Fair Load Distribution Among Cooperating Mobile Nodes

Karin Anna Hummel

Klagenfurt, 29.06.2009

karin.hummel@univie.ac.at
Distributed Systems Group
University of Vienna

Computer Science @UniVie

- 3 Institutes, 2 Research Labs
- Inst. of Distributed and Multimedia Systems
  – Multimedia Systems Group: W. Klas
  – Future Communication: K. Tutschku
  – Distributed Systems Group: G. Haring
- <Mobile Computing>, <Cooperative Systems>, etc. … *to come*

Methods (Future Comm., Distr. Systems): Analytical network (performance) modeling, simulation (distributed, mobility modeling), prototyping, measurements, user studies and statistical evaluation

Application Area – Mobile Distributed and Grid Computing

Grid Computing: How to integrate?

Austrian Grid: screening of 14 scientific research groups (2004; astrophysics, medicine, environmental research, etc.)

	Clients		Resources		
Mode	Observe and monitor	Control and delegate	Integrate as resource	Pure ad-hoc grid	No integration
No of groups	8	3	2	0	4

Intelligent Transport Systems / Car-to-Car (opportunistic networking)

Wireless Mesh Networks (limited router mobility, client mobility)

Spontaneous (ad-hoc) wireless networking: emergency and short-term community services (congresses)

Outline

- Motivation
  - Mobile distributed and grid computing
- Challenges of mobile distributed systems
  - Problems and approaches: self-organization and mobility-awareness
- Robust, decentralized load distribution
- Fairness
  - Fairness strategies – results for fault-free and faulty scenarios
- Unfair behavior and game theory
  - Investigations on consequences of TFT, gTFT, go by majority strategies in spite of unfair nodes
- Conclusions and current work

Challenges in Mobile Distributed Computing

Challenges

- Limited bandwidth, varying link quality
- Limited device capabilities, limited energy sources

→ Varying availability of resources (mobility, battery)

Goals for *computation* and *data dissemination*

- Robustness against disconnections
- Fairness among nodes

Achievement of goals without central controlling instances and local interactions only

- Self-organization of devices
- Mobility-awareness
- Energy-awareness
- (Context-awareness: supporting multi-homing, trust in free WMNs)

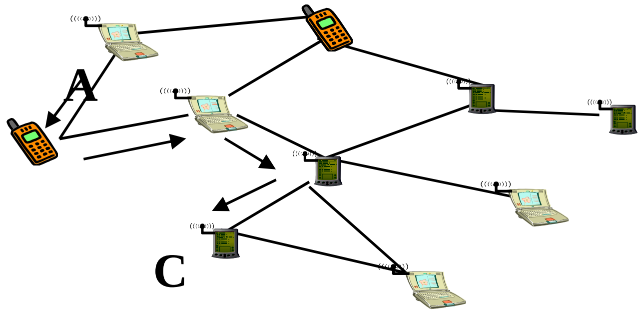

Mobility-awareness and Pathway Patterns

Create predictor for proactive data placement for mobile clients

- Based on relative frequencies of past movement
  - One month GPS taxi traces (Vienna taxi fleet)
- Assuming static information servers moving / replicating data
- Improvements to non-predictive strategy (*A*) possible (actual and predicted position *AP*) in terms of data distance *d* (hops)

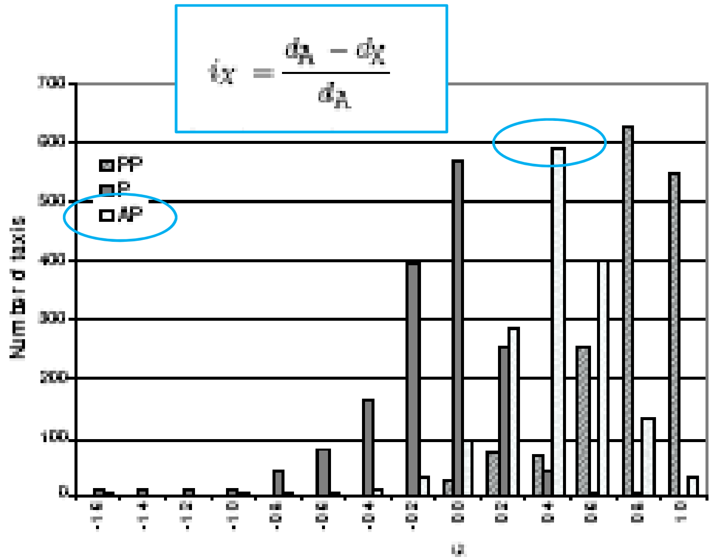

[Gossa08] J. Gossa, A. Janecek, K.A. Hummel, W. Gansterer, J.-M. Pierson. Proactive Replica Placement Using Mobility Prediction. MDM Workshops 2008

