Linkage liaisons: a scenario for the computational study of conflict and structure in multiagent systems

Malsch, Thomas; Paetow, Kai; Rovatsos, Michael

Veröffentlichungsversion / Published Version
Arbeitspapier / working paper

Linkage Liaisons –

A Scenario for the Computational Study of Conflict and Structure in Multiagent Systems

Thomas Malsch Kai Paetow Michael Rovatsos

Summary

Website presentation is becoming a crucial issue as more and more services of all different kinds are offered by more and more content providers to a rapidly growing audience on the Internet. In the face of accelerated competition and intransparency it is highly important for a service provider to present himself with an appealing and professionally designed website. To attract the attention of as many Internet surfers as possible and to raise their curiosity and interest, service providers will have to do everything to enhance their website's traffic and to transform casual visitors into habitual customers and partners. In order to attract more traffic, linkages directing the user from website to website is a crucial issue and building-up strategic linkage liaisons with other website owners on a B2B basis will be one of the future key elements of successful website presentation. Linkage liaisons between websites will give rise to emergent network structures (professional and non-professional) as well as to all sorts of minor and major linkage conflicts that arise among website owners (or their agents) in the course of incrementally engaging in building up new linkages and deleting old ones. Hence, linkage liaisons look like a particularly suitable scenario for designing, implementing, and testing models for real world applications and for scientific research purposes of comparing interactionism with social systems approaches in sociology. We assume that liaison models differ from more conventional DAI applications not only because they operate on a permanent basis (cf. computational ecology) and that they produce, reproduce and modify an "emergent" social structure while resolving (or not resolving!) conflictive episodes. The presumed advantages of the proposed model, both from a software engineering and from a sociological perspective, should be seen in the way that a linkage networks would have to meet the demands of shifting from a highly efficient and cost saving routine mode into a resource consumptive conflict mode and back again into a new routine mode after the conflict has been settled. In other words: the system should be able to learn from conflicts and, in doing so, an agent society would operate like a self-organising mechanism or a selfsustaining general purpose infrastructure.

Contents

-

1 Introduction

Website presentation is becoming a crucial issue as more and more services of all different kinds are offered by more and more content providers to a rapidly growing audience on the Internet. In the face of accelerated competition and intransparency it is highly important for a service provider to present himself with an appealing and professionally designed website. To attract the attention of as many Internet surfers as possible and to raise their curiosity and interest, service providers will have to do everything to enhance their website's traffic and to transform casual visitors into habitual customers and partners. But what makes a website presentation attractive and what can one do to improve a presentation? A new market segment is emerging in the service sectors for professional website consultants: website design and, more recently, website rating. In response to the growing demands, many thousands of websites designers and marketing consultants are on their way to help all those who have a trade to ply on the Internet and WWW. However, a professional website presentation alone will not do. In order to attract more Internet traffic – and traffic means business! – it would be useful to know more about the surfing behaviour of Internet users in general and specifically about the routes they select. Here, market research studies for website ratings and rankings might be revealing.

However, empirical information will not be all that revealing. It would largely confirm what we already know, namely that users, apart from using search engines as the most common point of entry to the WWW, also will look for interesting links they can find on those websites they have been directed to by their search machine: [Rogers, 2000] By means of website linkages which connect one site with another one, users will autonomously be capable of surfing from one website to another to find and select similar or related content elsewhere. Hence, linkages directing the user from website to website is a crucial issue and building-up strategic linkage liaisons with other website owners on a B2B basis will be one of the future key elements of successful website presentation. Linkage liaisons between websites will give rise to emergent network structures (professional and non-professional) as well as to all sorts of minor and major linkage conflicts that arise among website owners (or their agents) in the course of incrementally engaging in building up new linkages and deleting old ones.

Hence, linkage liaisons look like a particularly suitable scenario for designing, implementing, and testing CONSTRUCT (Conflict and Structure) models for real

world applications and for scientific research purposes of comparing interactionism with social systems approaches in sociology. The linkage liaison scenario is specifically attractive for CONSTRUCTNETWORK models out of the following six reasons:

1. It is a quasi natural scenario of website rating and a real life process occurring in front of our eyes in the WWW and Internet, and hence far more realistic than any sort of website ranking done by a jury of ”artificial” experts and published in the media. Instead, the ”real” experts are the website owners themselves who, by the very fact of having decided to link their website to another one, de facto have given a positive rating.

2. In the linkage liaison scenario an agent is defined as an active website representing and acting on behalf of its owner. To begin with, specific agent roles for web-search or Internet ”trawling” or for rating and ranking are not needed. Thus, unnecessary complications can be avoided in a prototypical implementation.

3. The basic prerogative of agents consists of attracting as much attention and traffic as possible by means of exchanging linkages with influential partners and, hence, to improve their public standing. This seems to be perfectly in line with the practical behaviour of real empirical website owners.

4. In a short-term perspective we seem to be faced with a clear-cut engineering task: As linkage negotiations (or linkage conflicts) are specified as local events or as episodic ”encounters” (or pseudo ”face-to-face” interactions) which terminate in a yes/no decision, they should be, at least in principle, computable.

5. In a global or long-term perspective, however, the scenario confronts the designers with a quasi endless process of reproduction and permanent structural modulation, i.e. with a ”living”, self-evolving network of linkages – and one of the design problems will consist of ”synchronising” the different time-scales of short term conflicts and long-term structuration.

6. Moreover, as a ”real life” phenomenon under empirical sociological observation, linkage liaisons on the Internet offer the opportunity to run computational simulation experiments of interesting sociological issues like power

distribution, social unequality, democratic control, coalition formation and the evolution of ”structures” like frames or generalised media of interaction.

Excursion 1: A conflict scenario in a subcultural network community

Before we begin to outline how to derive ”structure from conflict” bottom-up, it might be useful to give an illustration of a real life scenario of network communities. The illustration is, as all ideas in this paper, still very preliminary. So considerable social research work (qualitative case studies based on interviews, documents and other materials, quantitative statistical analyses) will have to be invested to gain a profounder empirical understanding of real life linkage liaisons on the Internet as a basis to improve our models and architectures. Generally speaking, a real life linkage scenario could be conceived as a loosely coupled network community with ill-defined system boundaries and fluctuating participants embedded in the communication infrastructures of WWW and Internet. To illustrate the idea of what ”structural change” could mean, let us take as an example a non-professional community of young people organising open air (Goa) parties[1] by making use of informal acquaintance connections via the Internet and WWW. Perhaps in the beginning just a handful of non-professional individuals who share a specific commitment to subcultural values (e.g. postmaterialistic values of ”flower power”) have linked their websites to organise and celebrate parties and performances and to have fun together. As the parties attract more and more people and money, a more professional management is required to sustain what turns out to have become a flourishing business with growing scales of investment and profit. Hence, a gradual shift takes place that may give rise to conflicts between ”traditionalists” (who stick to a subcultural life style of spontaneous anarchy) and ”professionals” (who wish introduce new professional standards of website design, business promotion, obligations and membership).

