PROJECT FINAL REPORT

Grant Agreement number: 213339

Project acronym: ALLOW

Project title: Adaptable Pervasive Flows

Funding Scheme: STREP

Period covered: from 01.02.2008 to 31.07.2011

Name of the scientific representative of the project's co-ordinator[1], Title and Organisation:

Prof. Dr. Kurt Rothermel

Institute of Parallel and Distributed Systems

University of Stuttgart

Tel: +49 711 685 88448

Fax: +49 711 685 88424

E-mail: Kurt.Rothermel@ipvs.uni-stuttgart.de

Project website address: http://www.allow-project.eu/

[1] Usually the contact person of the coordinator as specified in Art. 8.1. of the Grant Agreement.

Contents

1 Final Publishable Summary Report

1.1 Executive Summary

The objective of the ALLOW project was to develop a new programming paradigm for human-oriented pervasive applications. This paradigm shall enable pervasive technical systems to adapt automatically and seamlessly to humans involved and embedded in them, explicitly supporting people in achieving well-defined goals in dynamically changing environments and contexts. Furthermore, it enables the integration of humans into pervasive business and working processes in an unobtrusive way such that the resulting environments are secure and trustworthy.

The key concept for realizing these objectives in the project is called *Adaptable Pervasive Flow* (APF). Many processes in real life are defined in terms of flows either implicitly or explicitly. A flow consists of a set of actions that is *glued together* by an execution plan in order to achieve a goal under a set of constraints. This resembles the well-known workflow concept. However, APFs represent a much broader concept that enables adaptable pervasive applications. They are situated in the real world, i.e. they are attached to entities like artifacts or people, moving with them through different contexts. While being carried along, they model the behavior intended for their entity and adapt the entity's environment to this behavior. Thus, when a mobile user carries a flow that specifies his prospective actions, the pervasive computing environment will be set up for him by the flow. The flow itself may also adapt to reflect changes. In order to achieve this, flows are context-aware. They can sense the context pertaining to their entity's current environment as well as the entity's actual activities. The key feature that drove the project is that a flow represents a *temporal model of a complex application*. By modeling activities and transitions (either explicitly by a designer or by learning from observing the respective process) knowledge is encoded in the flow that represents past, current and future actions.

New concepts for learning, modifying and exploiting this *flow knowledge* for different purposes are the major outcome of this project. New mechanisms, methods, and principles for adaptable flow-based pervasive applications were investigated, and the following set of new fundamental technologies has been developed for enabling such applications:

- **A Flow-based activity recognition system** that is able to construct flows from observed activities and can improve recognition accuracy notably by exploiting this flow knowledge.
- **A Flow-based context prediction system** that can forecast the future context of a user (e.g. location) more accurately than previous systems, enabling more proactive applications.
- **A system for robust navigation in pervasive flows** that exploits the flow-knowledge in order to minimize the errors caused by inaccurate sensor readings and uncertain activity information.
- **New approaches to security for human-centric flows** to help *people* fulfill their duties securely, safely and to a high standard in the presence of uncertainty and incomplete specifications.

- **New UI concepts for interactive flow-based applications** that enable user to find and interact with flows in their environment.
- **New flow modeling concepts** that allow for a much more flexible composition and execution of flows and for user-specific views on large flow-based applications.
- **New concepts for adapting and evolving flows** such that they match the current situation and allow the users to continue despite the fact that certain preconditions are not met.
- **New concepts for distributing flows in a pervasive environment** such that user-perceived quality requirements and energy constraints are met.

1.2 Summary of Project Context and Objectives

Human users are increasingly embedded in an environment consisting of uncountable computing devices and artifacts that provide various degrees of computing power and awareness. Current estimations state that ten years from now, there will be 7 billion people surrounded by 7 trillion wireless devices and sensors living on this planet. Pervasive applications are software systems that run in such environments in a massively distributed fashion and support mobile human users in their daily activities. This vision holds several scientific challenges ranging from adequate hardware architectures to communication protocols and aspects of software distribution. The most challenging question, however, is: *How can a pervasive application adapt to the user in order to support him/her in an unobtrusive way?* Ideally, the application shall run in the background unnoticed by the user and adapt to his actions. This requires new paradigms for programming such applications.

Today, the interaction with computer systems of any kind is still done explicitly and manually in most cases. However, this is not an option for pervasive systems that shall offer support to a wide spectrum of users that may not be capable of dealing with the specificities of computing devices, even more so as these devices may not even be visible. Moreover, the necessity of having to adjust some pervasive application explicitly represents a major burden, contradicting the initial intention of pervasive computing. The same is true if users have to adjust themselves to such an application. Therefore, pervasive applications must be able to adapt themselves to the user autonomously in order to be feasible. However, there is still a considerable lack of technologies for enabling this kind of adaptivity.

Between 2008 and 2011, the ALLOW project was developing *Adaptable Pervasive Flows* as a key enabling technology for pervasive applications. Our overall key objective can be described as follows:

> ***The objective of the project was to develop a set of fundamentally new concepts and technologies as a core part of a new programming paradigm for enabling human-oriented pervasive applications. This paradigm enables pervasive technical systems to adapt automatically and seamlessly to humans involved and embedded in them, explicitly supporting people in achieving well-defined goals in dynamically changing environments and contexts. Furthermore, the results of the ALLOW project enable the integration of***

humans into pervasive business and working processes in an unobtrusive, secure and trustworthy way.

This work consists of several research activities in different directions with the following fundamental objectives:

- **A Flow-augmented activity recognition system:** The challenge of enhancing context recognition technologies with flow awareness is to find ways of utilizing the additional information encoded into the flow as sequence of actions to improve activity recognition results. While the initial plan was to use flow information to improve semi-/unsupervised training technologies of data classification methods, the medical application scenario suggested (through the given constraints of sensor hardware) a more important goal: combining current activity recognition technologies with flow knowledge to get more and more precise context information out of less and less precise sensor data.
- **A Flow-based context prediction system:** The objective of flow-based context prediction was to investigate how the knowledge encoded in a flow can be exploited to predict the future context of the associated user more accurately than possible with existing technologies. This knowledge indicates the most likely next activities of the user which in turn are associated with specific changes in context (e.g. location, speed, temperature and many other things). Predicting future context is of central importance in *proactive applications* that shall adapt to relevant changes in a user's context *before they actually happen*. Current context prediction systems are not incorporating application knowledge similar to that offered by flows, and thus, have a relatively low performance.
- **A system for robust navigation in pervasive flows:** The goal of our work on robust flow navigation was to enable the execution of flows in a dynamic pervasive environment. Since flows are heavily based on context information as their primary input, they are vulnerable to many problems like inaccurate, uncertain, missing or false context information, and these occur frequently in pervasive systems. Increasing the robustness of (flow-based) pervasive applications is a vital step towards unobtrusive pervasive systems since the active involvement of the users in dealing with errors can be greatly reduced. Existing context-aware systems mostly assume that received context information is correct which is unrealistic.
- **New approaches to security for human-centric flows:** Most approaches, to pervasive systems security, apply and extend traditional access control models. Such approaches fail to adequately address the needs of real-world scenarios where human activities play a central role in the operation of the system. The goal of this work has been to develop greater flexibility in security specification and enforcement that enables flow-based systems to cope with unexpected situations in a controlled manner even when specifications are incomplete. The key to the approach is to delegate greater responsibility and obligations on humans involved in flows and to allow security policies to be overridden while enforcing additional obligations.

- **New UI concepts for interactive flow-based applications:** Current workflow user interface concepts force users to follow a prescribed plan of action, whether or not the plane makes sense in the user's real world context. With a focus on human-centric flows, the goal of the interaction work of ALLOW has been to develop an interaction model and interactive system for flexibly guiding people's physical work activities, by helping them understand which actions are relevant and make sense in their current situation. As people and associated flows move throughout the environment, interfaces need to adapt while taking into account placement of people and devices, and interaction requirements of flows. Thus the gist of this work has been to develop adaptable, situated user interfaces for flow-driven applications, and to evaluate the resulting interfaces with respect to the impact on real world task execution.
- **New flow modeling concepts:** The means todays workflow languages provide to model and execute flexible pervasive flows are limited. The goals of our work were twofold. We want to provide means to model and compose partial process knowledge depending on some context information, since some parts of pervasive flows are assumed to be predictable and therefore it can be modeled. However, the complete flow is determined at runtime since the flow parts are selected depending on context information. And the other goal was to provide user-specific views on large flow-based applications, which are necessary to foster user-specific support in flexible and pervasive environments. The developed concepts shall helping people organizing their work, since they have not found their way into state of the art workflow technology. Therefore, we understand a person's task list as a flow of tasks, which supports and allows to control the work a person does by linearizing and refining their tasks into an appropriate work plan.
- **New concepts for adapting and evolving flows:** The goal of our investigation on flow adaptation and evolution was to guarantee that flows are able to adapt in order to match the current situation and allow the users to continue despite the fact that certain preconditions are not met. In particular, this required developing novel theories and mechanisms for two types of flow adaptation: short-term adaptation (also simply called "adaptation"), where a single flow instance is adapted in order to fit its current context; and long-term adaptation (also called "evolution"), which takes into account the whole history of adapted flow instances and the measures of their respective performance in a given context, and use these information to evolve the flow model.
- **New concepts for distributing flows in a pervasive environment:** The goal of this work is to find the best computers and devices for executing flow-based applications in the mobile users' vicinity such that the user's perceived quality of services is maximized while the energy consumption on any mobile devices is minimized. It is very important to find and adapt good distributions for flow-based applications since they may interact with many services, sensors, devices and users to accomplish their goals. This induces communication and calculation that, if placed inadequately, can decrease the effectiveness, the efficiency and the user experience greatly. Today's workflow systems do not account for energy and quality-aware distribution as they are mostly designed for enterprise environments with different requirements.

These issues were investigated based on two primary application scenarios that posed different challenges and requirements. Based on real-world case studies, subsequent simulations and field tests, we investigated each of these topics and their possible combination and interactions.

1.3 Main Results of the Project

1.3.1 Flows and Flow-based Applications

The key concept that was investigated within the project in order to realize these objectives is called *Adaptable Pervasive Flow* (APF – also simply called "flow" hereafter). Many processes in real life are defined in terms of *flows* either implicitly or explicitly. A flow as we define it in the ALLOW project, is a computer-based model that essentially consists of a set of actions that is *glued together* by some kind of *plan* (or control flow) for their execution in order to achieve some *goal* under a set of *constraints*. This resembles the well-known workflow concept which has proven successful in various fields such as business workflows and service-oriented architectures (SOA).

However, flows represent a much broader concept that ideally fits the problem of creating adaptable pervasive applications. They are *situated in the real world*, i.e. they can be logically or physically attached to entities like artifacts or people, moving with them through different contexts. While they are carried along, they model the behavior intended for their entity and *adapt the entity's environment to this behavior*. Thus, when a mobile user carries a flow that specifies his prospective actions, the pervasive computing machinery can be set up for him by the flow, wherever he goes and whatever he may intend to do there. Furthermore, many background tasks (like documenting a work process) can be executed automatically, and e.g. warnings and recommendations can be given to the user in case she diverges from the seemingly best practice of executing a certain process. Since people may change their minds, and since artifacts and people may be subject to changes in the environment, the *flow itself may also adapt* to reflect such changes. This requires flows to be context-aware. They can sense the context pertaining to their entity's current environment as well as the entity's actual activities. A flow can then use this context information to verify whether it matches specific conditions for continuing or completing. If a mismatch is found, the flow can be automatically adapted to deal with the new conditions.

1.3.1.1 Flow Design Space

The basic assumption of the ALLOW project is that humans follow a *flow of activities* in daily life. Figure 1 depicts an example: If a nurse in a hospital takes a blood sample, then this is naturally composed of a set of real-world activities that have to be executed in sequence to complete this action: "prepare injection", "take blood", "label sample", "store sample", "document procedure". We call this series of activities that one can identify when observing the nurse the *real-world process* (or *process* in short). The process itself has no physical representation. It is simply defined by what the nurse actually does. Our general vision assumes that such a process can also be explicitly modeled as an electronic *flow* in an IT system such that the flow can support the nurse in what she is doing. This flow is modeled by a flow designer and executed by some server in parallel to the real-world process. The system that is responsible for the flow execution is also called *flow system* in the following.

In the example below, the *flow* contains the same activities (called *tasks*) as the *process* and three additional tasks that support the nurse in her action: one task that prints the label for the blood sample, another task that automatically displays the correct documentation form for this kind of action on the nurse's PDA and a final task that stores the data input by the nurse via her PDA in the hospital's administration system. This data can then be automatically transferred into the patient's health record – something that would otherwise have to be done manually by the nurse.

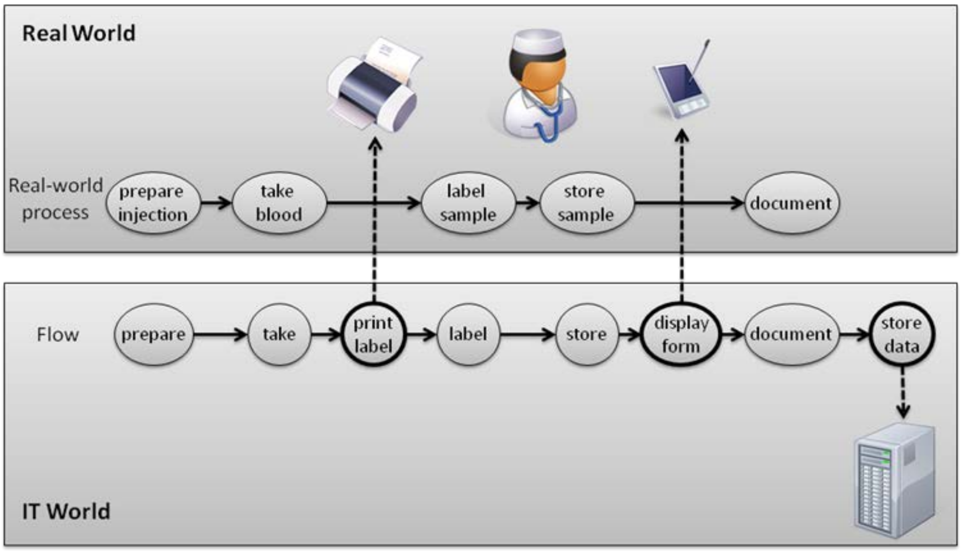

Figure 1 Sample flow-based application with an h-flow and an e-flow

Two general questions with respect to the concept of processes and flows are important:

1. **How rigidly is the flow defined?** Humans may have best-practice procedures for certain actions, but the actual order in which the individual activities are executed may vary between separate executions which is not necessarily an indication for an error. For example, the process in Figure 1 may also be perfectly valid if the sample is stored after the documentation. On the other hand, certain processes must follow strictly ordered tasks, e.g. due to safety or liability reasons.
2. **How does the flow interact with the user?** Since the flow is to be executed in parallel to the real-world process, it needs some means for synchronizing with this process such that, for example, the flow is really executing the task "label" when the nurse starts the "label sample" activity. The requirement of unobtrusiveness dictates that this should not be achieved via repeated explicit requests to the nurse or explicit direct input by the nurse. However, under some circumstances, the flow may require explicit input as will be seen in the examples below.

