Theory of Networked Computation?

Joan Feigenbaum
http://www.cs.yale.edu/homes/jf

FuDiCo; Bertinoro, Italy; June 2007

Future of Distributed Computing?

Conjecture: "Distributed-computing research" will be more like "networking research."

What characterizes "networking research"?

- "Sacrifice strict semantics for scalability."
  [Scott Shenker, PODC 2003]
- "Evolutionary fitness trumps elegance."
  [Jonathan Smith, Colloquium talk 2007]

Why is this Plausible?

- It's already happening; see PODC and EC proceedings of the last five years.
- Cultural trend: "networkization" of CS
- Funding trend: GENI, FIND, "Clean-Slate Design," … many 100s M$ worldwide
- Networks provide real-world examples of distributed computations.
- Intellectually compelling

Role of the theory community?

Elements of a "Theory of NC"

- Model(s) of computation
- General network-algorithmic techniques
- Algorithms for concrete problems of interest
- Lower-bound techniques, reductions
- Hardness results for concrete problems of interest

❖ Descriptive results and interpretation

Examples of "Networked Comp."

- Routing, congestion control, and other "network-layer" computations
- WWW search
- Auctions
- P2P file sharing
- Blogs, wikis, MySpace, and other "web-mediated communities"
- Yahoo! questions, ESP, del.icio.us, and other "human-aided computations"

5

Properties to Model

- Massive scale
- User self-interest
- Subnetwork autonomy
- Emergent behavior
- Extreme heterogeneity (of devices, uses, subnetworks, …)

❖ Results without convergence

❖ Agents, data, resources, and network conditions that change during computation

6

Outline

- **Theory of incentive-compatible IDR**
- What IDR has contributed to ToNC
- What IDR has not (yet?) contributed to ToNC
- Some ToNC-agenda items (technical and political)

7

Interdomain Routing

Establish routes between autonomous systems (ASes).

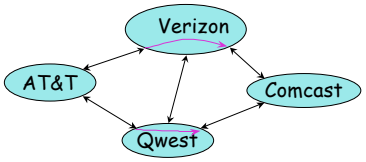

Currently done with the Border Gateway Protocol (BGP).

8

Why is Interdomain Routing Hard?

- Route choices are based on *local policies*.
- Autonomy: Policies are uncoordinated.
- Expressiveness: Policies are complex.

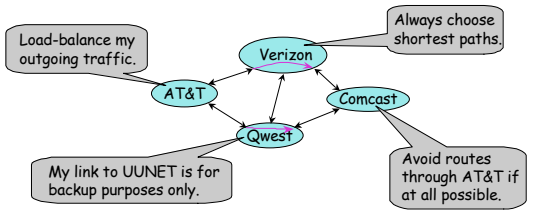

Desiderata (from Netw. Community)

- Distributed, adaptive route computation
- Destination-based forwarding; confluent tree $T_d = \{R_1, \ldots, R_n\}$ for each d
- Efficient use of time, space, and communication
- Loop-free routes, even in the presence of autonomous, potentially conflicting, routing decisions by ASes
  - ➔ Path-vector routing protocol

BGP Route Processing (1)

- The computation of a single node repeats the following:

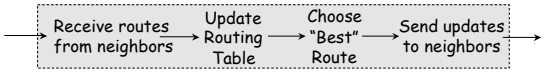

- Paths go through neighbors' choices, which enforces consistency.
- Decisions are made locally, which preserves autonomy.
- Uncoordinated policies can induce protocol oscillations. (Much recent work addresses BGP convergence.)
- Recently, private information, optimization, and incentive-compatibility have also been studied.

BGP Route Processing (2)