Robust, Decentralized Task Scheduler

Overview

- Based on distributed virtual shared memory
  - > Persistence of data, asynchronous communication
- Coordination based on distributed queues
- Workers decide autonomously when to take a job, considering:
  - > User policies
  - > Job requirements
  - > Current and predicted performance values
- Proactive Fault Tolerance (FT): redundant job execution to prevent job loss
- Reactive FT: handle system failures
- Very reliable nodes run critical tasks (e.g., FT services)

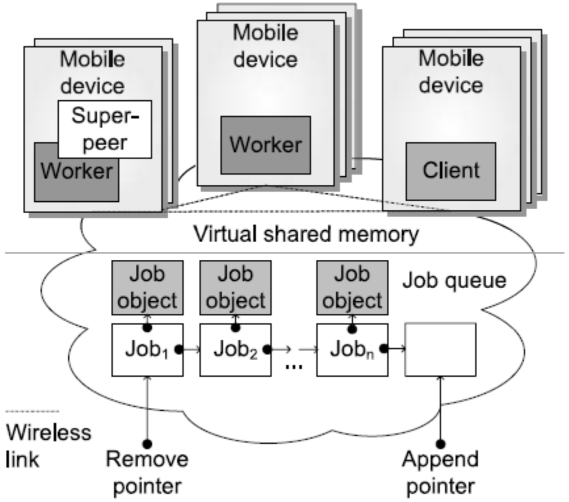

Prototype Implementation

- Distributed virtual shared memory (Corso; similar to JavaSpaces but distributed and objects own IDs)
- Condor Class Ads
- Globus compliant
- Java class loader