This opens the two-dimensional design space for flows that is depicted in Figure 2:

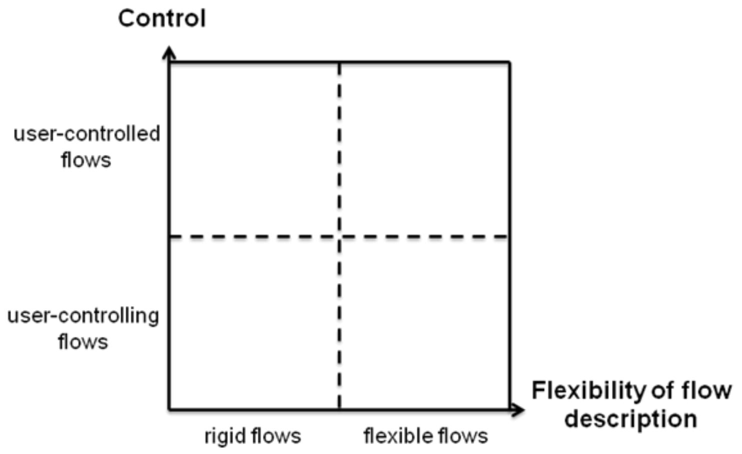

Figure 2 Two dimensional design space for flows

The horizontal axis expresses the flexibility of the flow model and has the following extremes.

- **Rigid flows** (cf. left side of Figure 3)**:** Flows may be modeled in a *prescriptive* way by specifying the complete control flow between the tasks. A prescriptive flow leaves no room for the human to deviate from the execution order defined in the flow. The flow in Figure 1 is an ordered sequence of tasks and, thus, represents a prescriptive flow.
- **Flexible flows** (cf. right side of Figure 3)**:** Flows may also be modeled in a *descriptive* fashion. A fully descriptive flow defines a set of tasks without explicit control flow. Instead, constraints may be specified like e.g. "the flow must be completed by the 'document' task", or "if the blood is taken, then the sample must eventually be labeled". Such a descriptive model can be very flexible and lets the human themselves choose the order in which to execute the tasks as long as they stay within the given constraints.

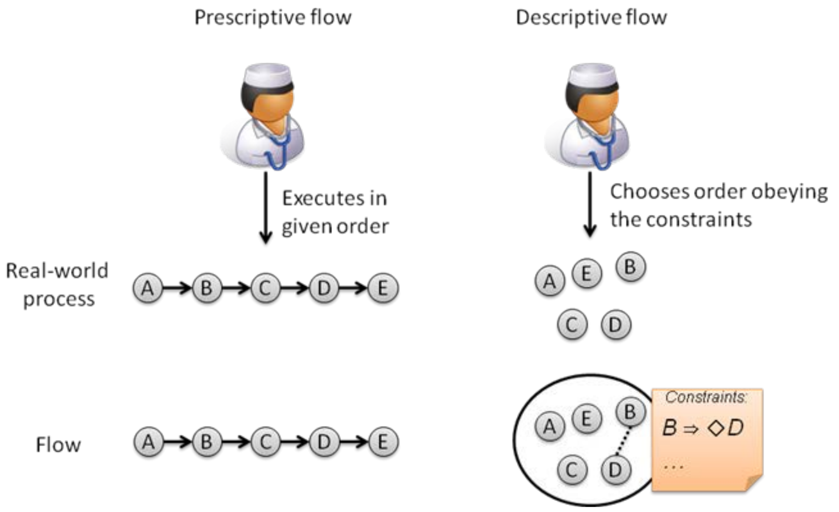

Figure 3 Extremes of the flexibility dimension

The vertical axis of the diagram in Figure 2 defines who has control over the execution of the process/flow. This axis has the following extremes:

- **User-controlling flows** (cf. left side of Figure 4)**:** In some flow-based applications, the flow tightly controls the user's actions. It explicitly tells the user to start a task and requires the user to report back when the task has been completed.
- **User-controlled flows** (cf. right side of Figure 4)**:** Other flows are entirely controlled by the user. In this case, activity sensing is used to sense which activity the user currently executes. Thus, using e.g. acceleration or sound sensors, the system can infer what the user is actually doing in the real-world process and relate that to the tasks in the flow such that the flow system can synchronize the execution of the flow with the actual process. A user-controlled flow has no explicit interaction with the user who is unaware of the flow.

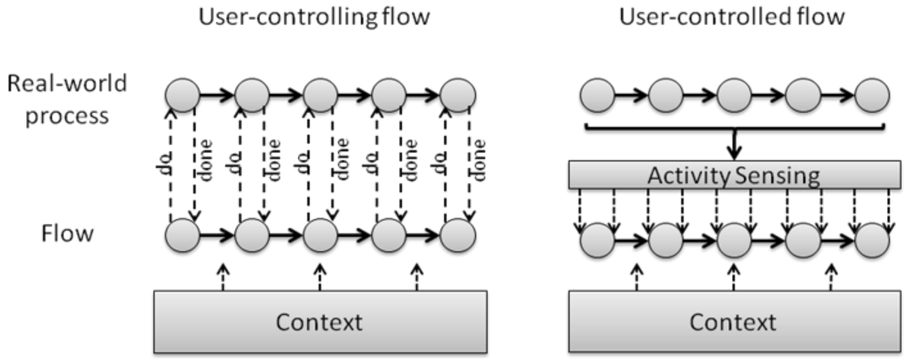

Figure 4 Extremes of the control dimension

1.3.1.2 Flow Categories Investigated in the Project

In the project, we investigated two representatives from the scale of possible flows sketched in Section 1.3.1.1:

- **Adaptive flows:** Adaptive flows are positioned in the lower left quadrant of the design space in Figure 2. They are rigidly defined and user-controlling. The rigid definition implies that these flows are vulnerable to changes in user behavior and to the dynamics of a pervasive environment. Typical application scenarios are found in business and work processes, e.g. in the area of logistics. In these scenarios, a well-defined process is modeled by a flow that explicitly interacts with electronic artifacts (servers, and services) but also with human users that are responsible for executing parts of the flow. The major challenge here is to adapt the flow in case the required execution conditions are not met due to dynamic changes. For example, if a required resource is not available (anymore), an adaptation system is responsible for modifying the flow such that it can be completed (e.g. using an alternative resource).
- **Human-centric flows:** Are positioned in the upper right quadrant of the design space. That is, they are flexible and user-controlled. These flows are designed to run in the background, typically having little direct interaction with the user. Some parts of such a flow may even be learned from the user behavior in an online fashion. Here, the user has the freedom of

choosing how to execute a certain process. Through flexible modeling concepts, learning mechanisms, the flow system arranges the execution to match the behavior of the human users and bother her with as little explicit interactions as possible. Activity recognition and intelligently filtered context information are the primary input to human-centric flows and replace expect user input a much as possible. The concept of flows is used in a number of ways to drive algorithms for context filtering, context prediction, and activity recognition below the actual application flow.

1.3.2 Human-Centric Flows

In this section, we summarize the main research efforts and results in the area of human-centric flows.

1.3.2.1 Activity Recognition

Motivation

Usually, activity recognition systems are set up in three stages:

- Assembling a sensor setup providing sufficient data for a given recognition task
- Collecting and annotating sensor data
- Selecting and training classification algorithms

Collecting the training data is a cumbersome task as it involves lots of manual, human work (mostly annotating the data stream).

The initial research goal was to reduce the human effort in this process by trying to replace the human made activity annotations with flow knowledge. However, during preparation of the case studies at the Mainkofen medical facility and the user interface study conducted with ULANC ([Altakrouri10]), we could not find a sensor setup fulfilling both requirements – delivering enough sensor data for activity recognition *and* being unobtrusive enough to fulfil the requirements of the medical environment. As some of the reasons for constraining the amount and type of sensor devices (e.g. limitations of agility) are not restricted to the deployment of pervasive flow systems to medical environments but will certainly occur on other scenes, the problem was considered worthy of further research within the allow project. Thus, we decided to neglect the research on semi-/unsupervised training combined with flow knowledge (as there are still other ways of handling this problem) in favour of enhancing otherwise insufficient sensor data with flow information to derive reliable context information from otherwise insufficient data.

Approach

As stated in the previous section, the sensor configuration used in the Mainkofen case study is not sufficient for detecting and recognizing the various tasks of the health care workflows. Every of the 5 sensor modalities used is to some extent able to provide some information about the current context of the user, but does in most cases not provide enough information to derive the current action with acceptable recognition rate. While some modalities have an overall relatively low error rate (like WiFi Position estimation) but do not provide exact data about the currently performed activity, others (like sound) are able to recognize some actions reliably while not being able to find others. Different ways of combining these modalities transform the raw sensor data stream into a sequence of possibly spotted activities the user might have performed, leaving out some of them and recognizing others that did not occur. While some optimization can be done on the level of sensor modality fusion which improves recognition results and additional statistical filtering can further reduce the amount of wrong classification results, still lots of “false positives” (i.e. recognized actions that were actually never performed) spoil the recognition results. To compensate for the shortcomings of the poor sensor data detail level, the flow knowledge is used as another filter to improve the ratio between correct classified actions and the wrongly recognized activities: the information about context and sequence of single activities can be used as additional probability measure to the information delivered by the sensor data classification.

Results

The sensor data collected in the Mainkofen mental hospital covered over all more than 100 recorded runs of the medical staff performing the daily morning routine – usually similar activities per day and patients with unpredictable deviations of the usualness. As expected, the first evaluations of the recorded data showed that, with the given amount of data (limited by the scenario environment constraints) it is not possible to recognize the actual activities of the test subject without finding large amounts of ‘false’ activities (i.e. recognizing actions the test subject did not perform), attempts on filtering the recognized activities also affected the recognition performance of actually performed activities. With the available amount of data, appropriate recognition of the workflow activities was not possible. The consequence on the research work was trying to find possible ways of using the flow models of the performed tasks as intelligent and adaptive filters for the context information: The descriptions of workflows can be interpreted as a set of consistency requirements for filtering recognized activities. Although the workflow might only de-/pre-scribe the most common way of performing a task, the flow and the current point in execution of the flow provides restrictions on current and future actions: the order of actions is often fixed or underlies other constraints independent of the actual execution of a workflow (e.g. a diabetic patient has his blood glucose level and blood pressure to be measured BEFORE having breakfast), other workflows provide special limitations (e.g. spacial constraints - a patient will not take a shower in his bedroom) etc. These constraints provided by the scenario specific flow model can be used to check recognized activities for probability and plausibility and thus eliminate recognized actions that are unlikely or impossible to have happened.

As [Muehlb11] suggests on a very simply application scenario, activity recognition can potentially profit by filtering spotted actions employing flow derived restrictions. Although the methodologies for combining sensor data classification results with flow data are not fully explored yet, even the recognition results of the very complex medical scenario of Mainkofen can be improved. However, flow data is still prone to suffer from similar limitations than other data sources for context recognition: The architecture for processing and combining flow data with other context data sources has to be custom designed and optimized for every specific application domain to enable adequate and acceptable results, and (just as faulty sensor data) flow models not sufficiently conforming to the way actual actions are committed in the real world can strongly influence activity recognition results in a negative way. Despite the discovered limitations, which are quite similar to other sensor modalities', flow information has been proven to be able to contribute valuable data to traditional activity recognition paradigms.

1.3.2.2 Context and Flow Navigation

The work in this area revolved around two basic questions:

1. How can we make flow navigation (and thus the execution of flow-based pervasive applications) more robust?
2. How can we exploit flows to generate more accurate context prediction, enabling proactive pervasive applications?

A common approach was used to answer both of these questions: We exploited the knowledge about the past and the future of an application encoded in a flow's structure. While this basic idea was the same, the approaches taken were different.

Robust Flow Navigation

The flows in the ALLOW project navigate mainly based on context data collected from the real world. Therefore, real-world effects that occur when context information is captured by sensing devices and categorized by context management and activity recognition systems can become a serious problem.

We investigated mechanisms for ensuring robust flow navigation under such effects [Wolf2010,Wolf2011]. Most of the context-aware systems proposed in literature assume that the information provided by these systems is correct. However, noise, cheap sensors, and classification errors produce data that can carry a high degree of uncertainty. Moreover, data can get lost due to hardware and software malfunction or simply because the raw data was too noisy to be categorized (false negatives). The same software may find patterns of known events while such an event actually never happened (false positives). Finally, the order of events may not be recognized correctly due to timing problems and race conditions (order errors). All of these effects happen regularly in real systems, and a flow that navigates based on context data will fail very quickly under these conditions.

With FlowCon [Wolf2010], FlexCon and FeVA [Wolf2011], we have developed a set of mechanisms that exploit the knowledge encoded within a flow in order to correct these effects before the context data is actually given to the flow. Flow knowledge is basically the knowledge about the temporal behavior of the flow that is encoded in the order of activities and in the transition between these activities. In a normal (non-flow-based) application, the temporal behavior is usually hidden in the program code, and it can be extremely difficult to infer what the application will do at a specific moment in time, even if one has full access to the code. In a flow-based application, this temporal behavior is explicitly defined in the transition system. The individual activities, the transitions and the conditions under which the transitions are taken are open and accessible. Therefore, it is much easier to infer the past, current and even the future behavior from the flow. We exploit this knowledge to reduce the uncertainty of incoming context data and to repair errors stemming from false positives, false negatives, and order errors.

FlowCon and FlexCon try to increase the confidence the flow can assume for single events that the context management and activity recognition system have detected. FlowCon is our first approach that focused on the more classical imperatively (hence strictly) modeled workflows, where each activities are executed in tight order. FlowCon extracts the dependencies from the flow to create the structure of a Bayesian Network and learns the dependencies between the events that drive the flow from previous executions. Using the Bayesian Network as additional source of information we gain confidence when processing the incoming context events. The FlexCon system is an advancement of FlowCon that focuses on more flexible and loosely coupled flows that allow the human user a more freely way of executing his tasks. However, those flows also potentially offer less information for flow navigation. FlexCon deals with this problem by using Dynamic Bayesian Networks to account for the changing execution of the flow. An important metric for the evaluation of FlowCon and FlexCon is the completion ratio of flows. This is the ratio of flows that is completed without getting stuck or ending in an error due to false navigation decisions. Without any assistance, a flow system that is faced with realistic uncertainty levels and noise only manages to complete up to 10% of all flows. With FlowCon, we observe a completion ratio or up to 90%. FlexCon performs worse due to the increased flexibility and dynamics in flow execution. But it still achieves between 20% and 80% flow completion.

FEvA resolves order-errors, false negatives and false positives. To do this, we compute a soft mapping of events to activities that is based on fuzzy logic. This mapping takes the structural information encoded in the flow as well as the flow's current execution state into account to find the most beneficial mapping of events. The mapping is adapted dynamically as further events are processed and the flow execution continues. We compared two flow management systems, one using the event mapping of FEvA and the other using a very simple event mapping. The results show that the robustness of the system using FEvA degrades much slower and more gracefully than for a system without FeVA, when exposed to rising numbers of false positives and out of order events. It even provides some graceful degradation when events are missed.

