Security and Cooperation in Wireless Networks

http://secowinet.epfl.ch/

A tutorial on game theory for wireless networks

Jean-Pierre Hubaux
EPFL

static games;
dynamic games;
repeated games;
strict and weak dominance;
Nash equilibrium;
Pareto optimality;
Subgame perfection;
…

Chapter outline

Brief introduction to Game Theory

- Discipline aiming at modeling situations in which actors have to make decisions which have mutual, **possibly conflicting**, consequences
- Classical applications: **economics**, but also politics and biology
- Example: should a company invest in a new plant, or enter a new market, considering that the **competition** *may* make similar moves?
- Most widespread kind of game: **non-cooperative** (meaning that the players do not attempt to find an agreement about their possible moves)

Classification of games

Non-cooperative	Cooperative
Static	Dynamic (repeated)
Strategic-form	Extensive-form
Perfect information	Imperfect information
Complete information	Incomplete information

Perfect info: each player knows the identity of other players and, for each of them, the payoff resulting of each strategy.

Complete info: each player can observe the action of each other player.

Cooperation in self-organized wireless networks

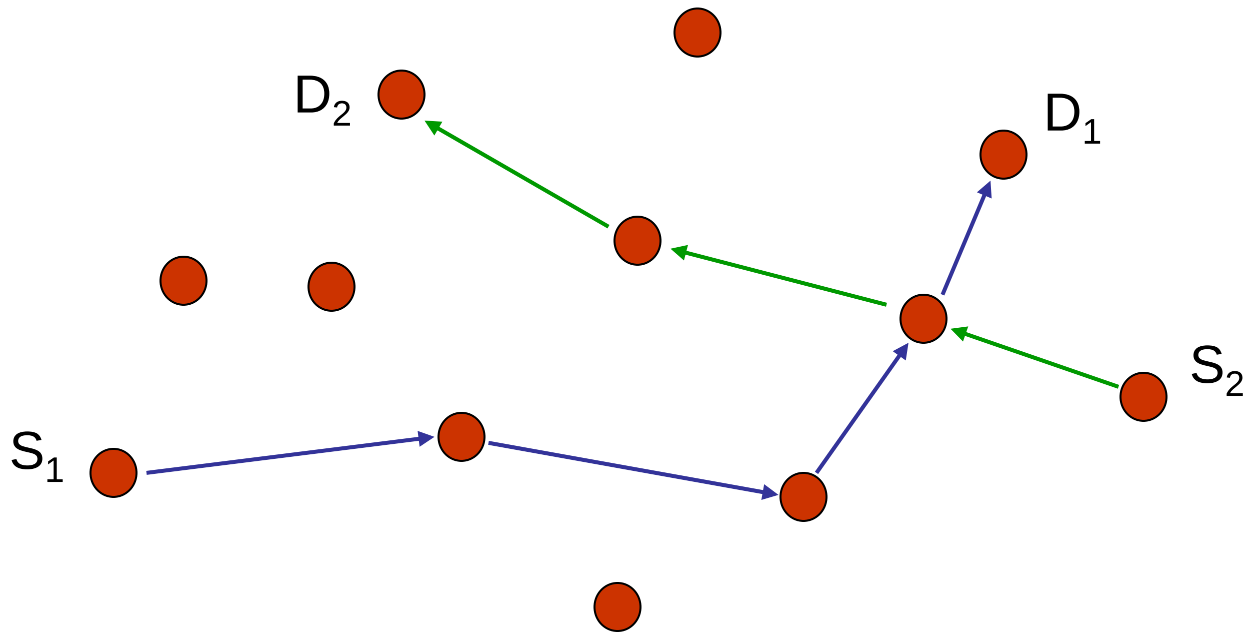

Usually, the devices are assumed to be cooperative.
But what if they are not?

