DESIGNING A CONTEXT-AWARE DISTRIBUTED SYSTEM – ASKING THE KEY QUESTIONS

Duncan Bates, Prof. Nigel Linge, Prof. Tim Ritchings & Tom Chrisp
Centre for Networking and Telecommunications Research, University of Salford
Salford, Manchester, M5 4WT.
UK

Abstract

Our previous research identified clear business benefits and the need for an intelligent information flow algorithm that could match information delivery requirements for mobile workers to the capability of the delivery networks. The primary mechanism in our model, which acts a delivery system to the algorithm, essentially involves the matching of the user's context against a previously documented workflow process model.

However, in designing such a system it is important to determine whether it can withstand all scenarios. In this paper, we set out a series of key questions that should be asked of the design of any context-aware distributed system. From these answers, we postulate how the model we have developed can be enhanced to accommodate them.

Key Words

User Context, Workflow Process Model

1. Introduction

The primary goal of our research is concerned with the derivation of an algorithm to allow the efficient delivery of information to a mobile user, based on the user's context. This need was identified from prior research carried out to develop a location aware application for field engineers in the utility industry [1].

In this paper the primary inputs to such an algorithm are examined, together with the key issues surrounding them, in order that the subsequent derivation of the algorithm is fully investigated beforehand.

The original question posed that set out the rationale behind the research was …

"If the volume of data exceeds the available bandwidth, then how do you prioritise the information flow such that the most important items are delivered first and in the most bandwidth efficient manner ?".

The suggestion is that there will always be a fundamental mismatch between the volume of content and the available bandwidth to deliver it to the users at a location convenient to them.

Therefore a new model was proposed (figure 1) for relating distributed information sources to a user's requirements and context. This had to be firstly defined as a workflow process model which not only described each stage of the tasks being undertaken but also contained links to distributed information sources and specified where and when each piece of information was required. Secondly, it was important to determine the complete user's context. User context was defined to include a user profile, location & time information, the reason why the given task is being performed and how that data is accessed. With the addition of the access device, this fits with the findings of Dey and Abowd [2] who examined a number of context-aware applications and concluded that a generalized definition of user context should contain elements of identity, activity, location and time. User profile breaks down further into the user details, user authority & user expertise. A third input is how much available bandwidth exists between the user and each distributed information source.

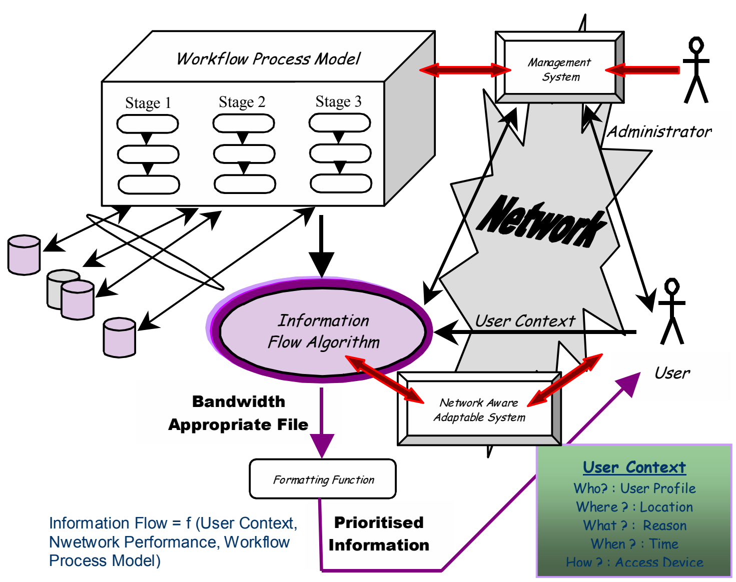

Figure 1: The Proposed Model

Once these three components are available and known, it is then possible to create an optimised information flow algorithm which can ensure a quality of delivery which determines the most relevant piece of information and delivers this to the user in a format that makes best use of the quality of their network connection. Hence, if a user has limited access capacity this algorithm would choose to send textual representations of critical information rather than graphical [3].

2. Research Rationale & Key Questions

Research Challenges

So the primary research challenge, alluded to in the question posed in the introduction, is the chosen delivery network might not always be able to deliver the required information in the highest quality. For this reason, data that is to be delivered is considered to be made up of a set of data choices. These range from the most bandwidth hungry choice, such as a video file, down to a simple text file. The set of data choices is considered to be roughly the same information content but expressed in different formats. So for example, given a network large enough in capacity, say for example a 10Mbps Ethernet, with low network loadings at the point of delivery, the top data choice would be available. However, what must also be taken into account is the user's access device. Given that that the user is accessing that data via a laptop for example, that top data choice, for example the video file, would be delivered. Clearly there are numerous other possibilities dependent on the type of delivery network, the network availability and the user's access device. This is why these components form part of the inputs into the optimised information flow algorithm.

