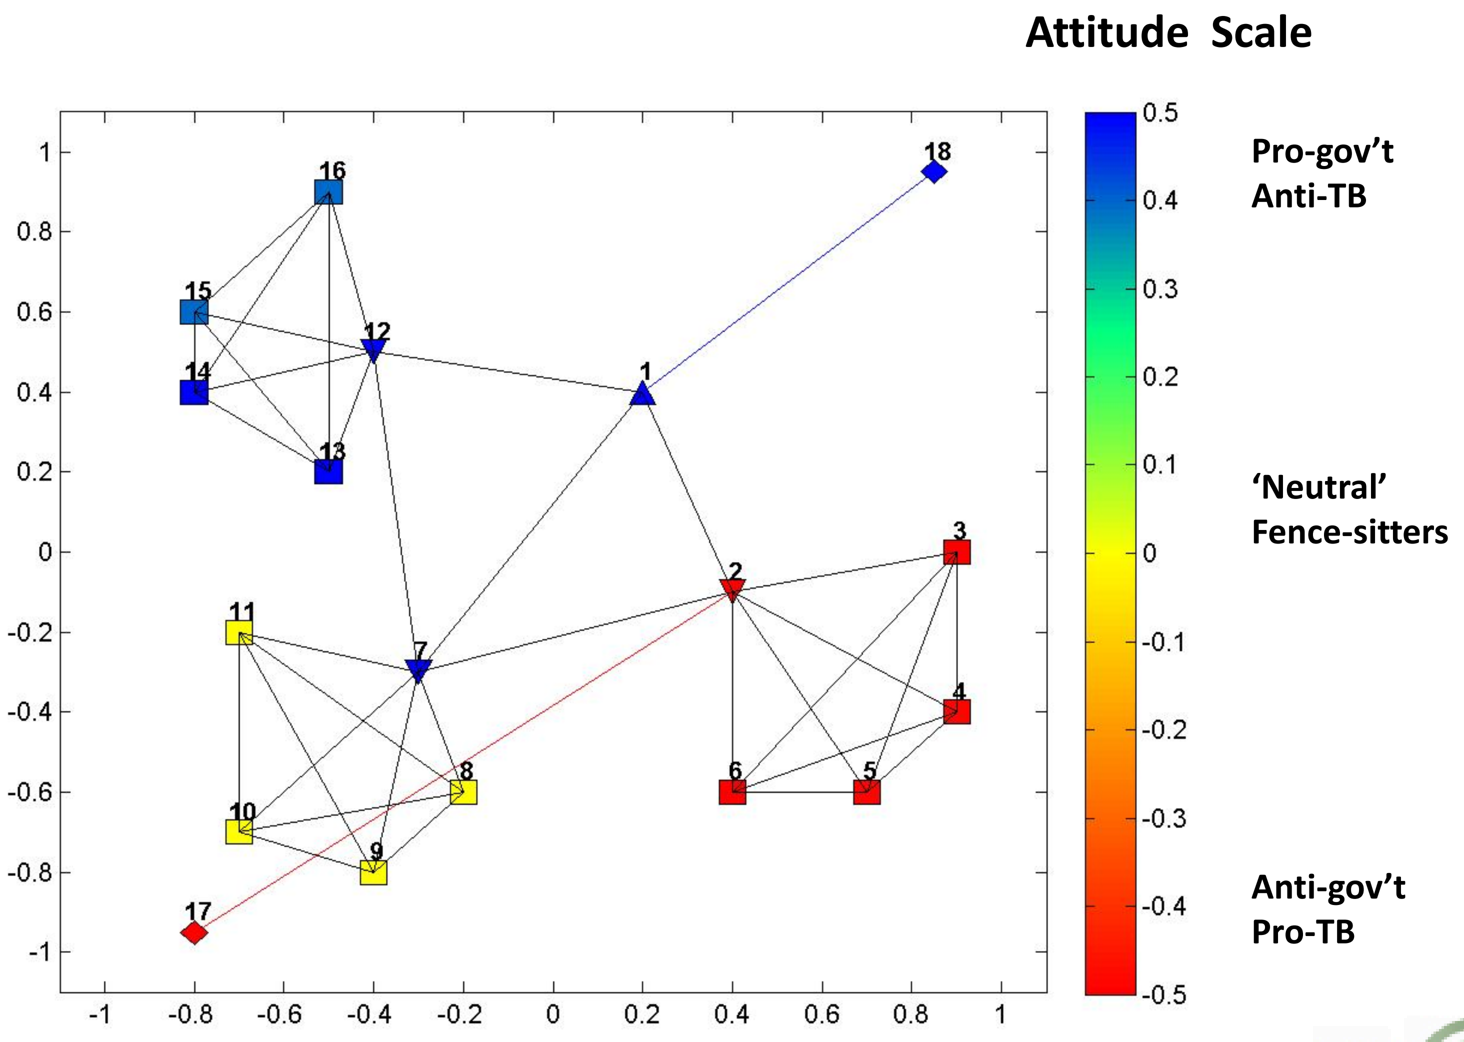

Interaction Simulation

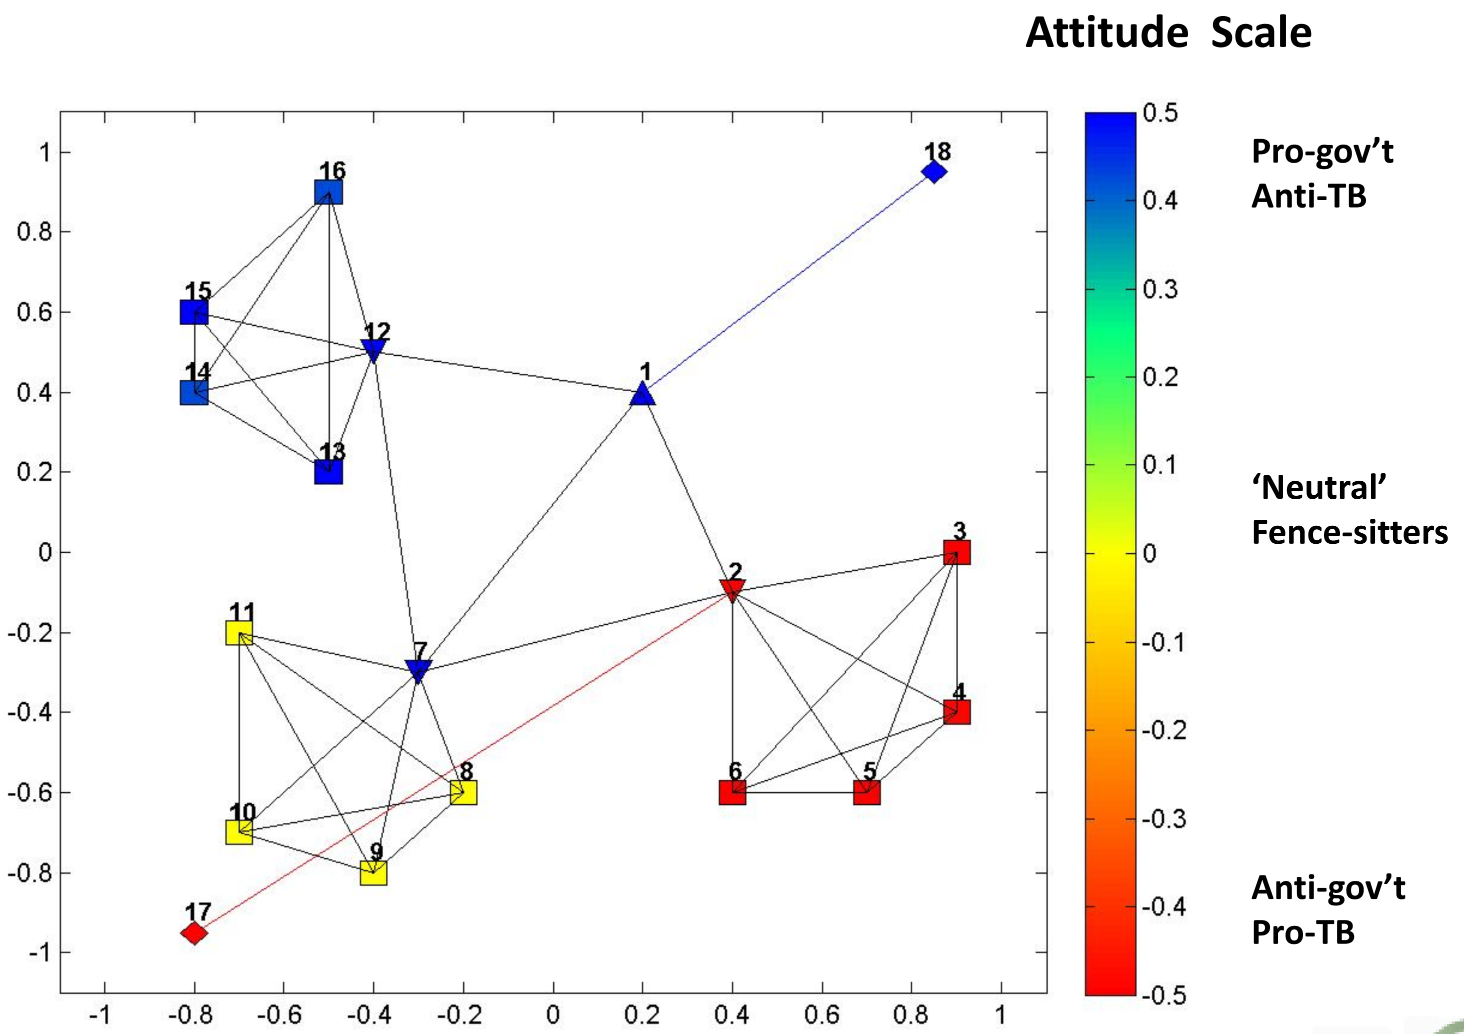

Interaction Simulation

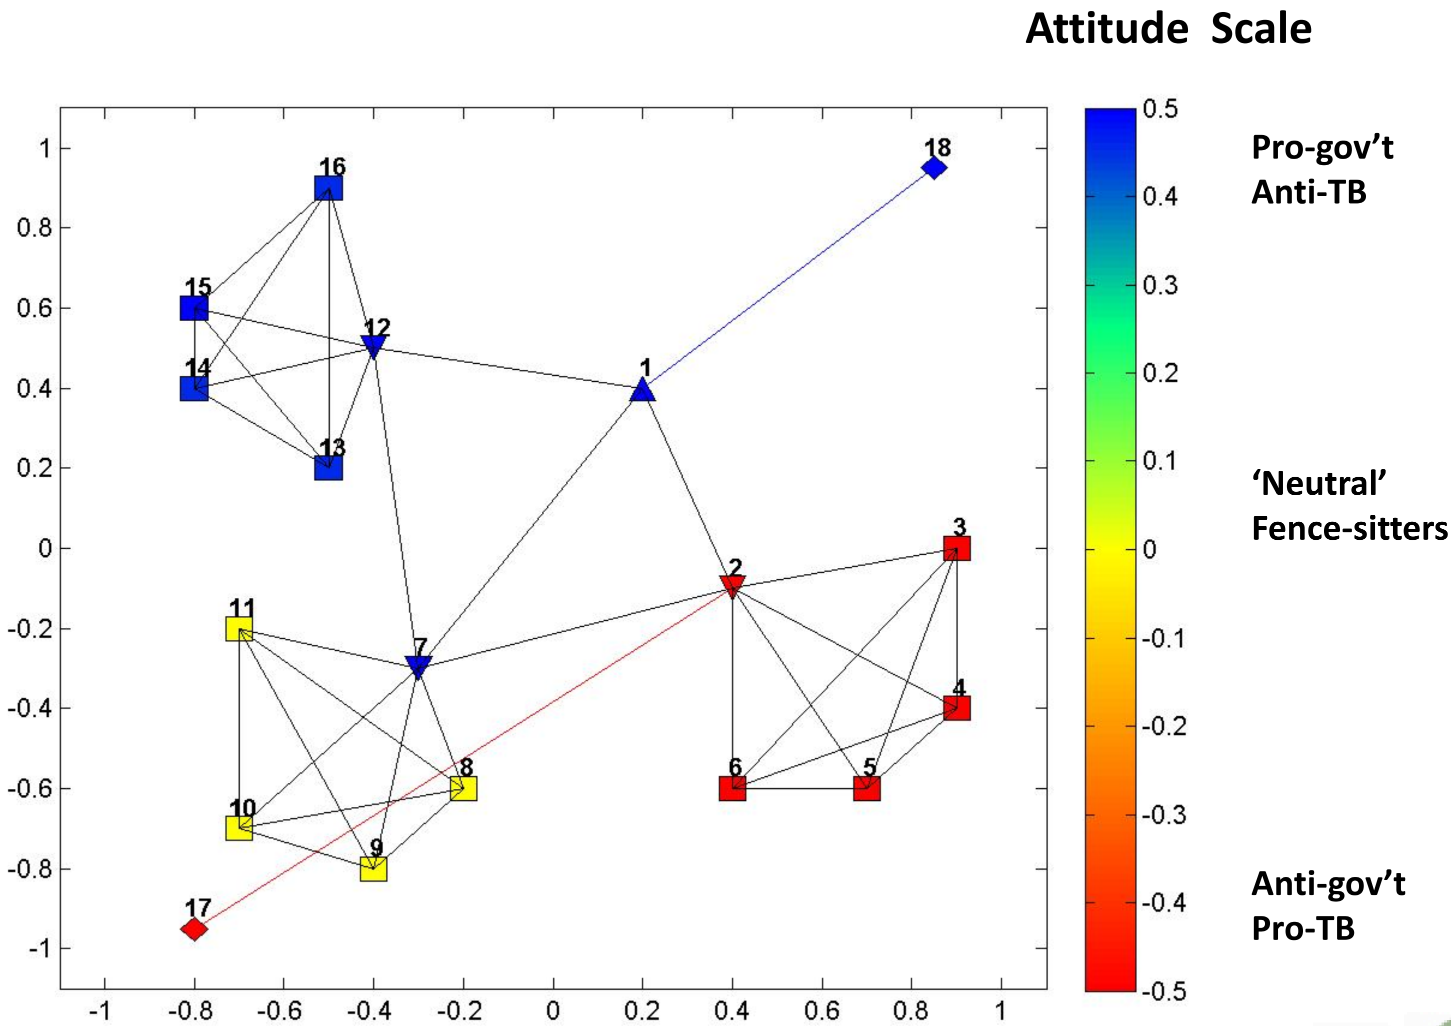

Interaction Simulation

Attitude Scale

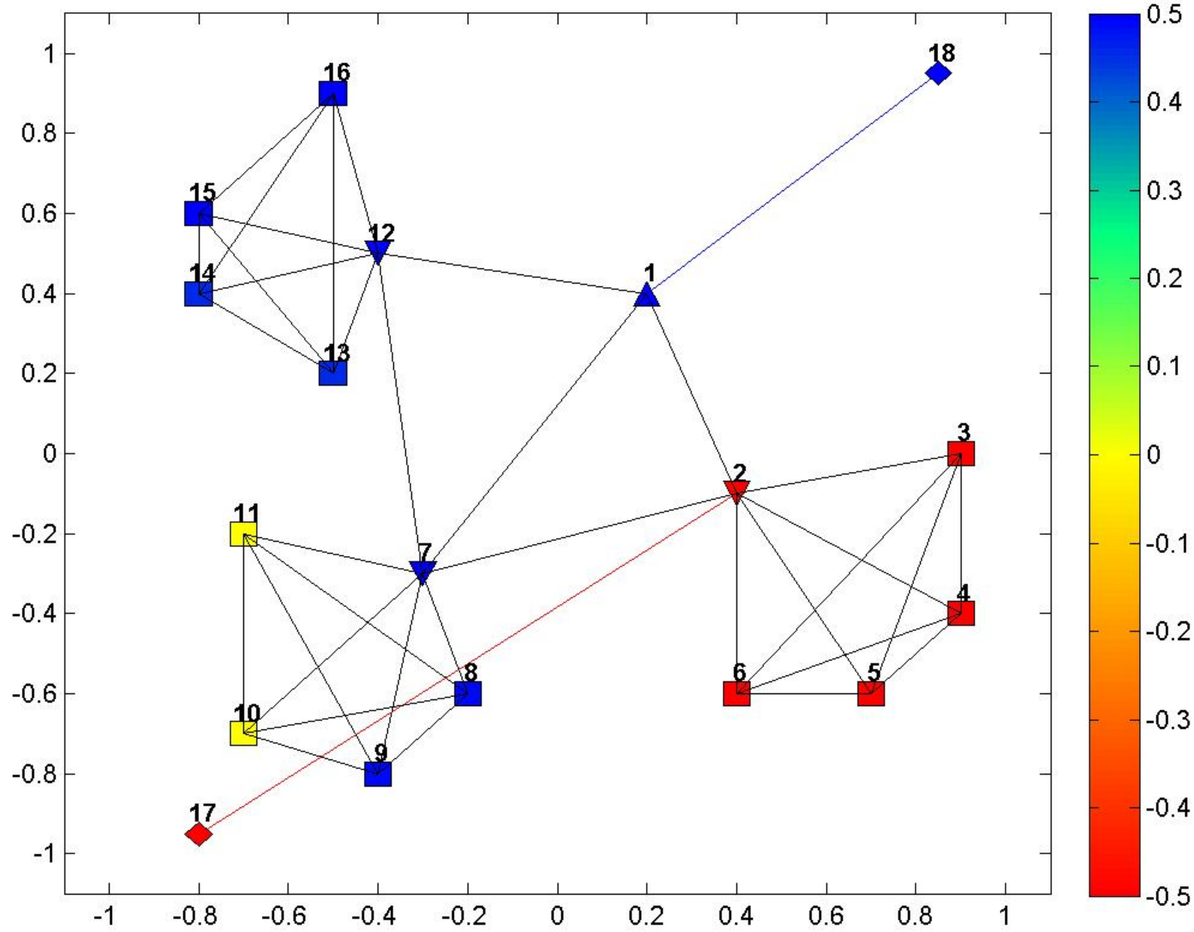

Pro-gov't
Anti-TB

'Neutral'
Fence-sitters

Anti-gov't
Pro-TB

Interaction Simulation

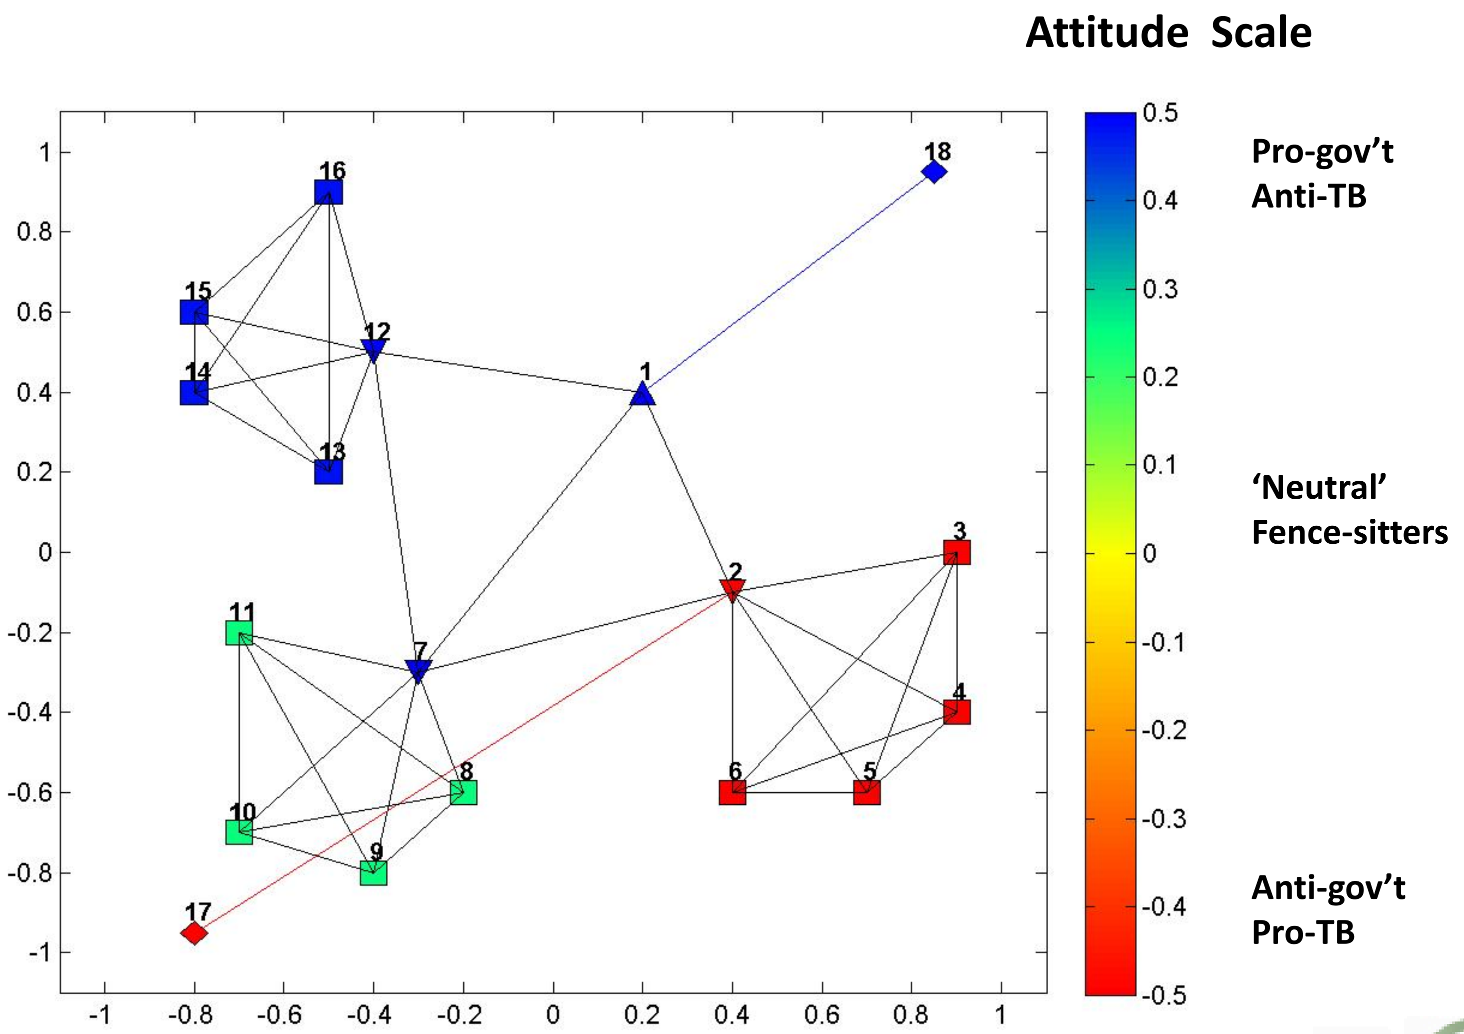

