



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A Collaboration Model for Community-Based Software Development with Social Machines

Citation for published version:

Murray-Rust, D, Scekcic, O, Truong, H-L, Robertson, D & Dustdar, S 2014, A Collaboration Model for Community-Based Software Development with Social Machines. in Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on . Institute of Electrical and Electronics Engineers (IEEE), pp. 84-93. <https://doi.org/10.4108/icst.collaboratecom.2014.257245>

Digital Object Identifier (DOI):

[10.4108/icst.collaboratecom.2014.257245](https://doi.org/10.4108/icst.collaboratecom.2014.257245)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



A Collaboration Model for Community-Based Software Development with Social Machines

Dave Murray-Rust*, Ognjen Scekić†, Hong-Linh Truong†, Dave Robertson* and Schahram Dustdar†

* Centre for Intelligent Systems and Applications,
School of Informatics, University of Edinburgh, UK
Email: d.murray-rust | dr @inf.ed.ac.uk

† Distributed Systems Group, Vienna University of Technology, Austria
Email: oscekić | truong | dustdar @dsg.tuwien.ac.at

Abstract—Today’s crowdsourcing systems are predominantly used for processing independent tasks with simplistic coordination. As such, they offer limited support for handling complex, intellectually and organizationally challenging labour types, such as software development. In order to support crowdsourcing of the software development processes, the system needs to enact coordination mechanisms which integrate human creativity with machine support. While workflows can be used to handle highly-structured and predictable labour processes, they are less suitable for software development methodologies where unpredictability is an unavoidable part the process. This is especially true in phases of requirement elicitation and feature development, when both the client and development communities change with time. In this paper we present models and techniques for coordination of human workers in crowdsourced software development environments. The techniques augment the existing Social Compute Unit (SCU) concept—a general framework for management of ad-hoc human worker teams—with versatile coordination protocols expressed in the Lightweight Social Calculus (LSC). This approach allows us to combine coordination and quality constraints with dynamic assessments of software-user’s desires, while dynamically choosing appropriate software development coordination models.

I. INTRODUCTION

Most social computing systems today are based around patterns of work that can be predictably modelled before execution, such as translation, bug discovery, image tagging [1, 2]. However, there are many cases where where a traditional workflow approach is too rigid to address the dynamic and unpredictable nature of the tasks at hand, and more flexible crowd working systems must be developed [3]. One example of such dynamic systems is the field of *social machines*—systems where computers carry out the bookkeeping so that humans can concentrate on the creative work [4]. This viewpoint can be used to model and produce a diverse class of systems, spanning task-oriented (Wikipedia) to generic (Twitter); scientific or humanitarian (GalaxyZoo, Ushahidi) to social (Instagram) [5, 6]. In these systems, interactions between computational intelligence and human creativity are deeply woven into the system, making it difficult to draw a clear line between the human and digital parts, separate their analysis and manage coordination.

Creating a social machine requires understanding of individual and group human behaviour alongside technical expertise, and a view of the system as an interconnected whole containing both human and computational elements. Further-

more, social machines must respond to the exigencies of unfolding situations, requiring human creativity in the face of unpredictability. In such cases it is important not to over-regulate participating humans, but to let them play an active role in shaping the collaboration during runtime. This includes leveraging human creativity and embracing the uncertainty that comes with it. On the other hand, it is often necessary to impose certain coordination and quality constraints for these collaborations in order to manage them. The constraints delimit the decision space within which the humans are allowed to self-organize.

Recently, a number of human computation frameworks supporting complex collaboration patterns were proposed (Section V). They mostly build upon conventional crowdsourcing platforms offering a process management layer capable of enacting complex workflows. While these systems represent important steps on the road to building complex social machines, in cases where unpredictability is inherent to the labour process and we cannot know all of the system requirements in advance, a different approach is needed.

In this paper we present models and techniques for coordination of human workers (software users and developers) in crowdsourced software development environments. They support dynamic bootstrapping and adaptation of social machines: using one social machine to generate/alter another one, thus allowing for flexible, community-driven development. This is a fundamental novelty, allowing more human influence on the execution of a computation.

The introduced concept augments the *Social Compute Unit* (SCU, Section II-B)—a general framework for management of ad-hoc human worker teams—with versatile coordination protocols encoded in *Lightweight Social Coordination Calculus* (LSC, Section II-C). This combination allows us to design and model social machines oriented towards crowdsourcing software development. Coordination protocols provide high level organisation of activities around development (including planning and user assessment/feedback), while a set of coordination and quality constraints guide the assignment of workers to tasks. This allows the system as a whole to strike a balance between imposed constraints and creative freedom in the software development cycle. Concretely, this means that the proposed model is able to take into account the feedback from the user population and subsequently alter the process of the development of software artefacts. Although we focus on collaborative software development, the solution we present is

generally applicable to a class of similar problems; in particular, situations where a social machine is being developed—and hence the developers must react to the changing needs and behaviour of the community—using another social machine to crowdsource the development.

The remainder of the paper is structured as follows: In the continuation of the section we present the motivating scenario of community-influenced, collaborative software development. In Section II we first analyse how the presented scenario can be modelled in terms of social machines. We then introduce background concepts which we use to build our model: Feature Oriented Software Development (FOSD), the Social Compute Unit (SCU) and the Lightweight Social Calculus (LSC). In Section III we present the coordination model for the social machine employing the previously introduced background concepts. In Section IV a proof-of-concept implementation is presented and evaluated through simulation. Section VI concludes the paper.