Chapter outline

Example 1: The Forwarder's Dilemma

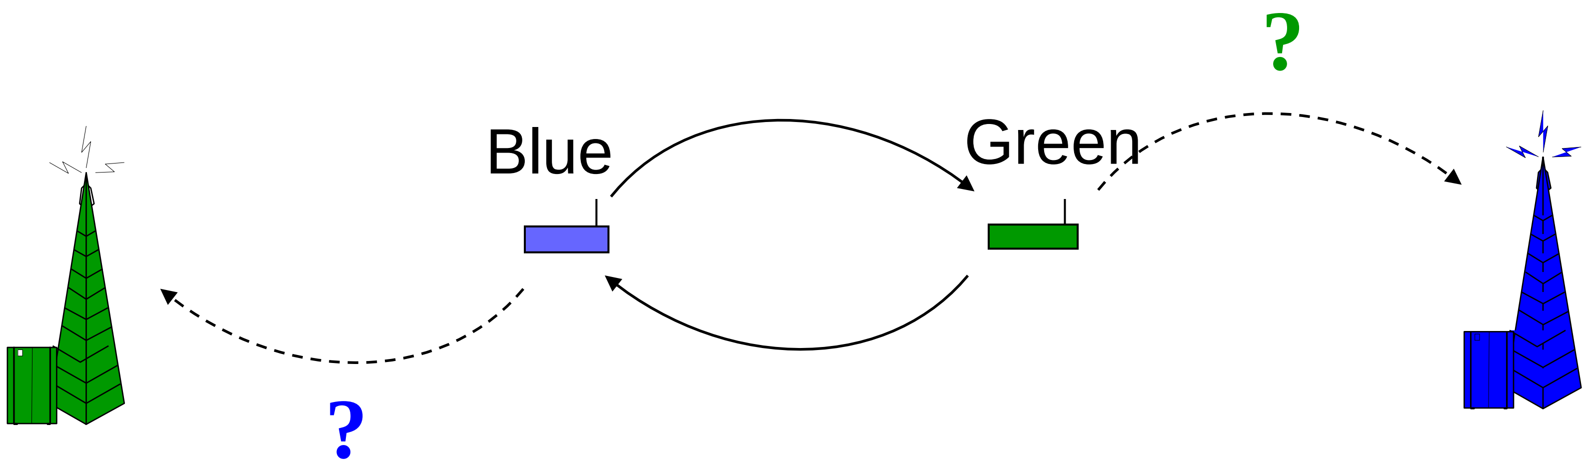

From a problem to a game

- users controlling the devices are ***rational*** = try to maximize their benefit
- game formulation: G = (P,S,U)
  - P: set of players
  - S: set of strategy functions
  - U: set of payoff functions →
    - Reward for packet reaching the destination: 1
    - Cost of packet forwarding: c (0 < c << 1)
- ***strategic-form*** representation

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 1)
Drop	(1, -c)	(0, 0)

Solving the Forwarder's Dilemma (1/2)

Strict dominance: strictly best strategy, for any strategy of the other player(s)

Strategy s_i strictly dominates if

$$u_i(s_i', s_{-i}) < u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}, \forall s_i' \in S_i$$

where: $u_i \in U$ payoff function of player *i*

$s_{-i} \in S_{-i}$ strategies of all players except player i

In Example 1, strategy Drop ***strictly dominates*** strategy Forward

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 1)
Drop	(1, -c)	(0, 0)

Solving the Forwarder's Dilemma (2/2)

Solution by iterative strict dominance:

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 1)
Drop	(1, -c)	(0, 0)

BUT

Drop ***strictly dominates*** Forward

Forward would result in a ***better outcome***

} Dilemma

Result: Tragedy of the commons ! (Hardin, 1968)

Example 2: The Joint Packet Forwarding Game

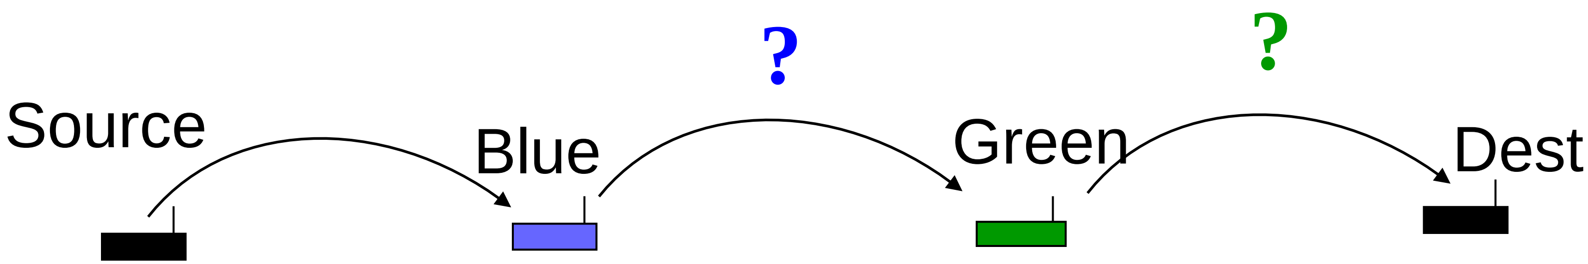

- Reward for packet reaching the destination: 1
- Cost of packet forwarding: c (0 < c << 1)

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 0)
Drop	(0, 0)	(0, 0)

No strictly dominated strategies !

