Purpose-driven, Self-growing Networks – a framework for enabling cognition in systems of systems

Bernd Bochow and Marc Emmelmann
Fraunhofer FOKUS
Berlin, Germany
Email: {bochow, emmelmann}@ieee.org

Abstract— **Purpose-driven, self-growing networks are a novel concept using their service or geographical extend to augment network capacity or operational constrains such as energy consumption. This paper builds upon the purpose-driven, self-growing paradigm: it contributes a first step towards a framework allowing a formal description and performance evaluation of such novel network functionality. The paper hereby shows how purpose-driven, self-growing networks can be described by their network lifecycle. The latter is decomposed into transition points and associated parameters, rules for transition between points, and associated costs. By applying backward reasoning, the paper contributes a mathematical description of the lifecycle and presents first ideas how such representation can be used for network optimization.**

I. INTRODUCTION

Various heterogeneous technologies will constitute the future of wireless networks. Such networks will provide coverage from femto-cells up to wide are coverage, will utilize licensed and unlicensed frequencies, will explore available white spaces in the spectrum, will have to support mobile and fixed devices, and will integrate single hop communication (on the wireless link) in addition to multi-hop relaying. In surplus, networks may potentially serve different purposes ranging from providing low bandwidth environmental sensing services, up to high bandwidth multimedia communication services. This potpourri raises several challenges for future wireless networks:

- How to improve the efficiency and sustainability of those coexisting networks?
- How to handle the increased complexity in network operation and management arising with the variety of coexisting technologies and network components?

The novel concept of purpose-driven, self-growing networking [1], [2] addresses those challenges. A self-growing network coexists, collaborates or integrates—potentially in symbiosis—with collocated networks utilizing their service or geographical extend to augment network capacity, or operational constrains such as energy consumption [3]. The self-growing process including network operation and management is realized by focused cognitive decision making controlling network and node reconfiguration. Depending on the ability of its network components, a self-growing network can autonomously and on demand switch between dedicated, generally pre-defined purposes. Examples for such purposes are providing sensing functionality (i.e. obtaining temperature and humidity information, usage of a given radio spectrum, general data forwarding within wireless sensor networks) or providing high bandwidth, low latency user communication. A cognitive engine associates the change in purpose along with providing a new or additional purpose with cost and benefit, such as projected energy consumption and remaining network or node lifetime. This allows to decide in advance on a change in purpose under operational (e.g. energy consumption) constrains as well as on the optimization of network efficiency and coexistence needed in consequence. As in addition to (human) management interaction, switching can be triggered by external events. Such event can include the indication of new QoS needs by user devices, or distributed sensors reporting on the network status and energy consumption of devices. Appropriate rule sets located within the decision engine can substantially increase the efficiency and sustainability of the network while reducing operation and management overhead.

In summary, a self-growing network utilizes state-of-the-art concepts and enablers to realize its evolution. Such facilitators include node and network reconfigurability, cognitive decision-making, and self-learning capacity. At the same time, a self-growing network can respond to exceptional operating conditions by applying these concepts and enablers to define a temporary purpose satisfying demands arising from the exceptional situation. A self-growing (derived from progressing to maturity) network thus can be seen as a managed autonomous network guided (from educating or raising) in its evolution by bounding rules. It is assumed, and has to be proven in specific scenarios, that this approach is more suitable for low-profile, resource-limited node and network architectures compared to envisaged full autonomic network solutions.

Hence, the self-growing paradigm evolves the state of the art as it incorporates both cooperative and autonomic aspects.

The research leading to these results has received funding from the European Community's Seventh Framework Program (FP7/2007-2013) under grant agreement CONSERN no 257542.

Cooperative behavior and problem solving is critical in the initial phase of self-growing, that is the small-scale network, as well as in the evolution to a larger scale network, able to serve different purposes of larger systems. Autonomic behavior is critical in the self-growing process where it drives decisions wether to join or leave cooperations.