A. Motivating Scenario

Developing software for a large user base with diverging interests can be challenging. As an illustrative example, let us consider the problem of developing a forum-like scientific platform—a scholarly social machine—to facilitate multidisciplinary cross-collaboration and sharing of results. This includes functionality such as: paper previews, comments, in-place formulae and data rendering, citation previews and bookmarking. While these are functionalities beneficial to all scientists, preferences for particular formats and services will likely differ among different sub-communities. For example, chemists and mathematicians will have different domain-specific requirements from the platform. Some examples are:

- Computer scientists need code syntax highlighting, LaTeX rendering, embedding of IEEEExplore and ACM DL citations. If any of these features is missing, the software is not useful to the community. However, they do not particularly care about chemistry-specific features.
- Chemists often use InChi strings to represent chemical formulae. If the software supports InChi, then chemists would also want support for compound lookup on PubChem, and visualisation with pyMol. Without these features, the platform does not help them particularly. On the other hand, syntax highlighting and IEEEExplore integration is not important.
- Individual scientists may be bothered by (lack of) certain features. For example, users may dislike being forced to use a LinkedIn account to log in, due to possibility of a third party accessing unpublished scientific findings.

Some of these features are orthogonal—code syntax highlighting and LaTeX rendering are both useful in their own right—while some are synergistic: a chemist might require both parsing InChi strings and PubChem lookup in order to carry out their particular workflow.

The complexity of developing such software lies in catering to the heterogeneous user needs, requiring numerous trade-offs




Fig. 1. Two connected social machines: i) the *development* social machine, where crowdsourced workers follow coordination protocols to create a software artefact; ii) the *target* social machine, where a community of practice forms around the software artefact created by i).

when deciding which features to implement. Furthermore, different sub-communities tend to change preferences regarding required or newly developed features during the development process which need to be taken into account. Finally, certain members of the scientific community (i.e., targeted users) may decide to take part in the development process themselves.

II. MODELLING COMMUNITY-BASED SOFTWARE DEVELOPMENT

The previously presented scenario is representative of many social machines, where a dynamic community forms around a particular (software) artefact. The population is likely to change, and feedback between the human participants and technological infrastructure can lead to changes in the purpose and direction of technological development.

This means that there are two *social machines*: i) the *target* social machine which includes the forum software and its community of users and ii) the *development* social machine, which is the software developers and their coordination architecture (Figure 1). We use the term *utility* to denote some metric for the benefits which a user (in the *target* social machine) derives from participation. While in principle this metric could include a multitude of components, within this paper we narrow our focus to treat utility as measuring how well the software’s feature set matches the user’s requirements.

The aim of the development social machine is to increase the overall utility of the user population, by creating features which match community needs and desires. The developers do not know ahead of time the true preferences of individuals, or the constitution of the community, and hence the effects of software changes on community behaviour are difficult to predict ahead of time.

Since crowd-labour is increasingly used for the development of software, we assume it is necessary to use development methodologies which split work into tasks that are amenable to crowdsourcing. This means that tasks have to be disassembled into simpler subtasks, and mapped to appropriate developers. The latter is itself a complex problem, as it also includes taking care of inter-task implementation dependencies. Hence, the *development* social machine must be able to i) assess user desires and preferences; ii) identify and prioritize features for development; iii) coordinate the development and deployment of these features; iv) organise these tasks over time with respect to a dynamically changing population and limited resources.

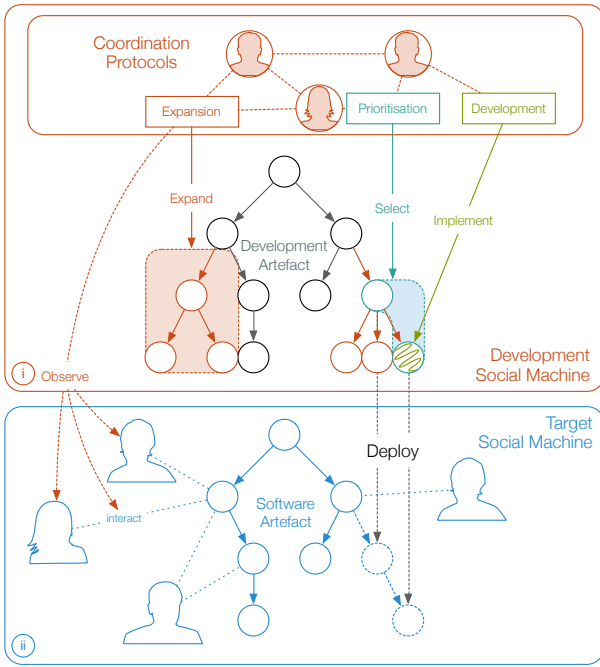


Fig. 2. Interaction between the development and target social machines, including development steps, developer interaction, user observation and community interaction.




Fig. 3. States and operations on a single node in the feature tree. Potential nodes are *expanded* into Possible nodes, which can be *evaluated* against user preferences, before being *implemented*.

These operations and their relation to the user population are outlined in Figure 2

A. Feature Trees for Artefact Development

A requirement of our model is the representation of the current state of the artefact under development—the *development artefact* in Figure 2. Since our example is based on software development, we use the *Feature-Oriented Software Development (FOSD)* paradigm, where software artefacts are represented as trees of *features*: “prominent or distinctive user-visible aspect, quality, or characteristic[s] of a software system” [7]. This representation is used so that development can be decomposed into small sets of related tasks that can be handled relatively independently, to aid collaborative creation of software artefacts.

