

A Discussion on Fundamental Approaches for the Engineering of Autonomic Systems

Peter Deussen³, Benoit Bruyère², Claire Fahy¹, Sven van der Meer¹, Maurice Israel², Edzard Höfig³, Dave Lewis⁴, Declan O’Sullivan⁴, Dominic Greenwood⁵, Alessandro Bassi⁶, Jacques Saraydaryan⁷

¹Telecommunications Software & Systems Group, Waterford, Ireland

²Thales Communications

³Fraunhofer Fokus, Berlin, Germany

⁴KDEG, Trinity College Dublin, Ireland

⁵Whitestein Technologies, Switzerland

⁶Hitachi Europe SAS, Valbonne, France

⁷ExaProtect, France

peter.deussen@fokus.fraunhofer.de

Abstract. An Autonomic System (AS) is one that delivers dramatic reductions in operational expenditure by implementing self-management features such as self-configuration, self-optimization, self-healing, and self-protection. Engineering autonomic features requires that service engineering is acutely responsive to ongoing knowledge of the operational changes experienced by deployed systems. Such knowledge requires on-going analysis of emergent system-to-system interactions, the emergent people-to-system interactions and, importantly, the changing interactions between those people. In this paper we introduce an approach to treating the network of interacting people and autonomic systems as a single dynamic system, which we term the Expanded Society, and a research agenda that will develop the approach.

1. Introduction

In this paper, we would like to introduce the seminal ideas of a new framework. This framework will eventually lead to the development of new engineering principles. These principles will allow Autonomic Systems to cooperate organically within instances of an Expanded Society, where for Expanded Society we mean the network of interacting people and autonomic systems as a single dynamic system. These enhanced Autonomic Systems will exhibit homeostatic properties, able to adapt autonomously to changes in their contextual and situational status on multiple levels (technology, business, usage, etc.) while maintaining overall stability. By viewing Autonomic Systems as an integral part of our future society, we look towards a “science of design” for those systems, beyond the existing system engineering principles for classical systems.

As stated in the abstract, Autonomic Systems promise to reduce operational expenditure by replacing human involvement in systems operation with automated

intelligent control loops that are governed by high-level operational rules. The engineering principles discussed in this paper will help to reduce the capital expenditure by providing guidelines for cost effective design and deployment of Autonomic Systems, which in turn will benefit from the introduced engineering rules by means of an extended life cycle (as compared to the rapid technological innovations and the almost on-demand service deployments).

The engineering of service software (Autonomic Systems) that this paper introduces addresses the engineering of associated autonomic features, while ensuring that these features are viable, i.e. the cost savings they deliver are maintained regardless of the frequent operational changes experienced over the lifetime of the service software. Engineering such viability cannot be achieved solely through traditionally static requirements analysis. Instead it requires that service engineering is acutely responsive to ongoing knowledge of the operational changes within the context of the deployed service software. Such knowledge requires on-going analysis of emergent system-to-system interactions, the emergent interactions between these systems and the people who use and manage them and, importantly, the changing interactions between those people. These person-to-person interactions drive the purposeful changes in the operational requirements of the systems while also responding to changes in the system-based interaction. To ensure that the analysis of person-to-person interactions is considered in close conjunction with those involving the systems directly, we treat the network of interacting people and autonomic systems as a single dynamic system which we term the *Expanded Society*.

The (predicted) constant flow of composed services requires methods to design interactions and tools to instantiate and deploy them in a viable way, i.e. enabling the continuous emergence of new services and managing the disappearance of obsolete ones, without destabilizing the operation of the existing ones. The vision of *Expanded Society* therefore encompasses the co-evolution of human organizations and technological systems as they aim to be self-managing at any point in time while also being homeostatic in this evolution, i.e. managing change and maintaining a stable identity while altering both form and function. Ultimately, Autonomic Systems engineered in this way will enable the emergence of a differentiated and highly structured landscape of durable, adaptive and successful communities comprising of system developers, system operators, service providers and end users.

This paper defines and discusses a research agenda that aims to define the engineering principles for Autonomic Systems and their collaborations, with regard to their embedding in the broader context of an Expanded Society. It begins by providing an overview in section 2 of related European research activities. In section 3, the approach is introduced, and core technical principles are then outlined in section 4. In section 5, an initial engineering methodology in support of the approach is described. Finally initial conclusions are presented in section 6.

