

Measuring and Analyzing Emerging Properties for Autonomic Collaboration Service Adaptation

Christoph Dorn, Hong-Linh Truong, and Schahram Dustdar*

Distributed Systems Group, Vienna University of Technology, Austria
{dorn,truong,dustdar}@infosys.tuwien.ac.at

Abstract. Dynamic collaboration environments in which team member utilize different pervasive collaboration services for their collaborative work pose many challenges for service adaptation. Given a team, the underlying collaboration services must fulfil the team's goal. Thus, it is not enough to adapt collaboration services to the context of an individual. One needs to understand the behavior of the team and the collaboration services in order to adapt these services. Though many research efforts aim at understanding team behavior at the human level, there is no such a framework that focuses on adapting collaboration services for teamwork.

In this paper, we introduce a set of novel metrics that characterizes emergent behavior of teams. We present a team analysis and adaptation framework (TAAF) which monitors diverse collaboration services, analyzes and provides relevant metrics for understanding dynamic teams and for continuous team and service adaptation. This paper also discusses how TAAF can be used to support self-management of collaboration services for collaborative teams.

1 Introduction

Recent advances in pervasive technologies have fostered the collaborative activities of knowledge workers across spatial, organizational, and professional boundaries [1,2]. Those activities are performed in a distributed and dynamic environment comprising of a variety of collaboration services used in different ways. In such an environment, pervasive collaboration services need to continuously adapt to the change of team context which is strongly dependent on the activities of team members. Existing autonomic adaptation approaches, however, concentrate on the adaptation of services to only the context of individuals [3,4]. Given a team of knowledge workers that utilize various collaboration services, a whole new level of complexity emerges when the adaptation needs to incorporate the behavior of the whole team. In this paper, we present a framework enabling adaptation of pervasive collaboration services based on a set of novel team metrics.

1.1 Motivation

Although, each member of a team uses collaboration services in a different way, the underlying services must fulfil the collaborative goal of the team. As a

* Part of this work was supported by the EU FP6 project inContext (IST-034718).

result, adapting collaboration services only to the context of an individual is not enough. We need to understand the behavior of the team in order to adapt these services. For example, understanding team execution phases will help reconfiguring service provisioning strategies: a pervasive document management service initially deployed for a small team should adapt its behavior when the team grows. Team resource usage patterns might reveal which services are relevant and should be selected for particular activities. This would significantly enhance current SOA-based approaches — such as in [3] — achieving self-adaptation. In short, if we need to support team-centric self-management of collaboration services, we must be able to understand the complex relations between the team and its utilization of services.

The complexity arising from dynamic teams operating in a heterogenous environment demands for a support framework to aid collaboration services adapting to emerging team behavior. To our best knowledge, there is no such a framework that focuses on adapting collaboration services for teamwork, though many research efforts aim at understanding team behavior at the human level [5]. Current scientific approaches to autonomic service adaptation focus on the system level and limited service consumer context [4,6]. Although research on context-aware systems [7] provides methodologies and frameworks to capture dynamic behavior, previous work consider merely the dynamics of individual humans. This lack of quantitative metrics and framework to provide data on emerging behavior motivates our work in this paper.

To tackle the above-mentioned issues, we apply the autonomic computing paradigm to the adaptation of collaboration services used in teamwork. Our ultimate goal is to develop a supporting software framework for *the adaptation of collaboration services for teamwork*. In our view, this requires a multi-disciplinary research effort where we need to combine research approaches from multiple domains, such as Computer Supported Cooperative Work (CSCW), context-awareness, autonomic computing, and SOA. However, providing such a framework for emergence-based autonomic adaptation is challenging. We need to characterize the behavior of team collaboration in terms of metrics that can be used by software, and to gather information from heterogeneous services and devices by means of monitoring. Based on that we can analyze the behavior of a team and its collaboration services to develop adaptation strategies.

1.2 Contributions

Our salient contributions of this paper are:

- A novel set of metrics characterizing emerging team behavior that can be used for service adaptation.
- An advanced set of analysis techniques for understanding emerging team behaviors in pervasive collaboration environments.
- The design and implementation of a novel framework for measuring and providing team metrics during runtime.

The work presented in this paper results in the *Team Analysis and Adaptation Framework (TAAF)*. To our best knowledge, it is the first attempt to

combine delivery of runtime characteristics and context awareness techniques to achieve emergence-based autonomic adaptation. This paper concentrates on the definition, monitoring and analysis of metrics for service adaptation. The subsequent steps of planning and execution to achieve adaptation inside the pervasive collaboration services remain outside the scope of this paper.

