

Konferenzbericht

GI/ITG/MMB/KuVS Fachgespräch zu “Performance Modeling of Self-Organizing Systems (PMSOS’06)”

Technical Co-Chairs

Hermann de Meer, Patrick Wüchner
Universität Passau
hermann.demeer@uni-passau.de
patrick.wuechner@uni-passau.de

Jens B. Schmitt
Technische Universität Kaiserslautern
jschmitt@informatik.uni-kl.de

Matthias Hollick
Technische Universität Darmstadt
matthias.hollick@kom.tu-darmstadt.de

September 21, 2006, Universität Passau

Abstract

Self-organization is expected to play a key architectural role for the future Internet as well as for large-scale pervasive computing systems such as wireless sensor networks. In conjunction with the “International Workshop on Self-Organizing Systems (IWSOS 2006)”, we organized the first “GI/ITG/MMB/KuVS Fachgespräch on Performance Modeling of Self-Organizing Systems (PMSOS’06)”. The event was hosted by the chair of Computer Networks & Communications of Universität Passau. To make the Fachgespräch as attractive as possible, we introduced and trialed a novel dialogue-intensifying format for discussions. In the aftermaths of the workshop, we have been able to collect evidence that the participants considered the workshop to be very vital and stimulating.

Motivation

The goals of this first Fachgespräch in the area of performance modeling of self-organizing systems were manifold. First, the discussions should provide the authors with high-quality feedback. Second, the discussions should be as intense and mind-opening as possible. Third, a novel conference format should be trialed to revitalize the discussion-instrument “Fachgespräch” in the MMB community; thus, increasing the attractiveness of this discussion format for MMB and KuVS organized technical meetings. The organizers were particularly motivated to attempt to establish the instrument of a Fachgespräch at the last meeting of the MMB Leitungsgremium and Fachexpertenkreis in Erlangen in March 2006. Special acknowledgements go to the MMB Chair Christoph Lindemann and all other supporters within the MMB working group. The MMB very much encouraged us to re-establish the formerly well-received format of the Fachgespräch.

The proposal of a novel and highly relevant topic fueled the re-birth of this instrument in combination with innovative

mechanisms to “self-organize” the discussion. Equally, the KuVS was supportive in accepting the Fachgespräch as a joint MMB/KuVS event, thus, consolidating the two well-established German bodies that are interested in the respective research areas.

“Self-Organizing” Fachgespräch

As one highlight of the workshop, a novel “self-organizing” way of running the workshop was trialed. It consisted of a court-like offense/defense process for each individual contribution. Before going into the details of this “self-organizing” format to run the workshop, we can summarize that this format worked out excellent. In an act of self-configuration all participants offered to serve in offending and defending the accepted contributions. The excellent organization of the event (e.g., with food supplies) supported the self-healing process of the participants after intense discussions. “Perturbations” introduced by participants also fueled the discussion in a positive sense. The discussion process was to some degree self-optimizing, because the offender/defender roles were dynamically supported by other participants, if necessary. Last but not least, the discussion proved to be self-protecting, because the offender did not directly clash with the author, but entered a discourse with a much more neutral defender. The authors thus could obtain an impression on how others perceived the good as well as the critical sides of their contribution, before entering the discussions themselves.

In brief, the Co-Chairs introduced the following workshop format: each presenter got an initial timeslot of 25 minutes to present her/his research results. In combination with the paper, this talk served as input for the following discussion, which was structured in a court-like fashion. In particular, the papers had been distributed in electronic form to all participants well before the meeting. Also, each workshop participant was invited to select one paper to defend or to offend beforehand. The discussion following the author’s presentation was then started by the so-called “offender”, who should (provocatively) question the work and point to potential weaknesses. Next, the so-called “defender” was to rebut the offense as far as possible. Both, offender and defender were asked to back their discourse with evidence from the paper or the presentation. Please note that during this tribunal, the author’s role was purely passive. After this initial exchange of arguments, the remaining participants were allowed to join the discussion, as was the author. A last phase served to collect a collaborative “judgment” addressing perceived open issues or attractive future work areas.