2. Related activities within European research work

To date, EU funding in the area of Autonomic Systems has focused on research into Autonomic Communications, in which the authors of this paper have played a central

role. A workshop in June 2003 (“Communication Paradigms for 2020”¹) identified the need for: “an integrated technology platform for sensing, communicating, decision making and reacting in future networking and networked businesses” [6]. It was envisioned to be “centered around networking selfware – a novel approach to perform network control, management, middle-box communication etc. based on principles of high evolvability and generic network instrumentation”.

In May 2004, the European Commission initiated the Autonomic Communication Coordination Action (ACCA). ACCA led to the creation of the *Autonomic Communication Forum (ACF)* described in section 2.1, and fostered the creation of several consortia that blossomed into a number of important projects, which we see as milestones that helped us to form our vision:

- ANA – focusing on the development of a post-IP approach for networking. Research is carried out that aims for the replacement of static IP protocol stacks with variable and composable networking functions. Network nodes are envisioned as being grouped into so-called compartments, i.e. domains with a common set of functions, policies, mechanisms, etc.
- CASCADAS – researching on novel approaches for component based service architectures, utilizing autonomic functions like self-organization, self-management, and self-preservation. Declared goal is the creation of an open-source, autonomic communication toolkit easing development, deployment and maintenance of autonomic services. For further information see [1].
- BIONETS – addressing novel service composition and service life cycle management approaches by utilizing bio-inspired principles to achieve autonomic operation.
- HAGGLE – aiming to define an innovative networking architecture designed to enable communication in the presence of intermittent network connectivity, which exploits Autonomic and Opportunistic Communication, in particular in the absence of end-to-end communication infrastructure.

Broadly, these projects are all developing autonomic communication systems of different type. These systems will provide a source of variety when considering the requirements for engineering of autonomic systems. Furthermore, they are ready to be used to illuminate the engineering problems involved in assembling heterogeneous autonomic systems to form a cohesive, federated whole. The practical experiences of these projects already provide a valuable sounding board for early assessment of the practicability of our engineering methods and tools. Our work represents a progression beyond the state of the art being developed in these projects by moving from the development of specific autonomic solutions to the development of a systematic and repeatable methodology for developing ‘integrateable’ elements and composite autonomic systems in the future.

More specifically, the CASCADAS analysis of the ‘Connected Society’ provides an input to the broader analysis of our *Expanded Society*. Both CASCADAS and BIONETS are taking service-oriented approaches to constructing autonomic communication systems, the former focusing on a component model, and the latter on

¹ A first brainstorming meeting was organized in Brussels on 22 July 2003, as a first step toward the definition of a future FET proactive initiative, launched at the end of 2004.

service parameterizations, configuration, and composition. We propose to expand on this work by developing a fully semantic autonomic service component model that can evolve and adapt both through service composition and configuration, but also through decentralized application of policy-based management. It is this managed interaction of service-based and policy-based adaptation in a broader organizational setting that is coupled to the AS development process that is the added value and integrative differentiator of our proposed approach.

Recently, the ACF has re-focused towards creating the first international standards body for autonomic communications. Officially introduced at the 1st Autonomic Networking conference in Paris (September 25-29, 2006), the first ACF meeting of the new Board took place in Brussels on March 1st 2007. The set of ACF standards will be led by industry and supported by academia, and will focus on mechanisms required to ensure interoperability. The coherence provided by standards is supported through research and development. Research, led by academia and supported by industry, provides coherence by aligning current and future research projects to the ACF vision as realized by its standards activities. Joint development will cement the cooperation between industry and academia, and provide coherence through coordination of projects funded by organizations such as the EU and the NSF.

All of the authors of this paper are active in the ACF, as board members, co-chairs of working/expert groups or contributors. Our work, presented in this paper, can be seen as a core initiative with strong support by and from the ACF working/expert groups.

3. Our Approach

We are aiming to develop new *engineering principles for collaboration of Autonomic Systems*, which are meant to operate within instances of an Expanded Society. In other words, we are focusing on the optimization of collaboration aspects within composed systems of Autonomic Systems. As the MACE 2006 [2] workshop has shown [3][4][5][7], the collaboration and the Autonomic Systems themselves are beyond those of classical ICT systems.

The collaboration of ASs will exhibit *homeostatic properties*, able to adapt autonomously to changes in their contextual and situational status. These changes can occur on multiple levels (technology, business, usage, etc.) and the subsequent adaptation should maintain the overall stability of the AS.