Mean Belief of Network Over Time

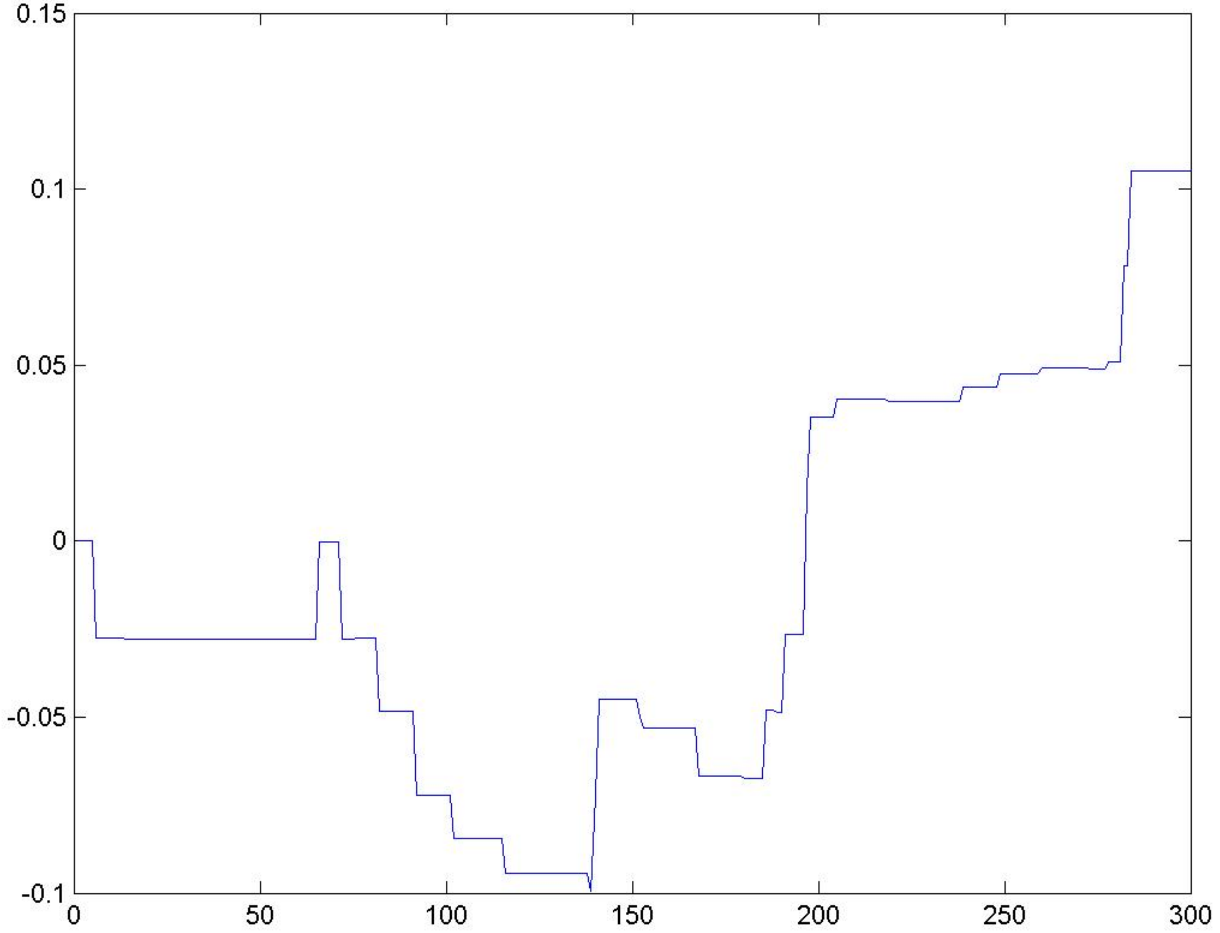

Properties of the Network Model

- Well defined first moment
  - Beliefs converge in expectation at a sufficiently large time in the future to a set of equilibrium beliefs:
    $$X^* = \lim_{k \to \infty} E[(X(t+k)|X(t)]$$
  - *X** is independent of X(t)
- In simulation the average belief of the network oscillates around the average equilibrium belief
- *X** can be calculated in $O(n^3)$

Network Science Center
West Point

The Game

- Complete information, symmetric
- 2 players control a set of stubborn agents, and then make connections to the mutable agents
- Payoff functions defined as a function of the expected mean belief of network (convex combination of the elements of X*)