Based on the requirements and possibilities in the scenario outlined in Section I-A, the feature tree in Figure 4 can be constructed¹. Here, broad classes of functionality, such as visual embedding of graphical objects are represented as

¹The tree was created using FeatureIDE. Details of assumed semantics can be found at http://www.itl.cs.uni-magdeburg.de/iti_db/research/featureide/

branches of the tree, with specific instances such as pyMol viewers forming sub-branches and leaves.

Feature trees can be used to represent the current state of software development; by labelling each node with a state, the team knows whether or not functionality for that feature has been implemented, and whether the conditions for implementing that functionality have been met. This forms a coordination artefact used by the development social machine, to organise construction and monitoring of the software artefact used in the target social machine. As shown in Figure 2, once features in the tree are implemented, the software artefact can be deployed to the user population.

Software development can be modelled as modifications to the feature tree: the re-labelling of nodes as new functionality is conceived of and implemented (before being deployed to the target artefact). In this paper, we use a simple state model (Figure 3), where each node is either: *i) Implemented*—code has already been created for this feature, and it is available to users; *ii) Potential*—the feature has been conceptualised and designed, but no code exists yet; *iii) Possible*—part of the universe of possible features, but one which is not currently under consideration for implementation².

Based on this representation, development can contain the following steps, or *development primitives*, which map to operations of the feature tree:

- 1) *Expansion* of the tree converts nodes from *possible* to *potential* by finding new features to implement. This might be through expert designers, co-creation or direct user solicitation.
- 2) *Evaluation* of community needs and their relation to individual features results in labelling nodes with some indication of how well the community will react. There are many ways to do this, including surveying the participants; public consultations; focus groups; monitoring of behaviour; and social media analysis.
- 3) *Prioritisation* of features to implement, which may be driven by the result of evaluations, voting by the population, investor demands, expert opinion etc. This decision may depend on which type of *costs* the controllers of the artefact would like to optimise (e.g., economic, temporal, social).
- 4) *Implementation* of the selected features, whether in-house, or crowdsourced, using some particular software design methodology. When implementing features, the constraints contained in the feature tree must be observed (e.g., mandatory features, alternative features).

Within our model, these tasks are carried out by assembling teams of crowd professionals—SCUs (Section II-B), capable of executing complex workflows. The formation of SCUs and coordination of their actions are carried out through the Coordination Model (CM), introduced in Section III, which allows flexible workflows that adapt to emerging situations.

B. Social Compute Unit (SCU)

An SCU [8] is a loosely-coupled virtual team of socially-connected experts with skills in the relevant domain. The SCU

²The *possible* state is largely a convenience for simulation.




Fig. 4. An example feature tree for a scientific forum software system.

is created upon request to solve a given task. It uses the crowdsourcing power of its members and their professional connectedness toward addressing the problem which triggered its creation and is dissolved upon problem resolution. The SCU is a programmable entity. This means that its various properties and functionalities (team assembly, task decomposition, runtime collaboration patterns, coordination, task aggregation) can be ‘programmed’ to support different types of human/machine collaborative efforts. For example, in [9] the authors show how the SCU can support well-defined business-processes, such as workflow patterns for IT incident management. However, SCU can also be used to perform looser collaboration patterns leaving space for human improvisation and creativity [10].

In this paper we use SCUs within the development social machine to execute tasks in the software development cycle, such as implementing a concrete software feature. Concretely, we build upon the particular SCU model presented in [11] and use it in the context of the encompassing social machine’s coordination model. Whenever a development primitive (from Section II-A) needs to be executed, a request with input parameters is sent to the SCU provisioning engine to form a team of developers/experts suitable for that particular primitive (task). The provisioning engine returns the closest-to-optimal matching subset of available developers, representing a SCU for that task.

The full list of available input parameters and descriptions

of the team formation algorithms are available in [11]. In this paper, however, we vary only the parameter named *job description set* (J), while assuming default values for the remaining parameters. J contains job descriptions for each subtask: $J = \{j_1, j_2, \dots, j_k\}$. A job description is a set of tuples $j_i = \{(t_1, q_1), (t_2, q_2), \dots, (t_m, q_m)\}$, where t_l is a skill type (e.g., ‘java developer’, ‘test engineer’) and $q_l = \{‘fair’, ‘good’, ‘verygood’\}$ is a fuzzy quality descriptor. The job description (t_l, q_l) specifies which skills a worker needs to possess in order to perform the subtask l successfully.

C. Lightweight Social Calculus (LSC)

LSC is an extension of LCC [12], which has been used to represent interaction in many systems [13]. LCC is a declarative, executable specification which can be communicated between agents at runtime; it is designed to give enough structure to manage fully distributed interactions by coordinating message passing and the roles which actors play, while leaving space for the actors to make their own decisions. LSC augments LCC with extensions designed to make it more amenable to mixed human-machine interactions; in practice, this means having language elements which cover user input, external computation or database lookup and storing knowledge and state.

An LSC protocol consists of a set of clauses; the head of each clause is a role specification, and the body a description of




Fig. 5. Example LSC clause from the meal organisation interaction model (slightly modified for clarity). An agent playing the role of *invitee* will wait for a message from a *confirmor* specifying the time and place for dinner; the values in the message for *Time* and *Place* are substituted in, and the agent then decides if it will *attend*, and sends back the appropriate message. It then resumes the role of *invitee* in case of alternate suggestions.

what an agent should do when playing that role (see example in Figure 5). The body contains message sending ($M \Rightarrow a(\text{role}, ID)$) and receiving ($M \Leftarrow a(\text{role}, ID)$), sequencing and choice (*then* and *or*), implication (*action* \leftarrow *condition*), the assumption of new roles ($a(\text{role}, ID)$) and any extra computation or conditions necessary.

