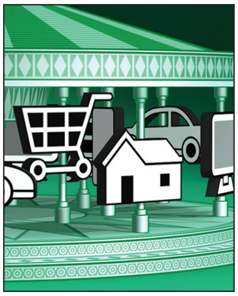

Programming Elasticity and Commitment in Dynamic Processes

Pablo Fernández • *Applied Software Engineering Group, University of Sevilla*

Hong-Linh Truong and Schahram Dustdar • *Distributed Systems Group, Vienna University of Technology*

Antonio Ruiz-Cortés • *Applied Software Engineering Group, University of Sevilla*

In the past, elasticity and commitment in business processes were underexplored. But as businesses increasingly exploit pay-per-use resources in the cloud for on-demand needs, elasticity and commitment have become important issues. Here, the authors discuss the value of using elastic resources and commitments to create more dynamic organizations that can easily balance the need to be adaptable and flexible, while also retaining a high level of manageability.

Availability, easy access to a huge number of software and human resources in the cloud (and from the crowds), and current market dynamics push organizations to scale efficiently (and save costs) but to be adaptable and meet on-demand customer needs. Such organizations search for intensive customizable solutions that could deliver specialized services while spotting and serving market opportunities through high adaptability. This leads to a clear tendency towards flexible process models that are highly reusable and adaptable, but with less manageability and control.[1,2]

A team-oriented, flexible process model[3] could potentially address these challenges for adaptable organizations. However, in this organizational shift, from the process perspective two goals collide:

- the need to be adaptable and manageable, and
- the fact that adaptation is boosted by keeping a flexible process system, but manageability usually is derived from a strict process environment.

One way to deal with this dilemma is to harmonize the concepts of *elasticity* and *commitments* among individuals, teams, and organizations as the key ingredients to support efficient management of resources, which in turn leads to more adaptable process models for organizations. Elasticity allows for dynamic on-demand changes, offering functions and associated costs and quality by leveraging existing resources onsite and on the cloud.[4] Commitment represents an explicit statement of settings (that is, the resources' compromised capabilities), objectives, and the compensation model over the potential outcomes delivered.[5] Obviously, both have a strong influence on composing and executing processes. Between them, commitment has a strong influence on elasticity, although this hasn't been well studied in dynamic processes. Here, we explore the relationships between elasticity and commitments to elaborate a list of research directions for taking elasticity into account and committing to the development of dynamic, elastic processes.

Elasticity and Commitment in Dynamic Process Management

Let's consider preventive maintenance as a prominent case where elasticity and commitments would play a crucial role in developing more flexible maintenance processes. Although the traditional maintenance model mostly comprises reactive actions as a consequence of failures, an evolved model of preventive processes could reduce the actual failure rate by means of a dynamic plan of preventive actions and data analytics from information provided by sensors in the field.

Figure 1 depicts a global overview of a preventive maintenance company that takes into account diverse resources with different profiles, including skills, experience, cost, or reputation; these resources are specialists who will perform certain tasks in preventive maintenance processes. In this example, let's assume there are three kinds of specialists: two kinds derived from different manufacturers of chillers (Chill-A or Chill-B) and electricians with a generic expertise. Different buildings (B1 to B4) include devices from different manufacturers and electrical installations.

From the commitments perspective, we foresee three different kinds of relationships in this scenario. First, at the resource level, each specialist would have (usually through a contract) a specific commitment to responsibilities (based on skills and availability), and the operational conditions could include some cost or incentives, depending on the resource's actual performance. Second, each team would have a commitment specifying a set of objectives to be addressed (such as an estimation of the resolution time), the resources allocated to solve the task, or the aggregated metrics derived from the team members' individual commitments. Last, at the organizational level, the maintenance provider has some commitment (again, typically within the context of a contracting process) over some specific services and guarantees with the provider's customers.

Consequently, once a resource or team establishes a commitment (the red document in Figure 1), it should include the expected profile and cost agreed, along with some possible compensations (penalties or rewards), based on arrival time (depicted as a gray round clock) or reparation time (depicted as a white round clock).