Proactiveness – Context Prediction

Navigating a flow based on current and past context information may be sufficient for many applications. However, if the flow has to be adapted in real time (i.e. while the user is waiting) to work around some problem (e.g. a missing resource), and if such problems may occur frequently (due to the mobility and dynamic of the entire system) then having reliable information about future context changes is vital for flow control. It enables the system to be proactive by recognizing and circumventing possible problems before they actually take effect. Therefore, we have investigated the area of *flow-based context prediction* to inform about expected changes in the user's context and support proactive control mechanisms.

The basic idea of flow-based context prediction is to exploit the flow knowledge as an additional source of information. This results in an improvement of the confidence in expected context changes which are observable in the real world. The flow reveals the states of a flow-based application and thus restricts the context changes to those that are typically occurring when the flow has arrived at a certain execution state. Following this idea, we have developed a context predictor that takes activity information (from a flow or some other source) and learns how these activities are connected with context changes [Föll2010]. This provides information about the context changes that are expected to happen each time the same flow activity is scheduled for execution. We represent this relationship as a probabilistic state transition system which is incrementally refined from the execution of flows at runtime. The state transition system describes the evolution of the user context, depending on the flow activity and previous user context from the user history. For context prediction, we traverse the state space of possible context changes to determine the most likely paths of future context occurrences. We distinguish between short-term and long-term prediction for which we have developed prediction algorithms that infer the series of the future context changes with highest probability. In our evaluation, we have shown that the inclusion of flow knowledge in the prediction model significantly improves the prediction accuracy, as classical predictors are limited by their application-agnostic view. The result is a system that can learn the behavior of context information more precisely than existing systems and is able to support more efficient proactive decisions due to the increased accuracy.

Moreover, we have developed PreCon [Föll2011], a novel approach to rendering flow-based context prediction more expressive than existing systems. PreCon allows time-dependent context prediction queries for multi-dimensional context (e.g. activity and location) and thus enables applications to express their queries for future context in a meaningful and flexible manner. PreCon applies methods of stochastic model checking (used e.g. for the verification of distributed communication protocols) to the analysis and prediction of human behavior. While classical model checking relies on fixed hand-crafted models of computer systems, our models are dynamic and learnt from traces of user context. In PreCon, user behavior is represented as Semi-Markov Chains (SMC) and temporal-logics is used as a query language. We extended existing model-checking techniques to deal with the online character of context predictions and to allow for continuous learning of SMCs. PreCon's query language provides a powerful means for applications to pose temporal queries for reachability and invariant properties of future context. Thus, PreCon goes far beyond existing approaches and represents a new class of context prediction systems that enable intelligent ubiquitous applications to take much more educated decisions. We evaluated PreCon based on a real-world case study in the area of healthcare using metrics from information retrieval and showed that it exhibits a good performance. Moreover, our evaluations yielded indications for choosing sensible parameters for different classes of applications. Overall, this represents a notable step forward from existing context prediction systems. We have shown that flow knowledge can be used to improve prediction capabilities. This means that flow-based applications have a great potential to be proactive and run in the background, unnoticed by the user. The feature of time-based queries provides applications and the flow system itself with a powerful means to pose meaningful queries.

1.3.2.3 Break-glass Security and Quality of Care

Motivation – Human-centric flows are subject to a great number of external events, which prompt people to act in previously unanticipated ways in order to achieve the goals of a workflow. For security this can manifest itself in different ways. For example:

- A request that ought to be permitted was not anticipated and thus not encoded in the policy, or an encoded security policy may be erroneously specified and thus omits permissible accesses (e.g. slow updating of security policy compared to changes in the organisation).
- It is not possible to anticipate and/or correctly machine-encode all possible emergency/exceptional situations under certain requests ought to be permitted.

These are real issues that that need to be addressed by practical pervasive flow systems that need to ensure that security policies still provide the necessary protection while not preventing people in fulfilling their duties. If security is too rigid, people will attempt to circumvent security altogether leaving the system even more exposed to security violations and attacks. This is likely to manifest itself particularly where workflow tasks do not involve computer resources at all but where there is a need to monitor who is doing what, to whom, under what circumstances, e.g. nurse X treating patient Y.

Currently, approaches in pervasive security systems are mostly oblivious to these issues, they typically adopt traditional access control techniques that are predicated on a *preventive* protection strategy (e.g. [Garcia10], [Hayat07]). This strategy is implemented by intercepting all access requests and granting only those that are entailed as permitted by the security policy. A crucial assumption underlying all instances of access control models is that security policies are, in a sense, complete and do indeed anticipate and precisely formulate not only which subjects but also when these subjects are to be granted their access requests. In other words all permissible requests are fully and correctly defined and will be granted when they are requested.

Another assumption made by most existing security models is that all requests deal with computer resources and thus they are indifferent to any physical access (such as nurse examining a patient) that is not modelled as a request to a flow system. In human-centric flows it is important that people are guided through what is permitted and what is not, in order to fulfill their duties and to ensure safety. In other words, it is people who are protected resources as well. This is especially true in medical scenarios, where for example a nurse, while aiding a patient, may expose herself to toxic drugs or infectious diseases. Therefore we believe that security needs to be pro-active and support people in their work duties, rather than simply wait for access requests or violations. Furthermore, security techniques have the potential to make people better aware of the security risks present in any situation and to encourage people to perform their tasks to higher standards, while providing informative feedback to the security system itself.

Problem Statement – We can summarise the motivating discussion through the following problem statements:

- What is an appropriate paradigm to enforce security policies for human-centric flows?
- How to manage violations that will inevitably happen? Notice that violations are not necessarily malicious -- people make forget, delay, reorder, delegate etc.
- How to guide and encourage people to perform their duties safely and to a high standard?

Approach **–** In order to address the stated problems, in this project we have decided to explore a different method for security enforcement. In particular, as the basis of our approach, we have investigated a *break-glass* approach to defining and enforcing human-centric security policies. The core idea is to allow people to *override* security policies that deny access. In break-glass people are empowered to make security decisions based on their assessment of the current situation. Our security policies focus on understanding whether people are competent and empowered to perform the tasks they require, rather than hard-coding strict permissions and denials. If they are, they can override policies but may have additional obligations imposed on them.

In order to manage security violations we have developed a new approach based on Teleo-Reactive (TR) processes from AI. Teleo-reactive processes effectively encode *reactive* behavior, which guides a human or agent to its goal, e.g. how to deal with security violations. They have powerful semantics that support continuously evaluated conditions and actions. In ALLOW TR agents are deployed on smartphone and the servers to monitor that the execution of flows follows security norms. If a violation is detected a TR agent will attempt to monitor and if necessary guide a user out of the violation.

Current workflow security models have focused on authorizing task executions. In ALLOW, we take a different approach and use flows as a way to guide people towards safe task executions. For example, the security model pro-actively injects additional safety tasks when a non-competent person is attempting to execute a task.

Results **–** There are three main results achieved within the scope of this project, each addressing one stated problem.

First, we have developed a flow security model based on the notion of break-glass security. As part of this development we have developed a novel security language that can make inferences over missing, conflicting, and partially trusted evidence and implemented it as logic program and can be used with any popular Prolog implementations.

Second, we have formulated new execution semantics for TR processes suited for event-driven workflow systems and implemented such TR system for Java and Android platform and shown how the systems can be used to deal with violations and break-glass obligations. The system has been made as open-source to others on the Ponder2 website for several months now.

Third, we have constructed a language to express *safety* policies, such as which tasks should and must be done while attempting to execute a certain flow task. As a demonstration we have shown how this language can be used to capture safety norms found in nursing practices.

1.3.2.4 Flows and User Interfaces

***Motivation* -** Traditional workflow systems are designed with a technology-centric view of the world: users are typically reduced to 'executing' tasks defined by a fixed plan of action or workflow (this view is exemplified for example by BPEL extensions BPEL4People and WSHumanTask). Making flows 'pervasive, i.e. bringing flows into the real world and extending their reach to complex work environments (e.g. a hospital) requires a completely different model, one that is steeped in an understanding of how people act, what motivates them and how they make decisions. In recent years a new theory of action has taken hold that emphasises the situated nature of action [Agre97, Hutchins95, Suchman87]. The situational theory posits that human action is a direct unmediated response to situations encountered in the course of routine activity, and that such action can be goal-attaining if the environment in which the activity takes place is sufficiently conducive [Johnston02, Johnston01]. For example, the actions of a nurse who cares for a patient are primarily shaped by what he or she perceives to be necessary in a given situation - and not by tasks prescribed by a mechanistic workflow. Yet as of now workflow technologies do not account for the situated nature of actions and for the fluid, minimally reflective actions of skilled actors engaged in routine tasks.

***Problem Statement* –** Flows are a novel paradigm for regulating the complex interactions that occur between humans situated in pervasive environments and pervasive services. The key problem here is to overcome the tension between the situated nature of work practices on the one hand and the prescriptive nature of flows on the other hand. This requires a novel interaction model that reinterprets flows as 'opportunities for actions' rather than as 'prescription to follow'. Thus our main task was to define a model for situated reactive guidance that helps people understand which actions are relevant and make sense in their current situation without forcing them to conform to a rigid plan of action. As flows and people move throughout the environment and as flows are adapted to fit the current context, the interface needs to adapt and make use of the most appropriate devices, taking into account placement of people and devices, and changing interaction requirements of flows.

Although theoretical system models that conform to the situated action theory have been proposed, such as 'situated information systems' [Johnston05, Waller09, Waller09b] and the 'situated choice support system' [Lederman11], these models have not seen concrete realization. Similarly, existing systems approaches using handheld and shared large-scale displays [Favela09] do not adequately support the situated and spatially distributed nature of hospital work.

***Approach* –** The gist of the user interface framework for pervasive flows is to transform a physical work environment into an interactive canvas for information which is used to provide actors with situated cues about the location, timing and nature of possible actions [Dix97]. These cues are driven by the flows that exist in the environment and are highly dynamic as they continuously adapt in concert with flows to reflect the current state of work. In concrete terms our approach towards realizing the situated action approach encompasses four parts:

Situated flows: we developed a novel concept of *situated flows*, which are flows that are located in time and space. This allows us to relate each flow activity with a particular physical location and a time, indicating where and when the activity should or must take place. The situated nature of flows is the foundation for turning a physical environment into an interactive space that makes visible opportunities for undertaking activities to achieve goals. In terms of user interface design, situated flows are the foundation for the interactive discovery of flows and activities by users.

Situated glyphs: Inspired by the work on medical visualization and recognizing the limitations of existing interaction approaches in hospitals environment we have developed a novel user interface paradigm for flows based on the notion of *situated glyphs*. Situated glyphs are visual representations of activities that are situated in place and time: they indicate when, where, which activities can or should take place in an environment. The situated glyph concept emerged from an experimental user interface study as Mainkofen Hospital focused on exploring intelligibility issues of situated information interfaces.

Micro-display networks: To realize situated glyph interface we developed *micro-display networks* composed of small-scale wireless embedded display units. Collections of micro-displays can be used as a physically dispersed, sparse, non-contiguous display space.

Interactive Flow Evolution: Workplace studies at Mainkofen and Hannover hospital uncovered that flow evolution cannot be fully automated in a hospital setting but must support the collaborative nature of established review procedures of medical and care procedures. In response to these insights, we have developed a new approach for *interactive flow evolution*, a user-driven learning process which translates *in-situ flow feedback* into *evolution hints* to compute *suggested flow evolutions*.

Results – Besides novel concepts and technologies described above the key results encompass technical and user experience aspects:

First, we developed three spatial distribution strategies for situated glyphs as follows:

2 *Activity centric* distribution strategy: This strategy places glyphs at the main location where an activity is to be performed.

3 *Entity centric* distribution strategy: This strategy places glyphs at or very near the main entity involved in the activity.

4 *Space Centric* distribution strategy: This strategy places glyphs in the observable space shared by multiple activities.

Second, through lab-based and workplace-based user studies we gained insights into the differences between the three spatial distribution strategies wrt the utility of situated glyphs and the quality of user experience. Some of the key insights are as follows:

5 *Mental Task Load and Information Utility*: in general, distribution strategies with higher number of glyphs fare better than a single centralized display.

6 *Task completion time and task errors*: a higher number of glyphs can lower the activities completion time. There is no correlation between the number of glyphs and error rate.

7 *Physical movement patterns*: there are marked difference between the strategies in the way users traverse a room when completing a distributed task.

8 *Attention management*: Situated glyphs require focused attention; they do not function as ambient display.

8.1.1 Adaptive Flows

In this section, we summarize the main research efforts and results in the area of adaptive flows.

8.1.1.1 Flow Modeling

Implementing pervasive scenarios using flow technology requires that flow technology must be able to react to highly dynamic changes in the environment. This means that the executed flows may look very different from one execution to another, even if the executions have the same goal. As a consequence flows cannot be modeled as a whole. Looking at the flows, however, one can recognize pieces that occur over and over again. Based on this observation we developed the concept of "flow fragments". Basically, the objective of flow fragments is to foster reuse of pieces of flow logic.

Furthermore, pervasive flows involve people in the execution. Several technologies like activity sensing operate on the activity flow of a respective person. This flow is called "person-centric flow". A person, however, participates within several flows at the same time and has to perform tasks derived from several flows at the same time. Thus, a person-centric flow is a flow consisting of tasks that are assigned to a single person and created by a variety of flows running within the environment. The person-centric flow provides a view on the tasks of a respective person. In order to support person-centric technologies like activity sensing a concept for generating person-centric flows has been developed.

Flow fragments

Flow fragments [EUL09,ELU10,Eberle10] are used to represent fragmentary process knowledge. Flow knowledge is assumed to be 'local'. This means that flow fragments represent fragmentary business knowledge for specific occasions. Flow fragments therefore capture modularized reusable flow logic.

Flow fragments can be composed either at design time or runtime into complete flows using scenario adequate integration techniques. Runtime decisions can be implemented within the flow fragment model itself using the constructs provided by the flow fragment modeling language. If the runtime cannot be modeled within the flow fragment, it can be implemented by the flow fragment composition technique by choosing a certain flow fragment to be composed with the running flow fragment.

Hence, the flow fragment approach is a hybrid approach, where runtime decisions are made either within the flow execution context in an imperative way or by the integration technique using flow composition.

The following list shows the contributions we provided in this area:

- Process fragment modeling language
- Integration techniques for flow fragments
- Recovery concepts for flow fragments

Figure 5 describes the execution steps that are processed during the execution of flow fragments. Firstly, a flow fragment gets selected and composed before the execution of the composed flow fragment can start. During the execution of this fragment other fragments can get selected and composed with the running process fragment over and over again.