Paper Structure: The remainder of this paper is organized as follows. Section 2 discusses team properties and adaptation. Section 3 introduces a set of novel team metrics. The design and implementation of TAAF is detailed in Section 4. We present experiments illustrating TAAF in Section 5. Related work is given in Section 6, followed by the conclusion and future work in Section 7.

2 Team Properties and Adaptation

Our objective is to support autonomic collaboration services. Figure 1 depicts a dynamic collaboration environment in which teams utilize different collaboration services for their collaborative work. Given the complexity of dynamic collaborations among team members, shown in the upper part of Figure 1, pervasive collaboration services (meeting scheduling, notification, document repository, etc.), shown in the lower part of Figure 1, should be adaptive.

Our approach is to define metrics characterizing teams to understand the dynamics services are confronted with. Based on that, we monitor team behavior, capturing the required data to determine and manage team metrics. This is the first step in the autonomic cycle focusing on observable complex relations between team members and their environment. We deliberately do not consider cognitive or psychological properties of the individual team members, as these properties do not constitute emerging team properties. Then, we analyze snapshot metrics and time-series to detect situation requiring adaptation. Threshold analysis, team lifetime phase detection, team comparison, and metric correlation are some supported techniques.

We have to consider several properties of the dynamic collaboration environment to understand how emerging team behavior affects service adaptation.

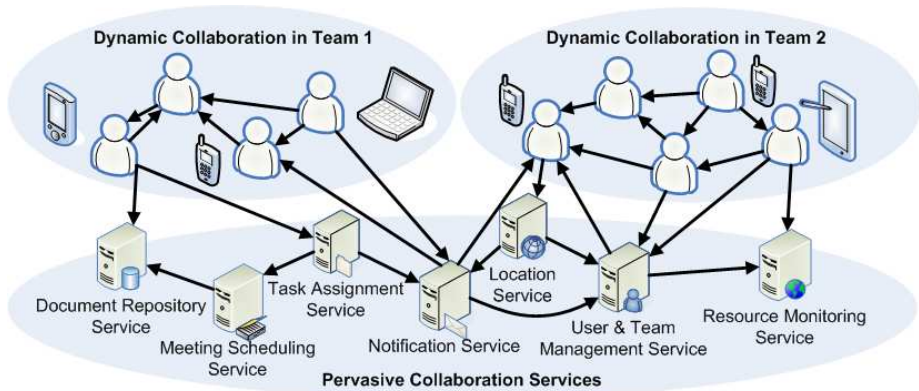


Fig. 1. Dynamic collaboration environment

Location. In dynamic collaboration environments, team members are distributed and mobile. They need to access services and communicate independent of their respective position. Furthermore, services provided by the team members' organizations are equally distributed. Example self adaptation strategies benefiting from analyzing the complete set of member movements are content distribution or service replication algorithms.

Organization. Team members originate from various organizations. Monitoring organizational structures and roles as well as dependencies across the whole team will enable the evaluation of the impact of each organization on the services and resources available to the team. For example, scheduling or communication services equipped with self-adaptive behavior can limit the team's dependency on a single organization.

Coordination. Members are coordinated based on goals, tasks, assignments, project-related roles, and skills. Analyzing individual properties is insufficient to understand their relevance in the overall team context. Collaboration services - such as a meeting scheduling service - that monitor coordinative metrics at a team level can self-adapt to select the most relevant meeting participants.

Interaction. Effective communication between distributed team members is vital. Analyzing scope (two members vs. the whole team) or type (synchronous vs. asynchronous) allows selecting the most suitable communication services - email, instant messaging, virtual conferences, mailing lists, or blogs. For example, as the team evolves, monitoring emerging interactions allows a communication recommendation service to continue recommending the most suitable form of communication.

Resources. Team members access a vast number of resources - distributed across a pervasive environment - from a multitude of devices. Rather than analyzing individual resource statistics, focusing on combined resource utilization at team level provides significant potential for adaptation. Example applications are resource monitoring services prioritizing the availability and reliability of the most vital resources.

By studying these properties, we develop and quantify metrics associated with teams. Specifically, we focus on metrics that describe emerging properties, arising from the relation between team members.