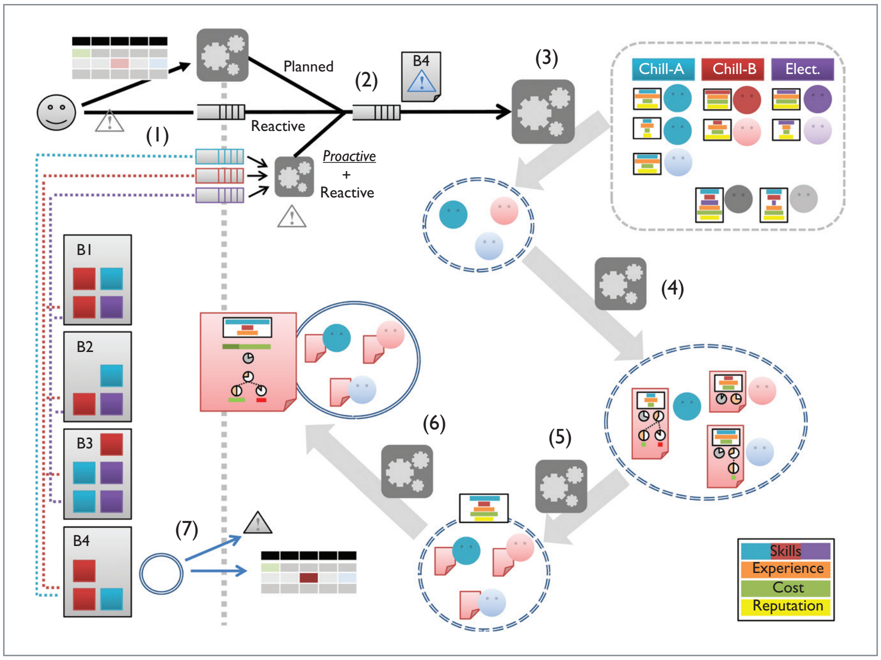

Figure 1. A scenario of dynamic processes in preventive maintenance. We enumerate the details of steps 1 through 7 in the main text.

Figure 1 shows a team's global cycle for different processes that reflect the role of elasticity and commitment. In the following, we enumerate these processes.

1. The company receives information from its customers from different sources. These can include sensors in the buildings send data that can trigger a reactive or proactive maintenance actions; explicit reactive action requests from the customer derived from an unexpected situation or unpredictable problem such as a random accident; and a plan of action agreed to with the customer for periodic maintenance routines. Each of these actions is framed in the context of specific commitments between the provider and each customer and therefore could involve different elasticity management.
2. A scheduling and selection mechanism identifies the most adequate set of (elastic) actions that should be triggered (predictive, proactive, or planned). This selection could be semi-automated and should take into account the different commitments between the provider with its customers and the provider and its resources.
3. The company develops a process to select the appropriate resources from its organization and from the cloud, taking into consideration resource availability (derived from the shift cycle, for example) and their skill sets, experience, and reputation. This process also evaluates possible complications or contingencies that could occur to design a team with a certain degree of elasticity. As an example of such a complication, teams could include an electrician, since a problem with electrical wiring is a typical orthogonal potential

source of failures that could be the root cause of the actual detected malfunction of a chilling device or sensor. Consequently, the teams on site could dynamically adapt to deal with the problem without the need to request a new resource specialized in electrical wiring.

4. For each team's resources, the company develops a process to define the appropriate commitment required to address the problem. The company bases this commitment on decomposing the global maintenance action to have a concrete set of specific actions that can depend on one another, conforming to an operational plan. In this context, the company should consider the team structure and dynamics along the individual commitments of resources, to observe a degree of elasticity that could deal with uncertainties. As a consequence, this process creates a concrete commitment for each resource that includes a precise definition of costs and other estimated performance outcomes, such as arrival or resolution times; additionally, in some cases, the company can consider penalties or rewards as performance incentives.
5. Based on the different resource commitments, the company calculates a global team profile. This profile aggregates the different skills and cost. In addition, it's important to analyze the team's coordination model, since it could affect the team's profile and, consequently, its performance.
6. Taking the team profile and the resource commitments as a starting point, the company then develops a team commitment. This commitment should be aligned with the global commitment of the preventive maintenance provider and the actual team profile.
7. The team starts to develop its duties, moving to the appropriate locations to perform the maintenance actions. Then two possible types of developments can occur: on the one hand, a team can develop the action straightforwardly and, consequently, when the action completes all resources return to headquarters; on the other hand, in the course of action, the team could face a problem that prevents them from successfully dealing with the action. These problems can raise new alerts to the headquarters that could trigger new immediate actions or be addressed in other planned actions eventually.