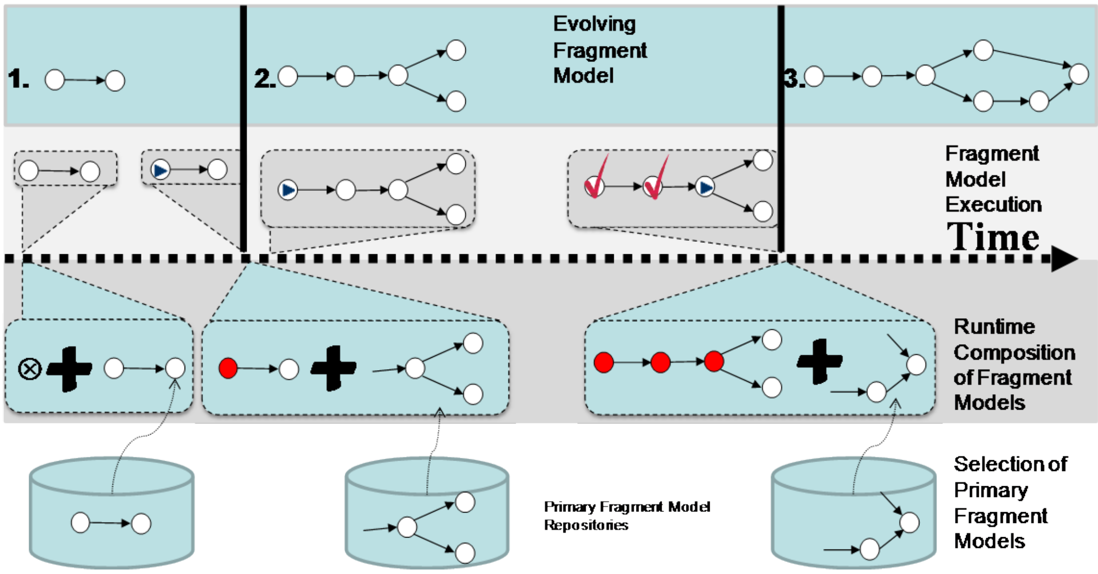

Figure 5: Execution of Flow fragments

Person-Centric Flows

The goal of this work has been to develop concepts helping people organizing their work, since such kind of concepts are not found in state of the art workflow technology. Our approach is based on the idea of a person-centric flow [Unger10a,Unger10b,UR10] summarized above. The basic idea behind a person-centric flow is to understand the set of tasks assigned to a person as a flow of tasks. A person-centric flow tries to simulate the behavior of people by linearizing and refining their tasks into an appropriate work plan. This is achieved by two major contributions: a concept for refinement of tasks and a concept of a task recommendation and control system.

- Refinement system: tasks within business processes are often modeled with a granularity which differs significantly from the granularity the tasks are actually executed by people. People divide tasks into smaller pieces which are more manageable for them or combine the execution of two or more tasks, because they appear very similar to them. Therefore, a system is introduced enabling the dynamic refinement of tasks. The system provides the following refinement patterns:
    - Split: splits one task into a set of sub-tasks.
    - Merge: combines one or more tasks into a single task.
- Recommendation and control system: people are very good in scheduling a small number of tasks. A growing number of tasks, however, lead to scheduling errors and an increasing cognitive load. A system that provides recommendations and monitors obligations can counteract this. The (refined) task list of a person is enriched with recommendations and obligations which are generated by domain-specific algorithms, extracted from top-level

business processes, or both. For this purpose we have developed a person-centric flow modeling language. The generated recommendations and obligations and the underlying set of tasks are mapped to flow model. Subsequently, the flow model can be used to guide a person based on the recommendations. Furthermore, the flow model can be used to check the compliance of the execution with the generated obligations.

8.1.1.2 Flow Adaptation and Evolution

In the pervasive applications foreseen in the project, the technical system needs to be able to continuously adapt to dynamically changing environments and contexts. Since flows are the key concept of the proposed approach for modeling the technical system, it is evident that the capability of flows themselves to adapt is mandatory for achieving the goals of the project.

In the project, we have hence addressed the problem of defining and developing a framework supporting the flow adaptation. Our goal was to have a framework able to cover different forms of flow adaptation, both in terms of different dimensions (horizontal and vertical adaptation) and of different time scales (short-term and long-term adaptation). More precisely, the framework should cover cases where the flow itself needs to be adapted, e.g., by changing its structure (horizontal dimension), as well as cases where the flow "components" needs to be adapted, e.g., by refining them, or by finding new implementations for some functionalities (vertical dimension). Also, the framework should cover short-term adaptation (also simply called "adaptation"), where a single flow instance is adapted in order to fit its current context, as well as long-term adaptation (also called "evolution"), which takes into account the whole history of adapted flow instances and the measures of their respective performance in a given context, and use these information to evolve the flow model. Finally, the framework should support the coordination of flow adaptation with the adaptation of other aspects, such as flow distribution, security policies, user interfaces, in order to allow for a coherent evolution of the whole pervasive application.

In order to solve the problem just described, in ALLOW we have developed novel theories and mechanisms supporting different forms of adaptation. The approach that we have followed in the investigation and development of these theories and mechanisms has been incremental, starting from built-in flow adaptation (where all the adaptability is explicitly implemented by the designer), moving then to a restricted form of adaptation suitable only for flow refinement (vertical adaptation), and finally addressing the more general case of horizontal adaptation and of flow evolution.

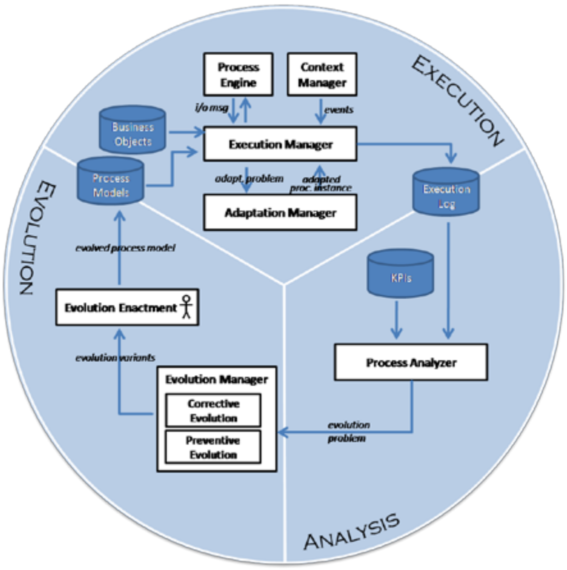

Figure 6 Flow Adaptation and Evolution Framework

In order to guarantee the overall coherence of the theories and mechanisms developed for the different forms of flow adaptation, we have proposed an overall framework for flow execution, adaptation, and evolution. This framework, depicted in Figure 6, defines relations and dependencies among all the different forms of adaptation, and places all of them in a coherent life-cycle for applications based on adaptive flows.

Finally, in order to evaluate the effectiveness and the efficiency of the developed mechanisms and of the proposed framework, we have adopted a real world scenario in the domain of logistics. More precisely, our reference scenario addresses car management in the automobile terminal of Bremen sea port, where nearly 2 million new vehicles are handled each year. By exploiting this scenario, we have studied the benefits of adopting the proposed approach.

According to the approach just described, different kinds of results concerning flow adaptation and evolution have been produced within the project. A first kind of results concerns the definition of novel theories and mechanisms; they cover in particular:

- **Built-in flow adaptation patterns** [Marconi09]. We propose a set of patterns, and the associated primitives for the Adaptive Pervasive Flow Language adopted in the project, that allow the designer to extend the flows with built-in mechanisms for adapting to changes in context and environment.
- **Vertical flow adaptation by on-demand flow creation** [Sirbu11]. For this adaptation we propose a mechanism for specifying and executing on-demand creation of sub-flows. These mechanisms can be exploited in those situations when flow composition may not succeed since adequate sub-flows or services are not present, but new sub-flows may be created using the knowledge that is present in the flow to be adapted and in the environment.
- **Horizontal flow adaptation by context-aware re-planning** [Bucchiarone11c]. We propose a planning-based technique able to modify the structure of the flow "on the fly" according to the

context-specific properties while satisfying the general flow goals and constraints. Indeed, re-planning techniques should provide more flexible mechanisms for horizontal flow adaptation than built-in context awareness.

- **Flow evolution** [Bucchiarone11a]. We propose a framework for evolving process models based on a history of process instance executions and adaptations. Our approach is context-driven. If the need to evolve the process model is detected, we analyze the relevant adapted process instances and look for recurring adaptation needs (i.e., the same constraint violation and system configuration). This allows us to construct and rank evolution variants which can handle the problematic context (corrective evolution). It also allows us to construct evolution variants which can prevent the adaptation need (preventive evolution).
- **Overall Adaptation** [Bucchiarone10]. The control of self-adaptive pervasive systems in usually performed through the cooperation of different autonomous components, each one dedicated to specific aspect of application execution. Each of these components usually is capable of self-adapting a particular part of the underlying pervasive system. To avoid problems like conflicts, race conditions and oscillations, we proposed a solution, based on a negotiation loop, that coordinates the different components and attempts to find an adaptation that satisfies all components' individual constraints and needs.

In addition to this, we have formalized and implemented the **framework** depicted in Figure 6 [Bucchiarone11a]. Finally, we have developed a **demonstrator** [Bucchiarone11b] that shows the framework in action, by exploiting our reference car logistics scenario

8.1.1.3 Flow Distribution

The area of flow distribution is concerned with the problem of finding and maintaining an effective and efficient distribution of a flow over the available computing resources. The basic idea of distributing flows is driven by the fact that flows need to maintain rich communication with their human users and with diverse services in their environment. Since the human users we regard are mobile, the available computing resources and the network connecting them it is not known a priori. Moreover, network bandwidth, stability and quality may be very diverse. Hence, effective mechanisms are required for placing different parts of a flow on those computing resources where they may be executed with the highest utility. We have investigated two important utility metrics in this project. The first one is *user-perceived latency*. This is basically the accumulated time that a user has to wait during the interaction with his flow. To allow users to operate unobstructed by any network effects (low bandwidth, high delay, etc) the activities of his flow have to be placed accordingly such that high-volume communication does not flow over low-quality network links. The second utility metric is *energy*. In the common scenarios regarded in the project, mobile users are equipped with mobile devices (e.g. off-the-shelf smart phones) that are part of the computing resources used for executing flows while the rest of the computing infrastructure are backend infrastructure servers. Preserving the energy of mobile devices is of primary importance since more energy consumption leads to shorter re-charge cycles, which in turn leads to more disturbances in the work process, e.g. requiring users to re-charge their mobile device in the middle of a shift. On the other hand, if the energy consumption is minimized to a known average value, cheaper devices with adequate battery capacity can be used, avoiding over-provisioning in terms of energy. Both effects lead to monetary benefits over the long run.

Minimizing Human Interaction Time

Flows represent compositions of various kinds of services that involve interaction with the user as well as require access to resources such as databases and information systems. Executing such flows in a distributed network of service providers, e.g., in domains such as logistics where resources are spread over multiple organizations and stakeholders, can require significant execution time for completion. This time is experienced by users as waiting time when interacting with such a flow and depends on two factors: First, on the runtime of services that need to be executed in order to provide a human with the desired information. Second, on the time required to transfer data between flow engines and service hosts participating in the execution of a flow. While the service runtime is usually bound by quality of service guarantees, the overhead of data communication is determined by the distribution of the flow activities to available workflow servers in the network and may create extensive delays due to limited bandwidth and significant propagation delay of network links. In order to reduce this overhead of data communication, we have developed an algorithm for distributing a flow over a set of workflow servers such that the interaction time experienced by humans is minimized [Hiesinger2011]. The algorithm takes a flow model and the specification of the network properties as input and creates a mapping of the flow onto the physical resources of the network. The central idea of the algorithm is to identify parts of the flow that share heavy communication load and create clusters of flow activities that are assigned to the same physical machine where the required resources and services are available. As the underlying problem is NP-hard, our algorithm is a based on a two-phase list-scheduling approach that defines a heuristic search through possible placements. In phase 1, a candidate solution is computed based on a reduced search space that only assumes network delays that are averaged over all links in the network where a required resource is available. As the real network delays may significantly differ from the assumed averages, this solution is then further refined and used as a starting point for optimization in phase 2. In phase 2 we refine the initial solution based on a hill-climbing algorithm and greedily alter the current placement to find better candidates that reduce the overall latency. As the distribution algorithm is scheduled before the actual deployment of the flow, the algorithm does not create additional costs that may affect the flow execution at run-time. We have shown in extensive simulation experiments that our algorithm improves the interaction time between humans and workflows by up to 80% compared to an approach in which the complete workflow is run on a single machine. As the integration of business systems and pervasive computing gains momentum, our approach thus achieves effective improvements in quality of service perceived by humans.

Minimizing Energy Consumption on Mobile Devices

At the same time, users rely on mobile devices to interact with flows virtually anywhere. While this enables unobtrusive interactions with the flow system, mobile devices have a constrained battery lifetime such that the energy consumption is a critical point for their usability. High energy consumption causes batteries to be drained fast, disrupting the business process and leaving users unable to continue their work. This renders the process less efficient and, thus, incurs higher costs. Therefore, preserving energy is an important goal with respect to seamless flow execution and cost reduction. However, extensive data transmission between user devices and flows running in the central infrastructure is required for each interaction with a mobile user. Consequently, this results in high energy costs since sending and receiving data are highly energy-intensive operations with current wireless communication technologies such as GPRS, UMTS or WiFi. As a response to these challenges of mobile flow technology, we have developed an algorithm for distributing fragments of a flow from the infrastructure to mobile user devices in order to reduce the energy costs on the mobile device [Fischer2011]. Distributing flows to mobile devices avoids transmissions of large volumes of data since this data is processed locally. However, shifting computational load on the mobile devices introduces additional costs for running the new tasks. Therefore, our algorithm constructs a cost graph that reflects the trade-off communication overhead vs. processing costs. The cost graph consists of the flow activities and the set of devices in the network. The edges in the graph connect flow activities with the device where the required functionality can be found. Also, the edges are weighted with the energy costs that arise from the two sources of energy: the data communication costs and the processing costs of services on the mobile device. A possible placement is given by the assignment of flow activities to network devices, such that the resulting costs are given by the edges that connect the different partitions in the graph. In order to find the best placement, we partition the cost graph using a minimum cut algorithm that guarantees minimum energy costs. In order to apply the minimum-cut approach for more than two devices, we have a proposed a graph transformation that creates a virtual device out of multiple devices and merges nodes and edges in the cost graph. The result of our extended algorithm is a partitioning of the cost graph, where each partition contains the workflow activities to be executed either in the infrastructure or on the mobile devices. Our evaluations show that this approach achieves average energy savings of 37% for GPRS and 32% for UMTS communication compared to the centralized infrastructure-based approach. In application domains that are heavily flow-driven (like logistics and health care) this represents a significant energy saving that allows for longer operation between two recharge cycles and thus for less distractions in the daily routine of the involved personnel. In scenarios where the available battery capacities already ensure distraction-free processes, our approach allows for downgrading to cheaper technology with less capacity. Both routes lead to significant monetary savings in areas where mobile devices are indispensable tools. Thus, our work represents an important step toward the cost-efficient and seamless integration of business processes and pervasive computing, enabling much more flexible flow-driven applications that involve mobile users.

8.1.2 Summary and Conclusions

In the ALLOW project, researchers from different areas of computer science joined forces to explore the new areas of flow-driven pervasive applications. Before this project started, the two areas of workflows and pervasive systems were separated. The fundamentally new approach of the project was to bring these two worlds together in order to 1. build more adaptive and more robust pervasive systems and 2. bring workflow-driven applications into the pervasive world in order to extend the sphere of workflow technology.