At first, the conflict will begin to develop on a low level of bilateral interactions (A versus B etc.). And it is not at all clear that there is a general underlying conflict between two different value orientations (commercial versus subcultural) lurging

[1]Other examples may be more compelling or suitable and should be selected out of a multitude of newsgroups for leisure and fun or for professional purposes or for hobbies, networks for business among enterprises (B2B) or networks of business to customer relations (B2C).[Kumar et al., 1999]

beneath a multitude of minor disputes. For an empirical sociological investigation it would be interesting to study how a network community gradually becomes aware of the more general dimensions and explicitly symbolises low level conflict as indicators of a collective (high level) conflict. Once a general awareness of a high level value conflict has developed within the community, the conflict may be defined as follows: The ”professionals” feel disturbed and inhibited in their business prospects by the ”traditionalists” who adhere to their ”postmaterialistic” practices of unreliability, nonchalance and neglect of commitments (e.g. A has ”forgotten” his promise to announce B’s next performance on his website, while B has faithfully advertised A’s new Jogi-tea sortiment). A is frustrated because the network propaganda had badly failed to spread and advertise his ”performance info”. As it seems, informality does not work satisfactorily and common subcultural values do not sufficiently frame and integrate the community’s identity any longer. The ”traditionalists”, on the other hand, have an innate aversion against formal rules and ”professional” restrictions because they identify themselves with openness and spontaneity and are suspicious of all kinds of business entrepreneurs who claim to be part of their community. Professionals are suspected of opportunistically abusing the subcultural community as a convenient market place for commercial interests. Traditionalists feel deeply threatened by excessive commercialisation and do not wish to be part of a formal network with strict reciprocity norms and binding obligations. In the course of conflict escalation the community network may split up into two different networks, i.e. most traditionalists delete their linkage liaisons to professionalists and vice versa. Thus, the conflict seems to be resolved by segregation.

At this point there are many more questions[2] waiting for an answer, but we shall concentrate on just one of them: What happens if both communities feel that finding a compromise is better than ”solving” the conflict by segregation? They will

[2]Which of the two subcommunities is going to survive? Will the professionalists loose their subcultural roots and their traditional ”customers”, will they be out-manoeuvered by bigger commercial players? Will the traditionalists be able to cultivate a niche market or will they starve from a loss of subcultural entrepreneurship? Perhaps the conflict parties are able to anticipate these dangers of segregation and find a way to coexist which is more beneficial with both sides. Indeed, segregation is just one possible outcome of conflict. Another way of coping with such conflicts could be seen in the emergence of proto-legal institutions for mediation and ”due process”, by giving the network a more formalised structure, or by making collective values more explicit and transforming them into norms. An implementation would humbly have to concentrate on just one of these aspects.

have to draw on institutions of dispute mediation and generally expected standards of argumentation. But what if such institutions do not exist? They will have to invent it. And when we look at behavioural norms in newsgroups (”netiquette”), it is not really unlike that more such institutions will be created in the course of future Internet developments.

2 Bottom-up instead of top-down?

We want to outline some basic principles for a dynamically evolving network of linkage liaisons with a limited number of agents which are free to say ”no” to each others prompts and suggestions. In the following sections a ”bottom-up” approach will be suggested for the construction of linkage liaisons via conflict communication. In order to avoid all unnecessary complications we will rigorously try to abstract from most of those technical features that would be needed to engineer a useful real life application: no scalability, no user interface, no mobile agents for web search, no communication facilities apart from those which are absolutely necessary to establish bilateral links, and agents equipped with as little knowledge and inferential power as necessary. In other words:

1. It is (unrealistically) presupposed that difficult problems like finding potential partners on the Internet (by matchmaking, agent mobility etc.) have already been successfully solved.

2. It is (unrealistically) presupposed that there are further mechanisms to support everyday communication and exchange among the website owners (since linkage liaisons have, of course, no end in themselves and serve as an infrastructure for some practical tasks).

3. And it is (unrealistically) assumed that agents know the preferences of their website owners or will be – at least – supplied by their owners with situated criteria or detailed instructions of how and when to engage (or disengage) in linkage liaisons. All these crucial issues will be excluded from our model.

In heroically neglecting most of these issues, we will try to consequentially keep track with a bottom-up approach: Apart from introducing a small set of predefined communication primitives, ”everything else” should be automatically generated by the system itself – or to put it more modestly: with as little help and input from the ”outside” (by a designer, a user, a databank, an external knowledge base etc.) as

possible. Instead, all of the more powerful sociological concepts (social structures) like expectations, frames, conventions, roles, commitments etc. should be produced by the artificial social system itself in a sort of self-referential (or recursive) process of ”emergent” self-production and self-modification. Of course we know that this is quite unfeasible and many good reasons and a lot of practical difficulties do indeed suggest that we should adopt a combined bottom-up and top-down approach. One of the most crushing theoretical objections against ”bottom-up” is that designing an evolutionary artificial social system out of a few communications or actions is self refuting because the designer cannot really abstain from a general ”top-down” vision of what his (!) system eventually should do. We suspect that this is as right as it is wrong. From a software engineering perspective such objections are correct because a pure bottom-up approach is technically unfeasible anyway. On the other hand, it is clear that system designers secretly dream the dream of an elegant, intelligent, bottom-up system design that solves a maximum of problems with just a minimum of predefined mechanisms.

But what about sociology? How would sociologists respond to the question? A bottom-up approach to system design appears to be in line with sociological theories of social action which suggest that social structures consist of nothing else but of repeated interaction and that structures cannot be adequately understood unless they are reconstructed from the bottom-up perspective of participant interaction as social practices.[3] Hence, for heuristic purposes we try to follow a bottom-up-design as faithfully and as rigorously as possible, hoping to gain some instructive insights from happily failing to push the approach to its limits.

[3]Of course, other sociological theories hold that communications are elements of social systems and as such they cannot be isolated and do not have an independent existence on their own. Following Luhmann [Luhmann, 1984, Luhmann, 1990], for instance, should we then have to construct an artificial social system from ”top down”? Not necessarily, says Luhmann, since system design is not limited to any specific approach and reconstructing a social system could begin almost arbitrarily from almost anywhere. But once we have decided how to begin, we are no longer free to proceed arbitrarily but must stick to our design decision with high fidelity for the sake of consistency and coherence, says Luhmann.