However, the sets of data choices do not exist in isolation. They are attached to a workflow process model, or more specifically the workflow processes within them, from here on in termed 'tasks'.

Workflow Process Model

Generally, as regards workflow process models or workflow management, there are a number of standards organisations that deal with the subject. Specifically there is the Workflow Management Coalition (WfMC), the Business Process Management Initiative (BPMI) and the Object Management Group (OMG). Of these the WfMC, with both IBM and Microsoft as funding members, is the primary body. In their Workflow Reference Model [4], the WfMC describe workflow management as follows

"*Its primary characteristic is the automation of processes involving combinations of human and machine-based activities, particularly those involving interaction with IT applications and tools.*"

In May 2002, the WfMC and BPMI announced their first joint standards meeting [5], so possibly in the future there may be some degree of co-operation & convergence between the standards. However, in the meantime, although it was felt that there ought to be an awareness of the workflow management standards, referred to above, a faster approach would be gained through a proprietary workflow process model. This model would be described using UML [6] [7], and would have a generic tree-like structure into which a specific set of workflow processes would be plugged.

The workflow process model can be described in terms of the discrete stages within a given worker's role. In order to define the workflow process model, the worker's role needs to be broken down into discrete stages. Furthermore those stages are then broken down into the tasks required to complete that stage of the work. Attached to some of these tasks are the sets of data choices that were referred to earlier. Those data choices may be distributed in that it may, for example, reside in different databases or data sources on different machines.

The structure of the workflow process model is described as a generic-specific structure, in that it is set up to be generic as regards the breakdown of the model into processes (process titles), stages and tasks. At the base level it allows specific workflow process models to be plugged into the generic structure, i.e. at the level of the set of tasks. Each 'generic task' has a redirect URL as one of its attributes / fields thus allowing the generic set of tasks to redirect to a real set of tasks. The real set might require extra database tables and thus extra classes in order to process the information from them. Similarly, extra 'specific' information may need to be delivered across to the workflow process model (server) from the management system (server).

As part of the research, we are examining a small number of case studies to describe their workflow processes. The project that preceded this work and the workflow processes were documented within the report that went with that project [8]. These have been anonymised and act as one of the case studies. Further work to get other people and/or companies on board that primarily employ mobile working practices is currently being investigated. Eventually enough case study material will allow the research to work towards a truly generic solution.

Expertise

Figure 2 shows a set of workflow processes for a communications worker. This diagram is a UML activity diagram and shows the activities linked by transitions and decision points (diamond symbols). The responsibilities for each of the given activities are designated by the so-called swimlanes (vertical dashed lines). So the communications worker, who in this case is described as a