This endeavor has resulted in a set of concepts and technologies covering activity recognition, flow modeling, context-awareness, user interfaces, security and adaptation and evolution. Fundamentally new concepts have been developed in all of these areas in a joint effort to enable Adaptable

Pervasive Flows. Aside from this direct result, the project has planted the idea of flow-driven systems firmly in each one of these areas, ensuring that this work will be carried on beyond the ALLOW project. This will eventually lead to a new way of thinking about pervasiveness and adaptivity.

8.1.3 List of References

[Altakrouri10] Altakouri, B., Kortuem, G., Gruenerbl, A., Kunze, K., Lukowicz, P.; Empirical Study of the Benefit of Activity Recognition for Mobile Phone Based Nursing Documentation, 14th Annual International Symposium on Wearable Computing; 2010 October 10-13; Seoul, South Korea.

[Bucchiarone10] A. Bucchiarone, A. Marconi, M. Pistore, S. Foll, K. Herrmann, C. Hiesinger, and S. Marinovic. An Overall Process for Self-Adaptive Pervasive Systems. In Proceedings of the Second International Conference on Adaptive and Self-adaptive Systems and Applications (ADAPTIVE 2010), 2010.

[Bucchiarone11a] A. Bucchiarone, A. Marconi, M. Pistore and A. Sirbu. A Context-Aware Framework for Business Processes Evolution. In Proceedings of 4th International Workshop on Evolutionary Business Processes (EVL-BP 2011) (to appear), 2011.

[Bucchiarone11b] A. Bucchiarone, S. Foll, K. Herrmann, M. Pistore and H. Raik. Adaptable Pervasive Flows: Towards a More Intelligent Environment. Demo at the 9th International Conference on Business Process Management (BPM 2011).

[Bucchiarone11c] A. Bucchiarone, R. Kazhamiakin, M. Pistore and H. Raik. Adaptation of Service-based Business Processes by Context-Aware Replanning. Submitted to the The IEEE International Conference on Service Oriented Computing & Applications (SOCA 2011).

[Marconi09] A. Marconi, M. Pistore, A. Sirbu, H. Eberle, F. Leymann, T. Unger: Enabling Adaptation of Pervasive Flows: Built-in Contextual Adaptation. ICSOC/ServiceWave 2009: 445-454

[Muehlb11] Michael Muehlbauer, Gernot Bahle, Paul Lukowicz, What Can an Arm Holster Worn Smart Phone Do for Activity Recognition? ISWC, pp.79-82, 2011 15th Annual International Symposium on Wearable Computers, 2011

[Sirbu11] A. Sirbu, A. Marconi, M. Pistore, H. Eberle, F. Leymann and T. Unger. Dynamic Composition of Pervasive Process Fragments. In Proceedings of the 9th IEEE International Conference on Web Services (ICWS), 2011.

[Hayat07] Z. Hayat, J. Reeve, and C. Boutle: Ubiquitous security for ubiquitous computing, Information Security Technical Report, Elsevier, 2007, 12, 172-178.

[Garcia10] O. Garcia-Morchin, and K. Wehrle, Modular Context-aware Access Control for Medical Sensor Networks, SACMAT, 2010.

[EUL09] Eberle, Hanna; Unger, Tobias; Leymann, Frank: Process Fragments. In: Meersman, R. (Hrsg); Dillon, T. (Hrsg); Herrero, P. (Hrsg): On the Move to Meaningful Internet Systems: OTM 2009, Part I.

[ELU10] Eberle, Hanna; Leymann, Frank; Unger, Tobias: Transactional Process Fragments - Recovery Strategies for Flexible Workflows with Process Fragments. In: Proceedings of APSCC 2010

[Eberle10] Eberle, Hanna; Leymann, Frank; Schleicher, Daniel; Schumm, David; Unger, Tobias: Process Fragment Composition Operations. In: Proceedings of APSCC 2010.

[Unger10a] Unger, Tobias; Eberle, Hanna; Leymann, Frank; Wagner, Sebastian: An Event-model for Constraint-based Person-centric Flows. In: Proceedings of the 2010 International Conference on Progress in Informatics and Computing (PIC-2010).

[Unger10b] Unger, Tobias; Eberle, Hanna; Leymann, Frank: Research challenges on person-centric flows. In: Gierds, Christian (Hrsg); Sürmeli, Jan (Hrsg): Proceedings of the 2nd Central-European Workshop on Services and their Composition, ZEUS 2010, Berlin, Germany, February 25--26, 2010.

[UR10] Unger, Tobias; Roller, Dieter: Applying Processes for User-driven Refinement of People Activities. In: Proceedings of the 14th IEEE International EDOC Conference (EDOC 2010).

[Agre97] P. E. Agre, Computation and Human Experience (Learning in Doing: Social, Cognitive and Computational Perspectives). Cambridge University Press, 1997.

[Dix97] A. Dix, Closing the Loop : Modelling action , perception and information, in AVI™96 - Advanced Visual Interfaces, 1997, pp. 20-28.

[Favela09] J. Favela, M. Tentori, D. Segura, and G. Berzunza, Adaptive awareness of hospital patient information through multiple sentient displays, International Journal of Ambient Computing and Intelligence, vol. 1, no. 1, pp. 27-38, 2009.

[Hutchins95] E. Hutchins, Cognition in the wild. MIT Press, 1995.

[Johnston05] R. B. Johnston, V. Waller, and S. K. Milton, Situated information systems for supporting routine activity in organisations, Int. J. Bus. Inf. Syst., vol. 1, no. 1/2, pp. 53-82, 2005.

[Johnston02] R. Johnston and S. Milton, The Foundation Role for Theories of Agency in Understanding Information Systems Design, Australasian Journal of Information Systems, vol. 10, no. 1, 2002.

[Johnston01] R. Johnston and S. Milton, The Significance of Intentionality for the Ontological Evaluation of Information Systems, in AMCIS 2001 Proceedings, 2001.

[Lederman11] R. Lederman and R. B. Johnston, Decision support or support for situated choice: lessons for system design from effective manual systems, European Journal of Information Systems, vol. 20, no. 5, pp. 510-528, Apr. 2011.

[Suchman87] L. Suchman, Plans and situated actions: The problem of human machine communication, Cambridge, Cambridge University Press, 1987.

[Waller09] V. Waller, Information systems in the wild: supporting activity in the world, Behaviour & Information Technology, vol. 28, no. 6, pp. 577-588, Nov. 2009.

[Waller09b] V. Waller and R. B. Johnston, Making ubiquitous computing available, Communications of the ACM, vol. 52, no. 10, p. 127, 2009.

[Hiesinger2011] Christian Hiesinger, Daniel Fischer, Stefan Foell, Klaus Herrmann, and Kurt Rothermel. Minimizing Human Interaction Time in Workflows. In Proceedings of the Sixth International Conference on Internet and Web Applications and Services (ICIW 2011), pages 22–28, March 2011.

[Fischer2011] Daniel Fischer, Stefan Föll, Klaus Herrmann, and Kurt Rothermel. Energy-efficient Workflow Distribution. In Proceedings of The Fifth International Conference on COMmunication System softWAre and middleware, COMSWARE, 2011

[Föll 2011] Stefan Föll, Klaus Herrmann, and Kurt Rothermel. Precon - Expressive context prediction using stochastic model checking. In Proceedings of the 8th International Conference on Ubiquitous Intelligence and Computing (UIC-2011), 2011.

[Föll 2010] Stefan Föll, Klaus Herrmann, and Christian Hiesinger. Flow-Based Context Prediction. In Proceedings of the 7th International Conference on Pervasive Services (ICPS 2010), Berlin, Germany, July 13-15, 2010. ACM, Juli 2010.

[Wolf 2010] Hannes Wolf, Klaus Herrmann, and Kurt Rothermel. Robustness in Context-Aware Mobile Computing. IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob'2010), 2010.

[Wolf 2011] Hannes Wolf, Jonas Palauro, and Klaus Herrmann. Fuzzy Event Assignment for Robust Context-Aware Workflows. In Proceedings of the Fourth International Conference on Dependability (DEPEND 2011), 2011.

8.2 Potential Impact

The final result of the project is a set of technologies enabling a new programming paradigm for adaptable pervasive applications based on the concept of flows as described above. The potential impact can be far-reaching. The provision of a solid conceptual foundation for human-centric, adaptive, pervasive applications is expected to have an impact in the following areas:

- **Strategic and economical:** The provision of a respective programming paradigm could open new application areas were the burden of manually configuring and interacting with a pervasive application were not acceptable before (e.g. health care and logistics). This can open new markets and help position Europe in terms of industrial developments and scientific research.

- **Social:** New application areas include, for example, home care and were severe social (as well as economical) problems could be mitigated by introducing a technology like ALLOW. If the technology functions without a need for direct intervention by the human user (in this case the elderly or disabled), this would allow for a better integration of these people into society.

Undoubtedly, information and communication technologies (ICT) proved to have a continuously increasing influence on the economy and the society in the past. Thereby, ICT has turned out to be a key tool for adjusting to changes at accelerating speeds and realizing higher economic growth.

Furthermore, almost every facet of our daily life has been touched by ICT already, and the number of embedded systems integrated into our surrounding environment is increasing day by day. The ALLOW project builds upon this trend in order to provide a solid technical foundation for the creation of human-oriented pervasive applications. These applications will provide seamless, natural, and intuitive support for many everyday tasks of people. Since such pervasive applications are frequently considered to form the next generation of ICT, there is no doubt that they will have significant impacts on both economy and society. The following details the most relevant economic and social impacts that can be expected from the technology enabled by the ALLOW project.

8.2.1 Economic Impacts

The technology enabled by the ALLOW project creates a number of economic opportunities on a European scale and on an individual and institutional level. In summary, the technology can strengthen the embedded systems industry as well as the software and services industry on a European scale. Furthermore, the ALLOW project can also help to establish a leading position for Europe in the realization of the next generation of computing systems and applications. On the individual and institutional level, it can help to provide thorough support for diverse business and work processes and, thus, it can reduce production and delivery cost of goods and services. Finally, the technology enabled by the ALLOW project can reduce the maintenance cost of future ICT, thereby allowing especially small and medium-sized enterprises (SME) to benefit from its advances.

Impacts on the European scale

In the past, the European economy was able to establish a very successful and strong electronic and embedded systems industry that contributed significantly to its wealth. At the present time, the worldwide electronics and embedded systems industry alone has an estimated market volume of approximately 250 billion Euros per annum. In addition to its high volume, the overall market is growing extremely fast, with average growth rates of more than 25% over the last years. In order to maintain a leading position in the area of electronics and embedded systems, it is very important for Europe to continuously maintain and improve its innovation capabilities.

The human-oriented pervasive applications enabled by the ALLOW project have the potential to serve as a sustainable stimulus for new innovations in this industry. As a result, they can help to strengthen Europe’s current leading position. This potential results from a combination of three different characteristics of the envisioned applications. Firstly, they will rely heavily on networked electronic and embedded systems that are invisibly integrated into many everyday objects. Secondly, the applications will be broadly applicable to various application domains as explained earlier. Thirdly, they will exhibit a comparatively high acceptance by users, since they will automatically adapt to changing user goals and they will support individual user tasks in a secure and reliable manner.

A similar argument can be made with respect to the software and services industry. In Europe, the ICT-related markets account for up to 8% of the gross domestic product (GDP) as well as for up to 6% of employment. With its estimated 20% share on ICT, the European software and services industry is responsible for a relevant part of Europe’s economic capacity. However, the software and services market is a fast-paced global market. Thus, in order to maintain Europe's current economic position, it is important to strengthen this industry as well.

The technology enabled by the ALLOW project has a great potential to make a significant contribution to the strength of the software and services industry, mainly for the following two reasons. 1. The human-oriented applications enabled by the ALLOW project will be based on the programming paradigm of Adaptable Pervasive Flows. In order to adapt to various environments, flow-based applications will rely on the presence of various services that can be composed in order to support a specific user task. This will require additional innovative services to develop the large number of novel application areas that can be supported through flow-based applications. 2. Flow-based applications will also stimulate a dramatic increase of the overall number of services utilized by companies as well as individuals which will broaden the overall market.

Finally, in order to maintain and improve Europe's market position with respect to the ICT industry as a whole, it is important to embrace and to control new technologies that bear significant economic potentials. As pervasive applications are frequently considered to form the next wave of computing applications, it is necessary to establish a leading position for Europe in this domain. Fundamental enabling technologies like Adaptable Pervasive Flows that exhibit a broad applicability and can stimulate a large number of different application areas, such a leadership position can be established at the early stages. In the fast-moving ICT world, an early technological leadership is usually crucial since the introduction of innovative technology quickly creates de facto industrial standards that are controlled by their inventors.

Individual and Institutional Impacts

Besides the previously mentioned economic impacts on the European scale, the technology enabled by the ALLOW project may also foster a number of positive effects on the individual and institutional level. Since flow-based pervasive applications can adapt automatically to the individual tasks of the persons working in them, they can facilitate diverse business and work processes. This enables users to perform their tasks more efficiently and it can minimize their cognitive load by reducing the amount of manual interaction. Furthermore, since this technology is able to automate a variety of tasks completely, its application can reduce the cost of many business and work processes. This might eventually reduce the overall production costs of goods by increasing the productivity of the involved work-force. Such productivity gains are particularly important for European enterprises since the cost of labor in Europe is relatively high, especially when compared to emerging economies like China and India. This potential cost reduction and productivity increase is additionally amplified by the fact that the proactive adaptation as enabled by Adaptable Pervasive Flows can significantly reduce associated management overheads.

Flow-based pervasive applications can also ease and improve the cooperation of different enterprises. Individual flows can compose spontaneously, and they may interact on a larger scale. Towards this end, a primary focus of the ALLOW project is the development of fundamental concepts to enable the composition of larger flows that are forming applications from smaller ones. Such composed flows can further the spontaneous goal-directed integration of diverse business and work processes. This significantly increases the flexibility of production chains. The importance of a higher flexibility is underlined by the fact that the production processes of many goods are becoming more complex and most goods must be produced in cooperation. Thus, this might potentiality improve the competitiveness of a broad variety of enterprises.

Especially the area of Small and Medium-sized Enterprises (SMEs) can benefit from the technology of Adaptable Pervasive Flows. Many small and medium-sized enterprises (SME) are not able to exploit technological advances that are targeted towards large-scale enterprises or that impose high financial hurdles. Flow-based technology can significantly reduce the management cost and maintenance overhead of future ICT applications. Since the technology can also be utilized to support small-scale business and work processes using low-cost devices, flow-based applications will not require high investments. Furthermore, by evolving automatically over time, they can significantly reduce the maintenance overhead introduced by today's applications.

8.2.2 Social Impacts

In addition to the economic impacts described above, the applications facilitated by the ALLOW project can also have significant social impacts. For instance, the developed technology can help in reducing the cost of Europe's health-care systems. Furthermore, it can also help to improve the support of elderly and disabled citizens in their own home environments while reducing the resulting supportive cost. Finally, this technology and its applications have the potential to reduce the digital divide resulting from today's ICT, and it can lead to an improved safety of today's working environments.