Weak dominance

Weak dominance: strictly better strategy for at least one opponent strategy

Strategy s'_i is weakly dominated by strategy s_i if

$$u_i(s_i', s_{-i}) \le u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}$$

with strict inequality for at least one s_{-i}

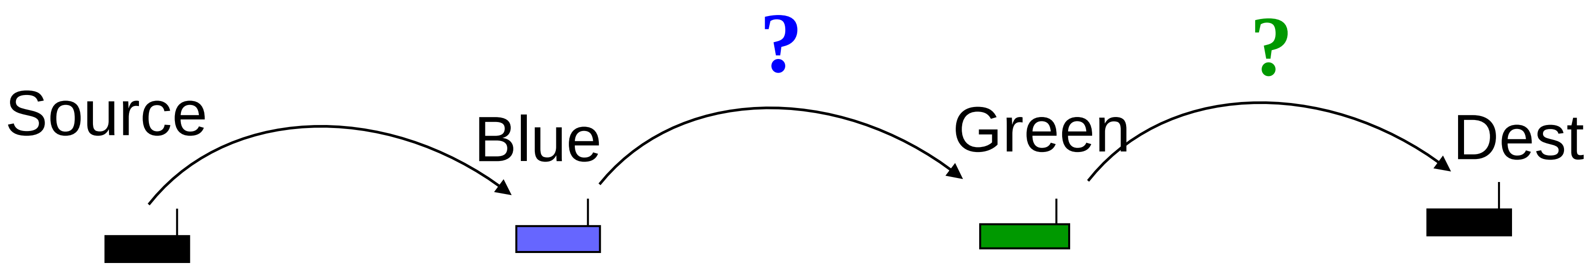

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 0)
Drop	(0, 0)	(0, 0)

Iterative weak dominance

BUT

The result of the iterative weak dominance is not unique in general !

Nash equilibrium (1/2)

Nash Equilibrium: no player can increase its payoff by deviating unilaterally

E1: The Forwarder's Dilemma

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 1)
Drop	(1, -c)	(0, 0)

E2: The Joint Packet Forwarding game

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 0)
Drop	(0, 0)	(0, 0)

Nash equilibrium (2/2)

Strategy profile s^* constitutes a **Nash equilibrium** if, for each player *i*,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \forall s_i \in S_i$$

where: $u_i \in U$ payoff function of player i

$s_i \in S_i$ strategy of player i

The **best response** of player *i* to the profile of strategies s_{-i} is a strategy s_i such that:

$$b_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i})$$

Nash Equilibrium = Mutual best responses

Caution! Many games have more than one Nash equilibrium

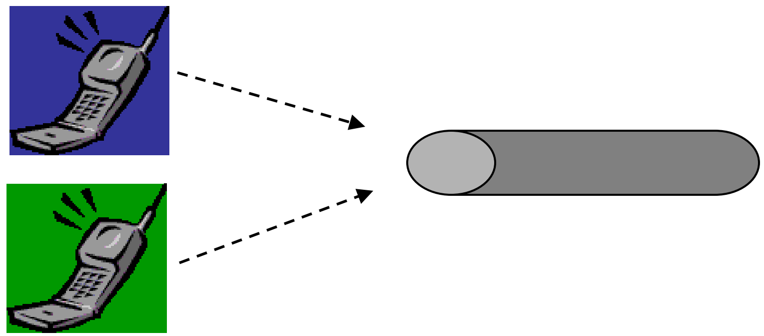

Time-division channel

Reward for successful transmission: 1

Cost of transmission: c
(0 < c << 1)

Blue \ Green	Quiet	Transmit
Quiet	(0, 0)	(0, 1-c)
Transmit	(1-c, 0)	(-c, -c)