Each agent’s interaction starts with a clause from a protocol, which is then repeatedly re-written in response to incoming events: incoming messages are matched against expected messages, role definitions are replaced with the body of matching clauses, values are substituted for variables and so on. As the interaction progresses, this *state tree* keeps a complete history of the agents actions and communications. This supports the creation of multi-agent institutions [14] where interaction is guided by shared protocols and a substrate which keeps track of state.

LSC is formal enough that it can be computationally manipulated, for example to synthesise new protocols [15]. It shares features with workflow languages—while providing more flexibility—and can be derived from e.g. BPEL4WS to create completely decentralised business workflows [16]. LSC has also been used in the creation of social machines by binding formal interaction models into natural interaction streams [17].

Within our model, LSC is used to model the development social machine, by specifying the interactions among developers, and between developers and the feature tree representing the state of the software artefact. It provides a means to create a formal representation of software development processes, allowing for computational coordination of their enactment, while providing more flexibility than a workflow would allow. LSC provides a bridge between low level operations—e.g. implementing a particular node—and high level concepts such as “agile methodologies”. By formalising the coordination protocols and making them first class objects, it is possible to share, modify, discover and rate individual protocols; by separating the protocol from the domain of application, it is possible to apply the same methodology to new domains. The flexibility of the language allows for sub-protocols to be chosen dynamically, so that development can be adapted in response to changing needs. Over time the system can build up a view of when each protocol is appropriate, and be able to assist with selection of protocols for novel situations.

III. COORDINATION MODEL

The *coordination model* represents the artefact regulating the interactions among social machines. It contains the following submodels, regulating different interaction aspects:

- **Data Submodel:** A formal data model used to represent the data that is processed and exchanged by social machines. It serves both as input and output for the social machine. In our example, the data model is represented by the feature tree representing the forum software. For the development social machine it serves to indicate the features to develop and dependencies; but also to track the progress of the development cycle. The resulting tree is then subsequently also used as the input of the target social machine for calculating the overall population utility, as well as to mark elicited features for future development cycles.
- **Quality-of-Service Submodel:** In essence, the development social machine is providing a software-development service to the target social machine. Therefore, we need a set of metrics to express the requested and measure the obtained quality of this service. In this paper, we adopt the metrics already provided by the SCU [11] to formulate requested QoS.
- **Interaction Submodel:** The coordination submodel contains a collection of LSC-encoded protocols managing interactions between social machines and their workers. The coordination submodel contains multiple possible protocols. A *metalevel protocol* is used to make real-time selection and enactment of an appropriate subset of concrete protocols, based on the current state of the coordination model, input from stakeholders or the current behaviour of the community interacting with the artefact. Selection could also include discovery of new protocols to use (e.g., as new development methodologies are introduced) as well as analysis of the historic performance of existing protocols in similar situations. The use of metaprotocols is crucial in order for the development social machine to be responsive to community requirements, and for it to adjust development trajectories accordingly.

Figure 6 illustrates the usage of the coordination model artefact for the scenario introduced in Section I-A. An iteration in the software development cycle starts by having an active LSC protocol send a request to the SCU Provisioning Engine (Figure 6, ①). The request contains the necessary QoS input parameters (described in Section II-B) for creation of multiple SCUs. Based on these parameters the SCU Provisioning Engine selects appropriate workers from the crowd of professionals (②), assembles and returns the SCUs. The newly created SCUs are passed the feature tree with nodes selected for implementation (③), finally constituting a functional development social machine. The development social machine starts performing the designated actions on the feature tree (④). The active LSC protocol from the interaction submodel takes care that the actions performed by different SCUs are properly ordered and repeated if necessary (e.g., due to failure, or insufficient quality). After the SCUs finish executing, the resulting feature tree is passed to the target social machine (⑤).




Fig. 6. Using the Coordination Model to support the community-based, collaborative software development scenario.

The modified tree is then evaluated by a function assessing the population utility (⑥). As explained earlier, this is a measure of target community’s satisfaction with the implemented features. Based on this value, and the new requests from the community, the metaprotocol can decide whether new development iterations are necessary, and if yes, which protocols to use (⑦). Depending on the new priorities, a different protocol can be chosen to control the development social machine in the new iteration. For example, for the *evaluate* action, new candidate features may be identified by a single SCU of experts, or by having multiple SCUs suggesting new features and then deciding by majority voting. Or, in case of a failure, we may decide to repeat the task with the same SCU, or escalate to a more reliable (and thus a more expensive) one.

In the following sections, we present a proof-of-concept implementation of this coordination model. We evaluate the implemented prototype by simulating a population and running a number of LSC protocols to showcase its functionality.

IV. IMPLEMENTATION AND EVALUATION

A. Prototype Implementation

In order to demonstrate the operation of the coordination model, we have implemented a simulation prototype which covers a subset of the conceptual model’s possible functionality. In the simulation, a pool of crowd workers participate in

improving the scientific forum software introduced in Section I-A. The (simulated) workers are managed by a system running various LSC protocols, representing different approaches to software development. This includes all of the task selection and implementation activities from Section II, as well as the team selection work discussed in Section II-B.