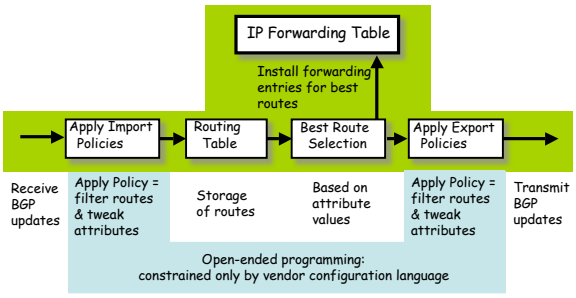

Protocol-Divergence Example

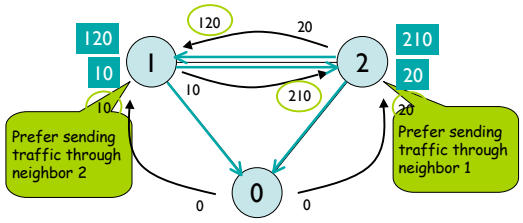

Dispute Cycles

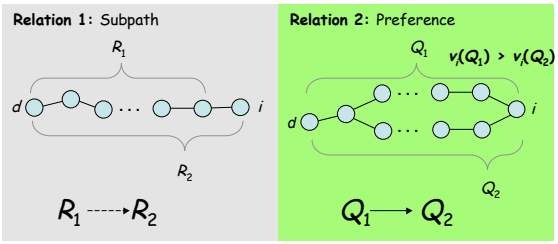

- Valuations **do not induce a dispute cycle** if there is no cycle formed by the above relations on all permitted paths in the network.
- No dispute cycle => robust BGP convergence [GSW02, GJR03]

Example of a Dispute Cycle

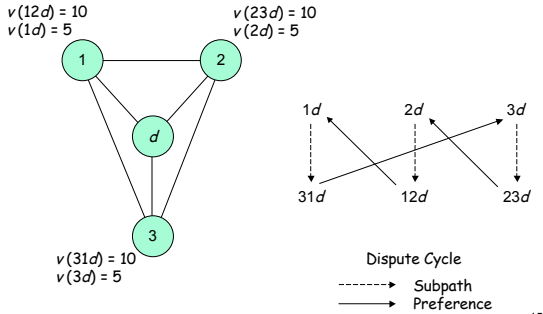

Gao-Rexford Framework (1)

Neighboring pairs of ASes have one of:

- a *customer-provider* relationship
  (One node is purchasing connectivity from the other node.)
- a *peering* relationship
  (Nodes have offered to carry each other's transit traffic, often to shortcut a longer route.)

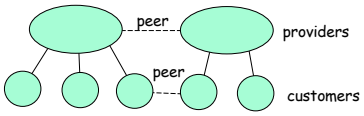

Gao-Rexford Framework (2)

- Global constraint: **no customer-provider cycles**
- Local *preference* and *scoping* constraints, which are consistent with Internet economics:

Preference Constraints

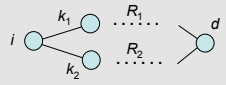

- If k_1 and k_2 are both customers, peers, or providers of i, then either ik_1R_1 or ik_2R_2 can be more valued at i.
- If one is a customer, prefer the route through it. If not, prefer the peer route.

Scoping Constraints

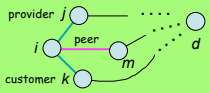

- Export customer routes to all neighbors and export all routes to customers.
- Export peer and provider routes to all customers only.

- Gao-Rexford conditions => BGP always converges [GR01]

17

Networking Community's Contributions

+ TCS-style theorems and proofs
+ Results that capture economic and engineering realities
- Seek stability but not optimality.
- May not properly incentivize ASes to follow the protocol.

18

Broadening the Effort

Seek an interdomain-routing protocol that is

- Computationally feasible
- Incentive-compatible
- Robustly scalable

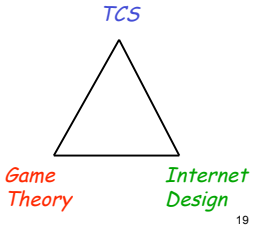

19

Economic Mechanism Design

Agents: 1, …, n
Strategies: s_1, …, s_n
Types: t_1, …, t_n
Actions: a_1, …, a_n
Outcome: o
Valuation functions: v_1, …, v_n
Payment functions: p_1, …, p_n
Utility functions: u_1, …, u_n

20

Classical, One-Round Mechanisms (1)

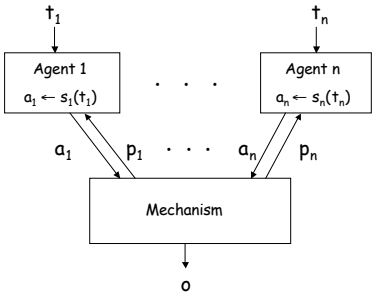

Classical, One-Round Mechanisms (2)

- Action vector $(a_1, \ldots, a_n)$ is "consistent with selfishness."
  - a_i maximizes $u_i(o, t_i) \equiv v_i(o, t_i) + p_i$.
  - Meaning of "maximize" depends on "solution concept," *e.g.*, NE, BNE, DSE, epNE, …
- Mechanism-design goal: $o(a_1, \ldots, a_n) \in G(t_1, \ldots, t_n)$
- Classical economic-MD question: For a given solution concept, which design goals can be achieved?