Terminologies and Notations: In the scope of this paper, *teamwork* is any work performed by a team to achieve a goal (defined by a Project). A *Team* consists of a set of members (more than one) engaging in teamwork, each *Member* being a human resource. Different members belong to different organizations, while a person can be a member of multiple teams. An *Organization* is a (legal) entity which defines the professional/employment background from which members engage in team work. Organizations can range in size from a single person, a dozen of people, to thousands of people. A *Project* consists of a goal to achieve, work steps specified to a certain degree of formalization, and constraints for achieving the goal. Teamwork consists of a set of *Activities* that describe the work actually performed by members to complete the project. By definition,

Table 1. Notations

Notation	Description
$team_i$	team i
m_i	member i
$size(team_i)$	the number of members assigned to $team_i$
a_i	activity i
$A(team_i)$	the number of activities executed by $team_i$
c_i	interaction i , a subclass of activity
r_i	resource i
org_i	organization i assigning members to a team
$ORG(team_i)$	the list of organizations involved in $team_i$
$l(m_i)$	location of member i
$ X $	number of elements in list or set X

Interactions are a subclass of activities, having multiple members involved, utilizing resources of type communication service, and being short-lived. A *Resource* is any computing, information, or communication service in pervasive environments that is used by team members in order to fulfil teamwork. Team members use resources to communicate, collaborate, and coordinate teamwork. Table 1 presents notations used in this paper.

3 Team Metrics

From the analysis of team properties we have developed a set of metrics characterizing relations between team members and collaboration services. Table 2 lists main metrics.

Team Location Entropy, $TLE(team_k)$, describes whether $team_k$ members' movements result in spatial clusters of workers or not, by determining the probability of all members being colocated.

$$TLE(team_k) = \frac{\sum_{i=1}^n \binom{l_i(team_k)}{2}}{\binom{size(team_k)}{2}} \tag{1}$$

where $l_i(team_k)$ is the number of members in $team_k$ at location l_i .

Table 2. Overview of main team metrics

Properties	Metrics
Location	Team Location Entropy $TLE(team_k)$ Team Mobility Entropy $TME(team_k)$
Organization	Organization Harmonic Mean $OM_h(team_k)$ Organization Arithmetic Mean $OM_a(team_k)$ Organization Membership Stability $OMS(team_k)$
Coordination	Team Size $size(team_i)$ Team Stability $TS(team_k)$ Activity Participation Harmonic Mean $AP_h(team_k)$ Activity Participation Arithmetic Mean $AP_a(team_k)$
Interaction	Interaction Participation Harmonic Mean $IP_h(team_k)$ Interaction Participation Arithmetic Mean $IP_a(team_k)$
Resource	Resource Access Harmonic Mean, $RA_h(team_k)$ Resource Access Arithmetic Mean, $RA_a(team_k)$ Resource Access Distribution, $RAD(team_k)$ Resource Utilization, $RU(team_k)$

Team Mobility Entropy, $TME(team_k)$, describes whether team members of $team_k$ relocate jointly or individually by determining the probability of the whole team being colocated before and after relocation.

$$TME(team_k) = \frac{\sum_{i=1}^n \sum_{j=1}^n \text{reloc}_{i,j}(team_k)}{\text{mob}(team_k)} \quad \forall i \neq j \quad (2)$$

where $\text{reloc}_{i,j}(team_k)$ determines the number of members in $team_k$ that have relocating from location l_i to location l_j and $\text{mob}(team_k)$ computes the total number of members in $team_k$ that have moved. Thus, members remaining at their location are not taken into consideration for calculating the TME.

Colocation and joint movements reveal tight interdependencies between members. TME and TLE specifically focus on the spatial relations between team members thus indicate the presence of similar needs and contexts. At the same time colocation and co-mobility reflect the complexity of providing communication and collaboration services. The higher TME and TLE values are, the less effort is required.

The *Organization Harmonic Mean*, $OM_h(team_k)$, is the harmonic mean of member count per organization within $team_k$ and is defined as:

$$OM_h(team_k) = \frac{\text{size}(team_k)}{\sum_{i=1}^{|\text{ORG}(team_k)|} \frac{1}{|\text{org}_i(team_k)|}} \quad \forall \text{org}_i \in \text{ORG}(team_k) \quad (3)$$

Organization Membership Stability, $OMS(team_k)$, is derived by observing changes in the number of participating organizations. Each joining or leaving organization, determined by function $\text{changeOrg}(\text{org}_i)$, results in a change of value 1.

$$OMS(team_k) = \frac{\sum_{i=1}^{|\text{ORG}(team_k)|} (\text{changeOrg}(\text{org}_i))}{|\text{ORG}(team_k)|} \quad (4)$$

Organization-related metrics provide an indicator of effort to provide services in a uniform manner. Multiple organizations in the same project may have, e.g., conflicting data representations forms or incompatible security policies.