As the ASs are designed for continuous operation, they will have to fulfill a variety of purposes that might be subject to change over their lifecycle. The purposes might even contradict each other from time to time. Achieving stability thus implies the capability to settle conflicts and to deal with/react to internal and external stimuli.

Autonomic Systems *assess the quality* of their own operation(s) constraint by flexible criteria and conditions. This assessment process will aid to exhibit the self-* properties, such as self-optimization, self-healing, self-protection, and self-management.

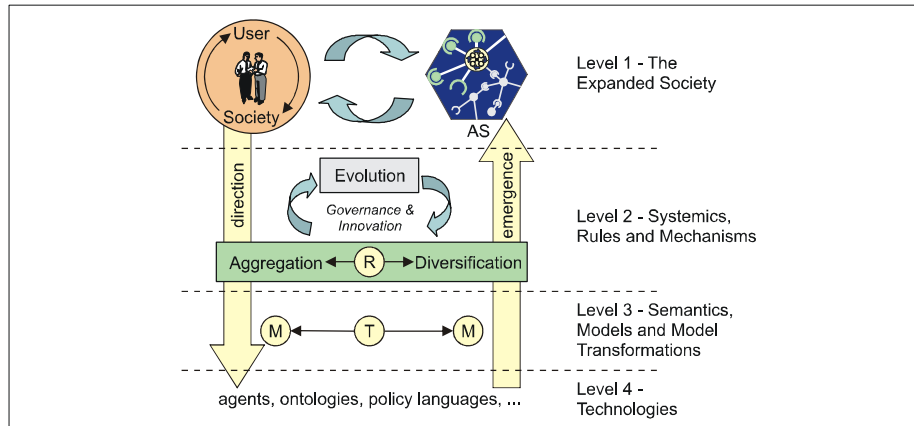


Fig. 1. Approach – Conceptual View of the identified Levels

The basis of our approach (in contrast to related work, such as IBM’s autonomic computing and HP’s utility computing) is to view Autonomic Systems as an *integral part of our future society*. We are aiming for an integrated solution (by means of engineering principles) on the level of society, not on the current best practice level that wants simply to “autonomize” systems (which are primarily designed as a single system or product). As shown in Figure 1, our approach integrates four different levels:

1. *Expanded Society* – studies the changing system-to-system, system-to-person and person-to-person interactions that emerge within a community of human users and Autonomic Systems. An instance of the Expanded Society could be a specific human organization or community of service providers and consumers.
2. *Systemic, Mechanisms, and Rules* – considers design principles and their conceptual foundation with respect to the aggregation and diversification as categories of principles. These categories are the further basis for the construction of Autonomic Systems. They may also help to form the mechanisms for role assignment and work distribution. The two main categories are *Evolution* (ongoing homeostatic system adaptation) and *Governance* (driving evolution to ultimately achieve goals of level 1).
3. *Semantic Models and Model Transformations* – uses models and model transformation as basic means to express the concepts of level 1 and 2. The models as well as the transformation needs to be enriched with semantic information, in order to allow for their continued extensibility reflecting the needs of the Expanded Society and considered technologies. This provides the basis for the knowledge which is created by progressively gaining information from the operation of the Autonomic Systems. The combination of all three aspects of this level will lead to “machine processable” models of various types, such as rules, ontologies, logics, UML models, etc.
4. *Interfaces to the technology* – defines the external interface that Autonomic System should provide in order to be compatible with our engineering principles and

consequently with each other. Candidates under investigation are agents, ontology, and policy languages; to name a few. Since our focus lies on engineering principles, we are not going to develop new technologies as such; rather re-interpret and re-combine existing ones.

As demonstrated in Figure 1, these levels are complemented by two conceptual flows: *Direction* and *Emergence*:

- *Direction*, a downward flow, from the level of Expanded Society towards the technological level of autonomic systems. This flow expresses the effort of the Expanded Society stakeholders to shape the behavior according to their needs and desires.
- *Emergence*, an upward flow, from the autonomic systems at the technological level. This flow expresses the influence of the situated and evolving behavior of the autonomic systems towards the Expanded Society. We assume that autonomic systems will develop global and local patterns out of their many point-wise evolution directions, whereas it is these patterns that will flow to the Expanded Society and its members.

This conceptual structure (the introduced layers) is intended to work in an autonomic way, needing principles close to the cybernetic ones governing smart robots, for example:

- prevent ongoing evolutionary adaptation from drifting into unintended or harmful directions as well as to optimize the collaboration of an individual Autonomic System (Governance), and
- steer the evolution into a direction that leads to continuously available, usable, and sustainable systems answering the society needs (*Innovation*).