The implementation uses the *scalsc* LSC library, with extensions to model feature trees, labour and team selection, and user populations³. Concretely, this comprises:

- 1) A population of simple software agents representing community members; this is simulated as a heterogeneous group of individuals, each with their own preferences about which features the community software should contain. The preferences are represented as scores for the presence of conjunctions or disjunctions of implemented feature-tree nodes. A typology approach is used, where archetypal users are defined for two classes (chemists and mathematicians), differing in their preferences for functionality. These users are sampled with multiplicative noise ($\mathcal{N}(1, 0.1)$) added to their preference scores to provide limited heterogeneity.
- 2) A feature tree representing the current state of the software, as defined in Section III, following the example in Figure 4;

³Complete source code and installation instructions can be found at https://bitbucket.org/mo_seph/social-institutions

- 3) A simplified, idealised SCU model, where teams are formed in response to quality constraints, and perform tasks on the feature tree. In this simplified model, we assume that one worker is returned per task, with a skill set that exactly matches the quality constraints [11]. Workers have scores for four skills: *implementation*, *evaluation*, *prioritisation*, *design*, each of which ranges from 0..1.
- 4) A labour model, relating worker qualities to the time, cost and quality of carrying out primitive tree operations. This model has been designed to represent the issues at hand in a stylised manner, while having a reduced parameter set to help understand the model’s behaviour. Operations are assigned a basic cost C_o , which is then multiplied by the cost of the worker who is carrying it out; worker cost is the sum of the worker’s skill levels (S) raised to an exponent $k = 0.5$, so the complete cost of a worker w operation o is: $C(w, o) = C_o \sum s^k$ (for $s \in S$). Details of operator cost, time and implementation specifics are in Table I⁴.

In order to effect changes to the feature tree, a set of LSC-based intra-unit- and meta-protocols are used to construct SCU according to quality metrics, and schedule them to carry out operations. Listing 1 shows an example high-level LSC protocol coordinating an agile development process: `form_scu` triggers the SCU formation, based on a set of required skills and the action to enact, while `do_task` controls the execution of the selected actions⁵. These protocols can be written as standard LSC [12], with a small set of extra predicates for forming teams and manipulating trees: `form_scu` and `do_task` mentioned above, as well as `current_tree` and `highest_priority`, which are demonstrated in Listing 1.

B. Scenarios

In order to illustrate the operation of the prototype, we run it under two contrasting scenarios, with three different *LSC-based workflows* imposed through appropriate LSC protocols. In the first scenario—*StablePopulation*, a population of 1000 members of the chemistry community (chemists) is simulated throughout the entire simulation runtime. In the second scenario—*DynamicPopulation* the initial population of 1000 chemists is replaced by 200 chemists and 800 mathematicians at timestep 20. This is a crude and stylised approach to representing a shift in user population, where the platform is adopted by a different user community, but it allows us to illustrate the prototype implementation’s behaviour.

The coordination models we use are loosely based on current or past practice in software development:

- The *traditional* model starts with a large public consultation, where most of the tree is explored and assessed before any implementation takes place. Actual feature implementation is then carried out by teams (i.e., SCUs) of average-skilled programmers, who have three attempts to implement any given node.

- The *escalation* model begins with the same initial public consultation, but is followed by a development process where initially an average (and cheap) developer attempts to implement each node. If that fails, a high quality, but more expensive, developer is found and brought in to finish the job. This is an example of a simple metaprotocol, allowing alternative development pathways to be chosen at runtime.
- The *agile* model defines a tight loop of evaluation and implementation, to allow development to respond to a changing set of user requirements.

C. Results and Discussion

Figure 7 shows the simulation outcomes when running the described scenarios and workflows. Under the *Stable* scenario, the *agile* workflow initially performs best, as development starts immediately. Over time, however, the *escalation* workflow achieves higher utility with fewer nodes due to a greater understanding of the complete feature tree. The *traditional* workflow is limited by the speed of its average-skilled developers, but utility does increase gradually. Under the *Dynamic* scenario, the initial behaviour is similar. However, when the population changes at timestep 20, the *agile* workflow is better able to adapt to the change, and create nodes which better represent the desires of the new population.

This demonstrates the utility of developing software using a flexible coordination model. When the user community changes, the coordination model can respond dynamically, by prioritising different features for implementation. The utility curve in the *Dynamic* scenario post population-change rises most sharply for the *agile* workflow, indicating its ability to responsively re-plan.

In contrast to conventional workflows, the LSC protocols are dynamic, and can be changed or “plugged in” during runtime. For example, the outcomes of a public consultation—as simulated here—can be used to influence the choice of protocols to be used subsequently, affecting the way that work is coordinated. This allows for a larger human influence on the execution of complex work processes.

The difference between the utility curves under the *traditional* and *escalation* workflows demonstrates the effects of bringing QoS constraints into the development protocol. The *traditional* workflow tends to be cheaper per unit time, using only low quality developers. However, the *escalation* workflow creates more nodes per unit cost, by being able to form SCUs with highly skilled workers when necessary to carry out difficult jobs. This also results in more nodes created per unit time, so the population utility rises more rapidly.

This demonstrates how a dynamic protocol can be responsive to population changes, and how integrating QoS constraints allows system designers to tailor development towards different goals. Taken together, these capabilities allow for dynamic, flexible protocols which can draw on a pool of cloud workers to create software artefacts in response to community needs.