Impacts on health care

At the present time, the expenditures on health care in the member countries of the European community account for an estimated 8.5% of the GDP and 9.3% of employment. Current estimates indicate that the cost of health care could rise up to 11.8% in 2030 even if the current growth rate levels off. This imposes an enormous stress on the European welfare systems in general. In addition to that, some European countries, such as Germany, are already experiencing a shortage of trained personnel in this sector. If this trend cannot be reversed, it will eventually result in a significant loss in quality of the medical care that is available to a majority of European citizens.

The technology and the applications enabled by the ALLOW project can help to reduce the cost of health care by automating many routine work processes. Such an automation would improve the efficiency of the processes which would in turn allow doctors, for instance, to spend more time with complex medical tasks such as diagnosis and treatment and spend less time doing administrative work or documentation. Due to the proactive and automatic adaptation enabled by the paradigm of Adaptable Pervasive Flows, pervasive applications can also reduce the duration of hospitalizations. Human-oriented pervasive applications can, for instance, automatically monitor the health conditions of patients at home or they can ensure that patients are applying appropriate dosages of medicine.

Impacts on elderly and disabled citizens

In addition to the issues resulting from increasing health care expenditures, the European welfare systems are also challenged by the fact that the society is rapidly growing older. According to current estimations, the average human life expectancy will experience an increase by almost five years between 2000 and 2050. Between 1998 and 2025 the part of the population classified as elderly will increase from 20% to 28%. The fastest growing group of the elderly population is the group of people aged over 80. This, in turn, will not only increase the cost of health care but it will also lead to additional costs for appropriate nursing homes.

The technology explored by the ALLOW project can mitigate such problems by enabling assistive applications that allow elderly citizens to stay in their own home environments for a longer period of time. A flow-based application can, for instance, be employed to coordinate home care activities conducted by different personnel such that the care is always tailored to the elderly person. Moreover, if the person's home is equipped with pervasive computing devices (probably common in future home entertainment systems, bathroom and kitchen devices) the automatic adaptation of this environment to the person's activities could enable many elderly people to lead an independent life much longer than today. This may offer a great relief for them and improve the overall attendance for elderly citizens. The same argument also applies to disabled people who could benefit from a higher degree of independence which could improve their overall quality of life.

Impacts on the digital divide

Another social challenge that is rapidly gaining importance for the European society is the digital divide, i.e. the fact that the benefits of many novel ICT technologies can only be leveraged by technically versed persons with comparatively high incomes. However, in order to maintain equal opportunities in Europe, it is necessary to provide a broad range of people with access to all services that support them in their daily life.

Pervasive applications that are built on top of the paradigm of Adaptable Pervasive Flows are a novel and supportive technology that can be made available to virtually all social classes. Since they can also be executed by inexpensive embedded systems, they do not necessitate high investments and, thus, this technology can also be used by citizens with lower incomes. Similarly, the envisioned applications do not require skilled or trained users that possess experience with computer systems as they are able to dynamically adapt to the needs and intended goals of users. As a result, flow-based technology will be able to reduce the growing digital divide.

Impacts on work environments

Another social impact that can be expected from the flow-based technology enabled by the ALLOW project concerns the safety of work environments. 4.7 million accidents happen in Europe per annum in the work environment. 5500 of these accidents are fatal. Besides the costs of human lives, it is estimated that work-related accidents are resulting in a loss of approximately 150 million working days per year. Thus, improving the prevention of accidents in working environments will save a considerable number of lives and it will also preserve precious working power. In addition to that, reducing the number of accidents will also reduce the cost for the medical treatment of injured workers. Many work-related accidents are caused by the fact that workers fail to recognize dangerous situations as a result of their rigidly defined environments.

Flow-based pervasive applications can help to identify such problems, and they can even adapt the environment to the workers instead of vice versa. Since the individual actions performed by a worker are part of a larger overarching work process, a flow-based application may be able to detect a dangerous situation earlier than a worker that focuses solely on the action at hand. Additionally, flow-based applications can exceed the horizon of one individual flow by cooperating with each other. As an example, consider two flows that are executing in parallel. If their ongoing execution would lead to a potentially dangerous situation, their underlying engines may detect this on the basis of their explicitly modeled states and actions. In order to mitigate this, they could either postpone the execution of one flow or they could even re-plan one of the flows such that the situation will not occur. Considering this, the flow-based pervasive applications enabled by the ALLOW project have the potential to reduce the number of work-related accidents leading to a higher productivity and lower health care costs.

8.3 Project Website and Contact Information

8.3.1 Project website

http://www.allow-project.eu/

8.3.2 Contact details

Prof. Dr. Kurt Rothermel (Project Manager)
Distributed Systems Group, IPVS University of Stuttgart
Email: kurt.rothermel@ipvs.uni-stuttgart.de
Phone: +49 711 7816 448
Fax: +49 711 7816 424

Dr. Klaus Herrmann (Coordination)
Distributed Systems Group, IPVS University of Stuttgart
Email: klaus.herrmann@ipvs.uni-stuttgart.de
Phone: +49 711 7816 425
Fax: +49 711 7816 424

9 Use and dissemination of foreground

Section A

The dissemination and exploitation plans aim at ensuring that the knowledge and results generated within the project are spread to all the relevant communities and can be taken up beyond the project.

In general, the project aims at disseminating results through publications, presentations, participations in scientific or industrial events, as well as organization of specific events, like workshops and demonstrations, and dissemination through joint projects with partners outside the consortium. The dissemination plan targets three kinds of audience: research, industry, and general public.

Dissemination to the scientific community

- *Publications in scientific journals and conferences* in the areas of service oriented software, workflows, pervasive computing, planning and reasoning, automated learning, contexts and

user interfaces, adaptation and self-organization in computing systems (see Table A1 and Table A2).
- *Organization of dedicated dissemination events*, such as workshops co-located with relevant conferences. These events have the objective to disseminate and demonstrate the results of ALLOW, as well as to compare ALLOW's results with related research results outside of the project (see Table A2).
- *Research partnerships*: ALLOW will disseminate the project results through research partnerships and, in particular, through the exploitation of the project results in other research projects. This includes bilateral collaborations between ALLOW and other existing projects, as well as the launch of new follow-up projects.
- *Website*: The ALLOW website contains a specific section dedicated to the scientific community, containing a description of the scientific results and research challenges, and a repository of relevant articles.

Dissemination and technology transfer towards industry
- *Demonstrations and open source code*: demonstrators based on the use cases addressed by ALLOW are exploited to showcase the project results to different audiences, but specifically to industry. Also, the tools developed in ALLOW are made available as open source to the industry and the research community.
- *Organization of workshops* dedicated to bring together researchers and practitioners from academy and industry to stimulate transfer of ideas and results (see Table A2).
- *Outreach actions towards specific companies*. Partners in the consortium have well-established collaborations with companies that are interested in the project results. These links are exploited, both at the partner level and at the consortium level, to disseminate and promote the project results, and to pave the way for possible adoptions of these results.
- *Website*: The ALLOW website will contain a specific section dedicated to Innovation & Technology Transfer, including a collection of the case studies, demonstrators, and open source resources.

Outreach the general public
- *Web site:* The ALLOW website will contain a specific section for the general public, with appropriate material, using adequate terminology in order to make the objective and results understandable to a large audience.
- *A brochure* will highlighting the ALLOW project and its objectives, as well as a CD presenting the results of the project. This will be an excellent means for disseminating information to the general public and the scientific community, e.g. at conferences and similar events.

TABLE A1: LIST OF SCIENTIFIC (PEER REVIEWED) PUBLICATIONS, STARTING WITH THE MOST IMPORTANT ONES										
NO.	Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers[2] (if available)	Is/Will open access[3] provided to this publication?
1	Horizontal and Vertical Combination of Multi-tenancy Patterns in Service-Oriented Applications	Ralph Mietzner, Frank Leymann, and, Tobias Unger	Enterprise Information Systems	Volume 5, Issue 1, February 2011	Taylor & Francis,	PA, USA	2011	pp. 59-77	http://portal.acm.org/citation.cfm?id=1923634	no
2	Smart Objects as Building Blocks for the Internet of Things	Gerd Kortuem, Fahim Kawsar, Vasughi Sundramoorth, and Daniel Fitton	IEEE Internet Computing, Special Issue of Internet of Things, 2010	Volume 14, Issue 1	IEEE		2010	pp. 44-51	http://ieeexplore.ieee.org/xpl/topAccessedArticles.jsp?punumber=4236	no
3	Customer-defined service level agreements for composite applications	Tobias Unger, Ralph Mietzner, and Frank Leymann	Enterprise Information Systems	Volume 3, Issue 3, 2009	Taylor & Francis,	PA, USA	2009	pp. 369-391	http://portal.acm.org/citation.cfm?id=1593602	no
4	Dynamic Composition of Pervasive Process Fragments	Adina Sirbu, Annapaola Marconi, Marco Pistore, Hanna Eberle, Frank	ICWS 2011		IEEE		2011	pp. 73-80		no

[2] A permanent identifier should be a persistent link to the published version full text if open access or abstract if article is pay per view) or to the final manuscript accepted for publication (link to article in repository).
[3] Open Access is defined as free of charge access for anyone via Internet. Please answer "yes" if the open access to the publication is already established and also if the embargo period for open access is not yet over but you intend to establish open access afterwards.

		Leymann and Tobias Unger								
5	How Much is Too Much: Exploring the Design Space for Situated Glyphs to Support Dynamic Work Environments	Fahim Kawsar, Jo Vermeulen, Kevin Smith, Kris Luyten and Gerd Kortuem	PERVASIVE 2011	Volume 6696	Springer, LNCS	San Francisco, CA, USA	2011	pp. 70-78	http://www.springerlink.com/content/d0k1w44175271h72/	no
6	Enforcement from the Inside: Improving Quality of Business in Process Management	H. Eberle, S. Föll, K. Herrmann, F. Leymann, A. Marconi, T. Unger, and H. Wolf	ICWS 2009		IEEE		2009	pp. 405-412	http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5175850	no
7	Enabling Adaptation of Pervasive Flows: Built-in Contextual Adaptation	Annapaola Marconi, Marco Pistore, Adina Sirbu, Frank Leymann, Hanna Eberle, and Tobias Unger	ICSOC 2009	Volume 5900	Springer, LNCS	Berlin, Heidelberg	2009	pp. 445-454	http://www.springerlink.com/content/8qn3xq57n3787ru4/	no
8	An Explorative Comparison of Magic Lens and Personal Projection for Interacting with Smart Objects	Fahim Kawsar, Enrico Rukzio, and Gerd Kortuem	MobileHCI 2010		ACM	New York, NY, USA	2010	pp. 157-160	http://portal.acm.org/citation.cfm?id=1851627	no
9	FlexCon – Robust Context Handling in Human-oriented Pervasive Flows	Hannes Wolf, Klaus Herrmann, Kurt Rothermel	CoopIS 2011		Springer, LNCS		2011			no
10	Robustness in Context-Aware Mobile Computing	Hannes Wolf, Klaus Herrmann, and Kurt Rothermel	WiMob 2010		IEEE		2010	pp. 46 - 53	http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5645026	no
11	Flow-Based Context Prediction	Stefan Föll, Klaus Herrmann, and Christian	ICPS 2010							no

		Hiesinger								
12	The Subprocess Spectrum	Oliver Kopp, Hanna Eberle, Frank Leymann and Tobias Unger	BPSC 2010	LNI 177	GI-Edition: Lecture Notes in Informatics	Leipzig	2010	pp. 267-279		no
13	Having Services "YourWay!" : Towards User-Centric Composition of Mobile Service	R. Kazhamiakin, P. Bertoli, M. Paolucci, M. Pistore, and M. Wagner	FIS 2009	Volume 5468	Springer, LNCS		2009	pp. 94-106	http://www.springerlink.com/content/6064607177472713/	no
14	Combining Different Multi-Tenancy Patterns in Service-Oriented Applications	Ralph Mietzner, Tobias Unger, Robert Tietze, and Frank Leymann	EDOC 2009		IEEE	Piscataway, NJ, USA	2009	pp. 131-140	http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5277698	no
15	Cafe: A Generic Configurable Customizable Composite Cloud Application Framework	Ralph Mietzner, Tobias Unger, and Frank Leymann	CoopIS 2009	Volume 5870	Springer, LNCS		2009	pp. 357-364	http://www.springerlink.com/content/49514881457uh24k/	no
16	Process Fragments	Hanna Eberle, Tobias Unger, and Frank Leymann	CoopIS 2009	Volume 5870	Springer, LNCS		2009	pp. 398-405	http://www.springerlink.com/content/w30w71ur7040p383/	no
17	Minimizing Human Interaction Time in Workflows	Christian Hiesinger, Daniel Fischer, Stefan Föll, Klaus Herrmann and Kurt Rothermel	ICIW 2011				2011			no
18	Supporting Interaction with the Internet of Things across Objects, Time and Space	Fahim Kawsar, Gerd Kortuem and Bashar	IoT 2010		IEEE	Tokyo	2010	pp. 1-8	http://dx.doi.org/10.1109/IOT.2010.5678441	no

		Altakrouri								
19	Energy-efficient Workflow Distribution	Daniel Fischer, Stefan Föll, Klaus Herrmann and Kurt Rothermel	COMSWARE 2011				2011			no
20	Fuzzy Event Assignment for Robust Context Aware Workflows	Hannes Wolf, Jonas Palauro, and Klaus Herrmann	DEPEND 2011				2011			no
21	Exploring the Design of Pay-Per-Use Objects in the Construction Domain	D. Fitton, V. Sundramoorthy, G. Kortuem, J. Brown, and C. Estratiou.	EuroSSC 2008		Springer-Verlag		2008	pp. 192-205	http://dx.doi.org/10.1007/978-3-540-88793-5_15	no
22	PreCon: Expressive Context Prediction using Stochastic Model Checking	Stefan Föll, Klaus Herrmann and Kurt Rothermel	UIC 2011				2011			no
23	Teleo-Reactive Policies for Managing Human-centric Pervasive Services	Srdjan Marinovic, Kevin Twidle, Naranker Dulay and Morris Sloman	CNSM 2010		IEEE		2010	pp. 80-87	http://dx.doi.org/10.1109/CNSM.2010.5691332	no
24	Rumpole: A Flexible Break-glass Access Control Model	Srdjan Marinovic, Robert Craven, Jiefei Ma, and Naranker Dulay	SACMAT 2011		ACM		2011	pp. 73-82	http://dx.doi.org/10.1145/1998441.1998453	no
25	Shinren: Non-monotonic Trust Management for Distributed Systems	Changyu Dong, and Naranker Dulay	IFIPTM 2010	Vol. 321/2010	SpringerLink		2010	pp. 125-140	http://dx.doi.org/10.1007/978-3-642-13446-3	yes
26	A Framework for Rule-based Dynamic Adaptation	Ivan Lanese, Antonio Bucchiarone, and Fabrizio	TGC 2010	Volume 6084	Springer, LNCS		2010	pp. 284–300	http://www.springerlink.com/content/c036j86588491148/	no