3 Defining a basic set of (inter)action primitives

First of all a set of primitives along the lines of KQML-like [Finin et al., 1997, KQML, 1999] speech acts[4] will be defined, each of which is framed by a sender and addressee $A, B, C, D, \ldots$ The agents A, B, C etc. can insert or delete links on their own website only and they are restricted to a "myopic" frog eyes' perspective and they share no global knowledge of the overall network structure. The communication primitives are given in the following list:

1. *insert link*(action: insert a link on ones own website to point out to another website: insertLink(A, B))
2. *delete link* (action: delete a link: deleteLink(A, B))
3. *check link* (action: checking a link on a website: check(A, B))
4. *reason* (action: inference of inserting or deleting a link: reason$(A,$ insertLink$(A, B))$)
5. *request* (communication: contains (inter)actions: request$(A, B,$ check$(A, C))$)
6. *accept* (communication: contains (inter)actions: accept$(A, B,$ insertLink$(A, C))$)
7. *reject* (communication: contains (inter)actions: reject$(A, B,$ deleteLink$(B, C))$)
8. *announce* (communication: contains future (inter)actions: announce$(A, B,$ object$(A, C,$ insertLink$(C, B)))$)
9. *tell* (communication: contains past (inter)actions: tell$(A, B,$ request$(A, C,$ insertLink$(C, B))$)

"insertLink" and "deleteLink" are indirect, asynchronous messages, since they can be received or observed by way of checking linkages on other websites. Thus,

[4]Note a decisive difference between our action and communication primitives and KQML-speech acts: while the latter use to trigger some kind of automated response in the recipient, the communication primitives suggested here do not impose a specific restriction to agent autonomy. Anything goes? If this is the case, will our scenario be chaotic and non-computational?

checking a website is equivalent to reading a binary-coded message from a mailbox or a blackboard: either there is a link to a particular website or there is none. In contrast, the action types ”check” and ”reason” are purely cognitive, i.e. unobservable by others. ”Reason” is a selectional operation based on information an agent has been told by another one or which he has gained by ”checking” another website. Both ”check” and ”reason” as such are unobservable by others (”arms-length relationships”: no sharing of personal knowledge). They are no communications or messages, because they does not leave any traces behind which could be observed or read by another agent later on[5]. However, ”check” and ”reason” can be read as the informational content of a message if included into one of the five direct (synchronous or asynchronous) message types ”request”, ”reject”, ”announce”, and ”tell”.

Note that the situational meaning of these primitives may be varying sharply according to different circumstances. For instance, ”announce” can express a hostile warning (announce$(A, B,$ deleteLink$(A, B))$) as well as a friendly offer (announce$(A, B,$ insertLink$(A, B))$). Furthermore, depending on the addressee's interpretation or point of view, a warning can occasionally be interpreted as good news because it will allow an agent, in the course of events, to get rid of some unwanted linkages. Or an offer can be interpreted as bad news because it means trouble if the addressee is not interested in exchanging links with the sender! This seems to complicate the model far beyond computational feasibility. And the question is how to implement technically useful and sociologically adequate restrictions or rules that help to channel an agents freedom of selecting whatever makes sense to him. From a human societies point of view, however, there is no need to worry too much about agents' internal or mental states, communicative selections, meaningful reactions etc. since it will be enough to observe how they observe each others' communications or interactions.

What we will have to find out in the following sections is whether it is (computationally and sociologically) feasible to model a more complex scenario of conflict processing which is based on not more than nine types of interaction. With our previously given definitions of (inter)action primitives we hopefully will be able to model (and eventually implement and simulate) the conflict dynamics and struc-

[5]Perhaps it is more realistic to assume that ”check” could be conceived as an indirect message to the website owner in case that owners are able to read the identity of visitors?

turation of a distributed network of loosely coupled liaisons among "active" websites. This includes the formation and disintegration of communities, of asymmetric power relationships, and of a diversity of mechanisms for conflict regulation, third party intervention and "due process".

4 Routine mode and conflict mode

Routine mode: In the routine mode we begin with a simple network of website-links with different patterns of references and relationships (symmetric, asymmetric, loosely coupled, strictly coupled etc.). For an illustration let us consider an agent A who wishes to join the network and inscribes on his own website a link to B's website (insertLink$)(A, B)$) while deleting a link to C (deleteLink(A, C)) whom he does or does not inform about the deletion, according to pure convenience, neglect and the like.[6] Then he proudly presents himself to B as someone who has beneficially linked to B's website (tell$(A, B,$ insertLink$(A, B))$). Whether B will be flattered, bored, pested or uninterested: A's offer is not just a constative information but also a performative message, namely an indirect or implicit request to exchange links, i.e. appeals to B to link his website to A's. B understands, i.e. we assume that he is able to distinguish between information and message, although, in our example, he does not display any overt reaction[7]. Later on A checks B's website (check(A, B)) and he finds no link referring to his own. A is not really disappointed (individual inhibition) because he had not actually expected to find himself acknowledged by B (who happens to be one of the more influential members of the network community), and so he continues to look for some other website-agent to exchange links with. Or it happens that he unexpectedly finds a positive answer (agree$(B, A,$ insertLink$(A, B))$) in his mailbox and a link on B's website ... and soon A expresses his gratitude by telling B that he is going to send a request to some of his own acquaintances (we assume that a function $acq(A)$

[6] It would conform to Goffman's interaction order[Goffman, 1983] to introduce an unquestioned convention that any action must be politely reported to whom it may concern. However, with regard to computational resources this would raise communication costs tremendously. From a sociological point of view we should distinguish between the conventions of face-to-face interaction and anonymous communication on the Internet.

[7] This comes quite close to Luhmann's definition of the operation called "Verstehen" as drawing a distinction between information and message. The concept may be interpreted to give complete autonomy to the agents and here is a crucial difference to KQML speech acts, which do not allow any message to remain unanswered.

returns a list of all agents agent A knows) to inscribe links to B's website:

$$\forall X \in acq(A).\text{tell}(A, B, \text{request}(A, X, \text{insertLink}(X, B))).$$

Meanwhile, many other agents inscribe and delete many other links, concurrently and at different times, and subsequently the network is in permanent evolutionary change. But is this the emergent social structure we are looking for? No, we are looking for the impact of conflicts on structure (and vice versa).

Conflict mode: An interactant A wishes to join the network and inscribes on his own website a link to B's website (insertLink(A, B)). Then he presents his link to B (tell$(A, B,$ insertLink$(A, B))$). B does not react neither explicitly nor implicitly, even after days A cannot find a respective link on B's website (check(A, B)). A doesn't give up and kindly but explicitly asks B to do him the favour of reciprocity (request$(A, B,$ insertLink$(B, A))$). At this point of interaction B (as an adherent to Goffman's rules of social intercourse) should be sufficiently motivated to answer A's request. But still B remains silent. A has to face the fact that B ignores his request which – in psychological terms – is more humiliating than a plain rejection. But neither rejection nor ignorance (reject$(B, A,$ request$(A, B,$ insertLink$(B, A)))$)[8] is a conflict. Instead, an interagent conflict begins to emerge as a conflict only when A declares his discontent in an unambiguous manner, either by asking B to give reasons for his rejection (request$(A, B,$ reason$(B,$ reject(insertLink$(B, A))))$) or by simply insisting on his request (i.e. repeating it), or, according to Luhmann's definition of a conflict communication, by explicitly rejecting B's rejection

$$\text{reject}(A, B, \text{reject}(B, A, \text{request}(A, B, \text{insertLink}(B, A))))$$

or by spelling out a warning (announce$(A, B,$ deleteLink$(A, B))$). We have an outspoken, meaningful dissent now – or a symbolically constituted social object. In other words: Both agents no longer can define the situation as a routine case but will have to realise that an established routine of networking has broken down

[8] We hopefully need not define an action "ignore" because we assume that this is implied by "reject". However, this issue is rather touchy and must be decided along with the question how to implement temporality (cf. Kai Lorentzen & Mathias Nickles on "Ordnung aus Chaos"). Do five minutes of silence have to be interpreted as "ignorance", hence as an "implicit rejection"? And what happens when a positive reply (BtellA(BinsertLinkA)) arrives seven minutes after A's original request and one minute after A has launched a counter-attack? He might, out of sheer impatience, just have spoilt the beginning of a wonderful friendship!

and they will probably follow up the coming events with growing awareness or concern! Or what is more likely, A's announcement just has no consequences and the conflict remains an episode and just dissolves.