Facts and Minutes

Performance modeling of self-organizing systems is an emerging discipline in computer science and information technology. Self-organization plays a key architectural role for the future Internet as well as for large-scale pervasive computing systems such as wireless sensor networks. Self-organization promises to enhance flexibility and evolvability of organically growing, large-scale distributed systems. Driven by randomness and feedback, the relation between cause and effect of these systems may appear chaotic: minor causes may have severe impact, whereas seemingly major causes may only have a small effect. Traditional mathematical models tend to be linear and, thus, are not suitable to model such chaotic systems. Models to be developed should be simple to remain scalable to

the huge number of entities in the systems under investigation and to be generally applicable.

The aim of the Fachgespräch was to investigate the methodology for modeling and evaluating the performance and reliability of complex self-organizing systems at the edge of chaos. To this end, the meeting brought together leading national and international experts, thus, creating a forum for knowledge exchange.

The workshop was held in the beautiful city of Passau on September 21, 2006. In the following, we give a brief overview of the event:

Jens Schmitt opened the Fachgespräch and introduced all participants to the special “self-organizing” format of the workshop. Six contributions were admitted, five from German first authors and one from an international author. The locality was introduced by Hermann de Meer, who also transferred the spirit of the discussion of the preceding IWSOS workshop. During the entire Fachgespräch the local organization related issues as well as some few “non-self-organizing” aspects such as timekeeping have been very well handled by Patrick Wüchner.

The technical sessions comprised six presentations, which covered a wide range of topics related to the performance modeling of self-organizing systems (see extended abstracts below for more details).

In particular, the feasibility of using performance modeling techniques from the domain of discrete-event systems for self-organizing systems was discussed. Important points raised in the discussion were the scalability of such mechanisms for the large state-space as well as the hope that the self-organizing communities can borrow and extend tools and methodologies of well-established related disciplines.

Randomized local optimization algorithms to allow the creation of self-organizing networks have been in the focus of the second talk. Here, the simplification and subsequent transformation to a problem class of much lower complexity led to promising results.

The talk that competed with the lunch break was on an extensive simulative study of the route forwarding behavior of flooding-based routing protocols in ad hoc networks. The insights obtained as well as initial steps to capture parts of the model in closed analytical form were identified as very promising. Possible next steps include especially the refinement of the analytical model.

The fidelity of simulation models and the generation of real implementation and simulation code using the same high-level description (SDL) were the core contribution of the next talk. Again, the importance of the method of simulation for obtaining insights in self-organizing systems was acknowledged. At the same time the discussion showed the limitations of simulation studies and made obvious that performance modeling is in fact an art as postulated by Raj Jain.

The following talk first departed from simulation and showed various extensions to network calculus and their respective applications in analytically modeling wireless sensor networks. To close the big circle opened by the author, the talk concluded with the intention to perform simulation studies with higher level of detail to back these nice analytical results. Here, especially the dynamicity of the methods and tools employed was subject of discussion.

After the second coffee break for the day, the last talk took an entirely different approach towards self-organizing systems. Inspired by self-organization in nature, the author contrasted engineered systems with their counterpart in nature to find similarities and dissimilarities. The instantiation of such systems as well as the limitations of this approach were discussed in depth.

Patrick Wüchner guided the wrap-up session, which got a very strong vote from all participants to plan a second Fachgespräch on the same topic in 2007. Also, a special issue in the PIK with open call was agreed upon for early 2008. Matthias Hollick concluded the Fachgespräch and, in the name of all participants, thanked the organizers for the excellent organization of the event.

Feedback from Participants

The feedback concerning the special format of the workshop among participants was indeed very positive. While the role of offender and defender were taken from very moderate to very critical, they succeeded in stimulating intense discussion.

Comments from the audience about the workshop in general and the discussion format in particular have been collected by e-mail:

Feedback about the contents of the workshop:

“It is important to capture the typical and often startling effects in self-organizing systems in the models to be developed without losing their scalability and without prejudicing the effects which can usually not be known beforehand.”

“The individual contributions covered a very wide and heterogeneous spectrum. Personally, I would love to see a more focused topic for the next workshop to yield a more homogeneous selection of papers.”

“PMSOS was on whole a learning experience for a student like me in the company of some esteemed researchers. It gave an insight on self-organizing system with respect to its properties and applications in various engineering disciplines.”

“The workshop showed that self-organization is an architectural principle as well as a phenomenon arising in large-scale complex systems. Performance modeling in that context is very challenging and analytical as well simulative approaches are required.”

“As a PhD student at the end of my first year, the workshop was of great help for me. I could figure out clearly what type of networking environments I should be caring about. I could also conclude the importance of the performance issues, and that I should consider them from the beginning while I am trying to define my security solutions for self-optimizing mobile networks.”