4. The Core Technical Principles

Following the introduced approach, this section discusses the technological objectives associated to each level of the approach. The following subsections provide a detailed overview of these objectives, starting with the integration of Autonomic Systems within the Expanded Society, followed by how to improve interoperability through semantic interactions, engineering direction through collaborative governance and emergence, and finalized by thoughts on how to create an autonomic process workflow and lifecycle methodology.

We need to advance the state of the art in composing and governing the behavior of (service-oriented) computational environments within which humans live and work, and in particular the associated communications ecosystems. To augment the previously stated objectives, we will need to develop a stratified methodology defining the dynamic, automated composition of workflows/processes from existing constituent services, or service bundles.

One approach is to employ reactive Finite State Machines (FSM) which offer stability through preconfigured static action plans that can be orchestrated dynamically to form complex FSM in the form of workflows. The advantage of the

FSM approach is the complete a-priori knowledge of the state the composed workflow process in relation to its current operational context.

Another, more proactive approach is to engineer goal-directed workflow/process compositions; directed both in terms of policy abstractions and technical constraint specification. Goals form conditions that an autonomous controller (e.g., autonomic manager) will attempt to achieve through the application of plans that in turn consist of sets of actions. Plans can be assembled dynamically and/or be made available through plan libraries.

The advantage of goal-oriented action planning over techniques such as FSM is the ability to provide more flexible and diverse options through the dynamic construction of plans from atomic actions, and indeed other plans, adapted specifically to meet stated goals. This approach involves views to define the environment, models to define the static characteristics and dynamic behavior of the environment, and contracts to exchange data between components in a standardized way. This encompasses the processes of analyzing business and system requirements, modeling the design, implementing a solution and then deploying the solution to constitute a lifecycle. Thus, our lifecycle methodology will codify this into a compelling and integrated design approach.

4.1. Integrating autonomic systems within the Expanded Society

We use the notion of the Expanded Society to understand the relationship between humans in their roles as operators, providers and consumers. In a second step, we apply it to a communication and service infrastructure that is formed from a collection of collaborating autonomic entities and systems. As well as being able to model the dynamics of these relationships at engineering time, it is necessary that these models can be processed by Autonomic Systems as part of their adaptive control loops. This requires the development of appropriate “machine processable” formalisms and tools for modeling the Expanded Society. Societal modeling extends and unifies a number of extant activities that are part of the engineering processes of systems we encounter today, including:

- Context awareness – as one of the key terms present in much of the literature concerning autonomic systems. It refers to the flow of information from a system’s environment into the system from which external conditions, such as usage context and user preferences/profiles, as well as business and technical goals and constraints are made available to the system at run-time. The meaningful comprehension of these conditions by the autonomic system provides the basis for dynamic adaptation of system behavior in response to changes in these conditions. A composed system should also have the ability to reflect on itself and to introspect the parts it is constructed of; this is referred to as “self-awareness”.
- Requirements specifications – typically form the first step in system development, defining the expected properties and functions of the deployed system. In one respect, requirements specifications define the relationship between the system user/operator and the system’s customers in technical terms (e.g., as the set of potential services that can be provided to the customer in terms of constraints such as reliability, performance, and controllability). This relationship is often

considered as static, but in autonomic systems it must be expected to be dynamic, for example by users that define new service configurations.

In our view, the Expanded Society modeling aims to extend these typically unidirectional, static or semi-dynamic relations with multidirectional dynamic interactions. Defined mechanisms will allow autonomic systems to actively perceive their own operational context, and to automatically adapt their internal models of the surrounding environment. This will result in a relatively continuous reflection and update of requirements, purposes, and goals assessed against outcomes. It furthermore creates the opportunity to use internal models to describe offered functions and services to other actors in order to negotiate potential communication patterns and modes of collaboration. Active creation of transient communities would then be based on the collaborative construction of a common model of interaction between the actors that are going to form the community.

4.2. Collaborative Governance and Emergence

The governance of an autonomic system refers to the management decisions exerted on the autonomic control mechanisms that monitor the state of the autonomic system, analyses its behaviors, plans modification of behavior in accordance with goal conditions, and activates the required functionality to execute the plan. Although autonomic systems aim to remove human activity from this loop, it is through governance that humans can impose their requirements and intent. Typically this is performed at a high level of abstraction by defining goals to be achieved by the autonomic system, such as maximizing the value of a utility function. Governance is exercised using policies that typically use rules-based languages to specify goals and constraints in a declarative form.