Having an entire organization join or leave or having an unequal distribution of members across organizations has significant impact on the team's performance due to complex coordination and resource provisioning challenges.

*Team Stability*¹, $TS(team_k)$, is derived by observing changes in the number of team members. The sum of joining members and leaving members is determined by $\text{change}M_{\text{joint}}(team_k)$, respectively $\text{change}M_{\text{left}}(team_k)$.

$$TS(team_k) = \begin{cases} \frac{\text{change}M_{\text{joint}}(team_k)}{\text{size}(team_k)} & \text{if } \text{change}M_{\text{joint}} > \text{change}M_{\text{left}} \\ \frac{\text{change}M_{\text{joint}}(team_k)}{\text{size}(team_k) + \text{change}M_{\text{left}}} & \text{if } \text{change}M_{\text{joint}} = \text{change}M_{\text{left}} \\ \frac{\text{change}M_{\text{left}}(team_k)}{\text{size}(team_k) + \text{change}M_{\text{left}}(team_k)} & \text{if } \text{change}M_{\text{joint}} < \text{change}M_{\text{left}} \end{cases}$$

¹ Answers.com defines team stability as “the degree to which the membership of a team remains the same. Team stability can be defined in terms of length of time that the team members remain together”.

Team stability reflects membership dynamics within a team and provides, together with team size, insightful information for determining suitable resource allocation strategies.

Activity Participation Harmonic Mean, $AP_h(team_k)$, specifies the harmonic mean over all activity involvements and is defined by:

$$AP_h(team_k) = \frac{|A_k|}{\text{size}(team_k) * \sum_{i=1}^{|A_k|} \frac{1}{inv(a_i)}} \quad \forall a_i \in team_k \quad (5)$$

where $inv(a_i)$ returns the number of members involved in activity a_i .

Interaction Participation Harmonic Mean, $IP_h(team_k)$, specifies the harmonic mean over the cardinality of all interactions and is defined by:

$$IP_h(team_k) = \frac{|C_k|}{\text{size}(team_k) * \sum_{i=1}^{|C_k|} \frac{1}{card(c_i)}} \quad \forall c_i \in team_k \quad (6)$$

where $card(c_i)$ returns the cardinality (number of participants) of interaction c_i .

Whether interactions tend to include the whole team or just a small subset of members reflects the scope of required interoperability between the employed collaboration and communication services. The same property determines to which extent self-adaptation algorithms will affect the overall team.

Resource Access Harmonic Mean, $RA_h(team_k)$, is the harmonic mean of resource access by members of $team_k$.

$$RA_h(team_k) = \frac{|R_k|}{\sum_{i=1}^{|R_k|} \frac{1}{use(r_i)}} \quad \forall r_i \in team_k \quad (7)$$

where $use(r_k)$ returns the amount of times resource r_k is used within $team_k$.

Resource Access Distribution, $RAD(team_k)$, is the average reuse indication how often resources are accessed by different members on average in $team_k$.

$$RAD(team_k) = \frac{DRA(team_k)}{|R_k|} \quad (8)$$

where the *Distributed Resource Access* $DRA(team_k)$ sums up the total amount of times resources R_k are accessed within $team_k$ by different members m_i . The *Distinct Resource Access*, $DRA(team_k)$, is defined as:

$$DRA(team_k) = \sum_{i=1}^{|R_k|} use_{set}(r_i) \quad \forall r_i \in team_k \quad (9)$$

where $use_{set}(r_i)$ denotes the count of distinct members having used resource r_i .

In contrast to Resource Access, *Resource Utilization*, $RU(team_k)$ is the reuse indication how long the average resource has been reused.

$$RU(team_k) = \frac{\sum_{s=1}^{|r|} t(r_s)}{t * |r|} \quad (10)$$

where $t(r)$ indicates the duration in which resource r_s is used within interval t . Resource Access Distribution enables identification of the commonly used resources and services. Subsequently this metric facilitates focusing adaptation efforts on these significant team resources.

The above discussed metrics provide insight into the dynamic properties of teams. These metrics characterize emerging behavior arising from indirect and direct interaction, activity, communication, and resources in teams. While these metrics provide only a snapshot of the team’s status at a specific time, temporal analysis of these metrics can detect the effects of preceding adaptation efforts.