“The workshop was useful for me as I got the idea about network calculus from the talk given by Jens B. Schmitt and also the talk given by Matthias Hollick was quite beneficial for my future work.”

“Excellent!”

“The workshop showed that an in-depth understanding of the characteristics and properties of self-organizing systems is crucial (even if definitions can only capture part of the picture) but not agreed upon.”

“The workshop discussed various facets of the timely topic of performance modeling of self-organizing systems. With the advent of such systems (esp. in the area of communication networks), these issues are of utmost importance and will be vital to allow a sustainable deployment at large.”

“While the understanding of SOS by the engineering, especially computer networks, community is partial and too specific with regard to the “scientific” SOS community (physics, biological and social sciences), they offer a clear, targeted view of functionality, measurement and performance of (engineered) SOS that the scientific community can only vaguely qualify.

As such their contribution to understanding and possibly control of particular SOS models can be unique. I believe that more and broader efforts in reusing/customizing/restructuring of classical SOS in engineering domains will be fruitful for both communities.”

Feedback about the format of the workshop:

“The format led to very lively and fruitful discussions and while the offenses could be very provocative the atmosphere remained constructive and friendly. This was certainly also due to the moderate size of the discussion group. I think the format should be handled with care in larger and less personal discussion groups.”

“I did not like the format. I have to admit that it facilitated to start discussion, but this can also be achieved by other mechanisms. I could imagine a system, where two participants have to collect discussion items beforehand. Especially, the authors should be allowed to defend their work themselves.”

“Personally, I judge the outcome of the individual discussions at PMSOS to outperform even well-regarded conferences in this area. (1) The authors were forced to listen to the feedback (no “defense-wars”), (2) the small group brought the criticism to the point without endless discussion, and (3) the discussion was much more honest and open. Not to forget (4) it was a lot of fun.”

“As for the offense/defense format, ... it made the workshop more dynamic and encouraged active participation.”

“I very much like the offense/defense structure, and I would like to see this coupled with a list of discussion topics ... proposed by the participants (and not necessarily related to the presentations).”

“I found the discussion at the PMSOS very lively and fruitful. Therefore I want to encourage the organizers to keep the format.”

“The workshop was interesting with the fact that it gave the author some good keynotes from the defender’s view as well

as from the offender’s. It gave the presenter an opportunity to know how well his paper was interpreted by the participants along with their views and suggestions for further improvement of the paper or its approach.”

“... it gave all the participants the possibility to be a part of the discussion process ...”

“... was indeed a novel idea and worked well to result in some fruitful discussions ...”

“... the idea of offense/defense format ... creates a healthy environment because the author got authentic feedback about the weakness of his work ... and it is also encouraging for the author ...”

“The organization and atmosphere at the Fachgespräch was excellent and asks for a continuation of this event.”

Additional Information

To illustrate the wide range of covered topics, in the following, we present extended abstracts of the individual contributions. The full versions of the workshop submissions are available online as a joint technical report (<http://www.fmi.uni-passau.de/forschung/mip-berichte/MIP-0609.html>). A printed paperback copy of the technical report is available upon request. Also electronic versions of the presented slides are available to the public. Please see <http://www.iwsos.net.fmi.uni-passau.de/pmsos.html> for the corresponding information.

For the organizers,

Matthias Hollick (Matthias.Hollick@kom.tu-darmstadt.de).

**Discrete-Event System Performance
Modeling of Self-Organizing Systems**
Patrick Wüchner and Hermann de Meer
Chair of Computer Networks and Computer
Communications, Faculty of Computer Science and
Mathematics, University of Passau,
Innstr. 43, 94034 Passau, Germany
{Patrick.Wuechner, Hermann.DeMeer}@Uni-
Passau.de

Extended Abstract

Man-made systems, e.g., communication systems like the Internet, are getting larger and more and more complex. Such large and complex systems are inclined to self-organize – if intended or not. Thus, to set the basis for understanding and developing future complex systems, it is extremely important to understand the behavior of self-organizing systems.

Often, the modeling of systems already makes a fair contribution towards the understanding of these systems. Therefore, we are planning to discover the basic principles of self-organizing systems in a universal way by building abstract and general models, particularly performance models, of complex, self-organizing systems.