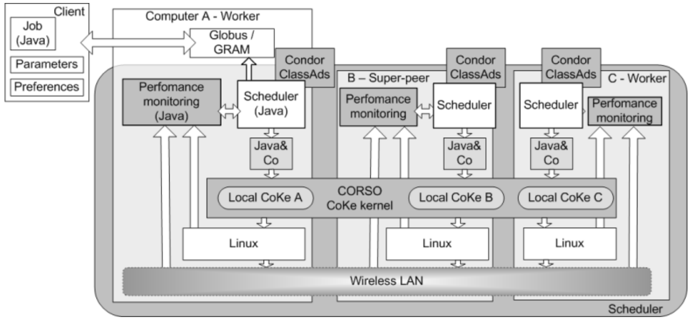

Speedup

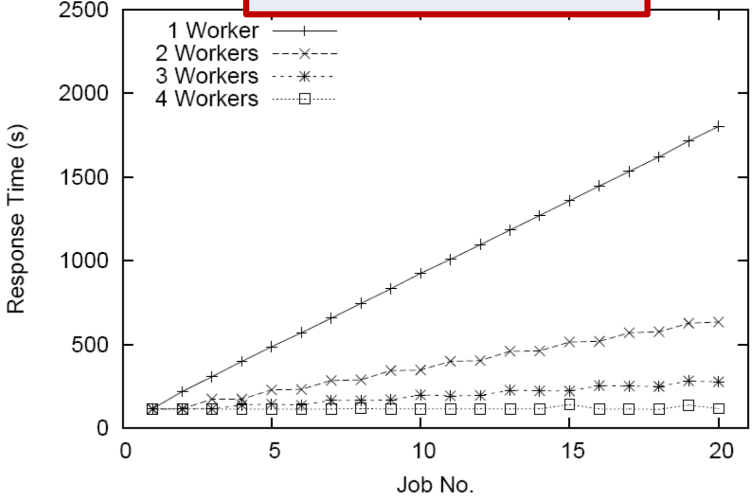

Intermittent Connectivity

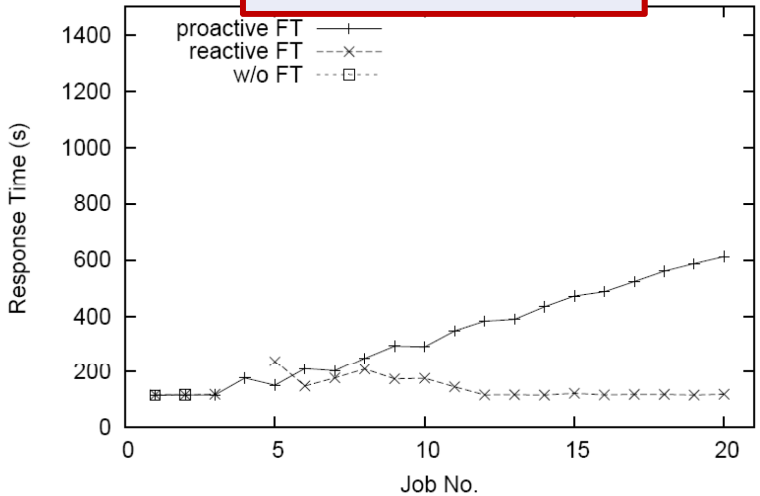

[HuJe07] K.A. Hummel, G. Jelleschitz. A Robust Decentralized Job Scheduling Approach for Mobile Peers in Ad-hoc Grids. In Proceedings CCGrid 2007

Fairness - Introduction

Assumption

- Devices with equal capabilities/contributions

Fairness means

- Balanced workload between devices
- Measured in terms of variance of processed jobs per node

Information dissemination

- Each device maintains a tree *T* where every node represents a device (to describe neighborhood)
  - > Root represents device itself
  - > Children represent group members
  - > Each node keeps track of recent performance data (self-descriptiveness)
- Each node periodically sends performance data to its group members in *T*
- Nodes aggregate values received

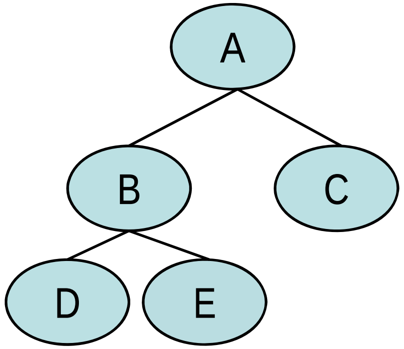

[Hum08a] K.A. Hummel, H. Meyer. Self-organizing fair job scheduling among mobile nodes. SASO Workshops 2008.

Gossip Among Group Members

Based on non-disjoint groups

- Assures spreading of decision throughout the system
- Avoids communication overhead (e.g., when compared to gossiping with all nodes)
- Group size n
- Example n = 3

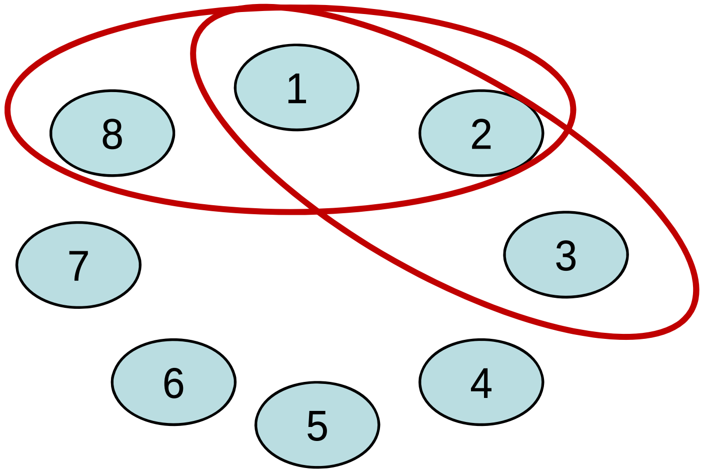

- Should provide a system structure allowing self-organization

Fairness - Strategies

Idea

- Decision whether to take or skip a job is based on chosen strategy
- Each strategy evaluates all performance values in *T* (utilization, battery, connectivity and expected future connectivity)

Classification

- Lazy strategy
  - *not best*: job is not taken, if at least one device in *T* is better
- Assiduous strategy
  - *worst*: job is not taken, if all devices in *T* are better
- Evaluation of average or majority
  - *Worse than average*: job is not taken, if average of devices in *T* is better
  - *Worse than majority*: job is not taken, if majority of devices in *T* are better
  - *Equal or worse than majority*: job is not taken, if majority of devices in *T* are equal or better

Deadlock prevention

- If job remains in queue for a defined time, job management without fairness is temporary activated, deactivate strategy

Experiments

Setup

- Discrete Event Simulation
- Job: RC5 key decoding application (24 bit key), requires 116.632 seconds on worker device
- 60 Jobs, one every 110 seconds

Experiment 1: Faultless Scenario

- Measurement of fairness (variance of number of processed jobs) and message overhead for varying group sizes