Example Payoffs

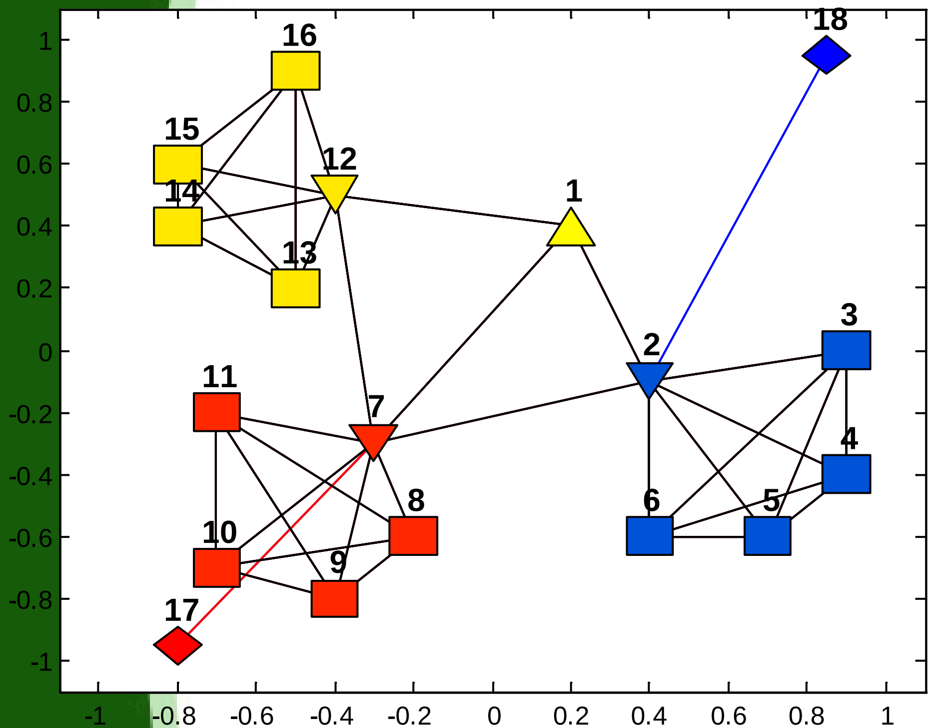

Mean Belief: -.0139

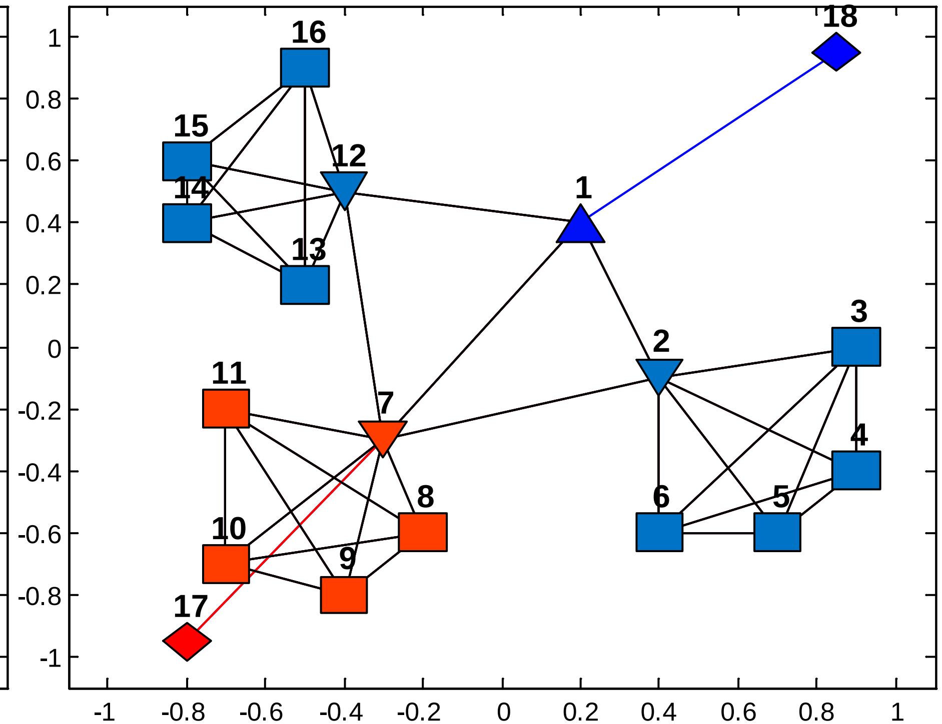

Mean Belief: .1545

Network Science Center
West Point

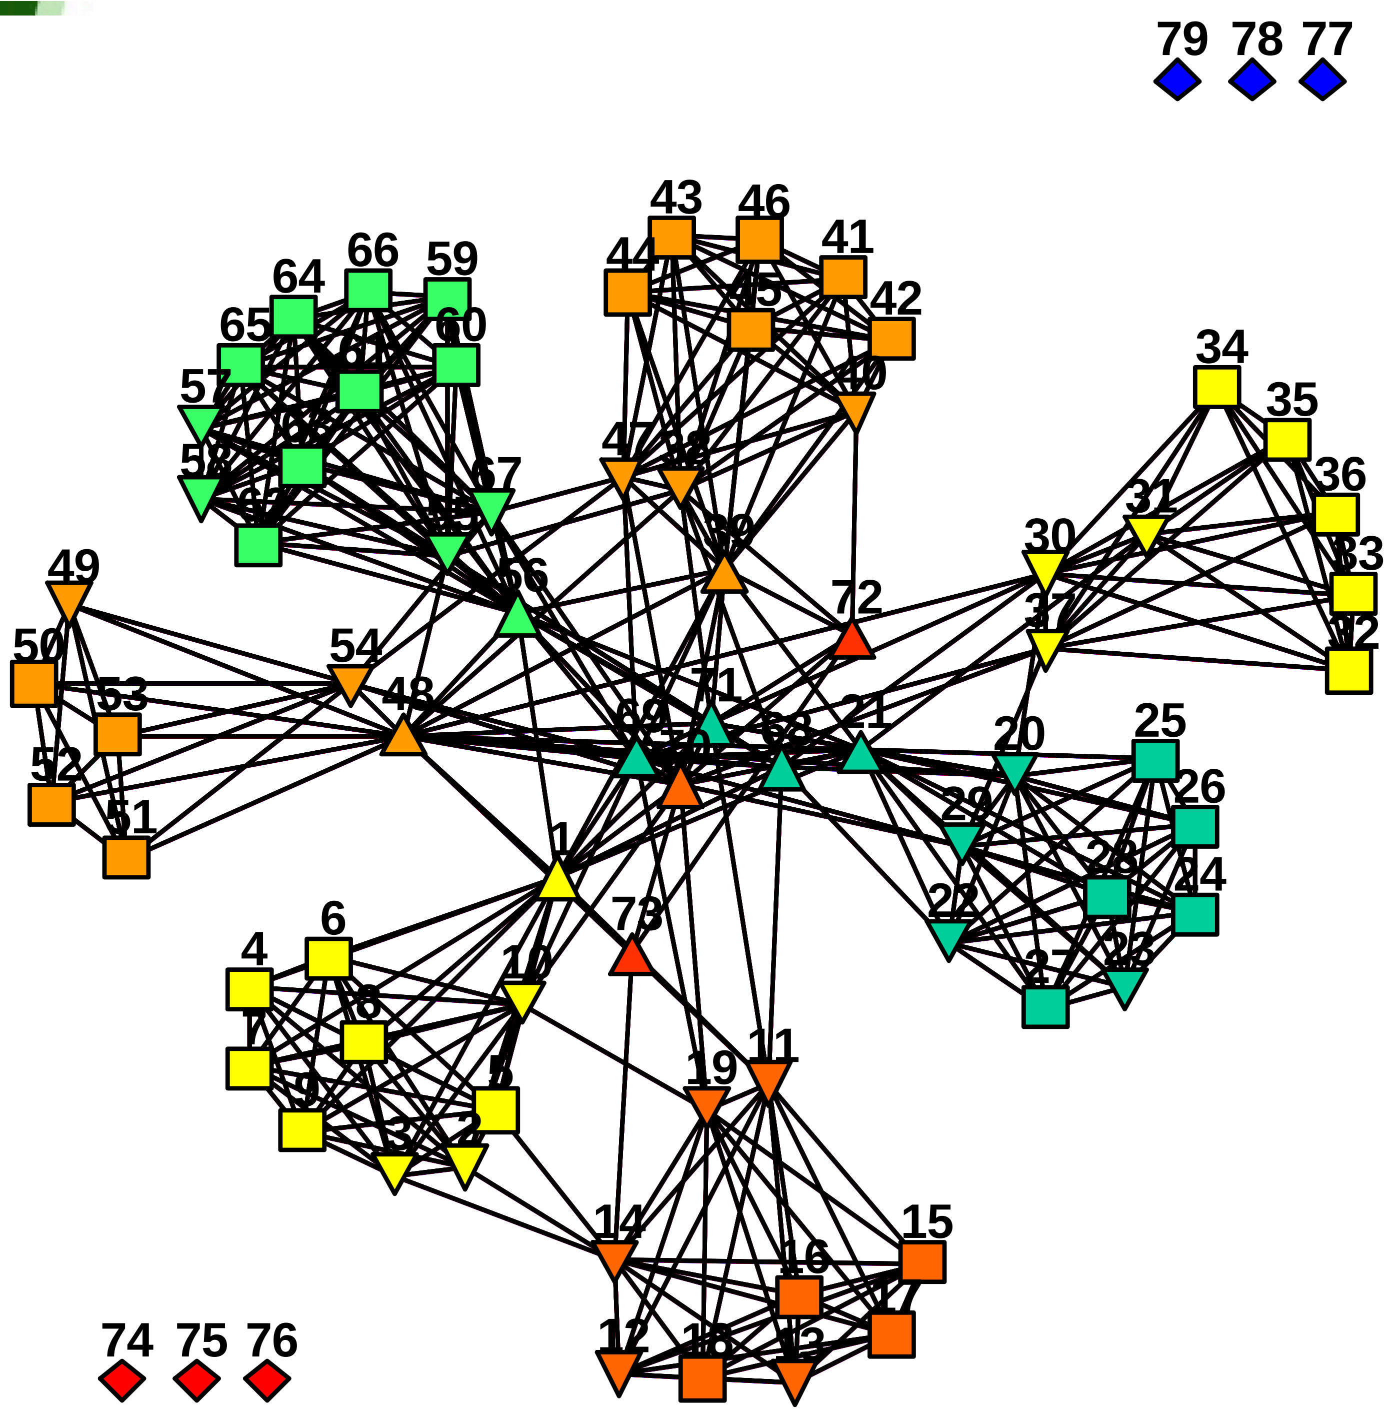

79 78 77
43 46 41 44 42
66 59 64 65 60
34 35 36
49
50 53 52 51
54 48
72 30 31 33
25 26 20
1 73
4 6 8 10
19 11
14 15 16
74 75 76

Finding Nash Equilibria

- Finding Nash Equilibria from a payoff matrix is straightforward
- However, given C connections for each player, it takes $O(n^{2C+3})$ to enumerate the payoff matrix
- Simulated Annealing runs in $O(n^3)$ to optimize a single player's strategy (vs MINLP)
- Applying the simulated annealing algorithm allows us to use Best Response Dynamics to find Pure Nash Equilibria in $O(n^3)$+
- The equilibria turn out to be highly robust to changes of the influence parameters

Network Science Center
West Point

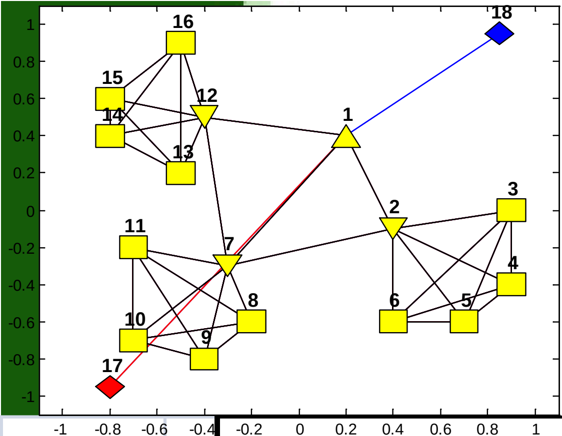

US Strategy \ TB Strategy	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0.000	0.119	0.295	0.295	0.295	0.295	0.052	0.283	0.283	0.283	0.283	0.119	0.295	0.295	0.295	0.295
2	-0.119	0.000	0.310	0.310	0.310	0.310	-0.046	0.240	0.240	0.240	0.240	0.000	0.237	0.237	0.237	0.237
3	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
4	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
5	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
6	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
7	-0.052	0.046	0.268	0.268	0.268	0.268	0.000	0.304	0.304	0.304	0.304	0.046	0.268	0.268	0.268	0.268
8	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
9	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
10	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
11	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
12	-0.119	0.000	0.237	0.237	0.237	0.237	-0.046	0.240	0.240	0.240	0.240	0.000	0.310	0.310	0.310	0.310
13	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000
14	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000
15	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000
16	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000

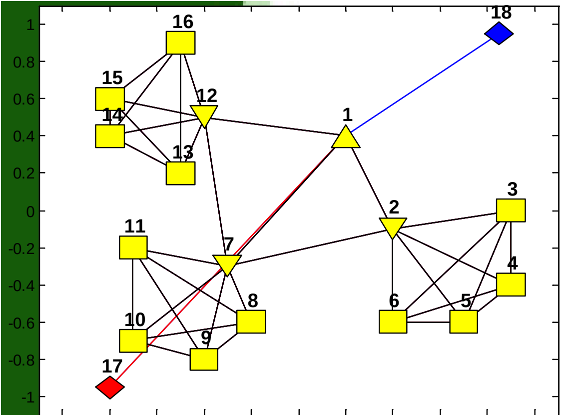

	TB Strategy															
US Strategy	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0.000	0.119	0.295	0.295	0.295	0.295	0.052	0.283	0.283	0.283	0.283	0.119	0.295	0.295	0.295	0.295
2	-0.119	0.000	0.310	0.310	0.310	0.310	-0.046	0.240	0.240	0.240	0.240	0.000	0.237	0.237	0.237	0.237
3	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
4	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
5	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
6	-0.295	-0.310	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.237	0.000	0.000	0.000	0.000
7	-0.052	0.046	0.268	0.268	0.268	0.268	0.000	0.304	0.304	0.304	0.304	0.046	0.268	0.268	0.268	0.268
8	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
9	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
10	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
11	-0.283	-0.240	0.004	0.004	0.004	0.004	-0.304	0.000	0.000	0.000	0.000	-0.240	0.004	0.004	0.004	0.004
12	-0.119	0.000	0.237	0.237	0.237	0.237	-0.046	0.240	0.240	0.240	0.240	0.000	0.310	0.310	0.310	0.310
13	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000
14	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000
15	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000
16	-0.295	-0.237	0.000	0.000	0.000	0.000	-0.268	-0.004	-0.004	-0.004	-0.004	-0.310	0.000	0.000	0.000	0.000