In contrast to existing approaches for autonomous, self-configuring, or self-managing networks [4]–[8], a self-growing network progresses along some rules of evolution along its lifecycle, following, for example, a predetermined progression from purposes requiring a lower level of complexity towards purposes requiring a higher level of complexity regarding reconfiguration and collaboration capacities. Thus, a self-growing network cannot freely evolve but is restricted towards an intended purpose and is driven by cost and benefit of a transition in purposes. Nevertheless, the degree of freedom to deviate from a planned lifecycle is a matter of the purpose of a self-growing network. In this scope, the optimal balance between the autonomic and cooperative paradigms may be different according to the purpose of the self-growing network. This will be reflected in the rules that govern the evolution of the network, favoring (and motivating) varying degrees of cooperation between the network elements.

This paper builds upon the concept of purpose-driven, self-growing networks and contributes a first approach on how to describe policies, rules, and decision making strategies.

II. PURPOSE-DRIVEN, SELF-GROWING NETWORKS

A. *Self-Growing via Cognition*

In technical terms, cognitive decision-making is generally understood as mimicking a human-like complex mental decision process. Implementations often rely on incremental and recursive reasoning and on inference processes closely associated with machine-learning strategies to refine rulesets and to resolve conflicts [9], [10]. Without restricting generality as well as potential implementation approaches, we here assume that decision-making for self-growing networks can be described by a set of rules that, when evaluated by one or more decision engines, can produce a choice from a set of potentially new purposes and ways to evolve. Within the architecture described, the decision engines realize the self-growing functionality of the network, while the rule sets available to decision-making describe the self-growing network's lifecycle. The latter may even be partitioned in parts describing characteristic evolvements of the system.

Giving one example, such (part of the) lifecycle may describe how a self-growing network evolves from a loosely coupled collection of independent networks that federate during the initial phase of operations forming a collaborative system by progressing through various states of integration. Sensor networks, femto- and macro-cels of different operators, and energy-constrained portable devices may join across an infrastructure network. As a result, the infrastructure network gains from utilizing portable devices and sensor network services as remote probes providing measurements allowing an optimization of the overall network's capacity or energy consumption.

B. *Decomposition of Lifecycle*

The network's lifecycle encompasses a self-determined or pre-planned path along a sequence of progression points that define (potentially temporary) stable points in the evolution of a self-growing network. Progression points can be associated with stable configurations of a network potentially providing different functionalities for a certain purpose of the network; the transition between them is described via rules. A lifecycle is defined as having one well-defined starting point and one or more potential end-points, as well as an arbitrary number of intermediate points, each of them defined by a progression point.

1) Progression Points & Attributes: Within the lifecycle, a progression point associates with a set of attributes. These attributes are described each by a non-empty set of parameters. A progression point is defined as being *measureable* if a set of metrics is made available for these parameters. An associated descriptive set of factors (i.e., values of parameters) then makes a progression point *well defined*. Since the transition from one progression point to the next along the lifecycle is measurable due to the change of parameters, it also implicitly describes the benefit (or cost) obtained from an evolutionary step as well as how well a network currently suits a given purpose.

Figure 1 shows a sample lifecycle comprising a sequence of progression points (A, B, C, D, G, H, I, K), transient progression points ($A.1, A.2$), and exceptional progression points ($E, E.1, E.2$).[1] Figure 2 illustrates for two progression and two transient points how the they are described by sets of parameters and their parameters.

In addition to linear progression, for a number of scenarios, a lifecycle may fork towards multiple potential target purposes. This is especially true for event-triggered progressions, where the type of the event determines the next purpose to enter (e.g., indication of limited battery power at mobile devices in conjunction with indicating an emergency situation). In this, the number of optional targets must be evaluated and all of them must be assessed and judged against the objective of the evolution. For example, a near term decision might have an impact on some long-term capacity, and the relevance of being able to realize a future configuration shall be judged against a short-term benefit. Certain progression points which are under "normal" operation not valid due to energy-constrained operation and limited battery power might become well feasible options for changes in network purposes due to an indicated emergency situation: short-term availability of "perfect communication" of all users might precedence the goal of achieving an extended network life-time.