We consider not only the governance of Autonomic Systems but also the emergence of new properties from the interaction of local system components within systems of autonomic systems combined with the human organizations that use and govern them, i.e., the Expanded Society. We are already aware that complex interactions among autonomic systems and their components can give rise to emergent coordinated patterns of activities. Using directed governance we can therefore provoke and guide emergence within boundaries to prevent the emergence of counterproductive behaviors while seeking, and moreover, pursuing and promoting constructive behaviors.

We must thus augment our understanding of what it means to govern different socially integrated autonomic systems as each will offer different properties, emergent potential and operational requirements based on the underlying nature of the services it provides, the nature of the human organization services and the means used to govern that provisioning. Governance of systems of autonomic systems that change frequently in their composition and are unconstrained in the sourcing of constituents pose a major challenge in forming a cohesive view of the governance semantics available across the whole system. Semantic-based expressions of autonomic system governance semantics, i.e., semantic policy languages, will play a key role in forming this cohesive model rapidly and accurately enough to be of operational use to human users and governors.

4.3. Semantic Annotation

Pragmatically, a society is conditioned to the knowledge shared between members in terms of their ability to communicate, to experiment with new strategies, to adapt links, and engineer methods to collaborate toward achieving common objectives. This implies that entities composing the society must converge on a common understanding of shared concepts, and to a certain extent also shared experiences and history. In our work, ontology-based semantic models will be used to bridge the gap in understanding between human society and the autonomic systems with which they interact. The challenge here is to represent knowledge relating to a service, policy, or domain through ontologically consistent semantic annotation such that an autonomic system can actively interpret and act on perceived state. The ontological models and semantic handling must be integrated within the computational systems architecture so that ASs can perform introspective inference on available knowledge.

Adding semantics to rule and policy expressions will mean the human-authored governance of ASs can employ ontological references to ensure consistency and provide a means for more convincing and powerful policy-based autonomic control and guidance and feedback to the governing users.

In addition, the semantic annotation of service descriptions enables more meaningful specification of their goals and effects. For instance, a composition goal defined as an interwoven mesh of concepts permits a more precise query of suitable services than is possible with a simple taxonomic classification. Ontology languages designed for the Semantic Web services (SAWSDL, OWL-S, WSMO) are well suited to this type of composition model that may be dynamically altered by multiple stakeholders over time. Not only are these changes immediately available for reasoning-based machine processing in development tools and runtime systems, but also the changes themselves are amenable to reasoned analysis as part of an engineering toolset for developing the Expanded Society.

We are aiming to create a core ontology describing the core modes and purpose of interactions between the user society and autonomic computation in the Expanded Society. Furthermore, we will derive the full set of semantic annotations and technology grounding mappings required to use this shared ontology for provisioning dynamic, policy-governed, shared service compositions.

5. Developing a Methodology for Engineering Principles

The most important step towards realizing our vision is to define a framework for engineering principles. The fundamental idea here is to combine the societal and economic operational context of autonomic systems in order to provide appropriate engineering guidelines and to define the boundaries in which these systems evolve. That means, engineering principles cannot be phrased in simple technical terms (for example as an instance of IBM's MAPE approach) but need to be put into the broader context of changing requirements and ongoing adaptation. In particular, the following issues need to be addressed:

- Which mechanisms are necessary to capture contextual and situational conditions that are needed to determine and to perform autonomic reactions, and what are appropriate reactions? How to validate and assess the effectiveness and appropriateness of those reactions? How to make systems pro-active and anticipate change?
- Which mechanisms are needed to enable autonomic system evolution in a sustainable way? How can evolution be controlled? Is it possible to achieve homeostasis by employing principles of self-organization?