Notional Best Response Dynamics Search Using Simulated Annealing

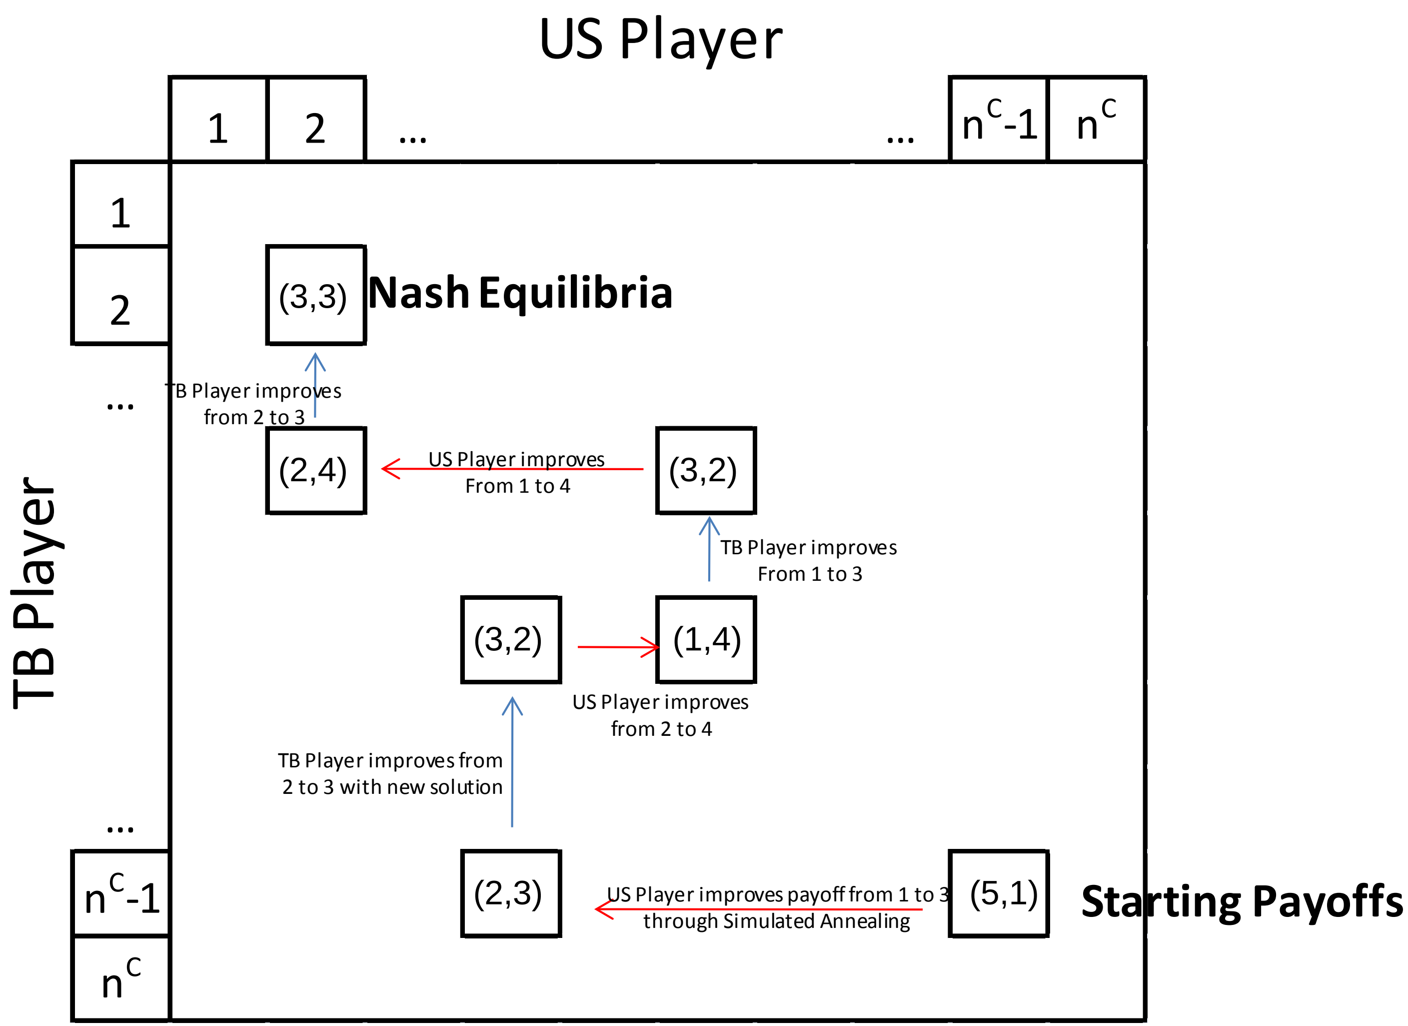

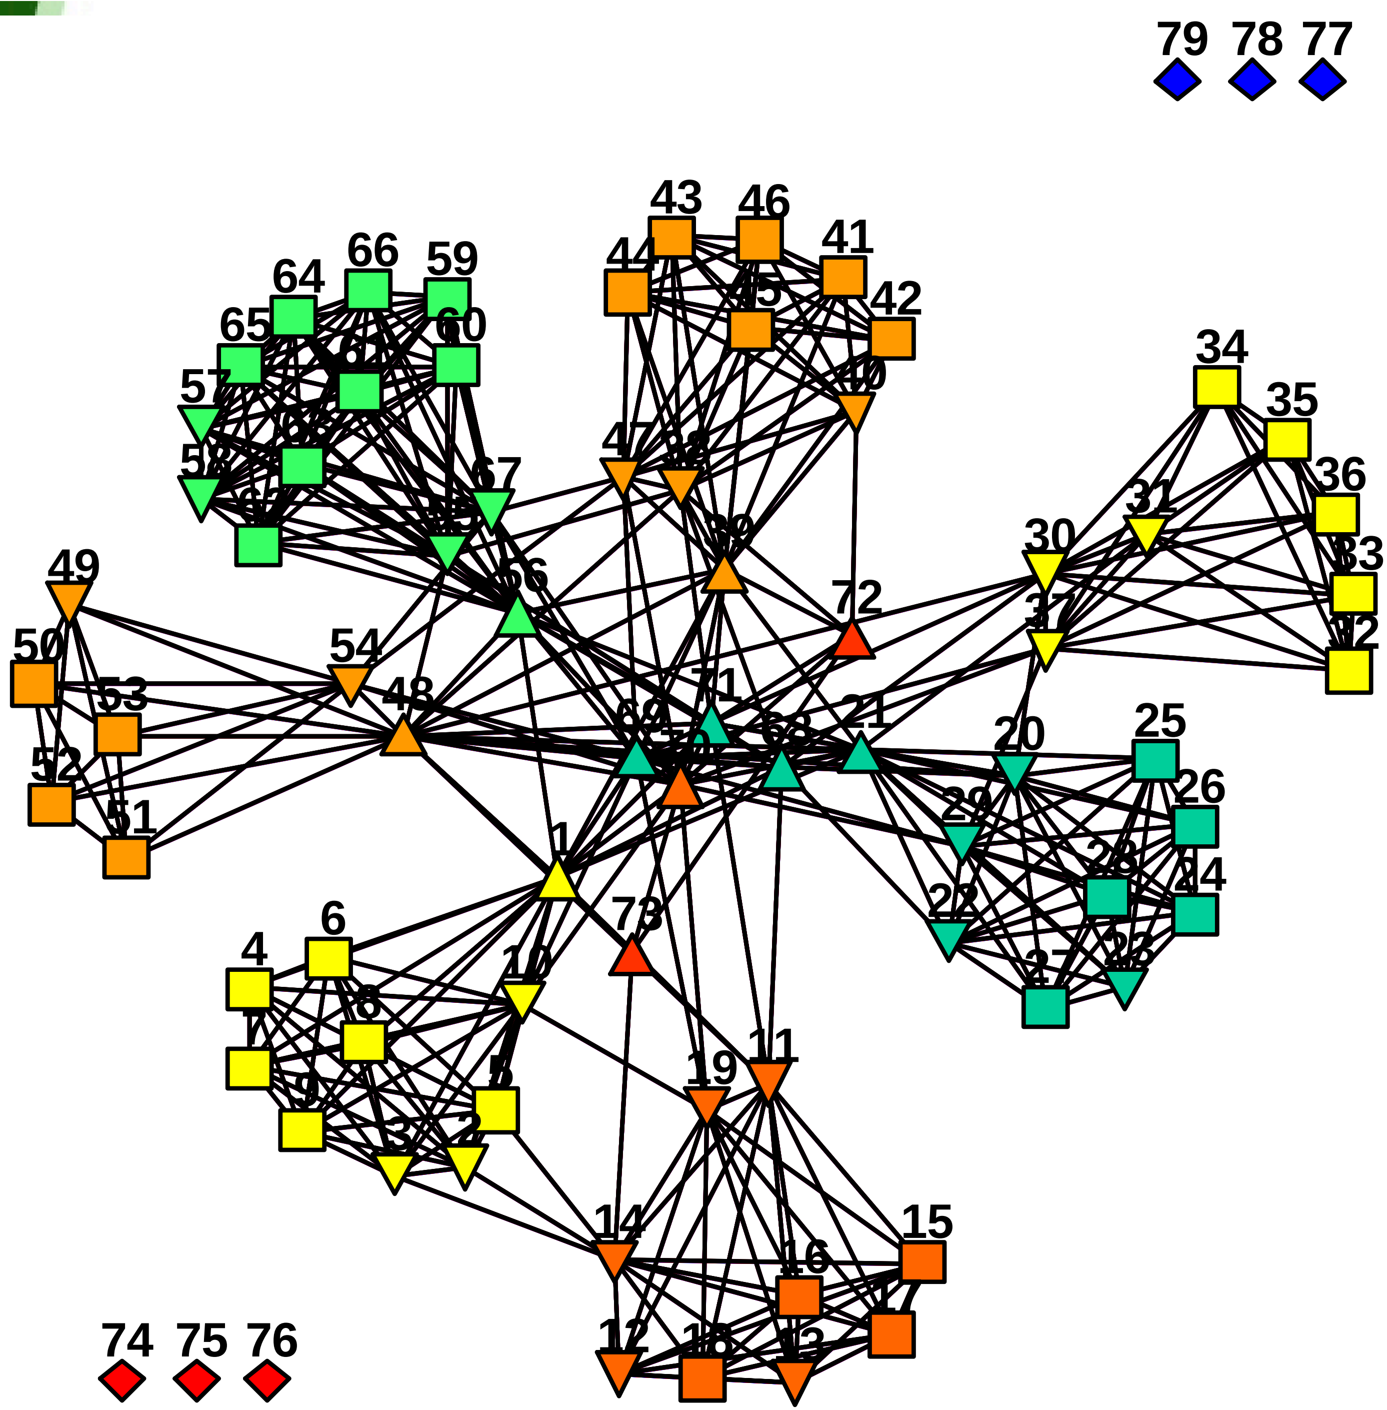
79
78
77
43
46
41
44
42
66
59
64
65
60
34
35
36
57
47
40
31
30
67
49
72
50
54
53
48
71
21
25
20
26
52
51
1
73
6
4
24
10
11
19
14
15
16
74
75
76