5 A simple interagent conflict and its dissolution

1. insertLink(A, B)
2. tell$(A, B,$ insertLink$(A, B))$
3. request$(A, B,$ insertLink$(B, A))$
4. reject$(B, A,$ request$(A, B,$ insertLink$(B, A)))$
5. request$(A, B,$ reason$(B,$ reject$(B, A,$ insertLink$(B, A))))$
6. reject$(A, B,$ reject$(A, B,$ request$(A, B,$ insertLink$(B, A))))$
   $\Rightarrow$ *double rejection: this is the conflict point!*
7. announce$(A, B,$ deleteLink$(A, B))$

Excursion 2: Emergent conventions?

More interesting, perhaps, is the case of an agent C raising an objection[9] against B (with whom he entertains a linkage liaison) exchanging also linkages with A. What happens after C has found out (by check or tell) that B has exchanged linkages with A, is that Cspells out a warning to delete his link to B, unless B deletes his link to A (or unless A refrains from realising his announcement):

$$\text{announce}(C, B, \neg\text{deleteLink}(B, A) \Rightarrow \text{deleteLink}(C, B)).$$

"If you don't break off (or refrain from taking up) linkage liaisons with A, I will break off my relationship with you!" This is somehow similar to a wedding announcement's function to prevent bigamy – the difference being, of course, that there is not law or convention that prohibits any linkage liaisons. In the contrary,

[9] We might consider to distinguish a "rejection" (against interactions directly concerning an agent) from an "objection" (against third-party interactions indirectly concerning an agent). But introducing a new primitive "object" would mean to raise unnecessary complications.

such kind of warning must be considered as an illegitimate interference into another agent's "liberty of linkage", definitely a rather modern concept of social life. However, later on we should take keen interest in demonstrating *how a convention that regulates (restricts and at the same time: enables!) linkage liaisons could emerge* or be learnt in the course of network evolution, either by deliberation or by adaptation.

Excursion 3: Global knowledge by self-design?

Similarly we must ask whether we must explicitly design a mechanism for automatically publishing or advertising any linkage proposal etc. on a blackboard? Or is our model scenario so charming because agents do have incomplete information about other networking (inter)activities? This comes close to pragmatically "situated" knowledge: Usually there is no pressing need and no motivation to acquire knowledge about something, which is of no immediate concern for oneself, even if it is not impossible to obtain the information. Although "check" is an activity that virtually enables any agent to gather information about the liaison structure in laboriously working his way from website to website, it is practically unfeasible to obtain a global overview of the entire structure because it would take too much effort to search through the endlessly ramifying branches of a search tree. Therefore, at some very early point of up-scaling, there will be an ever pressing need for "global knowledge" in form of yellow pages: Is this a designers task or is it feasible to implement a *collective learning process which enables an agent society to self-design its own blackboard?* So, allowing agents to have access to "global knowledge" is not the same as to provide them with the capacity to control "global knowledge". From the frog's eye perspective of a network participant it is only possible to understand one's immediate "acquaintance structure".

6 From interagent to collective conflicts

After a conflict has been explicitly declared by a double rejection and an announcement to delete a previously offered link, A has three options to continue the conflict communication on an interagent level and a fourth option to raise the conflict:

1. A may realise his warning and actually do what he had already announced before by actually deleting his link to B's website again (deleteLink(A, B)).

This may be called an "entropic" option because its consequence is network destructuration.

2. A may do nothing. Warning need not be followed by sanction. A zone of uncertainty[10] is established and the conflict impends – or as time goes by, evolves out of existence. There remains, however, an interesting side-effect which deserves closer investigation as A maintains his link to B but B does has no link referring to A's website: an uneven or asymmetric relationship!

3. A may continue to molest B with permanent requests hoping that B eventually will give in and finally agree (or B will take countermeasures to get rid of what he must conceive as a pest).

4. Or A may ask other agents to interfere on A's behalf and to persuade or convince B or to build up social pressure on him thus raising the conflict on a collective level.

A may turn to others (third party agents $X \in acq(A)$) for an intervention on A's behalf, ultimately to coerce B back on the right track. The conflict escalates... To insert some "social-world realism" into our scenario, it can be assumed that further steps of gradually up-scaling the conflict can be taken. On the basis of broadcasting his case to a broader public of network participants and to invite others to take collective action against B, A continues to escalate his conflict with B by including others (third party agents). By taking the following "indirect" measures, an interagent conflict will be gradually transformed into a collective conflict:

a. The first thing for A to do is to spell out a warning to B that A will bring the case to other networking agents:

$$\text{announce}(A, B, \forall X \in acq(A).\text{tell}(A, X, \text{reject}(A, B, \text{insertLink}(B, A))))$$

b. A publishes the issue at stake by actually telling other agents, his acquaintances, about his conflict with B:

$$\forall X \in acq(A).\text{tell}(A, X, \text{reject}(A, B, \text{insertLink}(B, A)))$$

[10] B might assume that she had been scared by an empty threat or unfounded menace – but she can never be sure. Those who have followed up the research work on commitments or persistent intentions (Cohen & Levesque 1990) will know how tricky it is to provide agents with an option of "de-committing" themselves in a feasible and implementable way!

c. Moreover A may want to bring his case to a broader public by asking others to broadcast[11] to their acquaintances whom he does not know: $\forall X \in acq(A)$, A issues the message

$$\text{request}(A, X, \forall Y \in acq(X).\text{tell}(X, Y, \text{reject}(B, A, \text{insertLink}(B, A))))$$

After having informed the public (not the public opinion: that comes later), A continues to mobilise other agents and to build up more collective pressure on B by asking the others to intervene on A's behalf (who become "supporters" by doing so) up to the point where A's supporters collectively exclude B from the network (more precisely: from their particular networking segment)

d. A asks others to intervene on his behalf and to persuade B to exchange links with A:

$$\forall X \in acq(A).\text{request}(A, X, \text{request}(X, B, (\text{insertLink}(B, A))))$$

e. A asks others to warn B to exclude him from the network, unless he does what they want him to do:

$$\forall X \in acq(A).\text{request}(A, X, \text{announce}(X, B, \text{deleteLink}(X, B)))$$

f. A's ultima ratio is to ask others to sanction B by practically excluding him from the network: $\forall X \in acq(A).\text{request}(A, X, \text{deleteLink}(X, B))$

g. The others follow A's request and exclude B from the network by deleting their respective linkages: $\forall X \in acq(A).\text{deleteLink}(X, B)$

Again, if the others do not follow A's suggestion and do not exclude B from the network (deleteLink(X, B)), the conflict is resolved in favour of B and to the disadvantage of A. However, their are more options impending. B will probably not

[11] Here we could introduce a more powerful communicative act called "broadcast": an agent requesting all her acquaintances to send a specific message to all their acquaintances to send a specific message to all their acquaintances ad infinitum until everybody is informed. To make such a compound-shortcut mechanism like "broadcast" emerge from collective experience alias from repetition, would it be convenient (and theoretically adequate?) to have a kind of connectionist (or neural-net) operator which detects repetitive communications and integrates them into communication structures alias more powerful pieces of program code?

remain calmly seated in his armchair and idly watch[12] A's activities in forming a coalition against him. Instead of allowing all this to happen, B will begin to mobilise his own cohorts in order to form a counter coalition. And in case that he is able to succeed and effectively counterbalance A's coalition forming activities, the situation will appear to be highly explosive and threatening to the entire network community.