		Montesi								
27	Distributed Orchestration of Pervasive Services	Leonardo Mostarda, Srdjan Marinovic, and Naranker Dulay	AINA 2010		IEEE		2010		http://doi.ieeecomputersociety.org/10.1109/AINA.2010.100	no
28	Collaborative Privacy Policy Authoring in a Social Networking Context	Ryan Wishart, Domenico Corapi, Srdjan Marinovic, and Morris Sloman	POLICY 2010		IEEE		2010	pp .1-8	http://dx.doi.org/10.1109/POLICY.2010.13	no
29	Teleo-Reactive Policies in Ponder2	Kevin Twidle, Srdjan Marinovic, and Naranker Dulay	POLICY 2010		IEEE		2010			no
30	Transactional Process Fragments - Recovery Strategies for Flexible Workflows with Process Fragments	Hanne Eberle, Frank Leymann, and Tobias Unger	APSCC 2010		IEEE		2010	pp. 250-257	http://dx.doi.org/10.1109/APSCC.2010.73	no
31	Process Fragment Composition Operations	Hanne Eberle, Frank Leymann, Daniel Schleicher, David Schumm, and Tobias Unger	APSCC 2010		IEEE		2010	pp. 157-163	http://dx.doi.org/10.1109/APSCC.2010.72	no
32	An Event-model for Constraint-based Person-centric Flows	Tobias Unger, Hanne Eberle, Frank Leymann, and Sebastian Wagner	PIC 2010		IEEE		2010	pp. 927 - 932	http://dx.doi.org/10.1109/PIC.2010.5687886	no
33	Implementation Architectures for Adaptive Workflow Management	Hanne Eberle, Frank Leymann, and	ADAPTIVE 2010				2010			no

		Tobias Unger									
34	An Overall Process for Self-Adaptive Pervasive Systems	Antonio Bucchiarone, Annapaola Marconi, Marco Pistore, Stefan Föll, Klaus Herrmann, Christian Hiesinger, and Srdjan Marinovic	ADAPTIVE 2010				2010				no
35	The benefit of activity recognition for mobile phone based nursing documentation: A Wizard-of-Oz study	Agnes Grunerbl, Paul Lukowicz, Kunze Kai, Bashar Altakouri, Gerd Kortuem	ISWC 2010		IEEE	Seoul	2010		http://dx.doi.org/10.1109/ISWC.2010.5665866		no
36	Adaptable Pervasive Flows - An Emerging Technology for Pervasive Adaptation	Herrmann, Kurt Rothermel, Gerd Kortuem, and Naranker Dulay	SASO 2008		IEEE		2008				no
37	A Context-Aware Framework for Business Processes Evolution	Antonio Bucchiarone, Annapaola Marconi, Marco Pistore and Adina Sirbu	EVL-BP 2011		IEEE		2011				no
38	Teleo-Reactive Workflows for Pervasive Healthcare	Srdjan Marinovic, Kevin Twidle, and Naranker Dulay	PerHealth 2010		IEEE		2010	pp. 316-321	http://dx.doi.org/10.1109/PERCOMW.2010.5470648		no
39	Modeling Dynamic Context Awareness for Situated	Hannes Wolf, Klaus	5th International				2009				no

	Workflows	Herrmann, and Kurt Rothermel	Workshop on Context Aware Mobile Systems, 2009.							
40	Retry Scopes to Enable Robust Workflow Execution in Pervasive Environments	Hanna Eberle, Oliver Kopp, Tobias Unger, and Fank Leymann	MONA+ 2009	Volume 6275	Springer, LNCS		2009	pp. 258-369	http://dx.doi.org/10.1007/978-3-642-16132-2_34	no
41	A formalisation of Adaptable Pervasive Flows	Antonio Bucchiarone, Alberto Lluch Lafuente, Annapaola Marconi, and Marco Pistore	WSFM 2009	Volume 6194	Springer, LNCS		2009	pp. 61–75	http://www.springerlink.com/content/14660rt24k442751/fulltext.pdf	no
42	Exploring the Application of Smart Objects in the Workplace	Daniel Fitton, Fahim Kawsar and Gerd Kortuem	Adjunct Proceedings UbiComp 2009		ACM		2009		http://www.ubicomp.org/ubicomp2009/programsDemos.shtml	no
43	Designing Interactions for Smart Objects with Flows	Fahim Kawsar and Gerd Kortuem	3rd International Workshop on Design and Integration Principle for Smart Objects (DIPSO 2009)		online proceedings		2009		http://eis.comp.lancs.ac.uk/workshops/dipso/dipso2009/	no
44	Exploring The Design of a Memory Model for Smart Objects	Daniel Fitton, Fahim Kawsar and Gerd Kortuem	Ambient Intelligence and Smart Environment	Vol 4	IOS Press		2009	pp. 33-38	http://dx.doi.org/10.3233/978-1-60750-056-8-33	no
45	Combining Crowd-based Sensing, Micro blogging and Activity Flow Models: A case study using soccer games	Kai Kunze and Paul Lukowicz	Pervasive 2011				2011			no

46	On BPMN Process Fragment Auto-Completion	Oliver Kopp, Frank Leymann, David Schumm, and Tobias Unger	ZEUS 2011	705	CEUR-WS.org, CEUR Workshop Proceedings		2011	pp. 58-64	http://ceur-ws.org/Vol-705/paper8.pdf	yes
47	Towards The Essential Flow Model	Oliver Kopp, Frank Leymann, Tobias Unger, and Sebastian Wagner	ZEUS 2011	705	CEUR-WS.org, CEUR Workshop Proceedings		2011	pp. 26-33	http://ceur-ws.org/Vol-705/paper4.pdf	yes
48	On Designing a People-oriented Constraint-based Workflow Language	Frank Leymann, Tobias Unger, and Sebastian Wagner	ZEUS 2010	563	CEUR-WS.org, CEUR Workshop Proceedings		2010	pp. 25-31	http://CEUR-WS.org/Vol-563/paper3.pdf	yes
49	Research Challenges on Person-centric Flows	Tobias Unger, Hanna Eberle, and Frank Leymann	ZEUS 2010	563	CEUR-WS.org, CEUR Workshop Proceedings		2010	pp. 97-104	http://CEUR-WS.org/Vol-563/paper12.pdf	yes
50	Collaboration Aspects of Human Tasks	Tobias Unger and Sebastian Wagner	CECPAW 2010	66	Springer, Lecture Notes in Business Information Processing		2010	pp. 579-590	http://dx.doi.org/10.1007/978-3-642-20511-8_53	no
51	Flow Driven Ambient Guidance	Gerd Kortuem, Fahim Kawsar, and Bashar Altakrouri	8th IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops),		IEEE	Mannheim	2010	pp. 796-799	http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5470544	no

			2010							
52	Designing Pervasive Interactions for Ambient Guidance with Situated Flows	Fahim Kawsar, Gerd Kortuem, and Bashar Altakrouri	4th International Workshop on Human Aspects in Ambient Intelligence (2010) WI-IAT 2010		IEEE	Toronto	2010	pp. 371-375	http://dx.doi.org/10.1109/WI-IAT.2010.119	no
53	Spin&Swing : Spatial Interaction with Orientation Aware Devices	Bashar Altakrouri, Fahim Kawsar, and Gerd Kortuem	Adjunct Proceeding, Eighth International Conference on Pervasive Computing (Pervasive 2010)		IEEE	Helsinki	2010		http://eprints.lancs.ac.uk/42495/	yes

TABLE A2: LIST OF DISSEMINATION ACTIVITIES								
NO.	Type of activities[4]	Main leader	Title	Date	Place	Type of audience[5]	Size of audience	Countries addressed
1	Conference	FBK,IPVS	BPM Conference 2011 – Demonstration	28 August – 2 September 2011	Clermont-Ferrand, France	Research	>50	
2	Presentation	ULANC	Pervasive 2011 – Demonstration	12-15 June 2011	San Francisco	Research, Industry	>100	
3	Workshop	ULANC	Workshop on Intelligibility and Control in Pervasive Computing – Organisation	12-15 June 2011	San Francisco, CA, USA	Research	<20	
4	Workshop	PASSAU	First Workshop on Hybrid Pervasive/Digital Inference at Pervasive 2011 – Organization	12 June 2011	San Francisco, US	Research	<50	
5	Presentation	FBK, IPVS	FET 11 Exhibition – Demonstration	4-6 May 2011	Budapest	Research, Industry	>50	
6	Workshop	Imperial	Privacy for Mobile Systems	May 2011	London, UK	Industry	<20	

[4] A drop down list allows choosing the dissemination activity: publications, conferences, workshops, web, press releases, flyers, articles published in the popular press, videos, media briefings, presentations, exhibitions, thesis, interviews, films, TV clips, posters, Other.
[5] A drop down list allows choosing the type of public: Scientific Community (higher education, Research), Industry, Civil Society, Policy makers, Medias ('multiple choices' is possible.

			with companies from GSMA.					
7	Presentation	Gerd Kortuem (ULANC)	Bell Labs – Invited talk	14 February 2011	Antwerp	Research	<100	
8	Presentation	Kurt Rothermel (USTUTT (IPVS))	NEC – Invited Talk	2011	Heidelberg, Germany	Industry	<50	
9	Presentation	Gerd Kortuem (ULANC)	Tokyo University – Invited talk	29 November 2010	Tokyo	Research	<30	
10	Presentation	Fahim Kawsar (ULANC)	Keio University – Invited talk	29 November 2010	Tokyo	Research	<40	
11	Presentation	Gerd Kortuem (ULANC)	Keio University – Invited talk	29 November 2010	Tokyo	Research	<40	
12	Workshop	ULANC	Workshop Digital Object Memories in the Internet of Things – Organisation	26 September 2010	Copenhagen	Research	<12	
13	Presentation	Naranker Dulay (Imperial)	PerAda Workshop on Privacy, Security and Trust, Rome – Invited talk	Nov 2010	Rome, Italy	Research	< 50	
14	Presentation	Kevin Twidle, Srdjan Marinovic (Imperial)	POLICY 2010 – Demonstration	July 2010	Fairfax, USA.			
15	Workshop	FBK, USTUTT	ALLOW-BIBA Research Center, Bilateral Meeting	10 June 2010	Bremen, Germany	Research	<12	
16	Presentation	ULANC	Pervasive 2010	20 May 2010	Helsinki	Research,	>100	

			– Demonstration			industry		
17	Workshop	Kevin Twidle (IMPERIAL)	FET-PerAda Meeting	March 2010	Brussels, Belgium			
18	Workshop	ULANC	Demonstration and Meeting with Veraz Inc	2 Feb 2 2010	Lancaster	industry	<10	
19	Presentation	Kurt Rothermel (USTUTT (IPVS))	PerCom 2010 – Keynote	2010	Mannheim, Germany	Research	>50	
20	Conference		International Conference on Service Oriented Computing ICSOC 2009	23-27 November 2009	Stockholm, Sweden	Research, Industry	>50	
21	Workshop		ALLOW Scientific Advisory Board (SAB) Meeting	16 November 2009	Stuttgart, Germany			
22	Conference		International Conference on Cooperative Information Systems (CoopIS 2009)	4-6 November 2009	Vilamoura, Algarve, Portugal	Research	>50	
23	Conference		International Conference on Ubiquitous Computing (UbiComp 2009)	29 September – 3 October 2009.	Orlando, US	Research, Industry	>50	
24	Workshop		International Workshop on Web Services and Formal Methods (WSFM 2009)	4-5 September 2009	Bologna, Italy		<50	
25	Conference		IEEE	1-4 September	Auckland,	Research	>50	

			International EDOC Conference (ECOC 2009)	2009	New Zealand			
26	Conference		IEEE International Conference on Web Services (ICWS 2009)	6-10 July 2009	Los Angeles, US		>50	
27	Conference		International Conference on Wireless Applications and Computing 2009	17-22 June 2009	Algarve, Portugal		>50	
28	Presentation	Gerd Kortuem (ULANC)	International Conference on Wireless Applications and Computing 2009 – Invited talk	17- 22 June 2009	Algarve, Portugal		< 40	
29	Presentation	Klaus Hermann	PerAda-SICSA Summer School – Invited talk	June 2009	Edinburgh			
30	Conference	Gerd Kortuem (ULANC)	Mobile and Ubiquitous Computing Track, PC Euro-Par 2009 – Program Chair	2009	Delft, The Netherlands			
31	Workshop	ALL	ALLOW Application Workshop	10 June 2008	Frankfurt, Germany	Industru	<10	
32	Flyer	ALL	Project flyer realized and printed.	2008	-	Research, Industry		
33	Web	ALL	Project web site on-line:	2008	-	Research, Industry		

			http://www.allow-project.eu/					

Section B
Part B1

No applications for patents, trademarks, registered designs, etc. have been filed within the project.

TABLE B1: LIST OF APPLICATIONS FOR PATENTS, TRADEMARKS, REGISTERED DESIGNS, ETC.					
Type of IP Rights[6]:	Confidential Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Application reference(s) (e.g. EP123456)	Subject or title of application	Applicant (s) (as on the application)

[6] A drop down list allows choosing the type of IP rights: Patents, Trademarks, Registered designs, Utility models, Others.

Part B2

Type of Exploitable Foreground[7]	Description of exploitable foreground	Confidential Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Exploitable product(s) or measure(s)	Sector(s) of application[8]	Timetable, commercial or any other use	Patents or other IPR exploitation (licences)	Owner & Other Beneficiary(s) involved
General advancement of knowledge	Adaptation and evolution concept	YES	-	Techniques and tools	J63 - Information service activities	2010-2011	Open source	FBK (owner)
General advancement of knowledge	Beagle Security Language	No	-	Techniques and tools	J62.0.1 - Computer programming activities	2010-2011	Open source	Imperial (owner)
General advancement of knowledge	Teleo-Reactive Policy Framework	No	-	Techniques and tools	J62.0.1 - Computer programming activities	2010-2011	Open source	Imperial (owner)
General advancement of knowledge	Flow-driven adaptable interaction concept	YES	-	Techniques and tools	J63 - Information service activities	2010-2011	Open source	ULANC (owner)
General advancement of knowledge	Flexible scientific workflows modeling and execution	YES	-	Techniques and tools	J63 - Information service activities	2011-2012	Open source	USTUTT IAAS (owner)
General advancement of knowledge	Adaptable coordinated service compositions	YES	-	Techniques and tools	J63 - Information service activities	2011-2012	Open source	USTUTT IAAS (owner)
General advancement of knowledge	Context prediction concepts	YES	-	Concepts and mechanisms	J63 - Information service activities	2011-2013	-	USTUTT(IPVS) (owner)
General advancement of	Robust flow navigation	YES	-	Concepts and mechanisms	J63 - Information service activities	2011-2013	-	USTUTT(IPVS) (owner)

[19] A drop down list allows choosing the type of foreground: General advancement of knowledge, Commercial exploitation of R&D results, Exploitation of R&D results via standards, exploitation of results through EU policies, exploitation of results through (social) innovation.