---

[1] Note, that the lifecycle diagram as drawn in Figure 1 intentionally resembles a state diagram, emphasizing that a sequence in time of progression points also may be understood as a sequence of network configuration states. In this diagram, initial, final, and exceptional states, such as those representing transient progression points, are marked by double outlined circles.

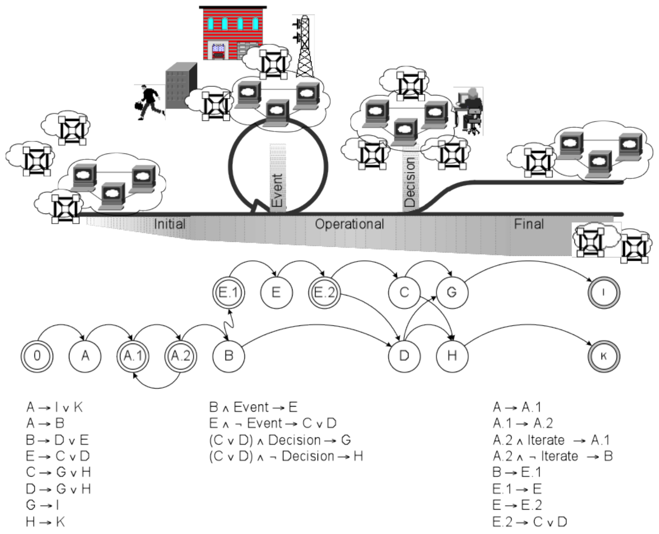

Fig. 1. Lifecycle example for a self-growing network.

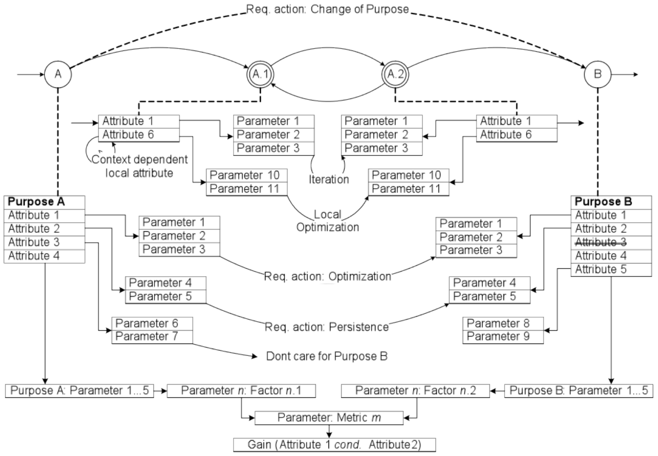

Fig. 2. Evolution between progression points of a self-growing network and associated benefit.

Such possible transitions are expressed by rules. The parameters associated with progression points may be used for deriving a metric which allows for such an evaluation showing if the network residing at a given progression point better suits its purpose (e.g. guaranteeing a minimal quality of service).

2) Rules Governing the Transition between Progression Points: How a self-growing network evolves between progression points shall be defined by suitable rules. Hence, the evolution of a self-growing network through its lifecycle can be described by a non-empty set of rules. Applying a rule may cause a change in attributes or parameters when commuting between progression points. The benefit of applying a certain rule is measurable, given that both the starting point and the endpoint of a transition between progression points are well-defined and are measurable. Rules might be static (e.g., known a-priori), volatile or dynamic (e.g., computed), or persistent (e.g., self-learned).

Backward reasoning is applied as a suitable mathematical inference method to describe the progression options of a self-growing network. Figure 3 provides an example for the top-level rules of the lifecycle example shown in Figure 2. Note that in an inference engine implementation, firing a rule in a timely manner implicitly requires an external trigger, or programmed request not shown here, except for transitional progression points which are assumed to be self-triggered.