Experiment 2: Faulty Scenario

- Disconnections simulated via timeline (e.g., due to mobility)
- Assumed one faulty node
- Fixed group size of 10
- Measurement of fairness
  - > error free case
  - > no fault tolerance
  - > proactive fault tolerance
  - > reactive fault tolerance

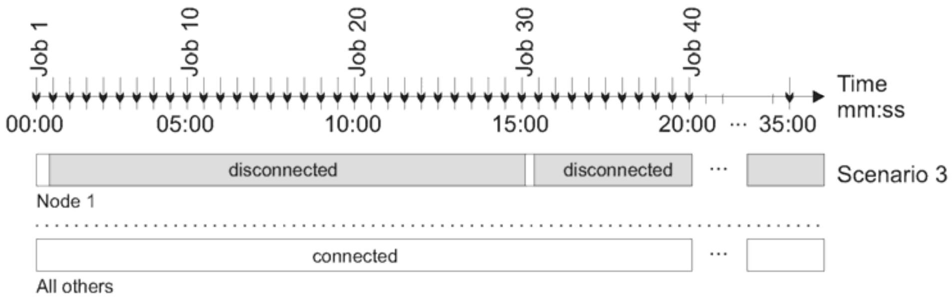

Results - Experiment 1 (Fault Free Scenario)

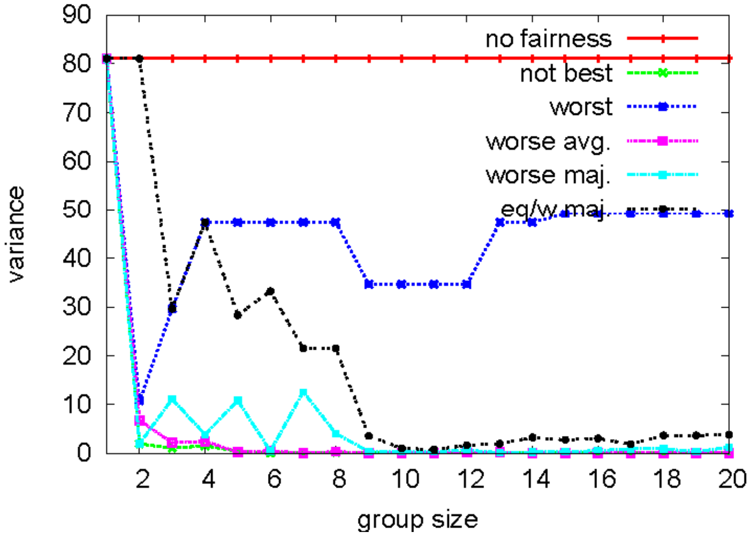

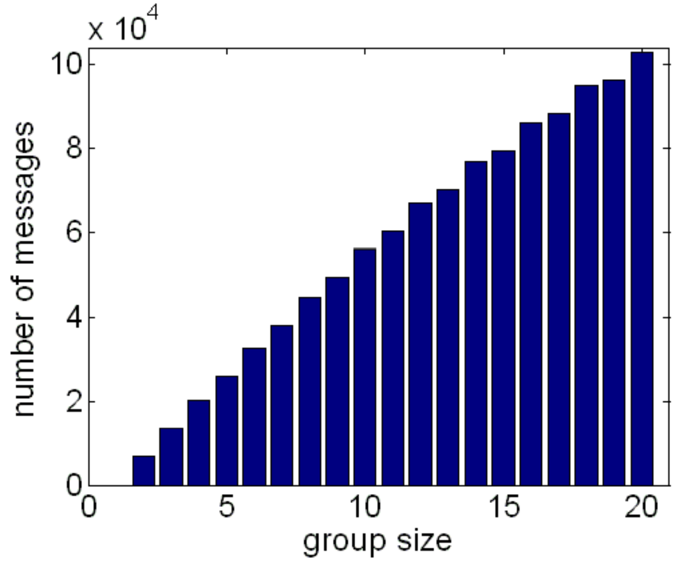

Observations

– Strategy *not best* outperforms other strategies with respect to fairness (incl. deadlock prevention)

Results - Experiment 2 (Faulty Scenario, n=10)

	no error	no FT	proactive FT	reactive FT
no fair.	85.26	17.64	203.67	61.57
not best	0	0.11	0.63	0.11
worst	36.42	28.42	91.95	38.11
worse avg.	0	0.11	0.11	0.11
worse maj.	0.32	0.32	0.99	0.32
eq/w maj.	0.73	3.16	2.85	3.16

Fairness: variances of number of processed jobs

Observations

- Phenomenon: No FT reduces variance due to overloaded nodes going offline
- Proactive FT increases variance due to redundant job exec.
- Worse avg. – not best – worse maj. – eq./w. maj. – worst – no fairn.