There is no strictly dominating strategy

There are two Nash equilibria

Mixed strategy Nash equilibrium

p: probability of transmit for Blue

q: probability of transmit for Green

$$u_{blue} = p(1-q)(1-c) - pqc = p(1-c-q)$$
$$u_{green} = q(1-c-p)$$

objectives

- Blue: choose p to maximize u_{blue}
- Green: choose q to maximize u_{green}

$$p = 1-c, \; q = 1-c$$

is a Nash equilibrium

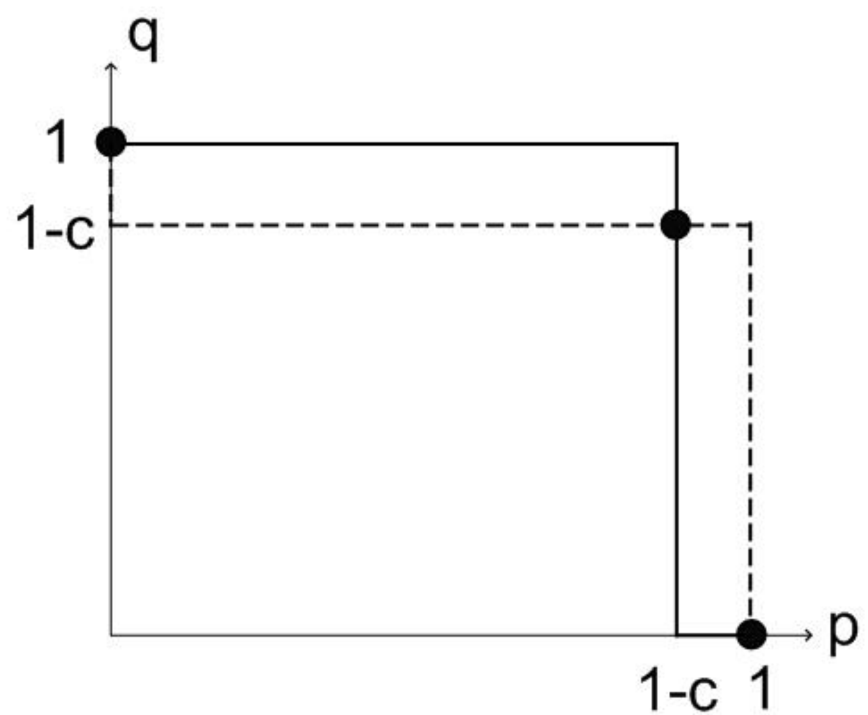

Example 4: The Jamming game

transmitter 

jammer 

two channels: C_1 and C_2

transmitter:
- reward for successful transmission: 1
- loss for jammed transmission: -1

jammer:
- reward for successful jamming: 1
- loss for missed jamming: -1

Blue \ Green	C_1	C_2
C_1	(-1, 1)	(1, -1)
C_2	(1, -1)	(-1, 1)

There is no pure-strategy Nash equilibrium

$$p = \frac{1}{2}, \ \ q = \frac{1}{2} \text{ is a Nash equilibrium}$$

p: probability of transmit on C_1 for Blue

q: probability of transmit on C_1 for Green

Theorem by Nash, 1950

Theorem:
Every finite strategic-form game has a mixed-strategy Nash equilibrium.

Efficiency of Nash equilibria

E2: The Joint Packet Forwarding game

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 0)
Drop	(0, 0)	(0, 0)

How to choose between several Nash equilibria ?

Pareto-optimality: A strategy profile is Pareto-optimal if it is not possible to increase the payoff of any player without decreasing the payoff of another player.

How to study Nash equilibria ?

Properties of Nash equilibria to investigate:

- uniqueness
- efficiency (Pareto-optimality)
- emergence (dynamic games, agreements)

Chapter outline

B.1 Introduction

B.2 Static games

B.3 Dynamic games

B.4 Repeated games

Extensive-form games

- usually to model sequential decisions
- game represented by a tree
- Example 3 modified: the **Sequential** Multiple Access game: Blue plays first, then Green plays.