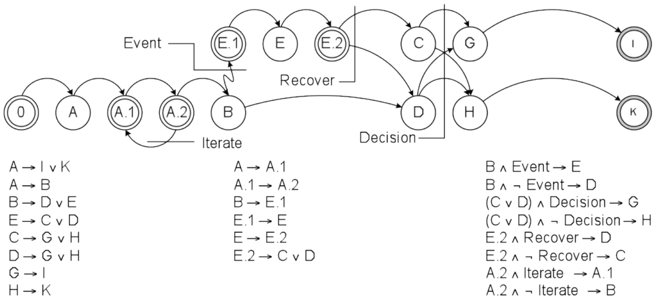

Fig. 3. Top-level rules for the lifecycle example given in Figure 1 highlighting decision points.

Accordingly, $B \wedge Event \rightarrow E$ here should be read as a conclusion: if B and Event then E. This implies that there exists a function that realizes the reconfiguration of the network into a configuration here described by E if it currently is in a configuration described by B, and an Event is detected and notified by some external function out of scope for this discussion.

Consequently, the cost (or price, or net utility, depending on the mathematical method to describe the effort required) to achieve this reconfiguration is herein denoted as $cost(B \wedge Event \rightarrow E)$ assuming some cost function that relies on externally defined metrics associated with parameters. This expression is synonymously used also as the cost of applying the corresponding rule, which allows comparing on the cost of certain alternatives in reconfiguring the network and may hence be used for optimizing the network configuration for a given purpose under cost constraints.

Figure 2 illustrates how attributes and parameters correspond between different well-defined progression points, and how the benefit of changing between purposes can be measured by factorizing parameters of an attribute set and applying common metrics to parameters.

C. Network Optimization

The sequence of progression points defines the lifecycle of a self-growing network and the set of rules defines how it evolves through this lifecycle. In that, a well-defined and measurable progression point might associate with a dedicated purpose of the network. This property of a self-growing network thus allows factorizing the transition between distinct purposes of the network. Accordingly, to compare the values of metrics associated with adjacent progression points provides a way to define and measure the cost or benefit of a transition between purposes. The system can be extended to support mapping and evaluation between attributes with different weights using properly designed rules. Fuzzy logic modeling can potentially facilitate this process since it is well suited for capturing complex non-boolean requirements. Metric comparison can also be interpreted as a way to evaluate a given rule set in terms of cost and benefit. Given that parameters may have multiple metrics, and given that a metric may apply to one or several parameters at a time, the approach is sufficiently flexible to enable the evaluation of the benefit of a certain network configuration at any time in the lifecycle of a self-growing network.

Clearly, the benefit of a purpose change cannot be determined for attributes that cannot be parameterized (e.g., Attribute 4 in Figure 2), if there exists no metric for a parameter, or if metrics exist but are not comparable between purposes. Simply speaking, a "before/after" comparison and a categorization in terms of "is more than" or "is less than" must be possible to measure the benefit. It is not necessary that parameters must have numerical factors in this. Measurability and comparability—potentially applying transformations to achieve comparability across different metrics—are sufficient, which can be achieved by an initial classification or fuzzification step.

It must be noted here that a progression point might be transient (i.e., is not a well-defined purpose). Under certain conditions the set of attributes describing a progression point cannot be associated with parameters (i.e., a progression point has a non-empty set of attributes and an empty set of parameters). Although such a transient progression point might be needed as an intermediate to commute between two well-defined progression points (e.g., to describe the transition through an unstable state with, potentially, zero-time to cross), it is not wise to consider them as a valid (temporary) network configuration. This is due to the problem of determining cost or benefit of approaching or leaving a transient progression point. In order to consider such transient points in the optimization process, it may be feasible to associate a local set of attributes with a transient progression point, or with a sequence of transient progression points, which are only meaningful within a local context (c.f. Figure 2 and progression points $A.1$ and $A.2$ for an example).