[8] A drop down list allows choosing the type sector (NACE nomenclature) : http://ec.europa.eu/competition/mergers/cases/index/nace_all.html

Type of Exploitable Foreground[7]	Description of exploitable foreground	Confidenti al Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Exploitable product(s) or measure(s)	Sector(s) of application[8]	Timetable, commercial or any other use	Patents or other IPR exploitation (licences)	Owner & Other Beneficiary(s) involved
knowledge	principles							

Adaptation and evolution concept. This exploitation foreground facilitates the development of more agile and robust software systems, in particular following the Service Oriented Computing (SOC) paradigm. FBK has started exploiting this foreground within the EU project S-Cube, as a baseline technology for service adaptation and evolution; this project will guarantee further developments and wider testing of the concept developed within the ALLOW project. FBK is also evaluating the transfer of the techniques and tools developed within this foreground to a spin-off working in the area of SOC. The tools are protected by an open source license. Expected impact is on the quality of software, in particular for what concerns its robustness, its resilience, and its agility; a quantification is not possible at this stage.

Beagle Logic Programming Language. Represents a novel logic programming language that can be used to express inference over a set of discrete truth-values, as opposed to the classical one. It advances the state of the art in programming languages by its expressiveness. Imperial has developed this language within the ALLOW project to cope with security policies in distributed and pervasive environments where evidence is missing, conflicting or partial. The language can be exploited in AI for logic-based agents, and also in other fields of knowledge representation.

Teleo-Reactive Policy Framework. Within the ALLOW project we have built and implemented a new type goal-driven Teleo-Reactive policy framework and shown how such policies can be used for managing human-centric and human-driven processes. We have open-sourced our implementation on both desktop and mobile platforms. We believe that TR policies can be exploited in other management scenarios in pervasive domains.

Flow-driven adaptable interaction concept. This exploitation foreground facilitates the development of more flexible, context-aware workflow systems for highly dynamic human work environments. ULANC is starting to exploit this foreground for proposed EU projects. ULANC is also evaluating the applicability of these concepts for a planned cooperation with a local SME. The underlying software is protected by an open source license. Expected impact is on the quality and effectiveness of human work processes.

Flexible scientific workflow modeling and execution: Flexibility is a core characteristic of scientific workflows. It must be possible to adapt the workflow during execution; i.e. by changing properties of the workflow or by adding additional activities. On the other hand the workflow must react to changes in the execution environment; i.e. by adding additional computing resources. These kinds of flexibility can be implemented by the concept of flow fragments. IAAS has started to exploit the results of the ALLOW project in the area of flow fragments to support the flexible modeling and execution of scientific workflows within the Stuttgart Research Centre for Simulation Technology and Cluster of Excellence “Simulation Technology” (SimTech).

Adaptable Coordinated Service Compositions: Service-oriented computing (SOC) is based on the idea of composing self-contained services into a higher level (business) service in a loosely coupled manner. Service Composition Models are involved in a lifecycle, which comprises modeling, verification, execution, and monitoring phases. However, existing service composition models are very inflexible in the way that the composition cannot take place during runtime. To address this problem IAAS has started exploiting the concepts of flow fragments within the EU project S-Cube as technology for runtime services composition.

Context Prediction Concepts. Concepts for prediction context information pave the way towards proactive, adaptive software in general. The concepts developed in the project are an important foundation for the work in future scientific projects and cooperation with industry partners. Making context-aware systems proactive is a basic requirement in many sectors of communication and information systems both in academia and in the industry.

Robust flow navigation principles. Flow navigation today mostly assumes that the input data is correct. This technology helps dropping this assumption and enables flow-based systems to deal with uncertainties and errors. This is an important foreground that we will exploit in cooperation with our academic and industry partners in order to create more robust pervasive applications.

10 Report on societal implications

A General Information *(completed automatically when **Grant Agreement number** is entered.*

Grant Agreement Number:	213339
Title of Project:	Adaptable Pervasive Flows
Name and Title of Coordinator:	Prof. Dr. Kurt Rothermel

B Ethics

1. Did your project undergo an Ethics Review (and/or Screening)? • If Yes: have you described the progress of compliance with the relevant Ethics Review/Screening Requirements in the frame of the periodic/final project reports? Special Reminder: the progress of compliance with the Ethics Review/Screening Requirements should be described in the Period/Final Project Reports under the Section 3.2.2 *'Work Progress and Achievements'*	***No***
2. Please indicate whether your project involved any of the following issues (tick box) :	
RESEARCH ON HUMANS	
• Did the project involve children?	***No***
• Did the project involve patients?	***No***
• Did the project involve persons not able to give consent?	***No***
• Did the project involve adult healthy volunteers?	***Yes***
• Did the project involve Human genetic material?	***No***
• Did the project involve Human biological samples?	***No***
• Did the project involve Human data collection?	***Yes***
RESEARCH ON HUMAN EMBRYO/FOETUS	
• Did the project involve Human Embryos?	***No***
• Did the project involve Human Foetal Tissue / Cells?	***No***
• Did the project involve Human Embryonic Stem Cells (hESCs)?	***No***
• Did the project on human Embryonic Stem Cells involve cells in culture?	***No***
• Did the project on human Embryonic Stem Cells involve the derivation of cells from Embryos?	***No***
PRIVACY	
• Did the project involve processing of genetic information or personal data (eg. health, sexual lifestyle, ethnicity, political opinion, religious or philosophical conviction)?	***No***
• Did the project involve tracking the location or observation of people?	***Yes***
RESEARCH ON ANIMALS	
• Did the project involve research on animals?	***No***
• Were those animals transgenic small laboratory animals?	***No***
• Were those animals transgenic farm animals?	***No***
• Were those animals cloned farm animals?	***No***
• Were those animals non-human primates?	***No***
RESEARCH INVOLVING DEVELOPING COUNTRIES	
• Did the project involve the use of local resources (genetic, animal, plant etc)?	***No***

• Was the project of benefit to local community (capacity building, access to healthcare, education etc)?	***No***
DUAL USE	
• Research having direct military use	***No***
• Research having the potential for terrorist abuse	***No***

C Workforce Statistics

3. Workforce statistics for the project: Please indicate in the table below the number of people who worked on the project (on a headcount basis).

Type of Position	Number of Women	Number of Men
Scientific Coordinator	0	7
Work package leaders	0	5
Experienced researchers (i.e. PhD holders)	0	10
PhD Students	4	10
Other	0	0

4. How many additional researchers (in companies and universities) were recruited specifically for this project?	**4**
Of which, indicate the number of men:	3

D Gender Aspects

5. Did you carry out specific Gender Equality Actions under the project?	○ ✓	Yes No

6. Which of the following actions did you carry out and how effective were they?

		Not at all effective				Very effective
❑	Design and implement an equal opportunity policy	○	○	○	○	○
❑	Set targets to achieve a gender balance in the workforce	○	○	○	○	○
❑	Organise conferences and workshops on gender	○	○	○	○	○
❑	Actions to improve work-life balance	○	○	○	○	○
○	Other:	Each of the partner institutions has installed gender equality instruments independently of the project.				

7. Was there a gender dimension associated with the research content – i.e. wherever people were the focus of the research as, for example, consumers, users, patients or in trials, was the issue of gender considered and addressed?

- ○ Yes- please specify
- ✓ No

E Synergies with Science Education

8. Did your project involve working with students and/or school pupils (e.g. open days, participation in science festivals and events, prizes/competitions or joint projects)?

- ✓ Yes- please specify: University open days
- ○ No

9. Did the project generate any science education material (e.g. kits, websites, explanatory booklets, DVDs)?

- ○ Yes- please specify
- ✓ No

F Interdisciplinarity

10. Which disciplines (see list below) are involved in your project?

- ✓ Main discipline[9]: Computer Science
- ○ Associated discipline
- ○ Associated discipline:

G Engaging with Civil society and policy makers

11a Did your project engage with societal actors beyond the research community? *(if 'No', go to Question 14)*	○ ✓	Yes No

11b If yes, did you engage with citizens (citizens' panels / juries) or organised civil society (NGOs, patients' groups etc.)?

- ○ No
- ○ Yes- in determining what research should be performed
- ○ Yes - in implementing the research

[9] Insert number from list below (Frascati Manual).

○ Yes, in communicating /disseminating / using the results of the project		
11c In doing so, did your project involve actors whose role is mainly to organise the dialogue with citizens and organised civil society (e.g. professional mediator; communication company, science museums)?	○ ○	Yes No

12. Did you engage with government / public bodies or policy makers (including international organisations)

- ○ No
- ○ Yes- in framing the research agenda
- ○ Yes - in implementing the research agenda
- ○ Yes, in communicating /disseminating / using the results of the project

13a Will the project generate outputs (expertise or scientific advice) which could be used by policy makers?

- ○ Yes – as a **primary** objective (please indicate areas below- multiple answers possible)
- ○ Yes – as a **secondary** objective (please indicate areas below - multiple answer possible)
- ○ No

13b If Yes, in which fields?

Agriculture Audiovisual and Media Budget Competition Consumers Culture Customs Development Economic and Monetary Affairs Education, Training, Youth Employment and Social Affairs		Energy Enlargement Enterprise Environment External Relations External Trade Fisheries and Maritime Affairs Food Safety Foreign and Security Policy Fraud Humanitarian aid		Human rights Information Society Institutional affairs Internal Market Justice, freedom and security Public Health Regional Policy Research and Innovation Space Taxation Transport	

13c If Yes, at which level?

- ○ Local / regional levels
- ○ National level
- ○ European level
- ○ International level

H Use and dissemination

14. How many Articles were published/accepted for publication in peer-reviewed journals?	**3**
To how many of these is open access[10] provided?	
How many of these are published in open access journals?	
How many of these are published in open repositories?	
To how many of these is open access not provided?	**3**
Please check all applicable reasons for not providing open access:	
✓ publisher's licensing agreement would not permit publishing in a repository ❑ no suitable repository available ❑ no suitable open access journal available ❑ no funds available to publish in an open access journal ❑ lack of time and resources ❑ lack of information on open access ❑ other[11]: ……………	

15. How many new patent applications ('priority filings') have been made? *("Technologically unique": multiple applications for the same invention in different jurisdictions should be counted as just one application of grant).*	**0**

16. Indicate how many of the following Intellectual Property Rights were applied for (give number in each box).	Trademark	**0**
	Registered design	**0**
	Other	**0**

17. How many spin-off companies were created / are planned as a direct result of the project?	**0**
Indicate the approximate number of additional jobs in these companies:	

18. Please indicate whether your project has a potential impact on employment, in comparison with the situation before your project:

❑	Increase in employment, or	✓	In small & medium-sized enterprises
❑	Safeguard employment, or	❑	In large companies
❑	Decrease in employment,	❑	None of the above / not relevant to the project
✓	Difficult to estimate / not possible to quantify		

[10] Open Access is defined as free of charge access for anyone via Internet.
[11] For instance: classification for security project.

19. For your project partnership please estimate the employment effect resulting directly from your participation in Full Time Equivalent (*FTE = one person working fulltime for a year*) jobs:	*Indicate figure:*
Difficult to estimate / not possible to quantify	✓

I Media and Communication to the general public

20. As part of the project, were any of the beneficiaries professionals in communication or media relations?

○ Yes ✓ No

21. As part of the project, have any beneficiaries received professional media / communication training / advice to improve communication with the general public?

○ Yes ✓ No

22 Which of the following have been used to communicate information about your project to the general public, or have resulted from your project?

✓	Press Release	✓	Coverage in specialist press
❑	Media briefing	❑	Coverage in general (non-specialist) press
❑	TV coverage / report	❑	Coverage in national press
❑	Radio coverage / report	❑	Coverage in international press
✓	Brochures /posters / flyers	✓	Website for the general public / internet
❑	DVD /Film /Multimedia	✓	Event targeting general public (festival, conference, exhibition, science café)

23 In which languages are the information products for the general public produced?

❑	Language of the coordinator	✓	English
❑	Other language(s)		

***Question F-10*:** Classification of Scientific Disciplines according to the Frascati Manual 2002 (Proposed Standard Practice for Surveys on Research and Experimental Development, OECD 2002):

FIELDS OF SCIENCE AND TECHNOLOGY

1. NATURAL SCIENCES
1.1 Mathematics and computer sciences [mathematics and other allied fields: computer sciences and other allied subjects (software development only; hardware development should be classified in the engineering fields)]
1.2 Physical sciences (astronomy and space sciences, physics and other allied subjects)
1.3 Chemical sciences (chemistry, other allied subjects)
1.4 Earth and related environmental sciences (geology, geophysics, mineralogy, physical geography and other geosciences, meteorology and other atmospheric sciences including climatic research, oceanography, vulcanology, palaeoecology, other allied sciences)
1.5 Biological sciences (biology, botany, bacteriology, microbiology, zoology, entomology, genetics, biochemistry, biophysics, other allied sciences, excluding clinical and veterinary sciences)

2 ENGINEERING AND TECHNOLOGY

2.1 Civil engineering (architecture engineering, building science and engineering, construction engineering, municipal and structural engineering and other allied subjects)
2.2 Electrical engineering, electronics [electrical engineering, electronics, communication engineering and systems, computer engineering (hardware only) and other allied subjects]
2.3. Other engineering sciences (such as chemical, aeronautical and space, mechanical, metallurgical and materials engineering, and their specialised subdivisions; forest products; applied sciences such as geodesy, industrial chemistry, etc.; the science and technology of food production; specialised technologies of interdisciplinary fields, e.g. systems analysis, metallurgy, mining, textile technology and other applied subjects)

3. MEDICAL SCIENCES
3.1 Basic medicine (anatomy, cytology, physiology, genetics, pharmacy, pharmacology, toxicology, immunology and immunohaematology, clinical chemistry, clinical microbiology, pathology)
3.2 Clinical medicine (anaesthesiology, paediatrics, obstetrics and gynaecology, internal medicine, surgery, dentistry, neurology, psychiatry, radiology, therapeutics, otorhinolaryngology, ophthalmology)
3.3 Health sciences (public health services, social medicine, hygiene, nursing, epidemiology)

4. AGRICULTURAL SCIENCES
4.1 Agriculture, forestry, fisheries and allied sciences (agronomy, animal husbandry, fisheries, forestry, horticulture, other allied subjects)
4.2 Veterinary medicine

5. SOCIAL SCIENCES
5.1 Psychology
5.2 Economics
5.3 Educational sciences (education and training and other allied subjects)
5.4 Other social sciences [anthropology (social and cultural) and ethnology, demography, geography (human, economic and social), town and country planning, management, law, linguistics, political sciences, sociology, organisation and methods, miscellaneous social sciences and interdisciplinary , methodological and historical S1T activities relating to subjects in this group. Physical anthropology, physical geography and psychophysiology should normally be classified with the natural sciences].

6. HUMANITIES
6.1 History (history, prehistory and history, together with auxiliary historical disciplines such as archaeology, numismatics, palaeography, genealogy, etc.)
6.2 Languages and literature (ancient and modern)
6.3 Other humanities [philosophy (including the history of science and technology) arts, history of art, art criticism, painting, sculpture, musicology, dramatic art excluding artistic "research" of any kind, religion, theology, other fields and subjects pertaining to the humanities, methodological, historical and other S1T activities relating to the subjects in this group]