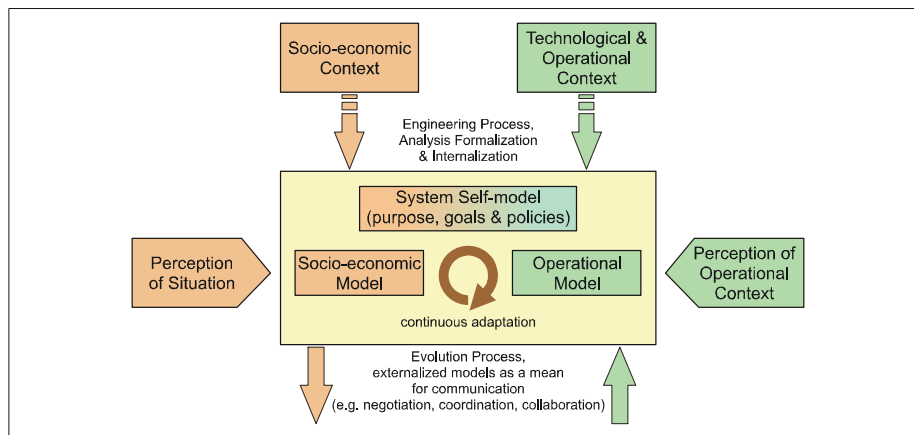


Fig. 2. High level engineering methodology

We have started assembling a comprehensive engineering methodology that departs from the analysis of the societal and economic context of an autonomic system, taking into account also the formative power of the deployed system, i.e. the creation of new communication and interaction patterns using this system, its interaction with systems that are already in place, and the emergence of new business opportunities and scenarios. This approach allows obtaining a more appropriate picture of the ongoing interaction and co-evolution of autonomic systems and the society using them.

New services do not emerge from the scratch, but appear in a landscape of already deployed and used technologies. Future systems will be based on seamless compositions of legacy systems and newly added components, services will be provided by orchestrations of highly dynamic processes. The definition and maintenance of proper trust relations will become one of the most important factors, as well as the ongoing validation of service level agreements and contractual defined requirements. The non-technical perspective developed in the first step will be extended by the incorporation of these (and others, more specific) technical issues.

Figure 2 exemplifies the key elements of the envisioned methodology starting from an analysis of the socio-economic, the technological and operational context of the system under design. Relevant factors of this complementary views will be formalized and thus be available not only as a description of engineering requirements, but also in the operational phase as a perceptive model frame to recognize situated and operational data (usage, collaborations, etc.). Complemented

by a system self-model defining the system purpose, restricting policies, or goals, these models are object of continuous updates in order to reflect changes in the environment of the system.

Models will have a paramount role also as instruments for communication (human-to-machine, or machine-to-machine). The exchange of models as part of internal “world-views” (formalized e.g. as ontologies) for various purposes (negotiation of collaboration contracts, invention of coordination patterns, establishment of trust/distrust relations, etc.) will become the driving mechanism of long-term evolution of systems and system interactions.

5.1. Methodology and Tools

Based on the conceptual view (Section 3 and Figure 1), we need to identify, specify, and build the methodology and tools to realize autonomic systems. It is necessary to experimentally apply the created artifacts to use cases to allow a quantitative and qualitative assessment of the principles as well as the developed methodology. In other words, the methodology (how to use) and the tools (engineering support by means of what to use) are always the joint outcome of a set of driving principles meeting the environmental constraints and boundary conditions of the Expanded Society.

The main focus of the methodology lies in combining different expertise, application domains, and business goals. The main focus of the tools lies in domain-specific standards, surrounding systems and application and vendor specific knowledge. This is done following the two main flows of information: Direction and Emergence (see Figure 1).

The overall insight about these two flows is that they form a dialectic pair, which means that they are in opposition to one another, but the result of their struggle is deeper than the balance of their respective intensities.

The precise nature of this relationship has to be investigated and further refined, but at least two examples already seem obvious to be mentioned here (the common trait being that both stem from a dialectic balance between *Governance*, being the function using the *Direction Flow*, and *Emergence*). The two examples are:

- Abstractly specified cooperation among Autonomic Systems: our approach strives to be decoupled from the internals of each collaborating AS. From the needs and forces in the Expanded Society however, some consequences that, while not uniquely specifying the cooperation, put some constraints on it. It will then be a Direction Flow that enforces these constraints, whereas an Emergence Flow will build up compliant collaborations starting from current configurations (available services, past history, known or learned evolution rules).
- Detection and reaction to threats: techniques of behavioral analysis would need to be applied in order to learn and evolve a characterization of abnormal or undesirable behavior, and to detect such behavior while it is being exhibited by an autonomic system. An Emergence flow will be fundamental to the learning and detection process, whereas the Direction flow will be the carrier of the explicit part of the definition of the normal and abnormal behavior.