This can be useful to describe transient network states where optimization takes place aside the main scope (e.g., in the scope of a neighboring network), resulting in an optimization within the main scope as a benefit for a cooperating network. Hence, the cost or benefit of progressing from a purpose to a transient progression point or vice-versa might still not be ascertainable, or might be meaningful only in the local context of the transient progression point(s). But the cost or benefit for progressing across a transient phase can be determined (c.f. Figure 2 and Figure 3 with respect to the transitions between A and B as well as B to E and E to C or D).

On the other hand, a transient progression point can be used to commute between contexts. That is, the cost or benefit for crossing transient progression points cannot be determined (e.g., due to a lack of comparability), but entering and leaving the transient phase both might be meaningful as the cost or benefit of leaving an initial context and entering a new context, although this might be expressed in terms of different factors and metrics (and might require human interpretation).

This illustrates how the cost of recovering from a reconfiguration necessary to handle an event in the near term has an impact on the benefit of a planned target purpose in the long-term (possibly affecting even the reachability of one of the pre-planned progression points) and hence on the optimization process. Resources consumed during an event then may require a different decision to optimize for the target purpose of a self-growing network. A potential approach to revise the networks lifecycle in a suitable way accordingly could be (among others) used to evaluate the cost of rules applied (and to estimate the cost of rules that must be applied in the future to reach the target purpose) and to reach a new balance between cost and benefit of purposes on the path towards a target purpose. Given the lifecycle depicted in Figure 3, the cost of applying a rule is illustrated as $cost(A \rightarrow B)$ to attain a new purpose B inferred from a purpose A under some external triggers, facts and conditions.

Considering the examples given previously, we can formalize the additional cost of handling an event by the network:

$$\begin{aligned} cost(Event \wedge Recover) &= cost(E \wedge Recover \rightarrow D) + \\ &\quad +cost(B \wedge Event \rightarrow E) + \\ &\quad -cost(B \wedge \neg Event \rightarrow D) \end{aligned}$$

while the minimum cost to reach purpose G is

$$cost(B \wedge \neg Event \rightarrow D) + cost(D \wedge Decision \rightarrow G)$$

The cost of approaching G and recovering first from the event-driven reconfiguration is

$$cost(B \wedge Event \rightarrow E) + cost(E \wedge Recover \wedge Decision \rightarrow G)$$

Omitting complete recovery and approaching a matched purpose (that probably only recovers partially since omitting D and applying C instead, which according to the rules given in Figure 3 is the only way to reach G without applying Recover) is

$$cost(B \wedge Event \rightarrow E) + cost(E \wedge \neg Recover \wedge Decision \rightarrow G)$$

which might be more beneficial in the long run.

Applying the evaluation method discussed above will allow comparing all potential evolutions in terms of cost and benefit on the basis of distinct attribute changes between purposes. It illustrates that the cognitive decision engine for purpose-driven, self-growing networks can feasibly described by sketching the networks life-cycle, and that such a life-cycle can easily be decomposed using progression and transient points in combination with rules (and associated costs) for each transition.

III. SUMMARY

This paper describes how the concept of "self-growing networking" can be identified as a novel type of network composed of (heterogeneous) network nodes and sub-networks that can cooperate and utilize their reconfiguration capacity to optimize on-demand for a dedicated (temporary) purpose, also augmenting capacity by associating with additional nodes, networks, services and functions in that. In this sense, the self-growing can be considered as building on paradigms such as: autonomicity and self-x capabilities, and cooperation and collaboration. Self-growing thereby allows to integrate a potpourri of heterogeneous network nodes to dynamically adapt to changes in users' needs and purposes of the network while optimizing for cost factors such as energy consumption.

The paper has shown that the cognitive decision engine required for enabling self-growing networking can be expressed by describing the networks life-cycle. One of the key contributions is the decomposition of this lifecycle into progression points and rule sets defining the progression between them. In addition, the paper shows how backward reasoning can be applied for describing the transition (rules) between progression points and how those rules can be associated with a cost function. This actually provides an initial framework for the purpose specific optimization of a self-growing network over its lifetime. The intelligence represented by this set of rules and costs allows an automated optimization of the network's configuration and hence highly degrades the remaining effort for managing and operating the network.