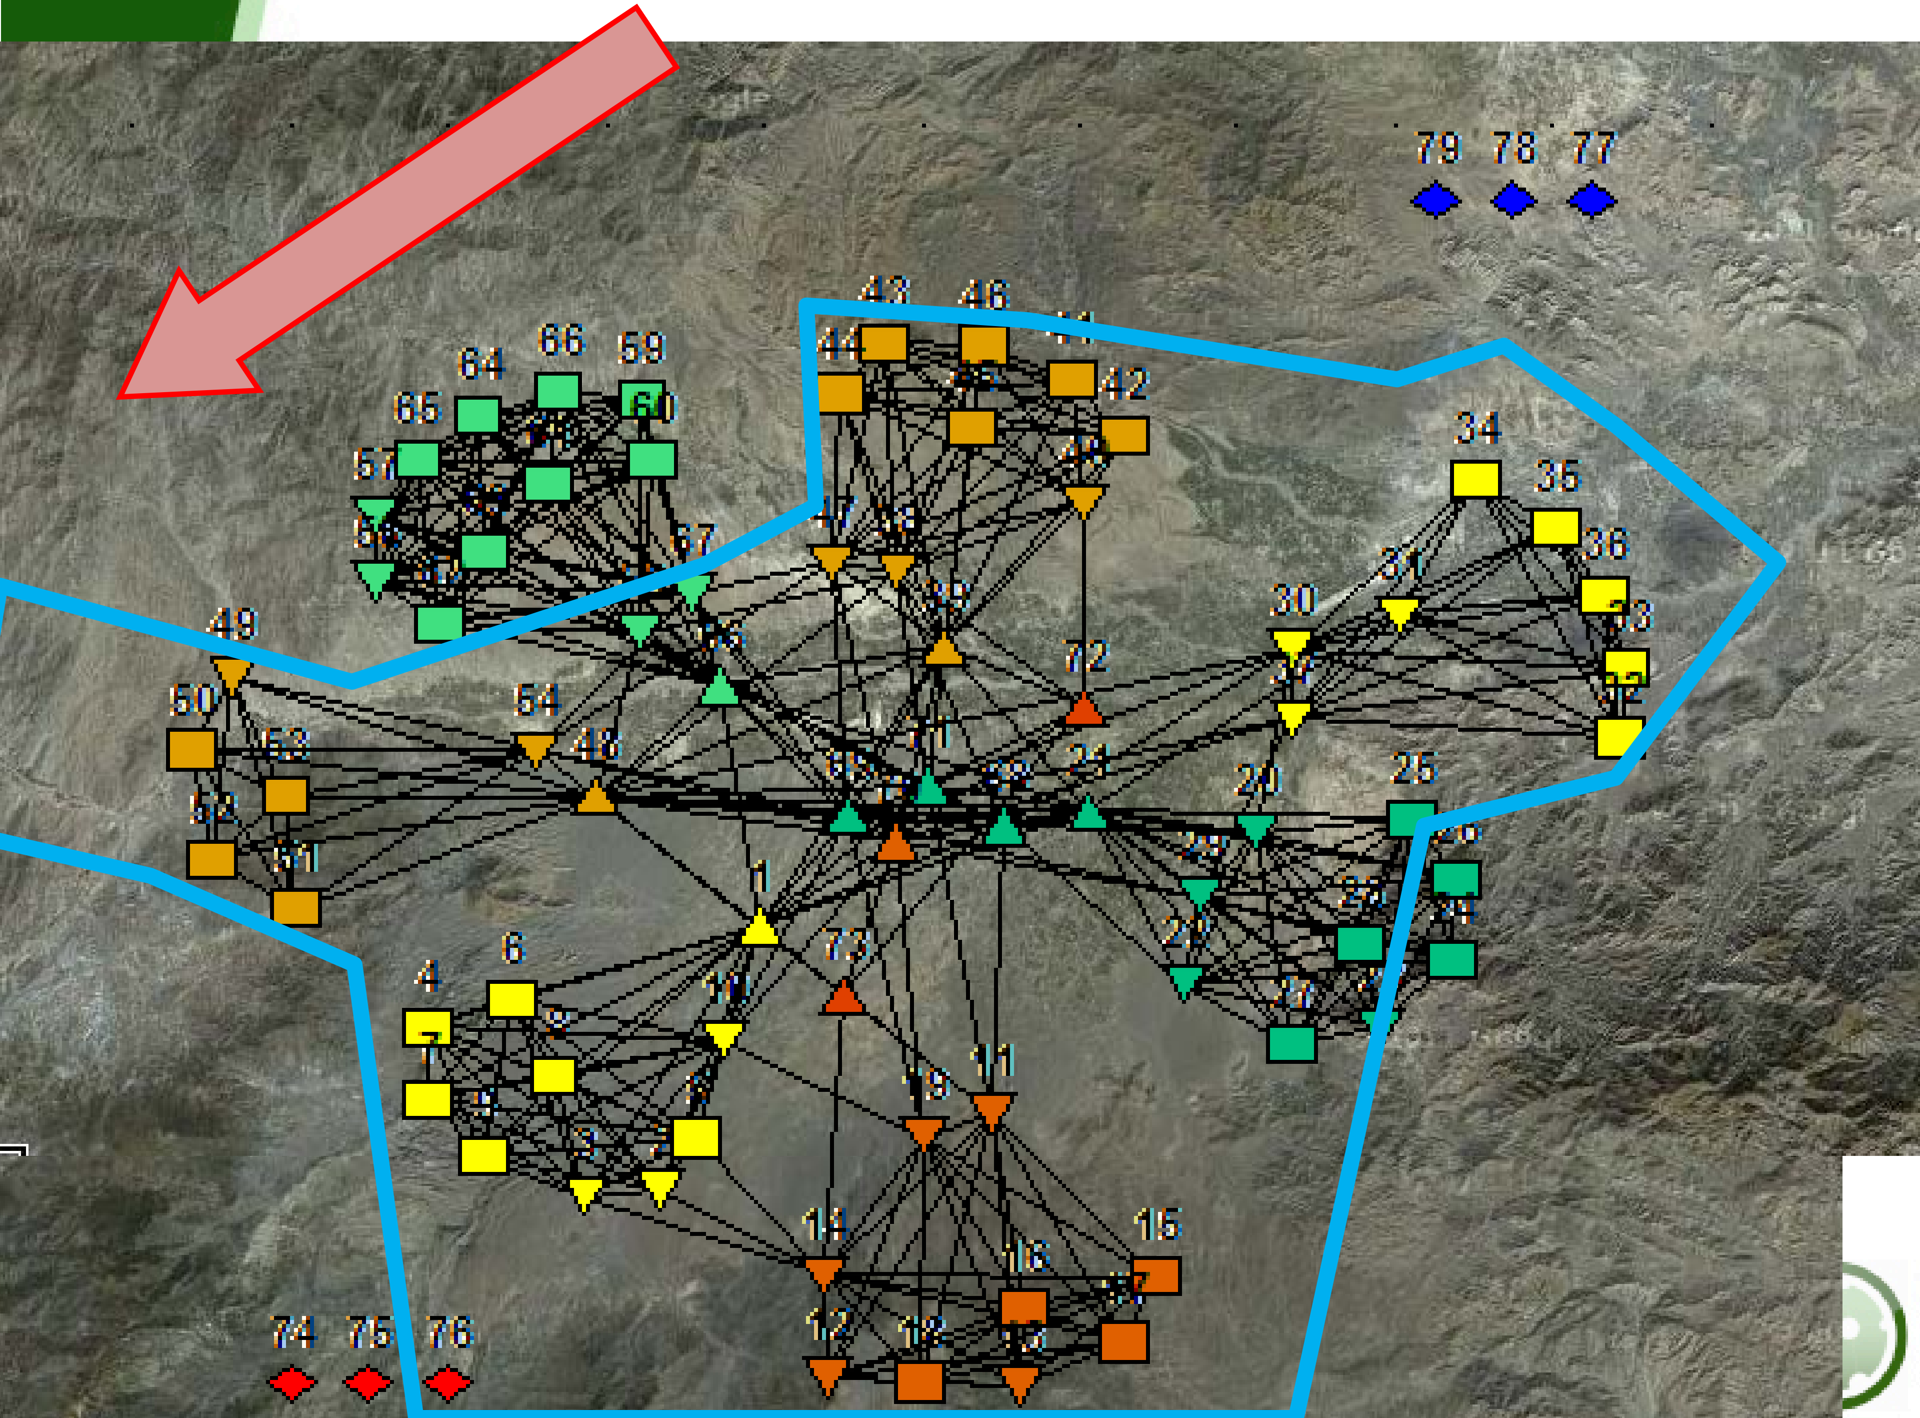

79
78
77
43
46
44
42
66
64
59
65
60
57
67
49
50
53
54
48
52
34
35
36
30
72
25
20
1
73
6
4
11
14
15
16
12
74
75
76