Solution Concepts

- $(a_1, \ldots, a_n)$ is a Nash Equilibrium (NE) if
  - $\forall_i$ and a_i' : $u_i(o(a_1, \ldots, a_i, \ldots, a_n), t_i) \geq u_i(o(a_1, \ldots, a_i', \ldots, a_n), t_i)$
  - Given other agents' actions, agent i is best off playing a_i.
- $(a_1, \ldots, a_n)$ is a Dominant-Strategy Equil. (DSE) if
  - $\forall_i$, a_i', and $(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$:
    $u_i(o(a_1, \ldots, a_i, \ldots, a_n), t_i) \geq u_i(o(a_1, \ldots, a_i', \ldots, a_n), t_i)$
  - Regardless of other agents' actions, agent i is best off playing a_i.
- $(s_1, \ldots, s_n)$ is an ex-post Nash Equil. (epNE) if
  - $\forall_i$, s_i', and $(t_1, \ldots, t_n)$:
    $u_i(o(s_1(t_1), \ldots, s_i(t_i), \ldots, s_n(t_n)), t_i) \geq u_i(o(s_1(t_1), \ldots, s_i'(t_i), \ldots, s_n(t_n)), t_i)$
  - Given that other agents follow the prescribed strategy, agent i is best off doing so, too, regardless of the other players' types.

Algorithmic Mechanism Design [NR01]

- Required polynomial-time o( ) and p( ).
- Focused on strategyproof, direct-revelation mechanisms.
- Put forth polynomial-time, strategyproof LCP-routing MD as a good abstraction of Internet routing.

VCG Mechanism for LCP [NR01]

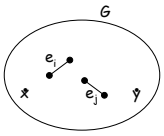

Agent i ≡ edge e_i

Type $t_i \equiv cost(e_i)$

Outcome o ≡ LCP from x to y (wrt reported costs)

$$p_i(o) = \begin{cases} 0 \text{ if } e_i \notin o \\ cost(o(G_{i<-\infty})) - cost(o(G_{i<-o})) \text{ if } e_i \in o \end{cases}$$

25

Moving Closer to Real IDR [FPSS02, SP04]

- Nodes (ASes), not edges, are the agents.
- All-source, LCP tree T_d to each destination d.
- No trusted center; ASes compute the routes themselves.
- Use BGP as an algorithmic substrate to preserve "evolutionary fitness" and encourage adoption.

⇒ BGP-compatible VCG mechanism for LCP trees that is incentive-compatible in epNE

26

Augment BGP Routing Tables

Dest.	cost	LCP and path prices				LCP cost
AS1		AS3	AS5	AS1		$c(i,1)$
	c_1	p^3_{i1}	p^5_{i1}			

1. LCPs are computed and advertised to neighbors.
2. Initially, all prices are set to ∞.
3. In the following stages, each node repeats:
   - Receive LCP costs and path prices from neighbors.
   - Recompute candidate prices; select lowest price.
   - Advertise updated prices to neighbors.

Final state: Node *i* has accurate p^k_{ij} values.

27

Performance of Algorithm

$$f = max_{i,j} \, || P(c; i, j) ||$$

$$f' = max_{i,j,k} \, || P^{-k}(c; i, j) ||$$

Theorem [FPSS02]:

This algorithm computes the VCG prices correctly, uses routing tables of size $O(nf)$ (a constant factor increase over BGP), and converges in at most $(f + f')$ stages (worst-case additive penalty of f' stages over the BGP convergence time).

28

"Optimality" wrt LCP is too narrow.

- More generally, seek $T_d = \{R_1, \ldots, R_n\}$ that maximizes $\Sigma_i\ v_i(R_i)$.
- Fully general v_i (routing policies)
- Next-hop policies: For all i and j,
  $v_i(ijk_3 \ldots k_m d) = v_i(ijl_3 \ldots l_t d)$
- Forbidden-set policies: For every source i, there is a set S_i of transit nodes such that $v_i(R_1) > v_i(R_2)$ if R_2 contains a node in S_i but R_1 doesn't.