Simulation modelling is used in the computational social sciences to explore theoretical ideas in the context of synthetic

⁴We acknowledge that the simulation will be sensitive to the parameter values chosen (especially k); the results we present are intended only to give an indication of capabilities, so no formal sensitivity analysis has been carried out.

⁵Further details on the prototype implementation, as well as the source code can be found at https://bitbucket.org/mo_seph/social-institutions

TABLE I. MODEL OF PRIMITIVE TREE OPERATIONS, SHOWING BASE COSTS, ASSUMPTIONS MADE AND IMPLEMENTATION DETAILS. e_{real} IS THE TRUE POPULATION UTILITY FOR A FEATURE TREE, AND s_x IS THE TEAM’S SKILL IN x .

Operation	Assumptions	Implementation	Cost, Time
Expand	Only children of <i>potential</i> nodes are considered. A better design process will create nodes which better match the population’s need	Order nodes by $e_{real}\mathcal{N}(1, (1 - s_{design}))$ and select the first.	2, 0.5
Evaluate	A better evaluation process will be closer to the true value population’s need	Label nodes with $e_{est} = e_{real}\mathcal{N}(1, (1 - s_{eval}))$.	0.5, 0.5
Prioritise	Better prioritisers order nodes more closely to their true evaluation order with respect to population’s need	Order nodes by $e_{est}\mathcal{N}(1, (1 - s_{prioritise}))$ and label with index.	0.01, 0.01
Implement	Select highest priority node; better implementers have more chance of success.	$P_{impl} = s_{implementation}$	1, 1

Listing 1. Example protocol used to coordinate “agile” development. An SCU is first formed to identify the next best node to implement. Then, another SCU is formed to implement that node. This sequence is then run in a tight loop to carry out responsive development.

```

a (agile (ExpQ, ImplQ), A) ::      %Agent role for doing "agile" development
                                %Create an SCU to expand the next best node
  form_scu (expand(1), [expansion(ExpQ)], ExpAssign ) then
  current_tree (InTree) then      %Get the current tree
  do_task (ExpAssign, InTree, ExpTree ) then %Carry out the expansion
  highest_priority (ExpTree, Next) then    %Find the best node to implement
                                %Form an SCU to implment it
  form_scu (implement, [implementation(ImplQ)], ImplAssign ) then
  do_task (ImplAssign, ExpTree, Result ) . %Carry out the implementation

```

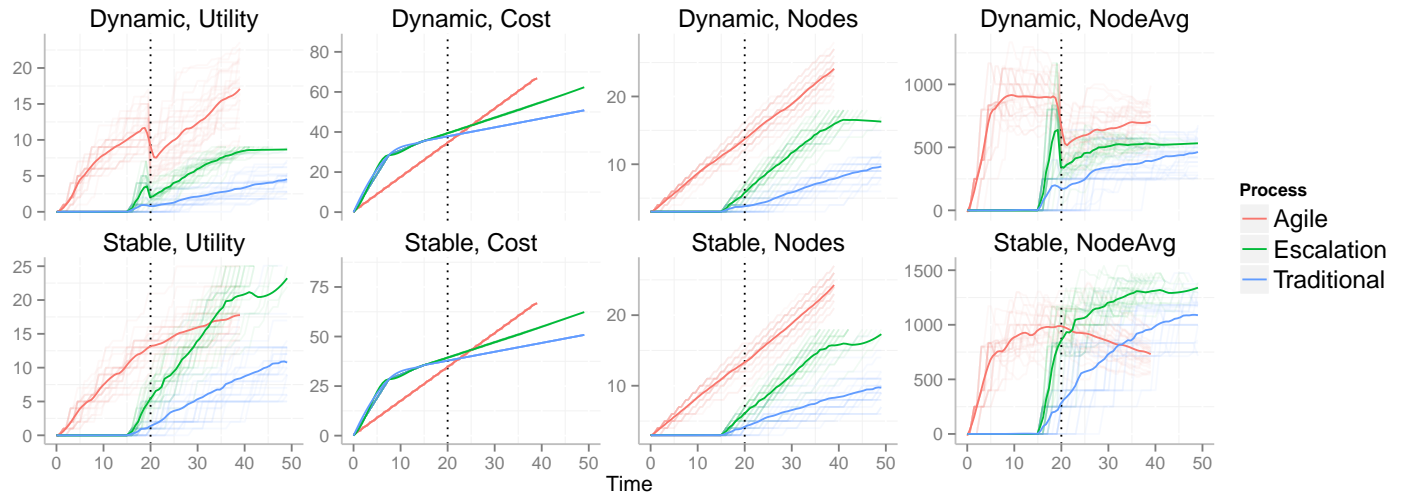


Fig. 7. Simulation outcomes for the two scenarios and three workflows. Faint lines represent individual runs ($n = 30$ in each condition), and solid lines are smoothed ensemble averages. *Utility* is the average utility of the feature tree across the population; *Cost* is the financial cost of developing the community-software; *Nodes* is the number of features (nodes) implemented; and *NodeAvg* is the average utility provided by each feature (an indication of the quality of community fit). The dotted vertical lines indicate the time where the *Dynamic* scenario undergoes a step change in the population composition.

populations, particularly where real studies would be impractical [18, 19]. Recently, this has been applied to crowdsourcing, in order to generalise results which otherwise would be tied to a particular situation [20]. However, simulation has the potential to play another role in this area, as developing a computational model of population behaviour can be used to “close the loop” and aid in the design of effective social machines [21].