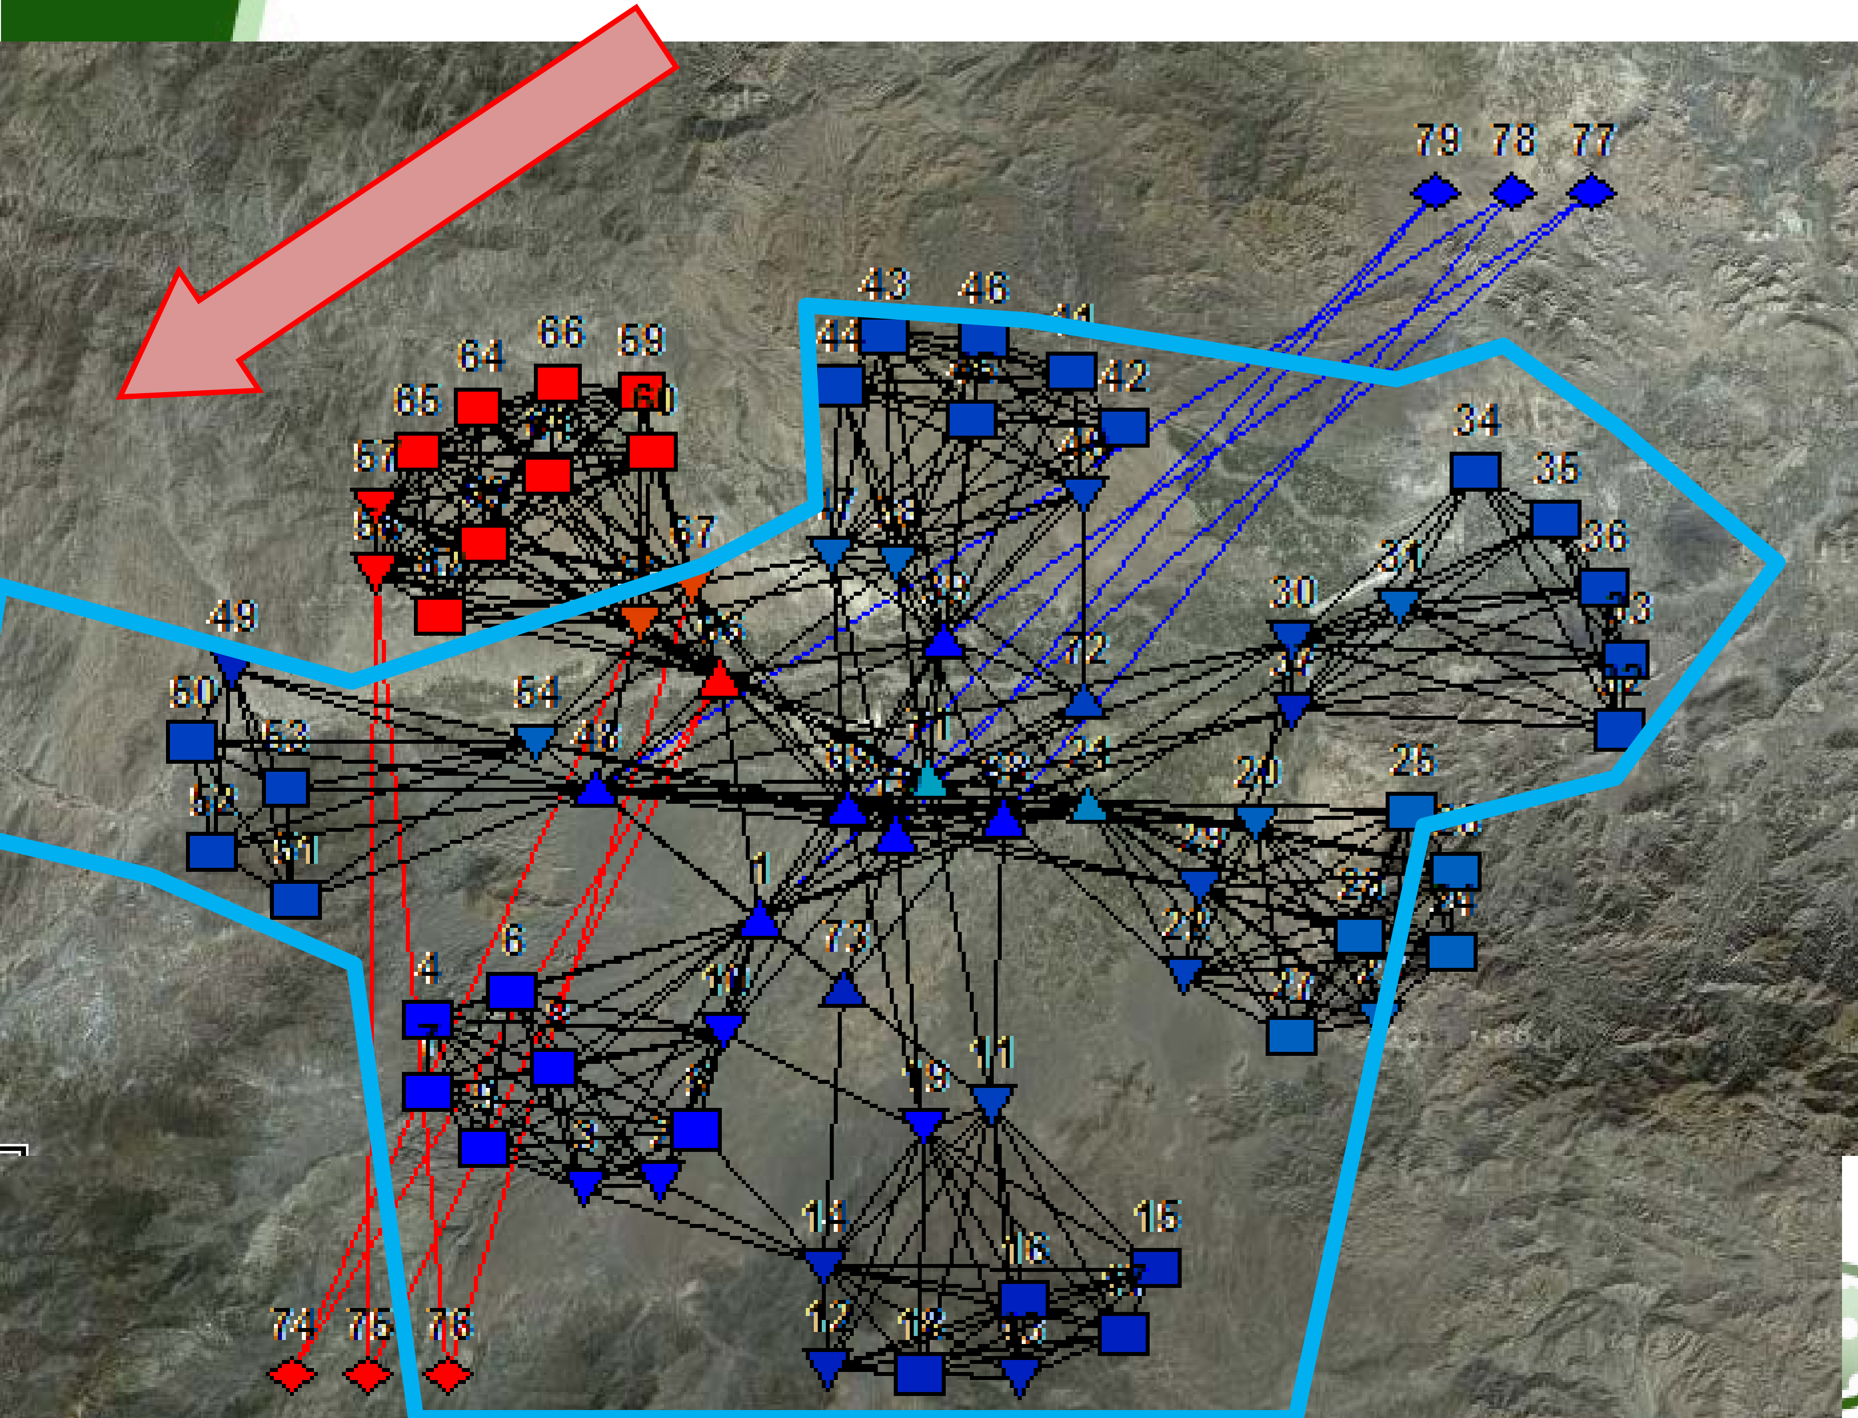

79
78
77
43
46
44
42
34
35
36
64
66
59
65
60
57
67
49
30
72
50
54
53
52
25
20
6
4
1
73
11
19
14
15
16
12
74
75
76

Thoughts

- The time to find Pure Nash Equilibria increases exponentially with C
- Simulated Annealing heuristic provides a method to conduct best response dynamics
- Current equilibrium payoff functions not completely realistic, but make problem tractable
- Inverse Optimization could be used to determine the enemy's payoff function based off of intelligence reports of villages