Autonomic systems have to take societal and economic context of their operation into account. Following this principle, the concept of purpose-orientation seems to be bound to play a key role within the methodology. The concept of purpose expresses what a system aims for and the reasons and rationales which make up the context for such an aim. Two major concepts are used to meaningfully and effectively represent purpose:

- Goals - descriptions of desirable states of the system within its surroundings and potentially of this very surroundings itself.
- Policies - structured sets of rules that give guidance, i.e., purpose, to the achievement of goals. They are typically specified as behavioral constraints for a system, or system components, in terms of their relationship with the environment.

Both concepts deal with expressing purpose. We expect that policies and goals will be fruitfully combined in different ways, particularly within the context of the dialectic Governance/Emergence interplay. Some examples are:

1. Business Policies and Management Goals. In this approach to Governance, the needs and wants from the Expanded Society are directly expressed as high-level policies. These policies are then used to specify management goals in terms of desirable system properties.
2. Business Goals and Management Policies. Alternatively, the Direction Flow can first express the Expanded Society desires as a set of goals, and then match each one of these goals to suitable management policies. Like in the previous example, multiple context-dependent management actions and techniques can be used while keeping the high-level business and management purpose stable.
3. Emergent Feedback Loop at the Purpose Level. In a first step, it is investigated how Governance actions can be changed in order to respond to a newly detected Emergence Flow. The second step is to create a closed (feedback and control) loop that modifies the purpose level triggered by an incoming Emergence Flow, i.e. but adding a new business goal “seize the opportunity” when favorable emergent behavior is detected or add the business goal “counter the threat” on detection of a dangerous emergent behavior.

5.2. Architectural Considerations

Methodology and tools need to be provided by means of an overall system architecture, which should be delivered in the form of a technical specification (which in turn will be submitted to the ACF). The first step of developing this architecture is to identify the various stakeholders by means of roles (users, network operators, content providers, society at large) and associating them with the technical environment (devices, sensors, actuators, networks). This will be realized in cooperation with the requirements from the considered use cases (see next subsection).

The second step is to identify and specify the interfaces which consume and/or provide services to an AS. These AS interfaces need to be capable to handle interactions required to command and control (Governance) and monitor (Emergence) an AS, based on business goals (from the Expanded Society).

Based on step 1 and 2, all capabilities and functions may be defined that allow for a federated AS framework (system with systems or system of systems). This framework identifies and details now all aspects enabling the Governance and Emergence paradigm for autonomic systems. Data flows and traffic description will complete this step specifying interactions between and within autonomic systems (i.e. components of an autonomic system).

The resulting technical architecture will be described as an instantiation of its components (hardware and software) logically expressed within the federated AS framework. The structure of the technical architecture will be defined by a software platform, including distribution and organization of intelligence (i.e. federation), management modules (i.e. registries, mediators) and networks (i.e. topologies). Particular care should be taken to enable traceability from the end-user requirements up to the defined technical components.

5.3. Modeling Considerations

The self-organization characteristic of an AS impacts the way dynamic service composition and the underlying system can be managed. Thus, the architecture will apply the following major modeling considerations

- Model the capabilities and constraints associated with a service and its possible environmental contexts (structure, syntax and semantics).
- Model the interaction between services (invocation and protocols).
- Model both the interaction between users and services and people among each other.

To maximize the benefit of self-organized management, two different main activities need to be modeled and the above introduced modeling considerations need to be applied:

- Communication between AS within an AS community to allow dynamic service and process composition, including the formation of new communities providing emergent services following a distinct purpose and
- Management of AS and AS communities incorporating different management models.

Three fundamental aspects have been identified:

- Model service interactions and evolution of services by modeling user needs (Expanded Society) within a domain and model the integrated functionality these services can offer to provide a solution.
- Model the deployment and management of the interactions specified by the business goals (Expanded Society) and the accompanied work flows.
- Model the discovery of services and their behavior that perform the tasks required (for example cooperation between AS and update of process compositions within an AS or an AS community).

Evolution, Emergence and Governance components will support these issues and their dependencies.

5.4. Security and Trust Considerations

Security and trust is deliberately not separated from the overall methodology. Instead, the engineering principles drive intrinsic consideration of the design of secure and trustworthy autonomic systems collaborations focusing on the support of confidentiality, integrity, and availability. Specifically, they will provide security within its homeostatic models that allow handling the following situations:

- An AS or AS community may operate without prior knowledge of other AS or AS communities;
- An AS or AS community may become malicious towards other AS or AS communities

Behavioral analysis will be used to support security and trust. Even if the AS is not previously known, it is possible to determine if its behavior is normal or not. It will be especially interesting to establish a behavioral pattern of all AS already met. In order to do so, the usual and previous AS activities are learnt and an AS surveyor compares on-the-fly present behavior with previously known behaviors. This behavioral analysis is decomposed to five functional processes: learning, detection, upgrade, evaluation and qualification. Learning describes the communications' behavior with AS partners and models this behavior. Detection analyses incoming events and qualifies them (as deviant or not compared to previous learnt referent profiles). Upgrade constantly updates the initially learnt behavior. Evaluation classifies these behavior evolutions, these deviances as normal or suspect evolutions. Finally, qualification aims at enriching detected suspect deviances.