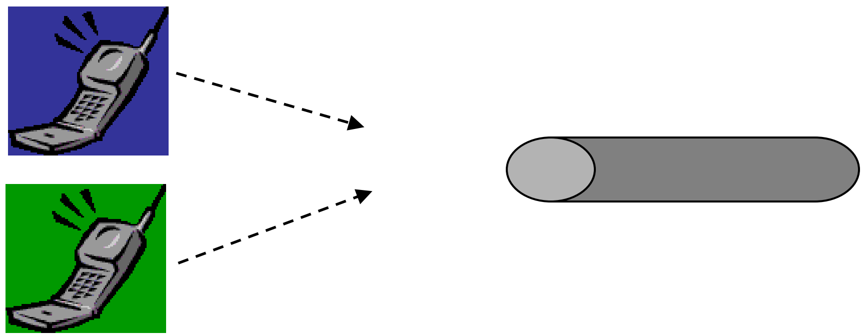

Time-division channel

Reward for successful transmission: 1

Cost of transmission: c
(0 < c << 1)

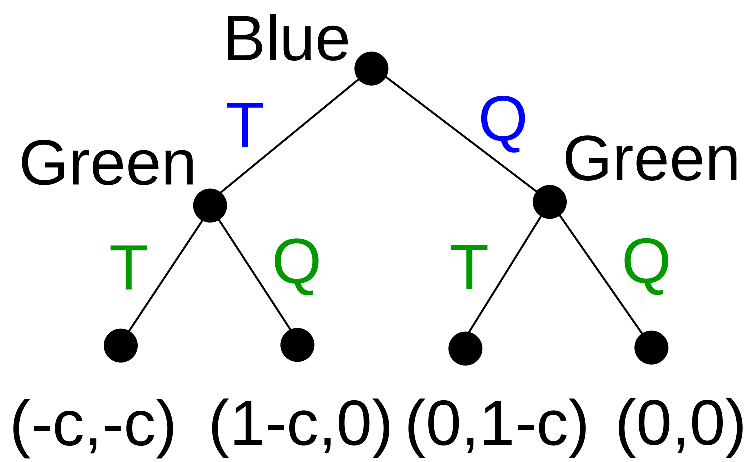

Strategies in dynamic games

- The strategy defines the moves for a player for every node in the game, even for those nodes that are not reached if the strategy is played.

strategies for Blue:
T, Q

strategies for Green:
TT, TQ, QT and QQ

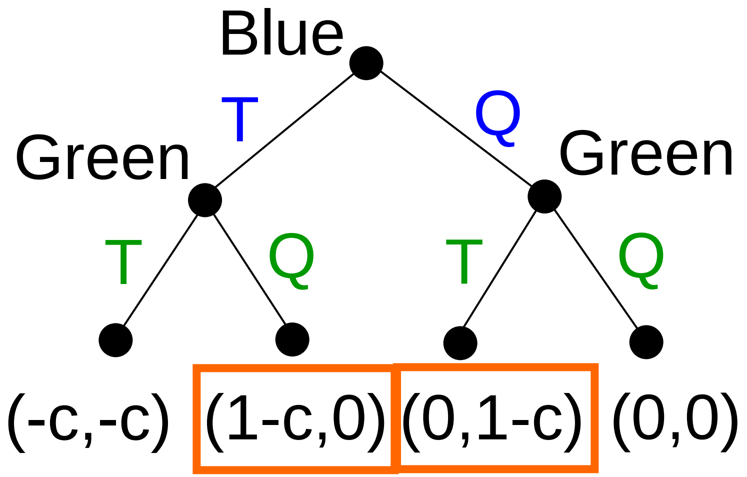

TQ means that player p2 transmits if p1 transmits and remains quiet if p1 remains quiet.

Backward induction

- Solve the game by reducing from the final stage
- Eliminates Nash equilibria that are *increadible threats*

Incredible threat: (Q, TT)

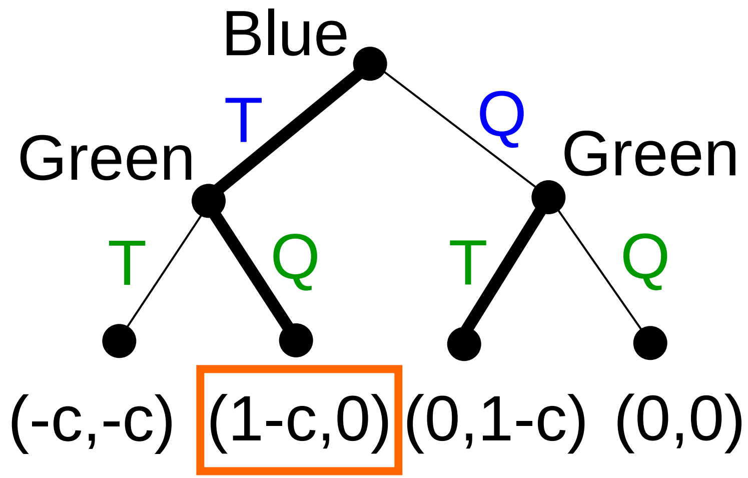

Backward induction solution: h={T, Q}

- Extends the notion of Nash equilibrium

One-deviation property: A strategy s_i conforms to the *one-deviation property* if there does not exist any node of the tree, in which a player i can gain by deviating from s_i and apply it otherwise.