4 The Team Analysis and Adaptation Framework

Figure 2 describes TAAF (Team Analysis and Adaptation Framework) which consists of middleware services and tools for monitoring and analyzing team metrics at runtime and utilizing metrics for service adaptations. The *Event Collection* gathers monitoring data related to teams from different collaboration services. Events will be pre-processed to extract the main relevant information which is the input for *Metric Calculation*. Metrics associated with teams are determined during runtime and the resulting metrics are stored into the *Team Data Store*. Based on that, *Metric Provisioning* provides, during runtime, metrics to *Team Analysis* tools and *Service Adaptation* components which require the metrics for adapting collaboration services.

4.1 Monitoring Team Behavior

Our work is focused on analyzing and managing the metrics reflecting properties and changes of teamwork. Thus, we have to cope with the complexity of diverse sources providing data required for team analysis. These data sources are collaboration services deployed on dynamic, heterogeneous hosts in a pervasive environment. In our work, these services are assumed to interact with TAAF




Fig. 2. TAAF architecture and data flow

Table 3. Examples of Event sources

Event Source	Deployment	Event type
Notification Service	Static	Interaction
User & Team Management service	Static	MembershipChangeEvents
Calendar Service	Static	ResourceAccessEvent
Activity & Project Service	Static	ActivityActionEvent
Context Provisioning Service	Mobile/Static	(Re)LocationEvent, InteractionEvent, ResourceAccessEvent
Document Repository Service	Mobile/Static	ResourceAccessEvent
Position Service	Mobile	Location

Table 4. Examples of events encapsulating team data

Events	Description	Frequency
MembershipChange	contains the list of user having joined or left the team	Weekly/Monthly
ResourceAccess	provide details on which member accesses which resource, optionally stating the duration of utilization	Hourly
ActivityAction	inform about members engaging in an activity	Hourly
Interaction	contain tuples of members communicating with each other	Minute/Hourly
Location	hold the current location of members	Minute/Hourly
Relocation	inform about the movement of members from one location to another location	Minute/Hourly

via a well-known interface. Such an assumption can be achieved via Web service standards. To obtain the data, TAAF relies on WS-Notification which is widely supported in pervasive environments, including constraint devices [8].

To describe data that TAAF can process, we have specified an XML schema for describing generic collaboration metadata and specific collaboration data. This schema allows the exchange of data relevant to teams by representing various types of data, such as a team identifier, an event identifier, URI of collaborative services, timestamp, and location information.

TAAF obtains relevant data from collaboration services based on push and publish/subscribe event delivery. Thus, TAAF can support various services with different capabilities. Dedicated collaboration services such as the *User & Team Management service* or the *Notification service* integrate collaboration sensors and provide a subscribe/notification interface. However, personal and highly dynamic sources, such as services on smartphones or PDAs equipped with location sensors, require the user to manually configure TAAF as a notification endpoint. This eliminates the challenge to locate and access volatile sensors for subscription and enables users to protect their privacy. Table 3 lists exemplary collaboration services from which TAAF retrieves relevant data whereas Table 4 presents examples of events provided by these collaboration services.

When the *Pre-Processing* component receives an event from the *Event Collection*, it verifies threshold values for event confidence and timestamp. Subsequently, the overall team structure is updated because TAAF needs to keep track of the overall team structure. The underlying team status is updated only once as a single event may be input to multiple metric calculations.

4.2 Analysis and Management of Team Metrics

Team metrics are determined during runtime. To achieve flexibility, metric calculation is performed within multiple subcomponents that register with the

Pre-Processing component; each subcomponent retrieves pre-processed data according to its subscription. However, in TAAF, tightly connected metrics — such as Team Location Entropy and Team Mobility Entropy — are determined in a joint fashion in order to improve performance. In addition, which metrics should be computed and the schedule of the computation can be defined in advance. Depending on configuration, metrics can be saved in the *Team Data Store*.

To facilitate the exchange of team metrics, we have defined an XML schema for representing metrics associated with teams. This representation can be used to describe various types of information such as metric identifier, and current and previous metric values. Listing 1 gives an excerpt of the metric XML schema.

```

1 <xs:complexType name="tMetric">
2   <xs:sequence>
3     <xs:element name="URI" type="xs:string" minOccurs="1" maxOccurs="1"/>
4     <xs:element name="CurrentValue" type="tValue" minOccurs="1" maxOccurs="1"/>
5     <xs:element name="History" type="tHistory" minOccurs="0" maxOccurs="1"/>
6   </xs:sequence>
7   <xs:attribute name="enabledHistory" type="xs:boolean"/>
8 </xs:complexType>
9 <xs:complexType name="tTeamSizeValue">
10  <xs:complexContent>
11    <xs:extension base="tValue">
12      <xs:sequence>
13        <xs:element name="Value" type="xs:positiveInteger"/>
14      </xs:sequence>
15    </xs:extension>
16  </xs:complexContent>
17 </xs:complexType>