As this work is the first step towards self-growing, purpose driven networking, upcoming work addresses the following key challenges: First, a classification of a minimal set of rules allowing network operation and growing satisfying users' and providers' needs is required. This should account for the possibility that the decision engine's cognition and self-learning approaches allow for dynamic evolvement of the ruleset over the network's lifetime. Second, approaches how to test (new) rule sets in a deployed network without influencing the existing behavior of distributed decision engines is a key challenge. Also, from a more practical point of view, a detailed specification of required functional elements within the network as well as protocols for exchanging rulesets and input parameters for the decision engine between network elements are required. Potential candidates for following a standard compliant approach, requiring the extension of existing protocols, are combinations of IPfix [11] and the IEEE 1900 protocol family of IEEE DYSPAN [12].

ACKNOWLEDGMENT

The authors would like to thank George Koudouridis from Huawei Technologies Sweden for many fruitful technical discussion his constructive feedback regarding this paper.

REFERENCES

[1] B. Bochow, M. Schuster, L. Thiem, and J. Tiemann, "A novel system paradigm for self growing wireless networks," 2010. [Online]. Available: http://publica.fraunhofer.de/documents/N-132320.html

[2] N. Alonistioti, A. Merentitis, M. Stamatelatos, E. Schulz, C. Zhou, G. Koudouridis, B. Bochow, M. Schuster, P. Demeester, P. Ballon, S. Delaere, M. Mueck, C. Drewes, L. V. der Perre, J. Declerck, T. Lewis, and I. Chochliouros, "Towards self-adaptable, scalable, dependable and energy efficient networks: The self-growing concept," in *UBICOMM 2010*, 2010.

[3] E. De Poorter, B. Latre, I. Moerman, and P. Demeester, "Symbiotic networks: Towards a new level of cooperation between wireless networks," *Wireless Personal Communications*, vol. 45, no. 4, pp. 479–495, Jun. 2008. [Online]. Available: http://dx.doi.org/10.1007/s11277-008-9490-5

[4] S. Milner, J. Llorca, and C. Davis, "Autonomous reconfiguration and control in directional mobile ad hoc networks," *Circuits and Systems Magazine, IEEE*, vol. 9, no. 2, pp. 10 –26, 2009.

[5] H. Hu, J. Zhang, X. Zheng, Y. Yang, and P. Wu, "Self-configuration and self-optimization for lte networks," *Communications Magazine, IEEE*, vol. 48, no. 2, pp. 94 –100, February 2010.

[6] S. Jeon and C. Ji, "Randomized and distributed self-configuration of wireless networks: Two-layer markov random fields and near-optimality," *Signal Processing, IEEE Transactions on*, vol. 58, no. 9, pp. 4859 –4870, 2010.

[7] H. He, Z. Zhu, and E. Makinen, "A neural network model to minimize the connected dominating set for self-configuration of wireless sensor networks," *Neural Networks, IEEE Transactions on*, vol. 20, no. 6, pp. 973 –982, 2009.

[8] D. Batista and N. da Fonseca, "A survey of self-adaptive grids," *Communications Magazine, IEEE*, vol. 48, no. 7, pp. 94 –100, 2010.

[9] D. Dubois and J.-L. Koning, "A decision engine based on rational aggregation of heuristic knowledge," *Decision Support Systems*, vol. 11, no. 4, pp. 337 – 361, 1994. [Online]. Available: http://www.sciencedirect.com/science/article/B6V8S-48MPR7Y-2C/2/483feba756c397276d4200e71ece2c7c

[10] T. Mitchell, "Machine learning," *Burr Ridge, IL: McGraw Hill*, 1997.

[11] RFC 5101, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information," January 2008. [Online]. Available: http://tools.ietf.org/html/rfc5101

[12] "IEEE DYSPAN Standards Committee Homepage." [Online]. Available: http://grouper.ieee.org/groups/scc41/