Subgame perfect equilibrium: A strategy profile s constitutes a subgame perfect equilibrium if the one-deviation property holds for every strategy s_i in s.

Finding subgame perfect equilibria using backward induction

Subgame perfect equilibria: (T, QT) and (T, QQ)

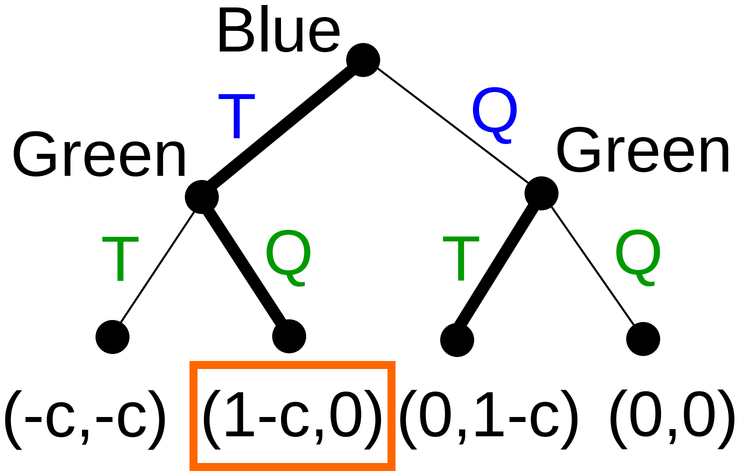

Stackelberg games have one leader and one or several followers

Chapter outline

Repeated games

- repeated interaction between the players (in stages)
- move: decision in one interaction
- strategy: defines how to choose the next move, given the previous moves
- history: the ordered set of moves in previous stages
  - most prominent games are history-1 games (players consider only the previous stage)
- initial move: the first move with no history
- finite-horizon vs. infinite-horizon games
- stages denoted by t (or k)

Utilities: Objectives in the repeated game

- finite-horizon vs. infinite-horizon games
- myopic vs. long-sighted repeated game

myopic: $\bar{u}_i = u_i(t+1)$

long-sighted finite: $\bar{u}_i = \sum_{t=0}^{T} u_i(t)$

long-sighted infinite: $\bar{u}_i = \sum_{t=0}^{\infty} u_i(t)$

payoff with discounting: $\bar{u}_i = \sum_{t=0}^{\infty} u_i(t) \cdot \omega^t$

$0 < \omega \leq 1$ is the discounting factor

Strategies in the repeated game

- usually, history-1 strategies, based on different inputs:
  - others' behavior: $m_i(t+1) = s_i\left[m_{-i}(t)\right]$
  - others' and own behavior: $m_i(t+1) = s_i\left[m_i(t), m_{-i}(t)\right]$
  - payoff: $m_i(t+1) = s_i\left[u_i(t)\right]$

Example strategies in the Forwarder's Dilemma:

Blue (t)	initial move	F	D	strategy name
Green (t+1)	F	F	F	AllC
	F	F	D	Tit-For-Tat (TFT)
	D	D	D	AllD
	F	D	F	Anti-TFT