One approach towards the modeling of systems in general is by describing them as discrete-event systems, i.e., as discrete-state, event-driven dynamic systems. While a huge set of modeling techniques were developed for discrete-event systems, we here concentrate on stochastic timed mathematical models that are useful for performance modeling. Within this class of modeling techniques several modeling formalisms have been developed. Some of the most prominent ones are Markov chains, queueing networks, and stochastic Petri nets.

Stochastic Petri nets seem to be the most promising modeling formalism in this group. They are especially capable of describing synchronization and concurrency. They might be evaluated either by transformation into Markov chains or more flexibly by discrete-event simulation. In the meantime, several additions to stochastic Petri nets have been developed to improve their modeling power and the modeling convenience. Also with respect to evaluation techniques several improvements were proposed. Among others, these improvements include simulation speed-up techniques (e.g., importance sampling, importance splitting/RESTART), largeness avoidance, state space truncation, fluid models, stiffness avoidance, and phase approximations.

Especially this plentifulness of methods makes it worthwhile to investigate if they can be used efficiently in the modeling process and evaluation of self-organizing systems.

We are still at the beginning of this investigation and not yet able to give a final answer to this question. What we can see already is that while some properties of self-organizing systems (e.g., autonomy, dynamics, energy usage) can be quite easily handled by modeling techniques developed for "classical" discrete-event system, the mapping of other properties is far harder to achieve, e.g., non-linearity, bifurcation, etc. Furthermore, the considered state-space based approaches tend to suffer from the largeness of the systems to be modeled. Here, it might be worthwhile to look for symmetries between the different components of the self-organizing system under study.

Besides identifying more properties of self-organizing systems that will have to be considered by a viable model, we plan to build a more detailed survey and classification of models already used by researchers dealing with self-organizing systems. From these models we might be able to learn how to extend the capabilities of DESs to increase their flexibility. Particularly, we will have a look at modeling and evaluation techniques for non-linear systems.

Self-Organizing, Evolving Networks for Performance Optimization

Peter Merz

Department of Computer Science,
University of Kaiserslautern, Germany
peter.merz@ieee.org

Extended Abstract

Self-organization plays an important role in the design of protocols/algorithms for computer networks, such as peer-to-peer overlays, ad-hoc mobile networks, and sensor networks. However, it is not trivial to design protocols or algorithms in a way that a desired behavior emerges.

By taking the point of view of optimization, we demonstrate how algorithms can be designed to optimize predefined performance metrics and hence exhibit desired emergent behavior. We show in the case of peer-to-peer overlay topologies that considering an approximation to the performance objective can yield fast convergence to a stable state of the network.

To demonstrate the feasibility of our approach, we consider the self-organization of peer-to-peer (p2p) topologies based on spanning trees. The choice for spanning trees is motivated by the fact that spanning trees are graphs with minimum number of edges to ensure connectivity and they are desired in applications such as multimedia streaming where any node can be content provider and wants to propagate its data by multicast to interested peers in the overlay. Ideally, the tree should have the property that routing data from any node to any other node is minimal with respect to message delay. Hence, in our example, our performance metric is routing efficiency. Finding a single shared tree with minimum routing cost is an NP-hard combinatorial optimization problem. Therefore, we propose to apply local search techniques from combinatorial optimization to derive rules for the entities (peers) in the self-organizing system. By letting the peers optimize their situation by greedily selecting parents to improve their routing cost, the system as a whole is enabled to converge to a state with minimum routing cost. However, due to the fact that the minimum routing cost spanning tree problem is NP-hard, a distributed local search for this problem can only be realized with a high message overhead and strong cooperation among the peers. Hence, we propose to optimize the degree-constrained shortest path tree problem instead, which provides an approximation to the NP-hard problem.

In the algorithm, each node/peer performs an optimization step in each round in order to improve the tree. An optimization step consists of probing for two different types of tree transformations (called moves). In the first move, a node (i) tries to connect to a new parent (q). If the node q accepts further children, and if the distance from node i to the root is reduced, the move is accepted. In the second move, the situation is similar. If node q already has the maximum number of children, node q may decide to accept the new child i and redirect its child j with maximum distance to itself. Hence, the distance to the root of one node (i) is reduced, possibly at the cost of the other (node j). In the first move, one edges in the tree is exchanged with another edge, while in the second move, two edges are exchanged. This way, the tree is optimized over time until the system converges to a stable state (assuming there is no churn), where there is no improving move and hence a local optimum is reached (or Nash equilibrium in the game theoretic sense).