We are planning to apply a metrics approach consisting of evaluation of two properties: usual AS activities and usual AS system states. These metrics are relevant to evaluate AS behaviors according to the introduced security metrics: confidentiality, integrity and availability. When the metric evaluation has been learnt, thresholds are created and every time a deviance is observed, an alert is triggered and reported to analyst.

These five processes will then be used within the two functional flows of Governance and Emergence. Governance uses innovation to create relations between AS and governs the maintenance of them. The Governance provides the means to translate behavioral alerts into recommendations for the AS: it could mean exclude one AS from the community as we do not trust it anymore, it could mean recommend AS to change their security parameters to adapt to the situation. Emergence provides information on the change of an AS or an AS community. As part of Emergence, we find it necessary to provide means to survey the behavior of communities of Autonomic Systems and detect abnormal behaviors that could be considered as security issue(s) that affect the trust we put on individual AS. The objective is to be able then to take into account these alerts in order to drive the collaboration of AS and the business processes deployed by the evolution.

6. Conclusion

In this paper, we have introduced our approach on the development of engineering principles for Autonomic Systems. This approach was developed as a result of our experiences within a number of European research projects, workshop, the ACF as well as preparations for the European Framework Programme 7. We are proposing to treat the network of interacting people and communities (society) on one hand and autonomic systems on the other hand as a single dynamic system, which we introduce as the *Expanded Society*. Thus, the discussed observations and proposed concept within this paper encompass the co-evolution of human organizations and technological systems. As the title of this paper suggests, we intend to open a discussion on engineering principles. Consequently, this paper describes concepts and ideas, which we would hope will advance the field and allow us to agree on the fundamental objectives addressing the various challenges. We will follow on with our work and hope that others will join this discussion within the ACF.

References

- 1 A. Manzalini, F. Zambonelli: *Towards Autonomic and Situation-Aware Communication Services: the CASCADAS Vision*. in Proc. IEEE Workshop on Distributed Intelligent Systems (DIS2006), Prague, Czech Republic, June 15 - 16, 2006
- 2 William Donnelly, Radu Popescu-Zeletin, John Strassner, Brendan Jennings, Sven van der Meer (Eds.): *Modelling Autonomic Communications Environments*. Multicon lecture notes, Vol. 2, October, 2006
- 3 Nazim Agoulmine, Sasitharan Balasubramaniam, Dmitri Botvich, John Strassner, Elyes Lehtihet, William Donnelly: *Challenges for Autonomic Network Management*. in Proc. of 1st IEEE International Workshop on Modelling Autonomic Communications Environments (MACE 2006), Dublin, Ireland, 25-26 October 2006
- 4 Sven van der Meer, William Donnelly, John Strassner, Brendan Jennings, Mícheál Ó Foghlú: *Emerging Principles of Autonomic Network Management*. in Proc. of 1st IEEE International Workshop on Modelling Autonomic Communications Environments (MACE 2006), Dublin, Ireland, 25-26 October 2006
- 5 Edzard Höfig, Björn Wüst, Borbála Katalin Benkö, Antoinetta Mannella, Marco Mamei, Elisabetta Di Nitto: *On Concepts for Autonomic Communication Elements*. in Proc. of 1st IEEE International Workshop on Modelling Autonomic Communications Environments (MACE 2006), Dublin, Ireland, 25-26 October 2006
- 6 Mikhail Smirnov: *Autonomic Communication - Communication Paradigms for 2020*. FET Brainstorming meeting, Brussels, 22 July 2003
<<http://cordis.europa.eu/ist/fet/comms.htm>>
- 7 Dave Lewis, Declan O'Sullivan, Kevin Feeney, John Keeney, Ruaidhri Power: *Ontology-based Engineering for Self-Managing Communications*. in Proc. of 1st IEEE International Workshop on Modelling Autonomic Communications Environments (MACE 2006), Dublin, Ireland, 25-26 October 2006