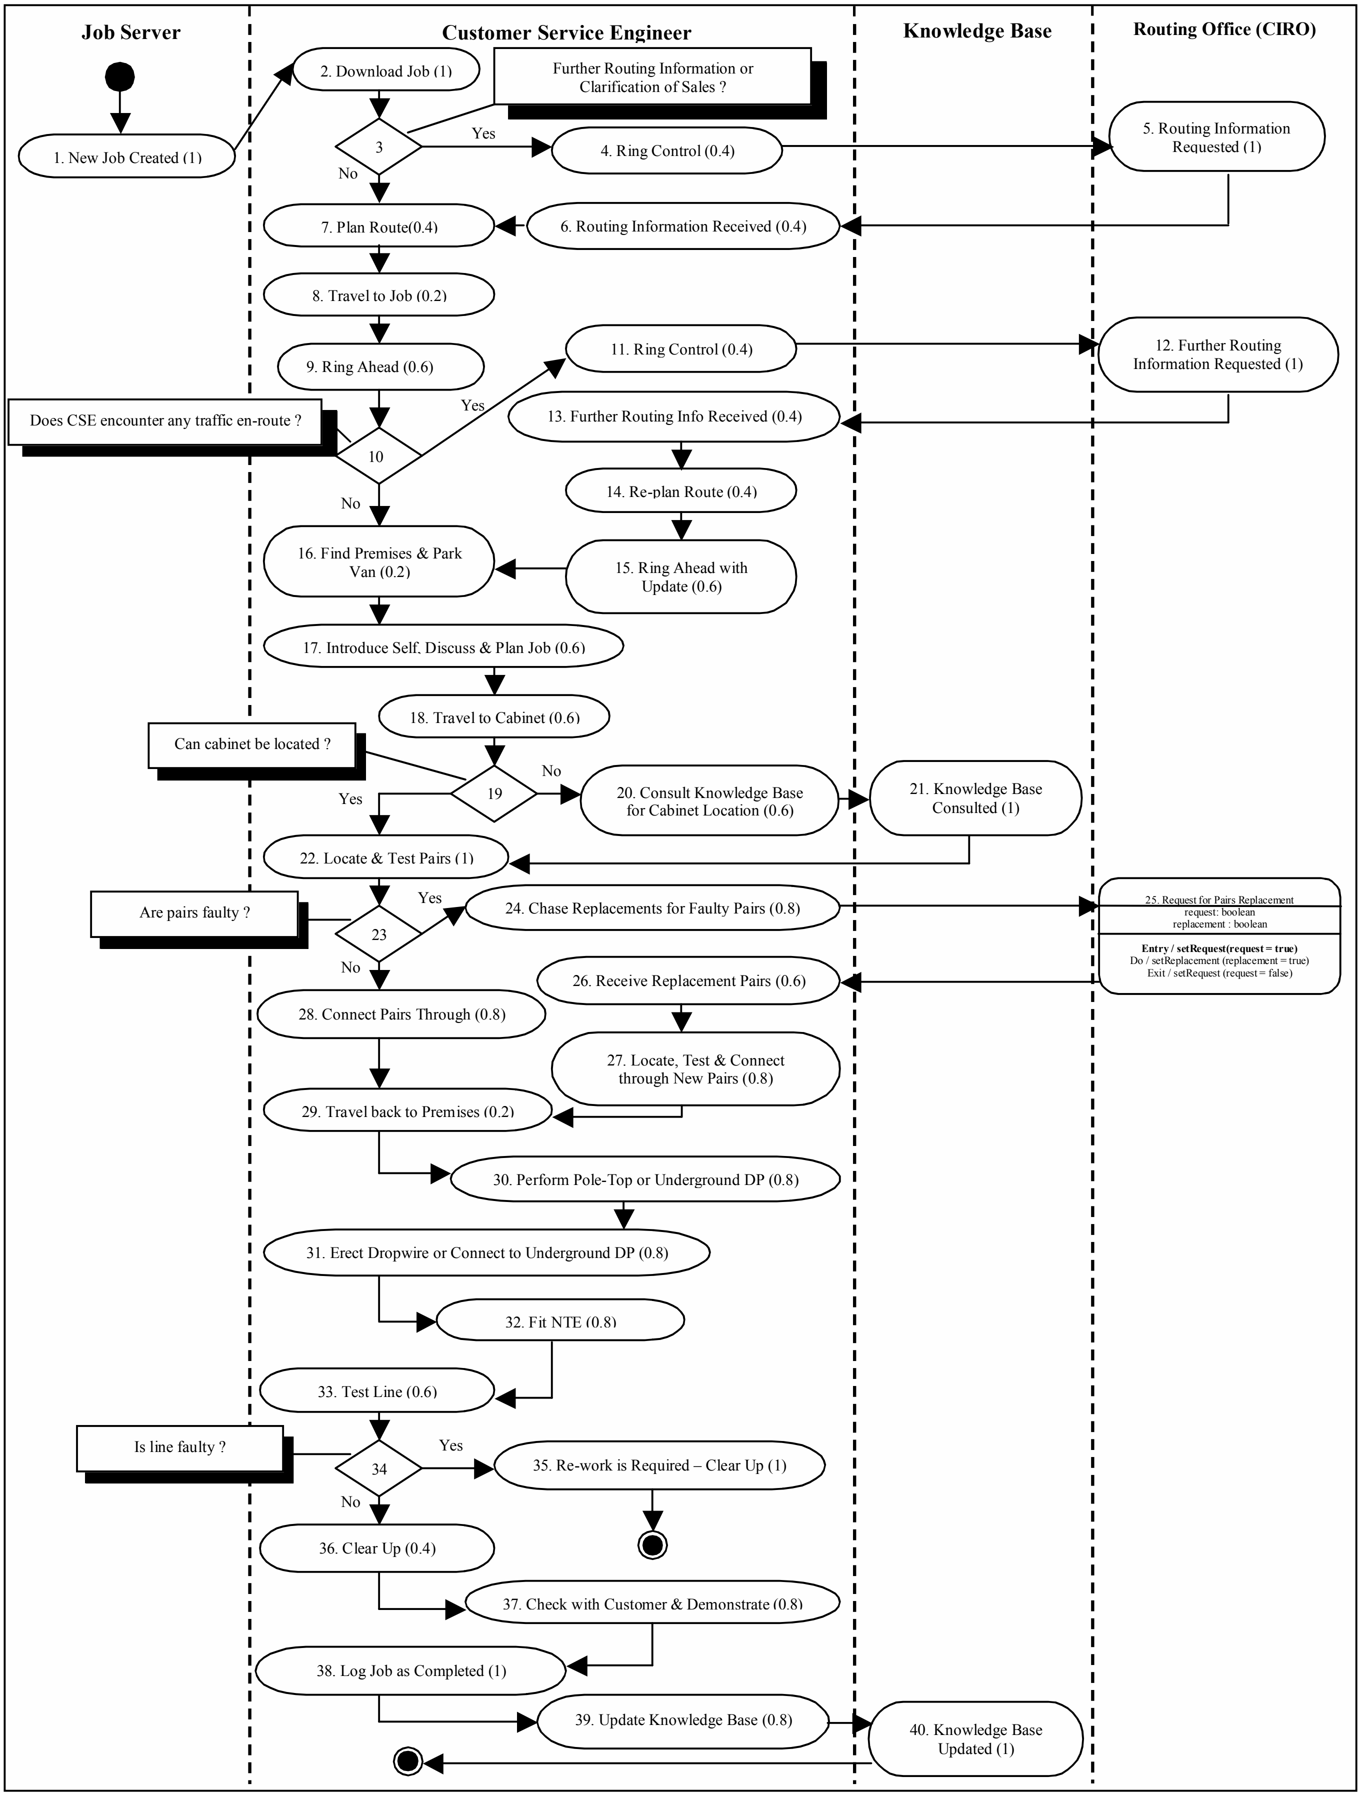

Figure 2 : Workflow Process Model Stage for a Communications Worker – User Expertise 0.2

customer service engineer, is responsible for most of those tasks, but he/she interacts with other people or systems, which collectively are known as actors.

A key attribute of user context is the expertise of the user because this has a bearing on the delivery of given tasks. A beginner on this particular job might need all these tasks delivering to him/her, but in contrast for the expert user this would become tiresome, and he/she might only require some of the tasks as prompts. In order to accommodate this into the model, not only must there be a value of user expertise, but also of task expertise against which to compare it.

There is a similar argument that can be made for user authority and task authority, although the comparison is the other way round (users need more authority than the task as against less expertise in order for the task to be delivered). However, the focus of this discussion will remain on expertise, and these ideas will be explored later.

J2EE Model

In order to develop a distributed application capable of delivering user context information across a network, drawing on disparate database sources getting, for example, job related information for that user and then match that information against a previously stored workflow process model for that user's role, it is likely that a multi-tier architecture needs to be considered. Moreover, since the user needs to be able to access the information remotely using a variety of devices, such as a mobile, PDA or laptop, then a web application seems the most reasonable way forward.

Our model was designed to meet the specifications of the J2EE platform [9]. The model is essentially divided into two and formed of a management system server and a workflow process model server, both J2EE compliant servers. The user logs on remotely to the management system, out in the field say. The user's context is then constructed and the user chooses one of the jobs previously assigned by an administrator. The essential user context & job information is then passed across to the workflow process model in order to be matched against the correct workflow process model of which the job is an instance. The specific tasks and any attached data are then returned to the user via the information flow algorithm. That is of course, if the algorithm determines that the task ought to be passed based on the user's authority & expertise. Any attached data is returned in the format appropriate to the user's access device and network availability at the time of delivery.