The interesting question is: What will happen next? In a constellation of two equally powerful coalitions opposing each other, there are at least three different ways to continue a conflict process at the collective or societal level:

1. Nothing will happen and, instead, the conflict cools out by way of annealing and will be soon forgotten (or stored in the collective memory of the network participants)

2. The conflict escalates into a "hot" war by someone[13] making his threat (announce) come true (deleteLink), thus possibly triggering a chain reaction in which the network community will be split up or differentiated[14] into two distinctive communities along the lines of previous coalition formation.

3. The conflict is stalemated because both coalitions are equally powerful and nobody wishes to run the incalculable risks of a hot conflict. At this point the conflict parties (or any other agent) can try (if it is not too late) to reframe conflict communication as a forum for "due process".

[12] To what extent do we have to assume that A and B are able to "observe" each others activities? A consistent "interactionist" perspective would have to restrict an agent's "local" knowledge about network activities only to those interactions in which he is directly participating or in which he is indirectly participating through other agents directly informing him about their interactions with third agents. In any case none of these agents should be endowed with "global" knowledge, i.e with a capability to monitor the entire network activities.

[13] Anyone? That would leave the coming events to all sorts of hazards which is certainly far too dangerous at a point of escalation where the whole network community is involved and potentially effected. So we would have to assume some kind of core agent (commander in chief) who acts on behalf of his coalition and controls switching from cold to hot war. And if so: do we have to predefine the role of a commander or is such a role to be conceived as the emergent outcome of previous interactions?

[14] Differentiation according to Luhmann's theoretical concept of segmentary or functional differentiation? Here we should follow up the question whether linkage conflicts might also foster functional differentiation along the lines of generalised media: e.g. one linkage network beginning to operate on a monetary code and the other on a code of subcultural value commitments. See below.

7 Forums for argumentation and mediation

Essentially both conflict parties will come to the point where they will either have to make their warnings come true or they will have to reframe the conflict setting into ”due process”, by inviting each other to renegotiate their case with the assistance of an impartial mediator. In the former case the network will probably undergo some drastic transformation and reorganise fundamentally by agents knitting new ties and linkages and breaking down old ones (perhaps along the lines of knowledge frames or generalised media?). However, since the scenario is embedded in a real world social context of emergent network communities on the Internet, the battlefield theatre that has been built up so far cannot really explode in a destructive fight because the utmost harm agents are capable of inflicting on each other is to break down interagent relationships. Indeed, as long as all agents have a realistic exit option, deleting a linkage liaison is just everyday business and does not appear to be a really harmful way of sanctioning someone. Furthermore, even huge linkage conflicts may be channelled in a more ”peaceful” way. In the latter case the scenario becomes far more complex because now reasons and arguments must be given and agents will have to disclose their motives: As long as the conflict parties do have complete freedom to take any time they need and since they do not have to act or react automatically or immediately, nor unconsciously, there will be no chain reaction triggered by any of the announcements, approvals or rejections. Agents are able to revise their previous declarations at any time. And they do not have to take action (insertLink, deleteLink, check, reason) just because they have announced to do so. Thus, an agent can suspend and even dispense any announcement at will. In such a highly contingent setting there is a need for implementing two facilities:

1. To cope with contingency – both conflict parties are unable to oversee (no global knowledge!) the consequences of ultimate conflict action (deleteLink) – agents will have to handle the coming sequence of interwoven interactions in a particularly cautious and thoughtful way. They must carefully reflect what to do next before taking action in a potentially explosive situation. Such ”contingency awareness” – which might grow through experience gained in a simulation experiment by repetition and accumulation, but must not be confused with global knowledge! – opens up new avenues for reframing a conflict scenario into a scenario of deliberation and argumentation (”due process”), and even while hostile interaction is already going on.

2. A ”due process” scenario (or forum) will have to reduce uncertainty and contingency by introducing expectation-based roles like (i.) impartial agents or mediators, (ii.) regulations and procedures for argumentation and deliberation, and (iii.) binding or obliging decisions underpinned with commitments and sanctions. With regard to our scenario of two emergent conflict coalitions we will assume that the conflict process has not entirely demolished all available (social, moral?) resources of the network community and that there are still enough impartial third party agents available who so far have declined all invitations to join any of the emergent conflict coalitions. These impartial agents may now play the role of mediators. When calling for a mediation process, the conflict parties as well as all other agents involved in the procedure ultimately commit themselves to the outcome of the decision. A due process scenario begins with the installation of a kind of proto-legal procedure and it terminates in a collective decision or a majority vote (a jury's verdict) by making use of argument and deliberation.

3. A conflict situation is transformed into due process when a conflict party (for whatever reasons) calls for a mediation procedure and the other agrees. However, in our scenario any such procedure or institution (mediators, argumentation, decision) does not exist so far. At this point it is just all too tempting to implement a predefined ”social structure” (off-line design of roles, regulations, decisions). However, this would be too easy. Our ambition is rather to predefine as few concepts as possible and find out instead how far we can get along with systematically applying the action primitive ”reason”. Up to this point we have deliberately avoided to make practical use of ”reasons” – with the exception of agent A asking B in our first example to give reasons for his rejection, a request which was ignored by B. Now we change gears by ”instructing” our agents to give and ask for reasons more frequently. This ”instruction” should be conceived as a collective learning process propelled by ”inhibition” or, more exactly, by creative attempts to overcome an inhibition[15] which agents have inflicted on each other. Agent societies should learn to differentiate between contexts (Goffman: frame [Goffman, 1974], situa-

[15]Changing gears from the (1) low level of pre-conflictive ”inserting” and ”deleting” linkage liaisons to the (2) medium level of ”conflict communication” and from there to the (3) high level of ”reasoning communication” implies growing resource consumption (runtime, processing capacity, memory, communication overhead etc.). Accordingly, agents should be able to switch between a ”routine mode” and a ”reflexive mode”.

tional context) in which it is appropriate to overcome a conflict by "reason communication" (argumentation and deliberation about motives, interests, consequences etc.) and contexts in which exchanging reasons would just end up in a mess of even more deeply entrenched hostility (hence a waste of computational and social resources).