Discussion

Task Scheduling Approach

- —Uses pro- and reactive fault tolerance mechanisms
- —Fully decentralized
- —All strategies improved fairness
  - > The strategy which takes jobs only, if the device has the best performance value in *T* leads to highest fairness
  - > Strategies evaluating the average or majority yield good results for higher group sizes

Future Work

- —Examine fairness between heterogeneous devices
- —Fairness w.r.t. other resources (e.g., energy)

On Game Theoretical Investigations

Terms

- Defecting/selfish: do not contribute resources
- Ever defecting: only defecting
- Cooperate: contribute resources

Strategies

- Tit For Tat (TFT)
- Generous TFT (g-TFT)
- Go By Majority (GBM)

Investigation of how fast selfishness is propagated among nodes

[Hum08b] K.A. Hummel, H. Meyer. On Properties of Game Theoretical Approaches to Balance Load Distribution in Mobile Grids . Int. Workshop on Self-organizing Systems, 2008.

Experiment Setup

Discrete event simulation

Setup

– 15 nodes playing TFT/g-TFT/GBM
– 5 nodes ever defecting
– group sizes 5, 10, 15 and 20

Scenario 1: Investigate strategies independently

– Propagation of selfishness among homogeneous strategies

Scenario 2: Investigate combinations

– Ex.: Propagation of selfishness with TFT + g-TFT

Results Scenario 1