Network Science Center
West Point

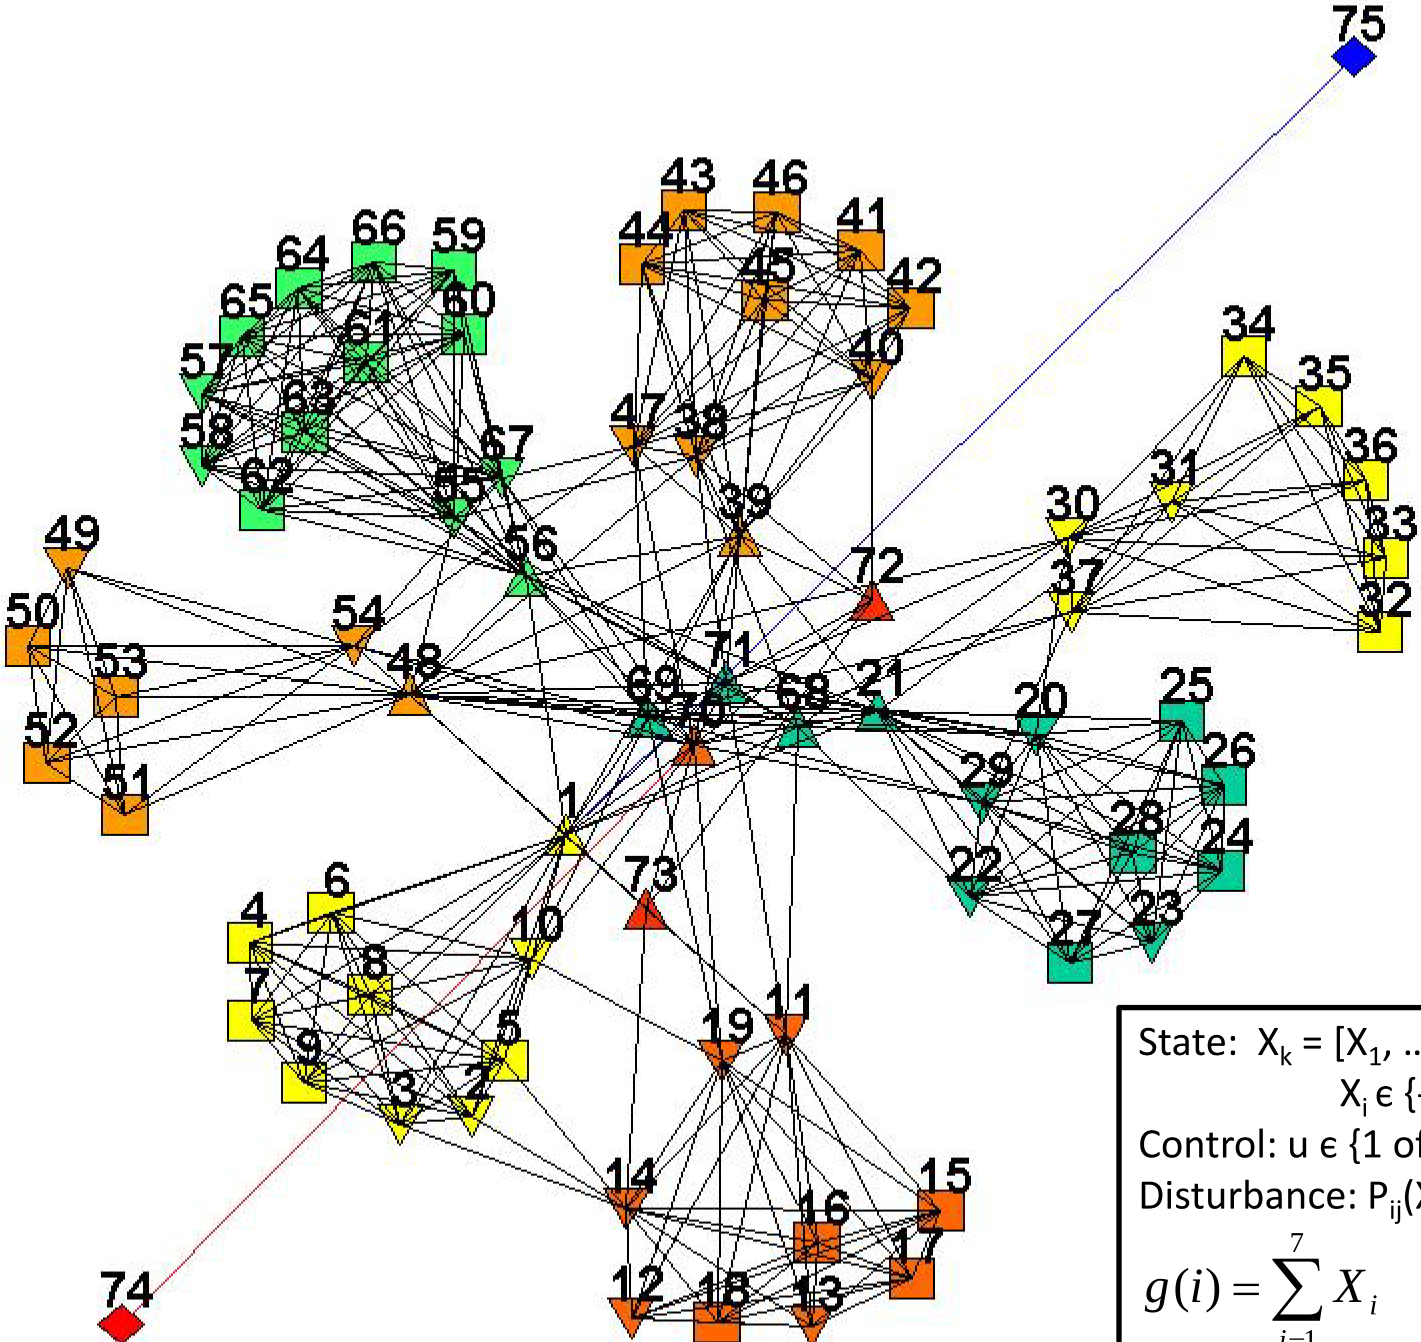

State: $X_k = [X_1, \ldots, X_7]$ $\quad |X_k| = 3^7$

$X_i \in \{-1,0,1\}$

Control: $u \in$ {1 of the 35 Inf. Nodes}

Disturbance: $P_{ij}(X_k,u)$

$$g(i) = \sum_{i=1}^{7} X_i$$

$$J(i) = \arg\max_{u}\left(g(i) + \alpha \sum_{j} P_{ij}(i,u) J(j)\right)$$

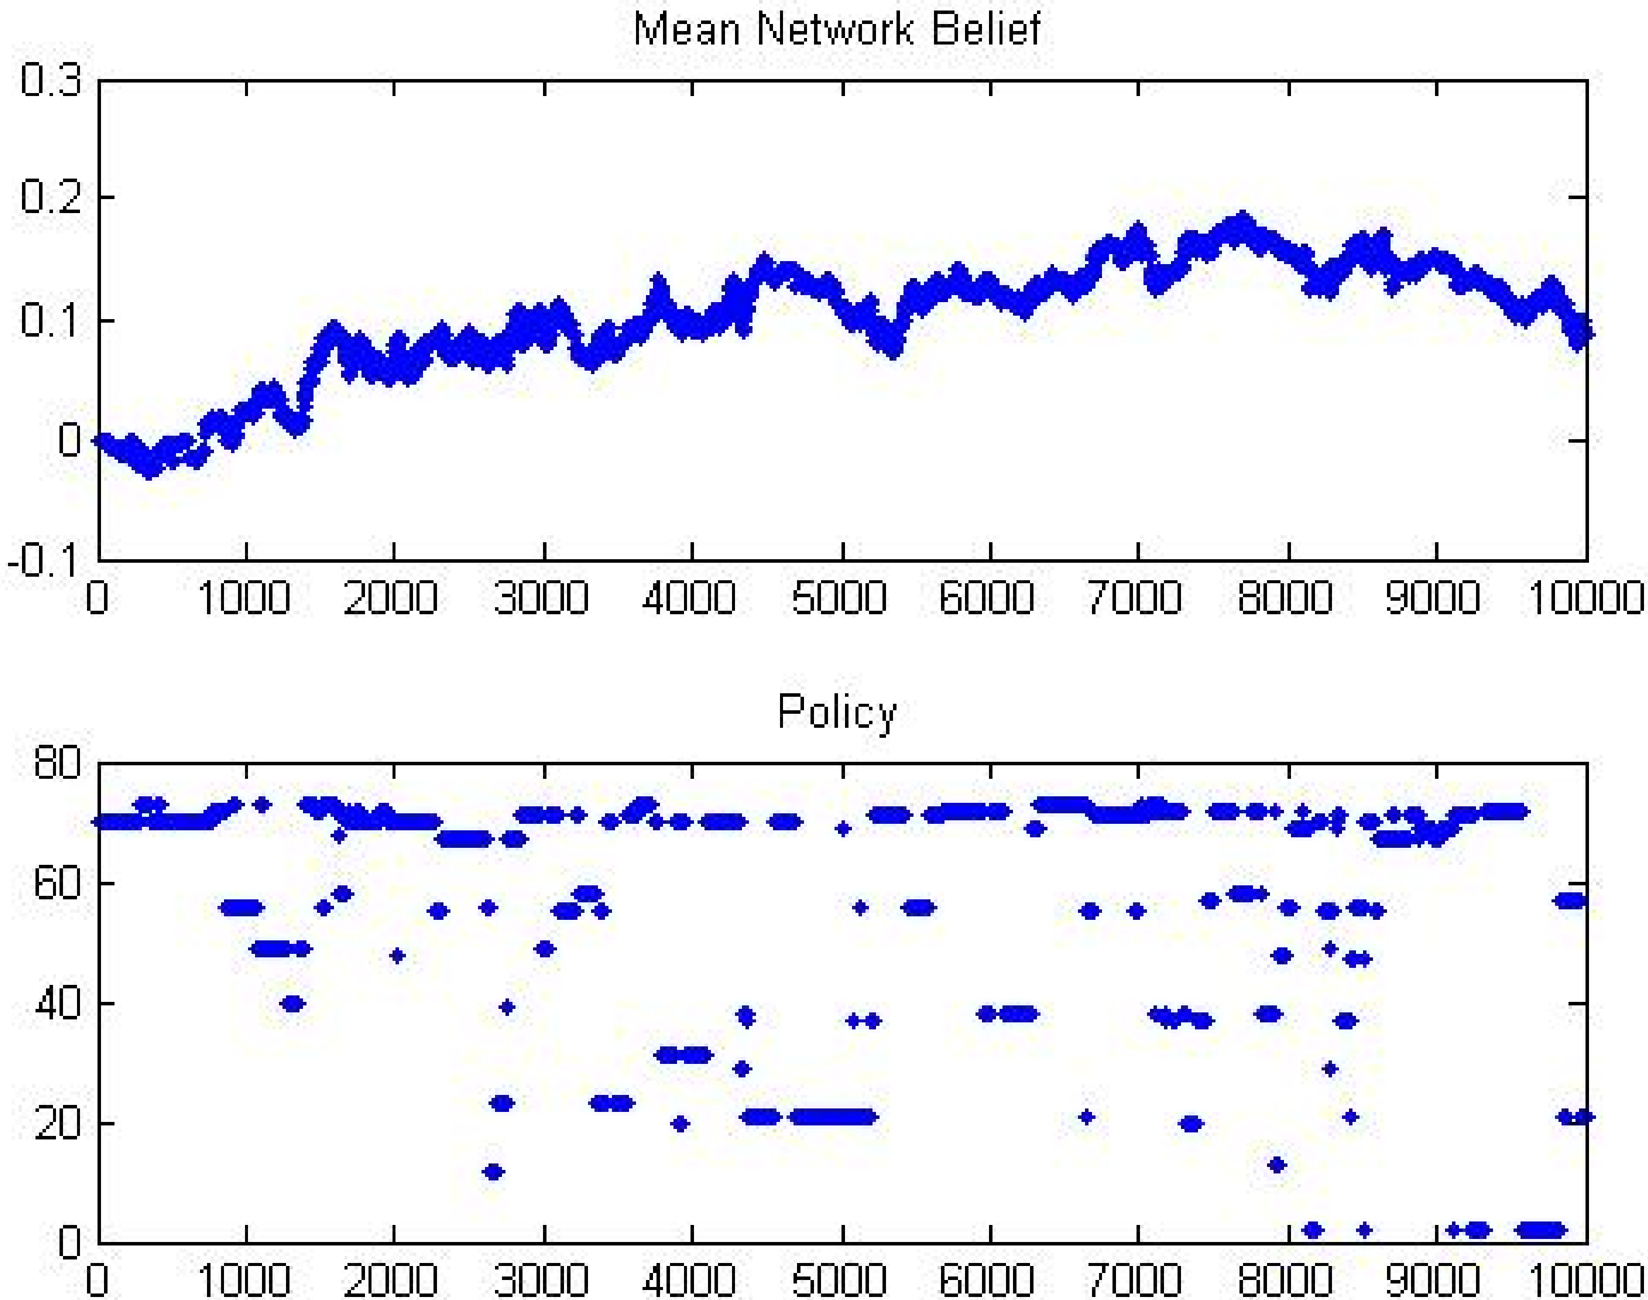
Mean Network Belief
0.3
0.2
0.1
0
-0.1
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000
Policy
80
60
40
20
0
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000

Conclusions

- Even with small amount of information, I can generate a set of stationary policies that do much better than a static strategy
- Once the network goes to a large positive belief, it generally stays there (self-sustaining)
- Flexibility and information have enormous utility in affecting the flow of information in this model

Network Science Center
West Point

Recommendations for Future Work

- Non-equilibrium payoff functions
- Different Strategy Options
- Partial information game
- Dynamic Networks

Network Science Center
West Point