In the following we will give a simple example of an agent C (already entertaining a linkage liaison with agent B) who opposes against B's announcement to exchange linkages with A:

a. C spells out a warning to B in saying "if you don't break off (or refrain from taking up) linkage liaisons with B, I will break off my linkage liaison with you":

$$\text{announce}(C, B, \neg\text{deleteLink}(B, A) \Rightarrow \text{deleteLink}(C, B)).$$

b. B ignores C's warning and now C informs his acquaintances and asks them to take collective action against B by excluding him from the network:

$$\forall X \in acq(C).\text{request}(C, X, \text{deleteLink}(X, B))$$

.

c. However, instead of supporting or rejecting C's case[16] (as in the previous example), acuaintance X will ask C to give reasons and justify why he does not want B to exchange links with A, i.e.

$$\forall X \in acq(C).\text{request}(X, C, \text{reason}(C, \text{request}(C, B, \text{deleteLink}(B, A))))$$

and they will also ask B to give his reasons.

d. C argues (1-) that A is providing the same kind of Internet service (evidence: website content "check") as C and that C would not have a fair chance to compete with a "big player" like A (evidence: A entertains "foreign" linkage liaisons with powerful partners outside the network):

$$\forall X \in acq(C).\text{tell}(C, X, \text{reason}(C, \text{deleteLink}(B, A), 1-).$$

[16]Note that our agents just have invented the notion of "impartiality"! Note also, however, that they would never be able to invent anything unless we write a program that enables them to do so. In other words: How is it possible to implement "deviant variation" or "creative invention" into a program as a nucleus for evolution or emergent institutions?

e. Meanwhile B sends a note to his acquaintances Y to inform them about his conflict with C and about X's request for reasons and invites them to a general discussion[17]:

$$\forall Y \in acq(B).\text{request}(B, Y, \text{request}(X, B, \text{reason}(B, \text{insertLink}(B, A))))$$

f. After acquaintances Y have accepted B's invitation, B addresses all who are involved now, including his prospective linkage liaison A, arguing (1+) that competition is better for the network community than having a "monopolist"; (3+) and that A is a gain for the network because he shares the network values (evidence: A's website content and that of some of his "foreign" liaisons):

$$\forall X \in acq(B) \cup acq(C).\text{tell}(B, X, \text{reason}(B, \text{insertLink}(B, A), \{1+, 3+\})).$$

g. Now everybody begins to send arguments pro and contra, the debate is confused and blurred by other agents trying to settle old scores, but before it all ends up in a chaos, D (from C's party) and M (from B's party) suggest to use D's website as a common address to get the reasoning process organised.[18]

h. The following discussion shows that a large majority is ready to follow B's reasons (1+,3+) because he argues from a general perspective of community values whereas C argues from a rather particular perspective of individual interests.[19]

[17] For the time being, B's invitation is merely incidental. From the perspective of a knowing outside observer (the designer) it looks as if B has just invented a central principle of legacy: rational argumentation under public control! How is it possible to institutionalise B's invention by simulating repetition, recursion, confirmation and condensation?

[18] D and M have just invented a blackboard system (with procedural rules like sequential argumentation) and what is more, they have enacted themselves as rudimentary mediators – a role that will probably raise more expectations in the course of events when a certain agent is asked again and again to be the mediator. In a large-scale scenario of many conflict forums occurring simultaneously, a temporal common website should be evolving into a permanent common website not "owned" by a single agent but by the entire network.

[19] Yet another invention! They have just created a distinction between acceptable (social) and unacceptable (individual) arguments: If the new distinction will be repeated, confirmed and specified in coming episodes, this might eventually also change the general critera for exchanging linkages.

i. Hence, the majority[20] invites A to exchange linkages with them and A becomes a de facto participant of the network.

Excursion 4: Emergent structures again

At this point we must halt and try to understand some of the heroic assumptions implied in our example. If we want to design an artificial social system which is able to generate (almost ex nihilo) an emergent structure of "due process" or reflexive mediation with differentiated roles (mediator, jury), procedures of argumentation (openness, evidence, sequentiality), standards of argumentation (normative justification instead of individual benefit), public obligation (commitment to majority decision) etc., B's suggestion (as in steps e. and f.) to include and henceforth publicly address all other agents (instead of just privately addressing a collection of acquainted agents, as in steps b., c., d.) must become more than an episode. It must be confirmed and specified, reconfirmed and respecified again and again in a very time consuming (if not: endless!) process of recursion. In the footnotes to the previous example we have identified some crucial inventions (or: contingent variations) and the question is: How is it possible that such "inventions" will crystallise into social structure in subsequent conflict episodes, if contingent variation and discontinuity is much more likely than repetition and recondensation? Again, it is all too easy and all too tempting at this point just to implement a pre-given structure that invisibly shapes and controls the course of collective action and communication, thus directing the agent society onto a prescribed evolutionary track.

Take, for instance, the point at which the majority decides in favour of A and B against C(step g.): Our agents might as well have voted in favour of C by explicating and accepting his implicit argument that the network's general goal is to support subcultural mutuality and self-organisation among small enterprises ("small is beautiful") and keep "big players" like A out of the community; or that A's commitment to the community values is doubtful because he has too many "foreign affairs" outside the network (evidence: A has many linkage liaisons with influential websites outside the network community). Moreover, the rationality of an argument may change over time: What is good for the network community at one

[20]The next step would be to institutionalise the majority principle according to which the minority is obliged to accept (not to act against or obstruct) majority decisions.

time, namely openness to new "members"[21] with rapid growth to gain more influence in a turbulent environment (evidence: ones own network is much smaller than other comparable ones), might be wholly threatening in a different situation where the network is too big and heterogeneous to maintain its integrity. Or we might envisage a situation which is framed by two different emergent belief systems or generalised media (on one side a small-is-beautiful communitarianism, and an aggressive commercial orientation dominated by big business on the other: cf. footnote 12).

Excursion 5: Performance ("fast motion") versus evolution ("slow motion")?

In the preceding episode of a conflict communication we have demonstrated in "fast motion" what is actually a very slow and time consuming process of evolutionary learning. What we would need instead is a "slow motion" computational simulation to make an artificial social system evolve in a sociologically more adequate manner. Hence, what we will need is a computational experiment running on a permanent basis instead of one of the usual short-term simulations which are able to hand out instant results to an impatiently waiting audience. Furthermore we need to elaborate a completely new concept of what may be called emergent "turning points" of learning or evolution: from individual invention to collective innovation, from collective innovation to institutionalisation and vice versa. From an engineering methods point of view such prospects do not look very appealing though, since software engineering methods are developed to obtain high efficiency and high speed at low costs. However, we contend that a very slow social system which is gradually building up and maintaining and modifying an emergent social structure could be efficient and fast from a "top down" perspective: If we consider a large society of many different network communities operating in parallel, these might be able to learn from each other at "fast motion" speed. While building a social structure from scratch usually takes ages, a network community adopting or

[21] Here is another issue which needs clarification: openness, closure, membership. Our example shows that, while the network community begins to engage in a collective discourse, the conflict is not simply about an agent A being adopted by and into a network community or not. Rather it reveals that the big story behind the little episode is about maintaining network identity when two different networks are merging, namely B's acquaintances and A's acquaintances. Or is it not a merger but just a one-point-liaison between two networks? But what is a network anyway?

importing proven concepts developed elsewhere would not necessarily have to wait for evolution. What we would need to elaborate, then, is a vision of two (or more) different time scales: a bottom level of high inertia which is slowly reproducing a social structures and a task-force level operating at high speed to cope with critical situations, and, this is the decisive point, which is able to do so because it can feed on resources provided by the slow-motion level. While the task-force level should meet engineering standards of efficiency, the bottom-line level should comply with engineering standards of robustness.