Now that we've identified the processes, let's consider some potential issues.

Challenges in Aligning Commitments and Elasticity Primitives

To provide a more dynamic process management, organizations face the need for elastic behaviors such as a dynamic team allocation to adapt in real time to the different situations that emerge. In addition, because organizations are normally tied to different commitments with their customers, they should respect the terms that usually include compensations associated with service-level objectives. Consequently, both elasticity and commitment management represent cornerstones that can support more dynamic processes. Unfortunately, current research in process management hasn't fully considered and taken into account elasticity and commitment (see the sidebar for related work).

Thus, here we summarize a list of challenging situations, along with their implications from the commitments (and elasticity) perspective.

- *The problem requires more capabilities so the team should scale in.* This situation typically occurs when there was a wrong estimation about the problem's cause or severity, and the allocated resources were insufficient, but the check onsite reveals the appropriate needs to solve the problem. In such cases, there's a need for new resources that, depending on the case, could come from the company itself, or from third-party organizations such as manufacturers or freelance specialists in the cloud. In any case, we need to create a new individual commitment with each of the new resources incorporated into the team and to modify the actual aggregated team profile to take into account the new capabilities and costs derived from the new members. Depending on the nature of the deviation, there could be a need for modifying the whole team's commitment. Another challenge is to adapt the coordination model to incorporate the new resources.
- *The problem requires fewer capabilities so the team can be scaled down.* As a complementary situation to the scale in, this is commonly derived from the fact that the problem was simpler than expected and there are underused or spare resources in the team that it can dismiss. As a result, the appropriate resources should be reassigned or returned to the resource pool. In this situation, the team's commitment remains unchanged, but the team profile can change based on the dismissed resources; the team can develop a possible recommitment with their remaining resources, for example, to include incentives.
- *While several teams are working on the field, there's a need to group them in order to deal with a joint problem.* This grouping could lead to different structures from a composed team consisting in a coordinated cluster of isolated teams to a merged team where all

Current Work in Elasticity and Commitment in Dynamic Processes

Traditionally, elasticity and commitments have been left out of the mainstream research area of business process management. The main reason is probably that so far, resources for business processes are dedicated through organizational or contract binding. However, as business processes exploit pay-per-use resources in the cloud for on-demand needs, elasticity and commitment become important issues.

Current approaches rely on a human-intensive set of processes to plan and adapt operations' different actions.[1] In the best cases, the current process management approaches rely on planning information systems with business-as-usual workflows that are rigid and assume a fixed organizational structure of teams and resources without an explicit sense of elasticity.[2] However, this approach usually boosts an anchoring that complicates the adaption to problems with a high level of variability that evolves over time at a high pace. As an example, the preventive maintenance industry faces such variability in medium-large scale scenarios where maintenance actions potentially involve a complex reality of highly interconnected devices that challenges the identification of actual malfunctioning elements with high accuracy.[3] Consequently, different kinds of specialists should typically interact and commit to work cooperatively to resolve the problem.

From a commitment perspective, in the typical business landscape a service provider (such as a predictive maintenance company) has formal commitments with its customers. Such commitments could contain some service-level objectives that are subject to penalties and rewards. From an intra-organizational point of view, providers rely on the commitments of their resources (usually contracts) that could include some incentive mechanisms based on the resource performance. It could also develop extraordinary liaisons with external resources (such as freelance specialists or manufacturers' resources). Consequently, to optimize its operations, the maintenance provider requires a harmonization between its commitments with customers and its potential resources.

In recent years, ongoing work has advanced the state of the art in this research direction. For example, Bikram Sengupta and his colleagues have explored the elastic Social Compute Unit (SCU) paradigm as an appropriate foundation for elastic business process management.[4] And Cristina Cabanillas and her colleagues have developed a dynamic resource allocation framework for processes.[5] Moreover, Carlos Müller and his colleagues showed that we can use agreement analysis[6] as a starting point to develop a formal commitment management framework.