```

Listing 1. Excerpt of metric specification

Metrics provide the fundamental data required for understanding teams, detecting correlations, and ultimately taking action to counter steer negative tendencies or amplify positive effects. Based on team metrics, we have developed various team analysis features which have been incorporated into a team analysis GUI that allows any user understanding the team metrics and their relevance to subsequent adaptation actions. Self-adaptive collaboration services subscribe at the *Team Analysis* component in order to receive notifications when metric analysis detects relevant metric values.

Threshold Analysis: detects metrics violating a predefined condition over a period of time. This analysis is used together with a notification mechanism to enable runtime reaction in critical situations.

Team Phase Analysis: evaluates general trends in a metric’s timeline that indicate several phases, such as project kick-off, execution, and completion phases. Duration and structure of phases provides insightful information for autonomic services making decisions on whether additional or available members and resources should be deployed or reduced.

Multi-team Analysis: compares metric timelines of different teams. With this analysis, we are able to observe teams over time and detect emerging differences. Similar team configurations, such as size and member distribution, can lead to

significantly different emergent behavior. Comparing the structure of two teams reveals how the same adaptation decisions — such as deploying or reducing resources — result in different outcomes. In addition, similar patterns in different teams can indicate the occurrence of team transformations.

Correlation Analysis: reveals correlation among multiple metrics. This analyzes relations between metrics, giving more meaning to individual metrics.

4.3 Prototype Implementation

We are currently implementing our framework based on Java. The *Team Data Store* is based on the eXist database². The following collaboration services are currently being integrated with TAAF: *User & Team Management Service*, *Context Provisioning Service*, *Calendar Service*, *Document Repository Service*, *Notification Service*, and *Activity & Project Service*.

The above-mentioned services are part of the Pervasive Collaboration Service Architecture (PCSA) deployed at multiple sites across Europe, including Vienna, Milan, Genoa, and Aachen, within the inContext³ project. Inside TAAF, we use OpenJMS⁴ to pass events between components. In addition to the existing message header information provided by JMS, we provide extended header fields for storing information on event type, team identifier, activity, user, and source thus enabling efficient intra-framework event selection. Of the metrics described in Section 3 we implemented all except Resource Utilization (RU). We clustered related metrics together such as arithmetic and harmonic mean, or size and stability. For visualizing metrics, we utilize the JFreeChart framework⁵. The user can select the number of desired metrics to be displayed at the same time (see Figure 3). The current prototype uses JMS queues for delivering metric updates. However, we are going to support this kind of update via WS-Notification.

5 Experiments

5.1 Testbed

Section 4.3 introduced the hosting environment for the pervasive collaboration services required for running our experiments; these services are based on the inContext's PCSA. The PCSA is currently being used for project developments only, therefore, we have not been able to obtain enough live data for our experiments. As our main goal is demonstrating how to exploit emerging behavior for autonomic adaptation, we simulate the emerging team behavior arising from a dynamic collaboration environment as depicted in Figure 1 (upper part). We implemented a team simulation based on the concepts introduced in Section 2 to achieve various emerging behavior.

² <http://exist.sourceforge.net/>

³ <http://www.in-context.eu/>

⁴ <http://openjms.sourceforge.net/>

⁵ <http://www.jfree.org/jfreechart/>

To simulate the team behavior, we adapted the model by Barabási and Albert [5] to create a scale free, directed, acyclical graph (DAG) of interdependent activities which are managed by the *Activity & Project Service*. In this DAG, the vertices represent activities and the edges represent the dependency between two activities. Each activity is associated with the following properties:

- Duration: indicates the amount of time required to complete this activity
- Location: indicates the location at which the activity is performed.
- Cost: specifies the cost associated with an activity.
- Priority: specifies the priority of an activity
- Activity status: is either *pending*, *available*, *work in progress*, or *completed*.

The *User & Team Management Service* is then enabled to assign each team member to an organization which provides a set of resources. Initially these resources are available only to members of that organization. *Calendar Service*, *Document Repository Service* are providing resources in the form of calendar entries and documents, respectively. During collaboration, these resources are shared between interacting members. The *Notification Service* provides communication in the form of instant messaging, SMS, and email. Organizations assign new members to a team or withdraw active members from the team. In addition, each member is able to spend a certain amount of time on an activity. Finally, the *Context Provisioning Service* provides details on member mobility.