29

Policy Consistency

Valuations are **policy consistent** iff, for all routes R_1 and R_2 (whose extensions are *not rejected*),

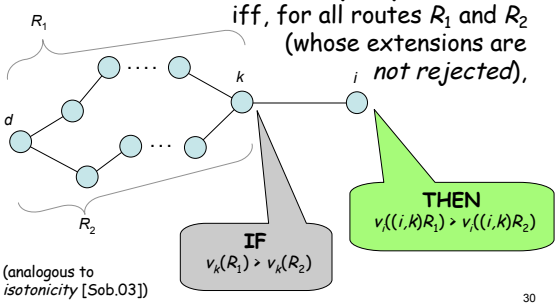

(analogous to *isotonicity* [Sob.03])

30

Requirements	Results	References
Lowest-cost routes	BGP-compatible VCG computation	FPSS02, SP04
Fully general routing policies	NP-hard even to approximate	FSS04
Next-hop routing policies	Centralized, ptime VCG computation. Not BGP-compatible	FSS04
Forbidden-set routing policies	NP-hard even to approximate	FKMS05
Gao-Rexford + policy consistency	BGP is inc.-comp. in collusion-proof epNE without payments.	FRS06, FSS07
Gao-Rexford + route verification	BGP is inc.-comp. in collusion-proof epNE without payments.	LSZ07

31

Outline

- Theory of incentive-compatible IDR
- **What IDR has contributed to ToNC**
- **What IDR has not (yet?) contributed to ToNC**
- Some ToNC-agenda items (technical and political)

32

Contributions

- Pushes the envelope on incentives in computation.
- Combines the relevant research areas (algorithms, networking, and economics) in a serious way.
- Helps explains why interdomain routing "works," despite the "proofs" that BGP is "wrong."
- Exemplifies "protocol-based algorithms design."

33

Unresolved Issues

- Computational modeling challenges, *e.g.,*
  - Results without convergence
  - AS graphs, policies, and loads that change during the computation
- Is epNE (or even "equilibrium") really a useful concept in networked computation?
- Are there any general techniques or insights here, or is IDR unique?
- Interaction with AS-graph formation

34

Talk Outline

- Theory of incentive-compatible IDR
- What IDR has contributed to ToNC
- What IDR has not (yet?) contributed to ToNC
- **Some ToNC-agenda items (technical and political)**

35

Integrating Diverse Theories

- At FuDiCo III, we've seen the influence of TCS, Game Theory, Distributed Computing, Networking, Cryptography, Model Checking, …
- If we try to combine all of the formalism and assumptions of these diverse fields, we will NOT be able to prove (or even state) meaningful theorems.

➢ **For each networked-computational problem, figure out what you need to be precise about and what you can fudge or ignore.**

36

TCS-Style Theorems Are Inadequate

- We will NEVER have a fixed IDR "instance." (Piatek made the same point about BitTorrent.) So what do "BGP-convergence" theorems mean?
- NC problems that are provably hard for *networking* reasons (*i.e.*, not because they're NP-hard) are few and far between.

➢ **Develop a complexity theory of networked computation: Relevant computational resources, "results without convergence," general algorithmic techniques, canonical hard problems and reductions, …**

37

Protocol-based Algorithm Design

- "BGP-compatible" algorithmic mechanisms can leverage the evolutionary fitness of today's IDR framework and would be easier to deploy than algorithmic mechanisms designed from scratch.
- Are there other pieces of the computational infrastructure that can be used in this way? Candidates: Search, keyword auctions, …

➢ **Consider the use of widely deployed, successful protocols as "computational substrates" for novel network algorithms.**

38

SPUR in Networked Computation

- Patterson says that Security, Privacy, Usability, and Reliability will be crucial for the success of 21st-century C&C. He's right!
- The "preventive approach" that dominated 20th-century research on security and privacy may be useless. Networks are popular precisely *because* they enable people and organizations to share information.

➢ **Consider after-the-fact accountability as an alternative to before-the-fact authorization. (Haeberlen made the same suggestion.) Use the financial world's approach to "commercial paper" as a starting point.**

39

Non-technical Challenges (1)

- They dream of GENI.
  - For us, it might be a nightmare.
  - It's not clear that the theory community will get ANY dedicated funding for ToNC. GENI and related programs may actually hurt us.
  - Is the funding situation better in Europe?
- We can be bigots, too.
  - None of the aforementioned IDR papers has appeared in a STOC or FOCS proceedings.
  - What would the (elites of the) European theory community think about ToNC?

40

Non-technical Challenges (2)

- We should be working with forward-looking companies.
  - The industrial-research culture that nurtured Griffin, Rexford, and other intellectual-boundary crossers is gone.
  - Is there a 21st-century alternative?

41

Conclusions

- Good opportunity to do novel theoretical work that has practical impact
- Multidisciplinarity is exciting but creates technical challenges.
- Further thoughts about ToNC, including the results of two NSF-sponsored workshops in 2006, can be found at http://www.cs.yale.edu/~jf/ToNC.html

42