8 Interagent (”micro”) and network (”macro”) arguments

Coming back from the distant vision of an artificial social system meeting sociological as well as software engineering requirements to the more immediate demands of designing a more simple multiagent system, we would need to elaborate further mechanisms and facilities to organise conflict forums, such as permanent blackboards (whiteboards) and institutionalised roles for mediators, specific calls for mediation and a closure procedure to prevent chaos and to organise the progress of collective inquiry, deliberation, argumentation, prompt conclusion, and to produce a binding outcome (majority decision). A majority decision has far-reaching implications. Agents must bear the consequences of autonomous action. As a social structure evolves, it would gradually reduce the consequences of individual action without restricting agent autonomy ex ante. Thus, indirectly, conflict parties just as all other agents involved in a conflict forum must commit themselves to the majority decision, whatever its outcome, which includes collective ”deleteLink” sanctions against anybody who acts against a majority decision. Again, such a tool would largely have to rely on the set of basic actions and communication primitives that have been defined above.

Quite difficult to handle, specifically from a computational perspective, are ”reasons” or arguments. Reasons seem to require a sophisticated semantic elaboration and must be disclosed and explicated and weighed in a complicated procedure of operationalisation and parametrisation. Going through a mediation and negotiation procedure conflict parties must reconstruct their conflict history, analyse the patterns of escalation, analyse their reasons and motives and suggest and test out different solutions. Agents therefore will have to be endowed with specific reasoning capacities beyond the scope of ”just” relinking and delinking themselves from

and to others. In order to avoid excessive semantics, we must try to find a rigorously simplified answer to this problem. In the course of forum evolution it might be convenient to include due process level arguments such as a specific ”order of interaction” (polite, civilised, obliging) or a more differentiated system of expectations and sanctions than just ”insertLink” and ”deleteLink”. Or it would perhaps be interesting to introduce a direct and more comprehensive website rating with criteria like professional appearance, information content, consumer attractiveness etc. But this looks far too complicated and would not add very much to the basic of an evolutionary structure. Therefore, an rigorously simplified reasoning semantics will not consist of more than three types of reasoning semantics both at a ”micro” and a ”macro” level.

At the ”micro” level of a linkage conflict an agent who is offered an exchange of links by another agent must be able to reason about linkage prospects: about the impacts on his personal business prospects (is the other agent a competitor or a promoter?), on bilateral dependencies (is the other agent influential or non-influential?), on value commitments and general belief systems (e.g. a kind of self-sustaining subcultural communitarism versus a strict commercial orientation – perhaps with an innate evolutionary tendency to transform the networks ”ideology” from communitarianism to commercialisation through linkage conflicts and shifting membership.[22] Again: Without rigorously reducing the number of ”content variables” or criteria to a maximum of two or three, we would easily run into a combinatorial explosion. So to make the scenario more tractable, we should introduce the following three types of interagent level reasons (which we have already drawn on in the previous example):

8.1 Interagent level arguments

1. support (promoter, competitor)[23]

[22] Prospects for evolutionary change: the ultimate criterion of commercialisation will be the introduction of ”linkage auctions”, which would render linkage liaisons more stable and less disputable because they have been paid for.

[23] Support: From reading A's self-description B will have to judge whether A will support his goals or must be regarded as a potential competitor. If this is too difficult to achieve, the input could be supplied directly from the user or by an interface agent. If the argument is symmetric in the eyes of both agents they would come to the same conclusion and their would be hardly any conflict. However, a support criterion is all but unambiguous and, moreover, an agent should be able to shift his concepts over time and circumstances, which gives rise to conflict and reorganisation.

2. influence (traffic or number of clicks, active links, references or passive links)[24]

3. value commitment (same, different)[25]

Interagent level arguments equip the agents with basic criteria for framing (evaluating) other websites together with an inferential capability to decide whether to propose (or reject) an exchange of linkages. By checking a website's content (in using, to begin with, an ordinal scale or low-level distinction, subcultural versus commercial orientation for instance) and influence (number of hits per day: provided such information is available), agents could reason about their individual advantages and disadvantages in case of establishing liaisons. (High traffic website agents will probably publish their traffic figures more readily in order to demonstrate their influencial power, while low-traffic websites will not.). Indirectly, agents could acquire information by checking the linkages from and to other websites and by relying on the wisdom of others (acquaintances) – a rather resource consuming procedure, as it seems.
The "micro" interagent level should be distinguished from a "macro" network community level with three different types of arguments (these have also been drawn on in the previous example):

8.2 Network level arguments

4. size of network (too small, too big, growing or shrinking?)[26]

[24]Influence: This is a quite straightforward argument because influence can be measured at cardinal-scale level. The problem arises from unreliable information sources. The only information which can be checked immediately is the list of active links, while passive links and number of clicks must be taken at face value. In principle, passive links could be followed up by checking the websites which are mentioned in the active-links list, but this could be quite time consuming. Clicks can not be checked at all, unless a website owner would allow foreign agents to monitor his private communication. . . However, a website community network could introduce a rule which would compel all website agents to subscribe to a standardised self-description. Making a self-description obligatory should be compatible with community conventions. Again, these should not be hardwired but emergent properties of network evolution.

[25]Value commitment: Is a website in line with the values and goals of the network community? This touches the issue of cultural "sameness" (fear of eroding community values, e.g. by excessive commercialisation, versus fear of marginalisation), an issue which is rife with dissent rather than bearing consensus.

[26]The question is: How do agents have access to "global knowledge" of how many agents are linked to the network, how has the network evolved in the past and where is its future development

5. power relations (symmetric, asymmetric, checks and balances)[27]

6. heterogeneity (same, different, differentiation/integration)[28]

On the level of collective reflection the ”macro” impact of a conflict on collective regulations of the network community can be analysed with regard to network size (is the network too large and, accordingly, social relationships too anonymous to support internal solidarity?), cultural integration (too foreign, too heterogeneous?), network influence onto the outside world (too marginal to influence the environment?), and democratic standards (asymmetric power relations?).

Introducing network level arguments, however, is incompatible with a bottom-up approach unless agents are endowed with ”global knowledge”. To reason about the advantages and disadvantages of a linkage liaisons and linkage conflicts from the overall perspective of general network prospects without global network knowledge is impossible. Global network knowledge would consist of an overall picture about the empirical network structures together with some rudimentary theoretical knowledge about the impact of different types of networks (star, circle, line) and their possible combinations, about meta-networks and network boundaries, about closure and saturation, and above all: about the risks and advantages of evolutionary network transformation and differentiation for the participants. Of course, in principle it is possible that all available local knowledge can be brought and puzzled together (based on ”tell”, ”check”, ”insertLink” operations) to form a comprehensive or synthetic picture of a network's liaison structure at a particular point in time. But is it realistic to implement such a resource consumptive method of building global knowledge? And could it be justified on the grounds of sociological adequacy?

---

going to, if we generally assume that their knowledge is reduced to what theys know about there acquaintances?