The data generator then simulates the invocation of our pervasive collaboration services. These collaboration services in turn deliver the actual events. When the simulated project begins, each member selects an *available* activity to work on, that is any activity which has all previous activities *completed*. An activity is *completed* once members have jointly spent enough time/effort to cover the activity's duration.

In each simulation round, we receive the set of collaboration events. Each member's selection results in an activity and location event. Additionally, for members engaging in the same activity at the same time, we receive an interaction event. Finally, interacting members utilize a subset of resources from their combined pool of resources, while members without interaction select resources only from their organization's resource pool. In both cases, services fire respective resource access events.

5.2 Examples for Emergence-Based Adaptation

For the adaptation example in this section, we created an activity graph of 200 nodes with activities spread across 10 different locations. The simulated team consists of 30 members from 4 organizations each providing 5 resources. Figure 3 presents an excerpt of the team analysis GUI visualizing the Activity Participation metric (Left) and Interaction Participation metric (Right) over the team's lifetime. Each graph includes harmonic and arithmetic mean. The meaning of the values are the same for both metrics: a value close to 1 indicates that (almost) all members participate in an activity, respectively an interaction, while a value close to 0 denotes a lack of collaboration as members work mostly alone on different assignments.




Fig. 3. Team Analysis GUI excerpt: The left graph provides Activity participation metrics while the right graph displays Interaction participation metrics for a team of 30 members working on 200 activities

Threshold analysis is a basic, albeit very useful technique enabling self-adaptive behavior. An exemplary project escalation service can utilize the subscription request in Listing 2 to receive alerts when the Activity Participation metric falls below 0.33. The simulated team crosses this threshold at the end of the kick-off phase (Figure 3 Left).

```

1 <subscription>
2   <type>
3     http://www.vitalab.tuwien.ac.at/projects/taaf/threshold_lowerbound
4   </type>
5   <teamuri>
6     http://www.vitalab.tuwien.ac.at/projects/taaf/teams#demoteam1
7   </teamuri>
8   <metricuri>
9     http://www.vitalab.tuwien.ac.at/projects/taaf#ActivityParticipationHMean
10  </metricuri>
11  <threshold>0.33</threshold>
12  <notificationendpoint>
13    ... [WS-Addressing Endpoint Reference] ...
14  </notificationendpoint>
15 </subscription>

```

Listing 2. Subscribing to threshold analysis

Other examples of potential self-adaptive behavior are:

Threshold Analysis: an autonomic content distribution service can decide to spawn extra distribution nodes when a team features decreasing team location entropy, or reduce the number of nodes when the team becomes more collocated.

Team Phase Analysis: a task scheduling service can ignore team instability at the beginning of a project, but starts to assign backup workers on critical tasks when the team remains instable during its execution phase.

Multi-team Analysis: a recommendation service can compare the effect of selecting the same communication service in different teams to adapt its selection strategy.

Correlation Analysis: Let us assume an inverse correlation of team location entropy and team interaction coverage. In this case, a meeting scheduler can suggest members from all locations to participate in a physical meeting.

6 Related Work

In our previous work, we analyzed teams with respect to impact on service requirements, resulting in a set of team forms and views [9]. Metrics associated with teams, however, have not been defined and quantified.

Scientists have invested great effort in providing concepts and tools for team based adaptation in the scope of context-aware devices and services. Vieira et al. [10] include interaction and organization aspects in their context ontologies but neglect emerging properties. Sterritt et al. [11] make the case for behavioral knowledge from which to compute metrics, but they remain at a general activity-focused level, not considering other teamwork aspects. Work on context gathering prior to these efforts generally focus on individual context neglecting team context altogether.

Current generic autonomic techniques and toolkits such as [4], [6], or [12] do not monitor the context of individual users, respectively limit monitoring to independent user properties such as location or device. De Wolf and Holvoet point out the potential of emergence for autonomic behavior [13] and also discuss the concept of emergence for engineering self-organizing systems [14]. They maintain, however, a pure system-centric view, applying emergent properties only to the autonomic system. In contrast, Bird et al. [15] apply email mining to discover emerging interaction patterns between users, but other major team properties are left aside. In a similar attempt, Valverde and Solé [16] investigate emerging self-organization in large open source social networks based on email repositories. However, such communities feature different characteristics compared to teams. TAAF specifically collects data about emerging properties from a wide range of sources and thus delivers more reliable and expressive data.