References

1. W.M.P. van der Aalst, A.H.M. ter Hofstede, and M. Weske, "Business Process Management: A Survey," LNCS 2678, Springer, 2003, pp 1–12.
2. M. Weske, *Business Process Management: Concepts, Languages, Architectures*, Springer, 2007.
3. Y.-T. Tsai, K.-S. Wang, and L.-C. Tsai, "A Study of Availability-Centered Preventive Maintenance for Multi-Component Systems," *Reliability Engineering & System Safety*, vol. 84, no. 3, 2004, pp. 261–270.
4. B. Sengupta et al., "Collective Problem Solving Using Social Compute Units," *Int'l J. Cooperative Information Systems*, vol. 22, no. 4, 2013.
5. C. Cabanillas et al., "Priority-Based Human Resource Allocation in Business Processes," LNCS 8274, Springer, 2013, pp. 374–388.
6. C. Müller et al., "Comprehensive Explanation of SLA Violations at Runtime," *IEEE Trans. Services Computing*, vol. 7, no. 2, 2014, pp. 168–183.

the different resources are joint in the same team. In both cases (merged or composition), there are important implications in the resultant team's profile and consequently, there can be changes in the team commitments. From the perspective of resource commitments, in the case of the composition, these could remain similar while the merge could involve important recommitments.

- *The problem is scattered and there's a need to decompose an initial team into lower-granularity subteams.* In this case, depending on the relationship of the resulting subteams, similar to the composition, this could result in coordinated clusters or isolated teams. In both cases, while the individual commitments could remain the same, new subteam commitments should be created based on the original team commitment.

Programming Elasticity and Commitment for Dynamic Processes

It's important to consider managing elasticity actions and the different commitments involved between organizations, teams, and individuals as cornerstones of the dynamic processes. Consequently, in order for the companies to evolve into more dynamic organizations, they should take into account these elements as programmable entities that could be managed in a more automated fashion. Next, we develop our proposed characterization of commitment model and elasticity primitives that point out the most promising areas that open discussion on a set of potential research challenges.

Elasticity Primitives

To tackle team dynamics, we propose using the Social Compute Unit (SCU) approach[6] that groups individual compute units/resources (or ICUs) as elastic service units that can scale up and down on demand. An SCU is a cloud-like virtual construct that exists only for the time required. It has a fundamental notion of computing power,

Table 1. Elasticity primitives.	
Primitive	**Description**
Create	The Social Compute Unit (SCU) is formed with the different individual compute units (ICUs).
Dissolve	The team is destroyed and ICUs are reassigned or dismissed.
Scale up	The SCU must grow by incorporating one or more new ICUs.
Scale down	The SCU must shrink to get rid of unused ICUs.
Merge	All ICUs from different SCUs are grouped together in new SCUs.
Split	An SCU is divided into smaller independent SCUs.
Compose	Several SCUs coordinate within the context of a bigger SCU that acts as the container.
Cluster	An existing SCU transforms into a composition of smaller SCUs.

where computing is performed through socially networked humans. These expert groups should exhibit well-defined competencies, letting systems automatically estimate the impact of additional experts in advance. SCUs help fulfill our basic requirements of measurability and elasticity. Consequently, in Table 1 we propose a set of elasticity primitives that represents a comprehensive set of actions that allow a flexible evolution of SCUs, to adapt to different situations and challenges.