The Repeated Forwarder's Dilemma

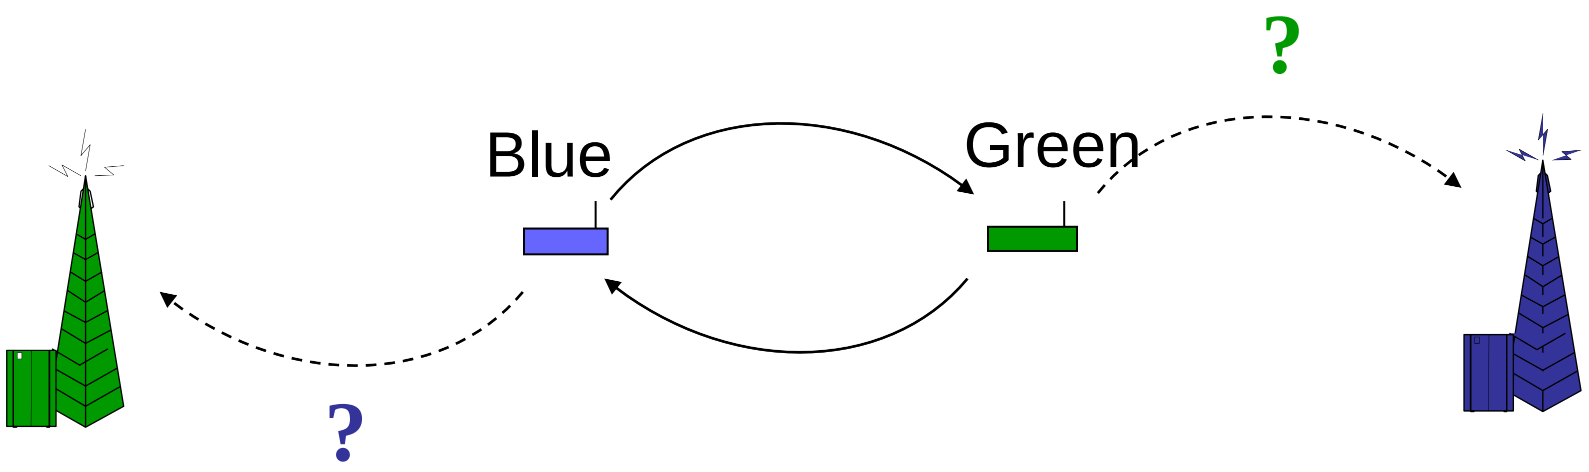

Blue \ Green	Forward	Drop
Forward	(1-c, 1-c)	(-c, 1)
Drop	(1, -c)	(0, 0)

stage payoff

Analysis of the Repeated Forwarder's Dilemma (1/3)

infinite game with discounting: $\bar{u}_i = \sum_{t=0}^{\infty} u_i(t) \cdot \omega^t$

Blue strategy	Green strategy	Blue payoff	Green payoff
AllD	AllD	0	0
AllD	TFT	1	-c
AllD	AllC	1/(1-ω)	-c/(1-ω)
AllC	AllC	(1-c)/(1-ω)	(1-c)/(1-ω)
AllC	TFT	(1-c)/(1-ω)	(1-c)/(1-ω)
TFT	TFT	(1-c)/(1-ω)	(1-c)/(1-ω)

Analysis of the Repeated Forwarder's Dilemma (2/3)

Blue strategy	Green strategy	Blue payoff	Green payoff
AllD	AllD	0	0
AllD	TFT	1	-c
AllD	AllC	$1/(1-\omega)$	$-c/(1-\omega)$
AllC	AllC	$(1-c)/(1-\omega)$	$(1-c)/(1-\omega)$
AllC	TFT	$(1-c)/(1-\omega)$	$(1-c)/(1-\omega)$
TFT	TFT	$(1-c)/(1-\omega)$	$(1-c)/(1-\omega)$

- AllC receives a high payoff with itself and TFT, but
- AllD exploits AllC
- AllD performs poor with itself
- TFT performs well with AllC and itself, and
- TFT retaliates the defection of AllD

TFT is the best strategy if ω is high !

Analysis of the Repeated Forwarder's Dilemma (3/3)

Blue strategy	Green strategy	Blue payoff	Green payoff
AllD	AllD	0	0
TFT	TFT	(1-c)/(1-ω)	(1-c)/(1-ω)

Theorem: In the Repeated Forwarder's Dilemma, if both players play AllD, it is a Nash equilibrium.

Theorem: In the Repeated Forwarder's Dilemma, both players playing TFT is a Nash equilibrium as well.

The Nash equilibrium s_{Blue} = TFT and s_{Green} = TFT is Pareto-optimal (but s_{Blue} = AllD and s_{Green} = AllD is not) !