TAAF differs from the above-mentioned work in many aspects as it explores emergent team properties for the self-adaptive collaboration services.

7 Conclusion and Future Work

Understanding and detecting emerging behavior, patterns, and transformations in teams ultimately enables team-centric self-adaptation of collaboration services. In this paper, we tackled issues related to team metrics, since runtime information on emerging team properties and team transformations is the key to service adaptation for pervasive collaboration environments. This has not been well addressed until now. We have presented a novel set of team metrics and described TAAF which is a framework for analyzing, managing and providing team metrics for service adaptation during runtime. TAAF can uncover associations between various metrics, notify collaboration services when thresholds are reached, visualize team life-time phases, and compare multiple teams, thus providing necessary features for achieving autonomic collaboration services.

Our future work includes the further development of metric monitoring and analysis parts of TAAF. Furthermore, we will concentrate our work on advanced service adaptation techniques for teamwork in pervasive environments.

References

1. Ferscha, A., Holzmann, C., Oppl, S.: Context awareness for group interaction support. *Mobility Management & Wireless Access Protocols*, 88–97 (2004)
2. Patterson, D.J., Ding, X., Noack, N.: Location by, for, and of crowds. In: *Proceedings of International Workshop on Location- and Context-Awareness (LoCA)*, pp. 186–203 (2006)
3. Fournier, D., Mokhtar, S.B., Georgantas, N., Issarny, V.: Towards ad hoc contextual services for pervasive computing. In: *MW4SOC 2006: Proceedings of the 1st workshop on Middleware for Service Oriented Computing (MW4SOC 2006)*, pp. 36–41. ACM Press, New York (2006)
4. Sterritt, R., Smyth, B., Bradley, M.: Pact: personal autonomic computing tools. In: *EASe Workshop at ECBS 2005*, pp. 519–527 (2005)
5. Barabasi, A., Albert, R.: Emergence of scaling in random networks. *Science* 286, 509–512 (1999)
6. Bigus, J.P., Schlosnagle, D.A., Pilgrim, J.R., Mills, W.N., Diao, Y.: Able: A toolkit for building multiagent autonomic systems. *IBM Systems Journal* 41(3) (2002)
7. Baldauf, M., Dustdar, S., Rosenberg, F.: A survey on context aware systems. *International Journal of Ad Hoc and Ubiquitous Computing* 2(4), 263–277 (2007)
8. Aiello, M., Dustdar, S.: Are our homes ready for services? a domotic infrastructure based on the web service stack. *Pervasive and Mobile Computing* (2008)
9. Dorn, C., Schall, D., Gombotz, R., Dustdar, S.: A view-based analysis of distributed and mobile teams. In: *Proceedings of the 5th International Workshop on Distributed and Mobile Collaboration (DMC 2007) at WETICE-2007*, IEEE Computer Society, Los Alamitos (2007)
10. Vieira, V., Tedesco, P.A., Salgado, A.C.: Towards an ontology for context representation in groupware. In: *Proceedings of the International Workshop on Groupware, CRIWG*, pp. 367–375 (2005)
11. Sterritt, R., Mulvenna, M.D., Lawrynovicz, A.: Dynamic and contextualised behavioural knowledge in autonomic communications. In: *Proceedings of the 1st International Workshop on Autonomic Communication, WAC*, pp. 217–228 (2004)
12. IBM: Autonomic computing toolkit: Developer’s guide (2004), <http://www128.ibm.com/developerworks/autonomic/books/fpy0mst.htm>
13. Wolf, T.D., Holvoet, T.: Emergence Versus Self-Organisation: Different Concepts but Promising When Combined. In: Brueckner, S., Di Marzo Serugendo, G., Karageorgos, A., Nagpal, R. (eds.) *Engineering Self Organising Systems: Methodologies and Applications* (2005)
14. Wolf, T.D., Holvoet, T.: Towards a methodology for engineering self-organising emergent systems. In: *Proceedings of the Int. Conference on Self-Organization and Adaptation of Multi-agent and Grid Systems*, pp. 18–34 (2005)
15. Bird, C., Gourley, A., Devanbu, P., Gertz, M., Swaminathan, A.: Mining email social networks. In: *MSR 2006: Proceedings of the 2006 international workshop on Mining software repositories*, pp. 137–143. ACM Press, New York (2006)
16. Valverde, S., Solé, R.V.: Self-organization and hierarchy in open source social networks. Technical report, DELIS - Dynamically Evolving, Large-Scale Information Systems (2006)