Elasticity Primitives in Action

Figure 2 shows a use case with a flow of actions that exemplify the dynamics of the different elasticity primitives in the case of the preventive maintenance domain. The use case starts in t_1 with two different preventive maintenance (P) tickets for building 1 (B1), requiring specialists in two types of devices (colored in blue and red). These two tickets trigger the creation in t_2 of an SCU conformed by three ICUs that go into B1 to develop the maintenance actions. In t_3, a new reactive task occurs derived from unexpected failure in B2 from a red device. This situation develops a split of the original SCU so that the specialist in-field is sent to B2 to attend the more urgent task and the previous planned preventive action is rescheduled in the future; meanwhile, the predictive maintenance actions in the blue device come to an end, dissolving the SCU. Once in B2, in t_4, the ICU realizes the situation's severity and alerts the headquarters (triggering a new reactive ticket) to scale up and form an SCU in t_5 with multiskilled specialists (colored in gray). In t_6, an unknown situation detected in B1 results in creating a new SCU with diverse specialists due to the uncertainty of the problem's source and, once the original problem in B2 is resolved, a merge results in a unique SCU in t_8 to deal with the problem in B1. After investigating in t_9, the problem's source is identified and results in a list of immediate actions that should be done not only in B1 but also in B2, and a planned action for future preventive maintenance. To develop these actions, in t_{10} the original SCU is transformed into a coordinated cluster composed by two SCUs that deal with the different types of actions in t_{11} with a progressive scale down (t_{13}) and dissolution (t_{17}) of the SCUs.

This elasticity management perspective represents an open context that raises research directions, such as how we can model and manage the profile; how to select the appropriate SCUs/ICUs; how to decide the right SCU coordination model; or how to manage elastic primitive triggering.

Commitments Model

We can address the semantics of the commitment from different perspectives; specifically, in the context of our scenario, we propose a commitment characterization in three key aspects:

- *Explicitly set the resource, team, or organization that will be involved in the task's accomplishment.* This should be derived from the potential capabilities of the resource profile. For computing resources, this could correspond with its configuration (such as a list of service property values); while for a human resource, this would be the explicit definition of skills and presumptions that will drive behavior, from a detailed plan of action to a declaration of capabilities used in the actual task.
- *Define a cost model for the resource's operation.* This model could vary in terms of complexity and dynamics, ranging from static pricing tables to a dynamic model that takes into account different metrics; consequently, this latter scenario would involve an additional requirement of monitoring operations that should be included in the overall cost.
- *List objectives that are expected to be derived from the task's development.* These statements could include explicit deliverables or could correspond with quality levels that are guaranteed. In any case, the actual under- or over-fulfilling of these goals could be framed within a compensation model, including potential penalties or rewards.[5]

Having an explicit declaration of the commitments within the resources and teams represents an important step towards strategically managing the organization as a whole. As the relationship between companies is