We demonstrate by simulations, that the proposed algorithm converges with a high progress rate and that it has the ability to find good approximations to the minimum routing cost spanning tree objective. Furthermore, we discuss some issues relevant for the analysis of such self-organizing, evolving networks. We consider the presented approach applicable to a wide range of scenarios with different objectives/performance measures in p2p, wireless ad-hoc, and sensor networks.

Modell zur Analyse des Routenfortschritts in Ad Hoc Netzen

Tronje Krop, Tobias Lange, Matthias Hollick und Ralf Steinmetz

Multimedia Kommunikation (KOM),
Technische Universität Darmstadt,
Merkstrasse 25, 64289 Darmstadt
{tronje.krop}@kom.tu-darmstadt.de

Erweiterte Kurzfassung

Mobile Ad hoc Netze unterstützen die spontane Kommunikation zwischen mobilen Teilnehmern. Die selbstorganisierende Bereitstellung der Verbindungen ermöglicht hierbei den Verzicht auf eine dedizierte Kommunikationsinfrastruktur. Die Leistungsfähigkeit der Verbindungen hängt stark von den verwendeten Basis Technologien ab. Insbesondere Protokolle zur Routenfindung und zur Koordination des Zugriffs auf das (drahtlose) Medium sind entscheidend für die Stabilität und Zuverlässigkeit der Verbindungen. Hierbei kommen häufig Mechanismen der Selbstorganisation zum Einsatz, welche Aussagen über globale Eigenschaften von Protokollen in Ad hoc Netzen schwierig machen.

In unserem Beitrag entwickeln wir ein Simulationsmodell zur Analyse der Leistungsfähigkeit des Routenfindungsprozesses in Ad hoc Protokollen. Schwerpunkt des Modells ist die Beschreibung der Effekte, die durch die Selbstorganisation der einzelnen Netzknoten bei der Weiterleitung von Routenanfragen hervorgerufen werden (basierend auf dem Fluten von Routenanfragen in einem drahtlosen Netz). Für die Bewertung führen wir die Konzepte Routendistanz, Routenfortschritt und Sprungdistanz als wesentliche Einflussfaktoren für Leistung und Lebenszeit von Routen ein. Erst das vereinfachte Modell für Broadcast Medienzugriffe und reaktive Routenfindung im statischen Kontext ermöglicht die notwendige Zahl an Routen für statistische signifikante Aussagen zu gewinnen. Im Gegensatz zu vielen analytischen Arbeiten modellieren wir dabei auch die Abhängigkeiten zwischen Links innerhalb einer Route.

Unsere Analyse zeigt, dass Routenfortschritt und Sprungdistanz eine deutliche Abhängigkeiten bezüglich der Routendistanz aufweisen. Gleichzeitig können weitere, emergente Effekte beobachtet werden. So zeigt der durchschnittliche Routenfortschritt abhängig von der Knotendichte und Distanz Schwingungen auf und, entgegen der allgemeinen Intuition, gibt es eine Symmetrie zwischen erstem und letztem Sprung in einer Route. Auch die systematische Erfassung des Verhältnisses von Routenfortschritt und Sprungdistanz zeigt ein charakteristisches Verhalten auf, welches mathematisch beschreibbar ist und als Ausgangspunkt weiterer Analysen dienen kann.

Performance Simulation of Distributed Embedded Self-Organizing Systems Modeled with SDL

P. Becker, R. Gotzhein, T. Kuhn
University of Kaiserslautern

{p_becker,goetzhein,kuhn}@informatik.uni-kl.de

Extended Abstract

The accuracy achievable in performance simulations depends on several factors. First, all resources influencing system performance must be simulated together. For instance, in a distributed embedded self-organizing system, both the hardware platform and the network are resource bottlenecks. Second, the behavior of the system under simulation is a decisive factor. Here, it is essential that the same binary code is used for simulation and production.

In order to improve the accuracy of performance assessments in both respects, we have developed *PartsSim*, a modular performance simulator for SDL design models:

- Core components of *PartsSim* are simulation control and event scheduler, which have been extracted from the network simulator *ns-2* and modified.
- Specialized simulator components cover communication network characteristics, in particular propagation model, medium arbitration, and network topology, and hardware platform. Currently, *PartsSim* supports models of wireless ad-hoc networks (WLAN, ZigBee), extracted and adapted from *ns-2*, and of MicaZ motes, taken from *Avrora*.
- The code of the system under simulation is generated from SDL design models, with the same compiler that is used to generate the production code.