The simulations presented here represent a highly simplified version of our conceptual model where intelligent computational machinery underpins human creative activity in the development of software artefacts for dynamic communities. The use of flexible process languages such as LSC means that the inter-unit protocols used here could be augmented to embody more refined development methodologies, with complex patterns of coordination where necessary. Similarly, at the metaprotocol level, there is space for a dynamic adjustment of the protocols and parameters chosen, in order to balance community and stakeholder demands against time and cost constraints. Since LSC is a first class protocol, the interactions specified can be exchanged, rated, discovered and modified both computationally or through human intervention. This can help create a better understanding of which methodologies work in which situations. At the intra-unit level, intelligent protocols could be used to more flexibly assign workers to sub-tasks, reacting to developing situations and changing requirements.

V. RELATED WORK

Social machines share common ground with other collective intelligence applications such as human computation and social computing (diagram in [5, p.2]). Many crowdsourcing systems can be seen as social machines. Of the existing commercial platforms, of particular relevance here are Topcoder⁶ and ODesk⁷, which use different mechanisms to organise diverse participants around software development. As crowdsourcing platforms are becoming widely used as research tools, a number of solutions appeared providing overlay abstractions offering more advanced workflow management and allowing users to perform more complex tasks/computations.

TurKit [22] is a library layered on top of Amazon’s Mechanical Turk offering an execution model (crash-and-rerun) which re-offers the same microtasks to the crowd until they are performed satisfactorily. The entire synchronisation, task splitting and aggregation is left entirely to the programmer. The inter-worker synchronisation is out of programmer’s reach. The only constraint that a programmer can specify is to explicitly prohibit certain workers to participate in the computations.

Jabberwocky’s [23] *ManReduce* collaboration model requires users to break down the task into appropriate *map* and *reduce* steps which can then be performed by a machine or by a set of humans workers. A number of *map* steps can be performed in sequence, followed by possibly multiple *reduce* steps. Human computations stops the execution until it is performed. While automating the coordination and execution management, Jabberwocky is limited to the MapReduce-like class of problems.

AutoMan [24] integrates the functionality of crowdsourced multiple-choice question answering into Scala programming language. The authors focus on automated management of answering quality. The answering follows a hardcoded workflow. Synchronisation and aggregation are centrally handled by the AutoMan library. The solution is of limited scope, targeting the designated labour type.

CrowdLang [25] brings in a number of novelties in comparison with the other systems, primarily with respect to the collaboration synthesis and synchronisation. It enables users to (visually) specify a hybrid machine-human workflow, by combining a number of generic collaborative patterns (e.g., iterative, contest, collection, divide-and-conquer), and to generate a number of similar workflows by differently recombining the constituent patterns, in order to generate a more efficient workflow. The use of human workflows also enables indirect encoding of inter-task dependencies.

To the best of our knowledge, at the moment of writing, CrowdLang and SCU are the only two systems offering execution of complex human-machine workflows. However, as explained before, both systems need to know the possible (sub-) workflows in advance. The coordination model presented in this paper complements the functionality offered by systems such as these two, by providing a higher-level coordination management layer.

VI. CONCLUSION

In this paper we introduced a novel coordination model for teams of workers performing creative or engineering tasks in complex collaborations. The coordination model augments the existing Social Compute Unit (SCU) concept with coordination protocols expressed using the Lightweight Social Calculus (LSC). The approach allows combining coordination and quality constraints with dynamic assessments of user-base requirements. In contrast to existing systems, our model does not impose strict workflows, but rather allows for the runtime protocol adaptations, potentially including human interventions. We evaluated our approach by implementing a prototype version of the coordination model for the exemplifying case-study and simulated its behaviour on a heterogeneous population of users, running different scenarios to demonstrate its effectiveness in delivering end-user utility, and illustrated responses to a dynamically changing population.

In summary, we have given a conceptual model for combining process models with crowdsourced teams to create software artefacts in support of dynamic communities. This formalisation paves the way for increased intelligence to be brought into crowdsourced software development, creating a more responsive, community-centred process.

ACKNOWLEDGEMENT

This work is partially supported by the EU FP7 Smart-Society project under grant 600854 and the EPSRC SociaM

⁶<http://www.topcoder.com/>

⁷<https://www.odesk.com/>

project under grant EP/J017728/1.