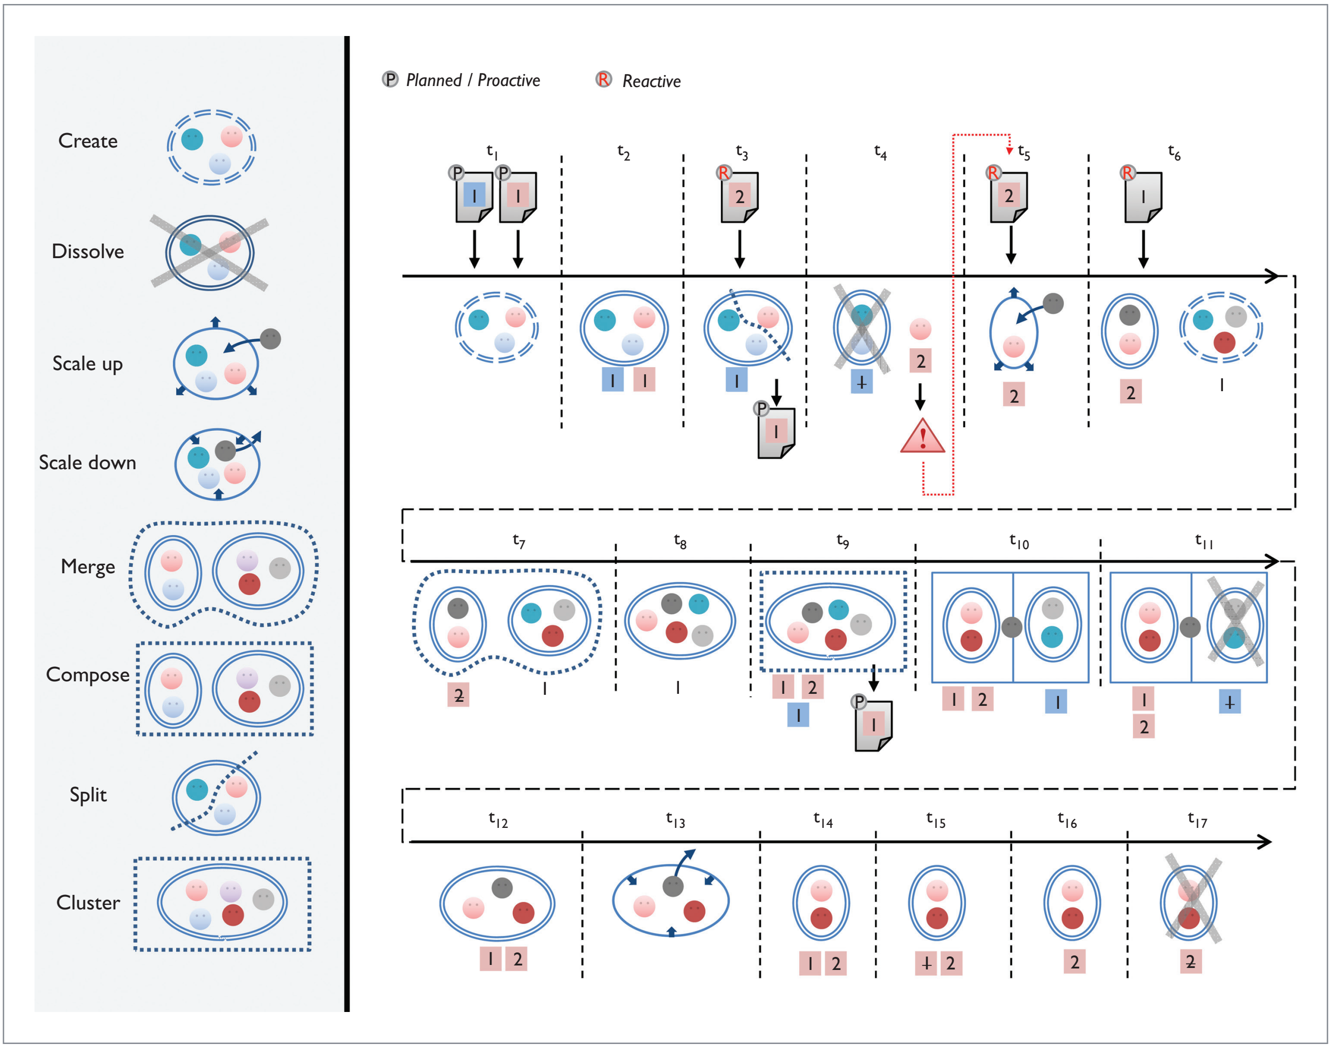

Figure 2. Elasticity use case. The flow of actions exemplify the dynamics of the different elasticity primitives in the case of the preventive maintenance domain.

also usually articulated in the context of formal commitments (contracts), we envision a scenario that would conform a global commitment ecosystem that's an appropriate characterization for a more intelligent and automated governance of the organization.

In addition, from our perspective it's important to highlight the strong influence between the commitments and elasticity primitives. On the one hand, commitments could be crucial to determine the most appropriate primitive to be triggered; as an example, we can foresee an organizational level of commitments with a certain customer that requires a more demanding scale-in policy, while other customers with a lower service-level agreement wouldn't require escalation at all. On the other hand, the primitive itself usually has a significant impact on the commitments involved. Specifically, Table 2 describes the different implications of the elastic primitives in terms of commitments. Because the triggering of primitives is derived from specific events, in some cases they could potentially affect the ICU commitments, while others affect the SCU commitment as a whole.

Commitment management challenges could include developing convenient formal languages to model commitments, analyzing the commitment fulfillment, or harmonizing the organization's commitments, SCU, and ICU. Moreover, it's important to note that the process management problem has been addressed traditionally from the perspective of a fixed organizational structure, and consequent approaches don't include the commitments as a first-level citizen that governs organizations' dynamics. From such a perspective, we foresee the challenge of process modeling and design to involve the elasticity and commitments and the runtime automation of such elements to articulate a more dynamic process infrastructure.

Table 2. Elasticity and commitment relationships.