Due to its modular structure, *PartsSim* can be modified and extended by exchanging simulator components and integrating further specialized simulators, respectively. For instance, by replacing the *Avrora* simulator, other types of microcontrollers may be supported.

First simulation experiments have provided clear evidence of the benefits of the increased accuracy. In these experiments, *PartsSim* has revealed severe timing problems that did not show up in simulations without platform delays and platform jitter.

Our future work will be focused on two aspects. First, we will use *PartsSim* to study the performance of networked systems in the ubiquitous computing domain that are based on MicaZ hardware. Second, we will incorporate the simulation of energy consumption.

Worst Case Modeling of Wireless Sensor Networks

Jens B. Schmitt

DISCO Lab, Computer Science Department,
TU Kaiserslautern

jschmitt@informatik.uni-kl.de

Extended Abstract

At the current state of affairs it is hard to obtain a predictable performance from self-organizing wireless sensor networks, not to mention performance guarantees. In particular, a widely accepted and established methodology for modeling the performance of wireless sensor networks is missing. In the last two years we have tried to make a step into the direction of an analytical framework for the performance modeling of wireless sensor networks based on the theory of network calculus, which we customized towards a so-called sensor network calculus. We believe the sensor network calculus to be especially useful for applications which have timing requirements. Examples for this class of applications are factory control, nuclear power plant

control, medical applications, and any alerting systems. In general, whenever the sensed input may necessitate immediate actions the sensor network calculus is the way to go, in our opinion. In the contribution for the Fachgespräch we summarize these activities and discuss the open issues for such an analytical framework to be widely accepted.

In particular, we have shown how to model important concepts in wireless sensor networks as for example the notion of a duty cycle using network calculus. Furthermore, we extended network calculus to cope with typical random deployments and topologies as often encountered in self-organizing wireless sensor networks. In an interesting application of the sensor network calculus to the problem of sink placement in a sensor field, we have shown how the analytical framework can guide us to make good decisions on the number and location of sinks. Interestingly, for this investigation we encountered a phase-transition effect: when increasing the number of sinks over a certain critical threshold then the delay distribution in the sensor field suddenly turned from unpleasant to favorable almost instantly. The nice thing is that the analytical model was able to capture that frequent effect of self-organizing systems despite being on an abstract level in order to remain scalable. As future work we perceive the stochastic extension of the sensor network calculus as well as further model sophistication to capture even more sensor network characteristics. Simulative validation of the analytical models is on the way.

Measures of Developing Stability and Perturbations in Ecologies of Semantic Web Services

Elpida S. Tzafestas

Institute of Communication and Computer Systems
NTUA Campus, Zographou 15773, GREECE
brensham@softlab.ece.ntua.gr

Extended Abstract

We are discussing the issue of performance of engineered self-organizing systems that belong to the class of virtual organizations where an entrepreneurial or other activity acts as an interface between a number of existing remote services or applications and a number of customers that interact with them in complex ways. Such virtual organizations are characterized by operational and performance criteria that are independent and even sometimes incompatible with those of individual components (external services/applications) involved.

Within the relative INFRAWEBs (www.infraweb.org) effort, we are developing a security reasoner, which is responsible for maintaining an image of the network of external semantic web services used by an application and which follows the artificial immune systems approach. This parallel overlay network that self-organizes and reaches an equilibrium can be said to have discovered the normal behavior of the target service network. The first performance criterion in this setup is that of *stability*. More specifically, we need to find one or more measures that have to be stable in time to ensure operability of the application. One design guideline for stability is to search *behavioral attractors* of the system, i.e. attractors for derived values and measures that express the functionality and

operational goals of the system, rather than purely structural network metrics. A second criterion for performance evaluation is the type of response to *perturbations*. If the overall system can regain stability after a while, this is an indication of an operational, “healthy” application. In such cases, the matter of study is the relation between the former and the new equilibria. We also need to devise and study counter-measures to specific perturbations, i.e. *counter-perturbations* that can make up for potentially harmful ones. Moreover, the virtual organization should be able to maintain by self-organization a correct image of the *developing equilibrium*, in the same way that a real biological system learns from experience.

Like in individual-based biological modeling, all the above criteria and proposed mechanisms are of qualitative and not quantitative value: a system may find the equilibrium predicted in simulation, but the exact parameters and “meaning” of the equilibrium will have to be assessed on-line.