REFERENCES

- [1] O. Tokarchuk, R. Cuel, and M. Zamarian, "Analyzing crowd labor and designing incentives for humans in the loop," *IEEE Internet Computing*, vol. 16, no. 5, pp. 45–51, 2012.
- [2] A. Doan, R. Ramakrishnan, and A. Y. Halevy, "Crowdsourcing systems on the World-Wide Web," *Communications of the ACM*, vol. 54, no. 4, p. 86, Apr. 2011. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1924421.1924442>
- [3] A. Kittur, J. V. Nickerson, M. Bernstein, E. Gerber, A. Shaw, J. Zimmerman, M. Lease, and J. Horton, "The future of crowd work," in *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, ser. CSCW '13. New York, NY, USA: ACM, 2013, pp. 1301–1318. [Online]. Available: <http://doi.acm.org/10.1145/2441776.2441923>
- [4] T. Berners-Lee and M. Fischetti, *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. Harper Information, 2000.
- [5] N. R. Shadbolt, D. A. Smith, E. Simperl, M. V. Kleek, Y. Yang, and W. Hall, "Towards a classification framework for social machines," in *SOCM2013: The Theory and Practice of Social Machines*. Association for Computing Machinery, 2013. [Online]. Available: <http://eprints.soton.ac.uk/350513/>
- [6] P. Smart, E. Simperl, and N. Shadbolt, "A Taxonomic Framework for Social Machines," in *Social Collective Intelligence: Combining the Powers of Humans and Machines to Build a Smarter Society*, D. Miorandi, V. Maltese, M. Rovatsos, A. Nijholt, and J. Stewart, Eds. Springer Berlin, In Press.
- [7] T. Thüm, C. Kästner, F. Benduhn, J. Meinicke, G. Saake, and T. Leich, "FeatureIDE: An Extensible Framework for Feature-Oriented Software Development," *Science of Computer Programming*, vol. 79, pp. 70–85, 2014.
- [8] S. Dustdar and K. Bhattacharya, "The Social Compute Unit," *Internet Computing, IEEE*, vol. 15, no. 3, pp. 64–69, 2011. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5755601
- [9] B. Sengupta, A. Jain, K. Bhattacharya, H.-L. Truong, and S. Dustdar, "Collective Problem Solving Using Social Compute Units," *International Journal of Cooperative Information Systems*, vol. 22, no. 4, 2013.
- [10] M. Rivani, H. L. Truong, and S. Dustdar, "On the elasticity of social compute units," in *CAiSE*, ser. Lecture Notes in Computer Science, M. Jarke, J. Mylopoulos, C. Quix, C. Rolland, Y. Manolopoulos, H. Mouratidis, and J. Horkoff, Eds., vol. 8484. Springer, 2014, pp. 364–378.
- [11] M. Z. C. Candra, H.-L. Truong, and S. Dustdar, "Provisioning Quality-aware Social Compute Units in the Cloud," in *11th International Conference on Service Oriented Computing (ICSOC 2013)*. Berlin, Germany, December 2-5: Springer, 2013.
- [12] D. Robertson, "A lightweight coordination calculus for agent systems," in *Declarative Agent Languages and Technologies II*, ser. Lecture Notes in Computer Science, J. Leite, A. Omicini, P. Torrioni, and p. Yolum, Eds. Springer Berlin Heidelberg, 2005, vol. 3476, pp. 183–197. [Online]. Available: http://dx.doi.org/10.1007/11493402_11
- [13] —, "Lightweight coordination calculus for agent systems: Retrospective and prospective," in *Declarative Agent Languages and Technologies IX*, ser. Lecture Notes in Computer Science, C. Sakama, S. Sardina, W. Vasconcelos, and M. Winikoff, Eds. Springer Berlin Heidelberg, 2012, vol. 7169, pp. 84–89. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-29113-5_7
- [14] M. d'Inverno, M. Luck, P. Noriega, J. a. Rodriguez-Aguilar, and C. Sierra, "Communicating open systems," *Artificial Intelligence*, vol. 186, pp. 38–94, Jul. 2012.
- [15] J. McGinnis, D. Robertson, and C. Walton, "Protocol synthesis with dialogue structure theory," in *Argumentation in Multi-Agent Systems*. Springer, 2006, pp. 199–216. [Online]. Available: http://link.springer.com/chapter/10.1007/11794578_13
- [16] L. Guo, D. Robertson, and Y. Chen-Burger, "Using multi-agent platform for pure decentralised business workflows," *Web Intelligence and Agent Systems*, vol. 6, no. 3, pp. 295–311, 2008. [Online]. Available: <http://iospress.metapress.com/index/E26236758K550221.pdf>
- [17] D. Murray-Rust and D. Robertson, "Lscitter: Building social machines by augmenting existing social networks with interaction models," in *Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion*, ser. WWW Companion '14. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2014, pp. 875–880. [Online]. Available: <http://dx.doi.org/10.1145/2567948.2578832>
- [18] N. Gilbert and K. Troitzsch, *Simulation for the social scientist*. McGraw-Hill International, 2005.
- [19] C. M. Macal and M. J. North, "Tutorial on agent-based modelling and simulation," *Journal of Simulation*, vol. 4, no. 3, pp. 151–162, 2010.
- [20] A. Bozzon, P. Fraternali, L. Galli, and R. Karam, "Modeling crowdsourcing scenarios in socially-enabled human computation applications," *Journal on Data Semantics*, pp. 1–20, 2013.
- [21] E. Kamar, S. Hacker, and E. Horvitz, "Combining human and machine intelligence in large-scale crowdsourcing," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2012, pp. 467–474.
- [22] G. Little, "Turkit: Tools for iterative tasks on mechanical turk," in *Visual Languages and Human-Centric Computing, 2009. VL/HCC 2009. IEEE Symposium on*, Sept 2009, pp. 252–253.
- [23] S. Ahmad, A. Battle, Z. Malkani, and S. Kamvar, "The jabberwocky programming environment for structured social computing," in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '11. New York, NY, USA: ACM, 2011, pp. 53–64. [Online]. Available: <http://doi.acm.org/10.1145/2047196.2047203>
- [24] D. W. Barowy, C. Curtsinger, E. D. Berger, and A. McGregor, "Automan: A platform for integrating human-based and digital computation," *ACM SIGPLAN Not.*, vol. 47, no. 10, pp. 639–654, 2012.
- [25] P. Minder and A. Bernstein, "Crowdlang: A programming language for the systematic exploration of human computation systems," in *Proc. of Social Informatics (SocInfo'12)*, 2012, pp. 124–137. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-35386-4_10