Primitive	SCU commitments	ICU commitments
Create	Single commitment created from the different ICU commitments.	Created from the ICU profile and target actions.
Dissolve	A fulfillment check over the commitment is developed, and global compensations are calculated.	Evaluating the ICU's performance over its commitments to update its profile (with an improved reputation, for example) and calculate the appropriate compensations.
Scale up	It can change if the reason of scaling was exogenous to the organization (such as a new customer request).	The new resources must commit and some reassignment of tasks can modify the commitment of the pre-existing ICU.
Scale down	It normally remains unaltered.	They normally remain unaltered.
Merge	A new commitment should be created by combining the different commitments of source SCUs.	A recommitment could be made based on the aggregation of responsibilities in the newly created SCU.
Split	New commitments must be made based on a distribution of responsibilities from the source SCU commitment.	A recommitment could be made based on the delimitation of responsibilities of the newly created SCUs.
Compose	A new commitment should be created, taking into account the coordination model and the commitments of the source SCUs that remain unaltered.	They normally remain unaltered.
Cluster	New commitments for each created SCU must be made based on the source SCU and the coordination model defined.	A recommitment could be made based on the delimitation of responsibilities of the newly created SCUs.

We need to see the importance of promoting the concepts of elastic resources and commitments, to create more dynamic organizations that can adapt their behavior while keeping a high level of manageability. Innovative management approaches that utilize elasticity and commitment could significantly improve the classical business models. Still, from a technical perspective, we must revisit concepts of process modeling, composition, and execution to be able to take into account elasticity and commitment models in Internet-scale resources.

Acknowledgments

This work was partially supported by the European Commission (FEDER), and the Spanish and Andalusian R&D&I programs through grants P12-TIC-1867 (COPAS), TIN2012-32273 (TAPAS), and TIC-5906 (THEOS).

References

1. W.M.P. Van der Aalst, M. Weske, and D. Grünbauer, "Case Handling: A New Paradigm for Business Process Support," *J. Data & Knowledge Eng.*, vol. 53, no. 2, 2005, pp. 129–162.
2. C. Dibrell, J.B. Craig, and D.O. Neubaum, "Linking the Formal Strategic Planning Process, Planning Flexibility, and Innovativeness to Firm Performance," *J. Business Research*, vol. 67, no. 9, 2014, pp. 2000–2007.
3. Y.L. Doz and M. Kosonen, "Embedding Strategic Agility: A Leadership Agenda for Accelerating Business Model Renewal," *Long-Range Planning*, vol. 43, nos. 2–3, 2010, pp. 370–382.
4. S. Dustdar et al., "Principles of Elastic Processes," *IEEE Internet Computing*, vol. 15, no. 5, 2011, pp. 66–71.
5. C. Müller et al., "Towards a Formal Specification of SLAs with Compensations," LNCS 8841, Springer, 2014, pp 295–312.
6. S. Dustdar and K. Bhattacharya, "The Social Compute Unit," *IEEE Internet Computing*, vol. 15, no. 3, 2011, pp. 64–69.

Pablo Fernández is an assistant professor in the Applied Software Engineering Group at the University of Sevilla. His work focuses on service agreements, automated trading, and business processes. Fernández has a PhD in computer science from the University of Sevilla. Contact him at pablofm@us.es; www.isa.us.es/members/pablo.fernandez.

Hong-Linh Truong is an assistant professor in the Distributed Systems Group at the Vienna University of Technology. His work focuses on service engineering analytics — in particular, cloud computing; service-oriented architectures and computing; Internet of Things; complex and elastic distributed systems; and context-aware computing. Truong has a PhD in computer science from the Vienna University of Technology. Contact him at truong@dsg.tuwien.ac.at; dsg.tuwien.ac.at/staff/truong/.

Schahram Dustdar is a full professor of computer science (informatics) and he heads the Distributed Systems Group at the Vienna University of Technology. His work focuses on Internet technologies. Dustdar is a member of the Academy Europeana, an ACM Distinguished Scientist, and recipient of the IBM Faculty Award 2012. Contact him at dustdar@dsg.tuwien.ac.at; dsg.tuwien.ac.at/.

Antonio Ruiz-Cortés is an accredited full professor and head of the Applied Software Engineering Group (ISA, www.isa.us.es) at the University of Sevilla. His current research focuses on service-oriented computing, software product lines, and business process management. Contact him at aruiz@us.es; www.isa.us.es/.