[27] The same question: How do agents know about the asymmetric power distribution within the network and from where do they get the preferences or criteria in favour of more or less symmetric or asymmetric, more democratic, egalitarian etc. network structure? Are such different preferences emergent properties and could they be subjected to open conflict or inquiry into inquiry?

[28] How much heterogeneity is a network able to support without breaking into different fractions of value communities or sub-societies?

9 Desiderata and open issues

As it seems, deriving a first paper version of a rather complex model of network formation and collective conflict processing out of a minimum of speech-act like communications, is not altogether impossible. However, several issues remain unresolved: How are agents able to reason on linkage exchange and coalition formation at the "local" interagent level; how and on what grounds do they decide whether to exchange a link or not, whether to join a coalition or not; how can they reason about the outcome of linkages, conflicts, and coalitions at the "global" level of a network; and what renders stability or continuation (i.e. structure) to any coalition or any linkage in the face of agent autonomy and contingent interaction? Moreover, how are linkage liaisons embedded within the empirical context of webside-owners in modern society, how can they respond to the needs of a modern society when agents look like tribal warlords calling their followers to the arms on the basis of kinship affiliation and face-to-face interaction? Does our scenario embody something like feudal commitments rather than modern social behaviour? On one hand, grounding conflict coalitions on "feudal" (or rather: communitarian) obligation for and trust in mutual assistance seems to be quite plausible with regard to modern subcultural communities rooted in today's reality of modern societies. On the other hand, when agents are endowed with the autonomous right to say no and to join and accept or to refuse and reject any commitment at any time, the scenario is far from any modern social reality too. This is so because we have not solved the basic dilemma of the bottom-up approach: generating emergent frames of interaction (i.e. normative structures, repeated patterns of interaction) is not warrented by strong agent autonomy. Granting autonomous decision making to the individual agent harshly collides with the fact that we must presuppose an autonomy which is tacitly controlled by procedures and standards for legitimate argumentation, power, influence, values. . .

From a software engineering point of view things do not seem to look much better, perhaps. Chaotic oscillation[29] at run-time behaviour is imminent unless we assume that our agents A and B dispose over some legitimate power to oblige others, un-

[29] Note that notions like "fluidity" and "variation" should not be easily excluded from further consideration because of alleged non-computability. Instead, they should be considered more seriously than this would be done from an engineering perspective. Something like "chaotic oscillation" (or at least a functional equivalent) should be considered particularly because interaction and communication are rather slow – not to speak of structural social change! – compared to cognitive operations.

less we assume predefined roles, some global knowledge at everybody's disposal etc. If not: What might increase the probability that other agents will readily follow A's suggestions? And why not B's? And who is the maverick anyway? A or B? If everybody is free to link him – or herself to others or not, there is no moral or legal obligation in the first place. If everybody could thwart anybody else's plans at any time, such prospects normally would not be considered as an invitation to coordinated action or socially structured behaviour anyway. The question, then, is how coordination is possible under the precondition that all agents are endowed with equal autonomy. Since unlimited agent autonomy in an open website-network scenario will yield unstructured linkage conflicts and produce chaotic behaviour (or social entropy), we will have to introduce (or impose) some sort of restrictions. Moreover, if each agent follows its own prerogative, it is unlikely that others should be willing to form a coalition with A, unless we assume some preconditioned mutuality (if A is ready to assist $C, D, E \ldots$ in their particular cases they will assist A in his case...). Moreover A will not begin to escalate a conflict if he does not have good reasons to believe that his case is going to be supported by a sufficient number of others. Again, on what grounds? Somehow, his case will have to be justified by some basic values, norms or principles or a volont generale which motivates potential supporters to join a coalition.

We assume that CONSTRUCT models or architectures differ from more conventional DAI applications not only because they operate on a permanent basis (cf. computational ecology) and that they produce, reproduce and modify an "emergent" social structure while resolving (or not resolving!) conflictive episodes. The presumed advantages of the proposed model, both from a software engineering and from a sociological perspective, should be seen in the way that a CONSTRUCTNETWORK would have to meet the demands of shifting from a highly efficient and cost saving routine mode into a resource consumptive conflict mode and back again into a new routine mode after the conflict has been settled. In other words: the system should be able to learn from conflicts and, in doing so, an agent society would operate like a self-organising mechanism or a non-trivial machine or a self-sustaining general purpose infrastructure. As a self-restoring and self-modifying infrastructure it must maintain a history file of its own conflict episodes. And it would need some neural-network like mechanisms to re-enforce frequently used patterns of interaction, conflict resolution, due process, "emergent" conventions etc. In order to built such a system, a first prototype should perhaps be restricted to

learning just a minimum of core conventions. For instance, an agent society could learn by conflict mediation that ”idiosyncratic” justifications at the interagent level are gradually replaced by ”generalised” justification at the network level – in analogy to the pragmatist intuition (in line with Mead and Dewey, cf. [Mead, 1930, pp. 412], [Habermas, 1992, pp. 224], [Honneth, 1998]) that ”particular” standards of justification are more and more replaced by ”universal” democratic standards.

References

[Finin et al., 1997] Finin, T., Labrou, Y., and Mayfield, J. (1997). KQML as an agent communication language. In Bradshaw, J., editor, *Software Agents*, pages 291–316. AAAI Press/The MIT Press.

[Goffman, 1974] Goffman, E. (1974). *Frame Analysis: An Essay on the Organisation of Experience*. Harper and Row, New York, NY. Reprinted 1990 by Northeastern University Press.

[Goffman, 1983] Goffman, E. (1983). The Interaction Order. *American Sociological Review*, 48:1–17.

[Habermas, 1992] Habermas, J. (1992). Individuierung durch Vergesellschaftung. In Habermas, J., editor, *Nachmetaphysisches Denken. Philosophische Aufsätze.*, pages 187–241. Suhrkamp, Frankfurt am Main.

[Honneth, 1998] Honneth, A. (1998). Democracy as Reflexive Cooperation: John Dewey and the Theory of Democracy Today. *Political Theory*, 26:763–783.

[KQML, 1999] KQML (1999). *The UMBC KQML Web*, http://www.cs.umbc.-edu/kqml/.

[Kumar et al., 1999] Kumar, R., Raghavan, P., Rajagopalan, S., and Tomkins, A. (1999). Trawling the web for emerging cyber-communities, http://ww8.org/w8-papers/4a-search-mining/trawling/trawling.html.

[Luhmann, 1984] Luhmann, N. (1984). *Soziale Systeme. Grundriß einer allgemeinen Theorie*. Suhrkamp, Frankfurt am Main, Germany.

[Luhmann, 1990] Luhmann, N. (1990). Anfang und ende: Probleme einer unterscheidung. In Luhmann, N. and Schorr, K. E., editors, *Zwischen Anfang und*

Ende: Fragen an die Pädagogik, pages 11–23. Suhrkamp, Frankfurt am Main, Germany.

[Mead, 1930] Mead, G. H. (1930). Philanthropy form the point of view of ethics. In Faris, E., Laune, F., and Todd, A. J., editors, *Intelligent Philanthropy*, pages 48–133. Chicago, IL.

[Rogers, 2000] Rogers, R. (2000). Playing with search engines, mediamatic, issue 9.3, http://www.mediamatic.nl:16080/magazine/magazine.html.