- 5 ever-defecting nodes
- group size (a) n=5, (b) n=10, (c) n=15 and (d) n=20

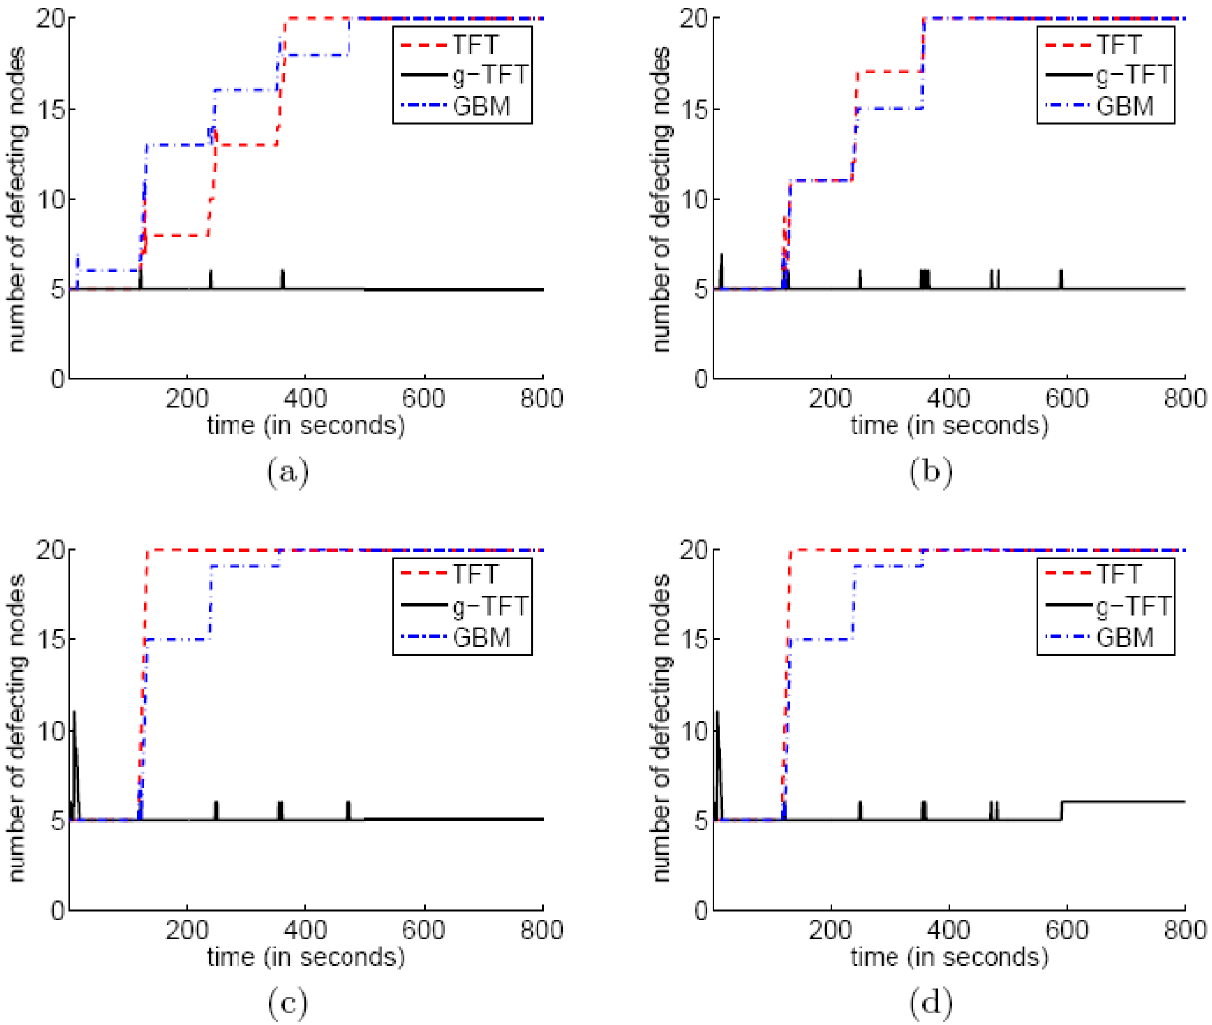

Results Scenario 2 (ongoing)

– 14 TFT and 1 g-TFT
– 5 ever defecting nodes

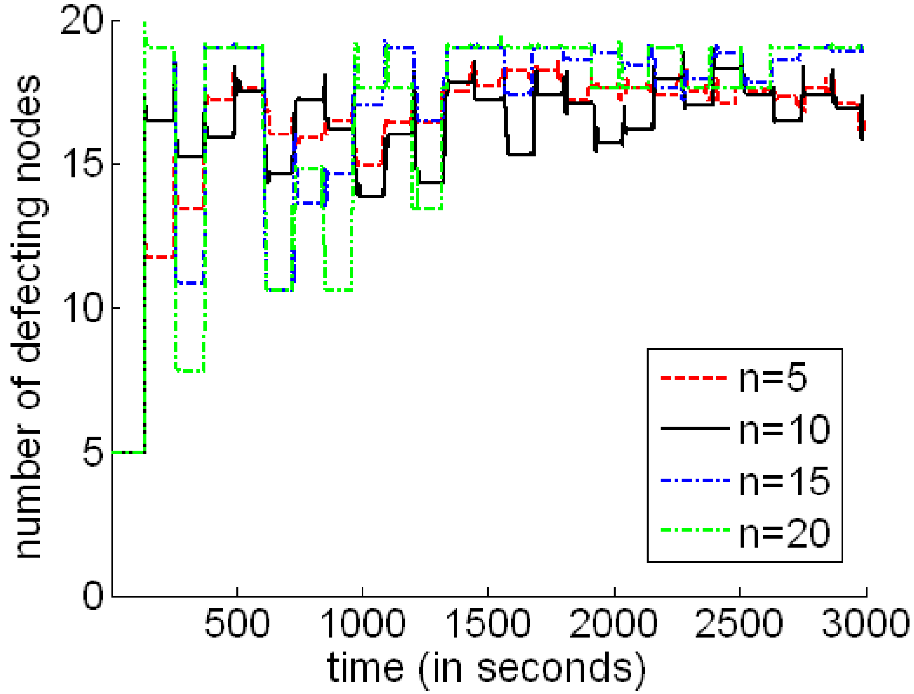

– Where should g-TFT nodes be placed to damp spreading of selfishness? How many do we need?

Discussion

Game Strategies

– We investigated the propagation of selfishness by means of three game strategies
– Selected non-stochastic strategies tend to propagation of network-wide selfishness
– First observations: Stochastic element and small group sizes damp down spreading of selfishness

Future Work

– combine game strategies with payoff function
– incorporate mobility

Conclusion

Improve mobile distributed computing performance by

- Mobility-awareness
- Self-organization of tasks / placement of data based on local interaction

Investigations

- Improve fairness by means of rules
- Spreading of selfishness in case of minimal information exchange and no explicitly modeled benefit for cooperation

Involvements in Related Networking Conf.

IWSOS 2009 Intern. WS on Self-Organizing Systems (December 2009, Zurich, CH)
http://www.iwsos2009.ethz.ch/
Full Paper submission deadline: July 10, 2009

eEnergy 2010 – University of Passau, D
Contact:karin.hummel@univie.ac.at