Experiment: Tournament by Axelrod, 1984

- any strategy can be submitted (history-X)
- strategies play the Repeated Prisoner's Dilemma (Repeated Forwarder's Dilemma) in pairs
- number of rounds is finite but unknown

↓

- TFT was the winner
- second round: TFT was the winner again

R. Axelrod *The Evolution of Cooperation*
Basic Books, 1984

An Example beyond Engineering

Country 1 \ Country 2	Reduce military investment	Increase military investment
Reduce military investment	(1, 1)	(-1, 2)
Increase military investment	(2, -1)	(0, 0)

Payoffs:
2: I have weaponry superior to the one of the opponent
1: We have equivalent weaponry and managed to reduce it on both sides
0: We have equivalent weaponry and did not managed to reduce it on both sides
-1: My opponent has weaponry that is superior to mine

Discussion on game theory

- Rationality
- Payoff function and cost
- Pricing and mechanism design (to promote desirable solutions)
- Infinite-horizon games and discounting
- Reputation
- Cooperative games
- Imperfect / incomplete information

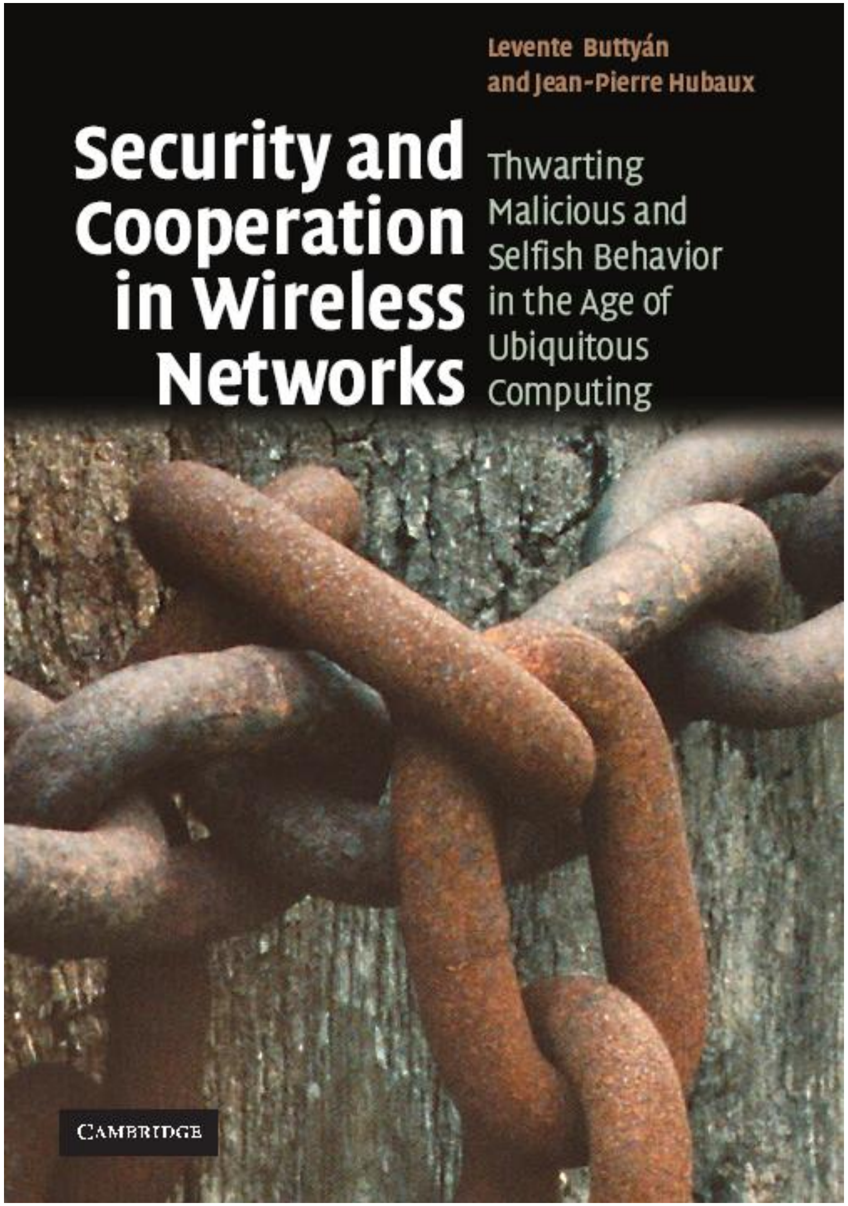

http://secowinet.epfl.ch

Conclusions

- Game theory can help modeling greedy behavior in wireless networks
- Discipline still in its infancy
- Alternative solutions
  - Ignore the problem
  - Build protocols in tamper-resistant hardware