A Set of Key Questions

At this stage, a number of pertinent questions were asked of the model in order to examine areas that required closer attention. These questions are listed below …

1) *How does a given stage of a workflow process model act under variable user expertise ?*

2) *What should the system do if the user deviates from the documented workflow process model or the model doesn't match the user's experience ?*

3) *How can the logic of the workflow process model be expressed to the user ?*

4) *Are there elements of user context that will be revealed through new case studies that are not yet considered important ?*

5) *Should some data always be delivered in the format of the highest data choice, and therefore would the user be prepared to wait beyond the chosen delivery time ?*

6) *Does user expertise have an effect on data choice as well as delivery of tasks ?*

7) *How does location actually have a bearing on the system ?*

Turning to the first question on that list, let us examine how a given set of workflow process tasks acts under different user expertises. In other words, how does the workflow process model for those tasks change when the delivery or non-delivery of tasks is taken into account ? The delivery of a given task is determined by the algorithm, and would seem to take the form of a simple comparison between user expertise & task expertise, as in the pseudo-code below.

```
if (User Expertise <= Task Expertise) {
        Deliver Task
}
```

Figures 2-4 illustrate a stage of a workflow process model that is shown under variable user expertise. The tasks have been arbitrarily assigned task expertises from 0 – 1 (0.2, 0.4, 0.6, 0.8 or 1) and the first diagram (figure 2) shows the stage for a beginner (user expertise : 0.2). Consequently all the tasks would be delivered to the user, because the user expertise is always less than or equal to the task expertise. However as user expertise increases, you would expect that the user would not want to see all the tasks, as this would become tedious and unnecessary. This is the case with figure 3, which shows the same stage under a user expertise of 0.8.

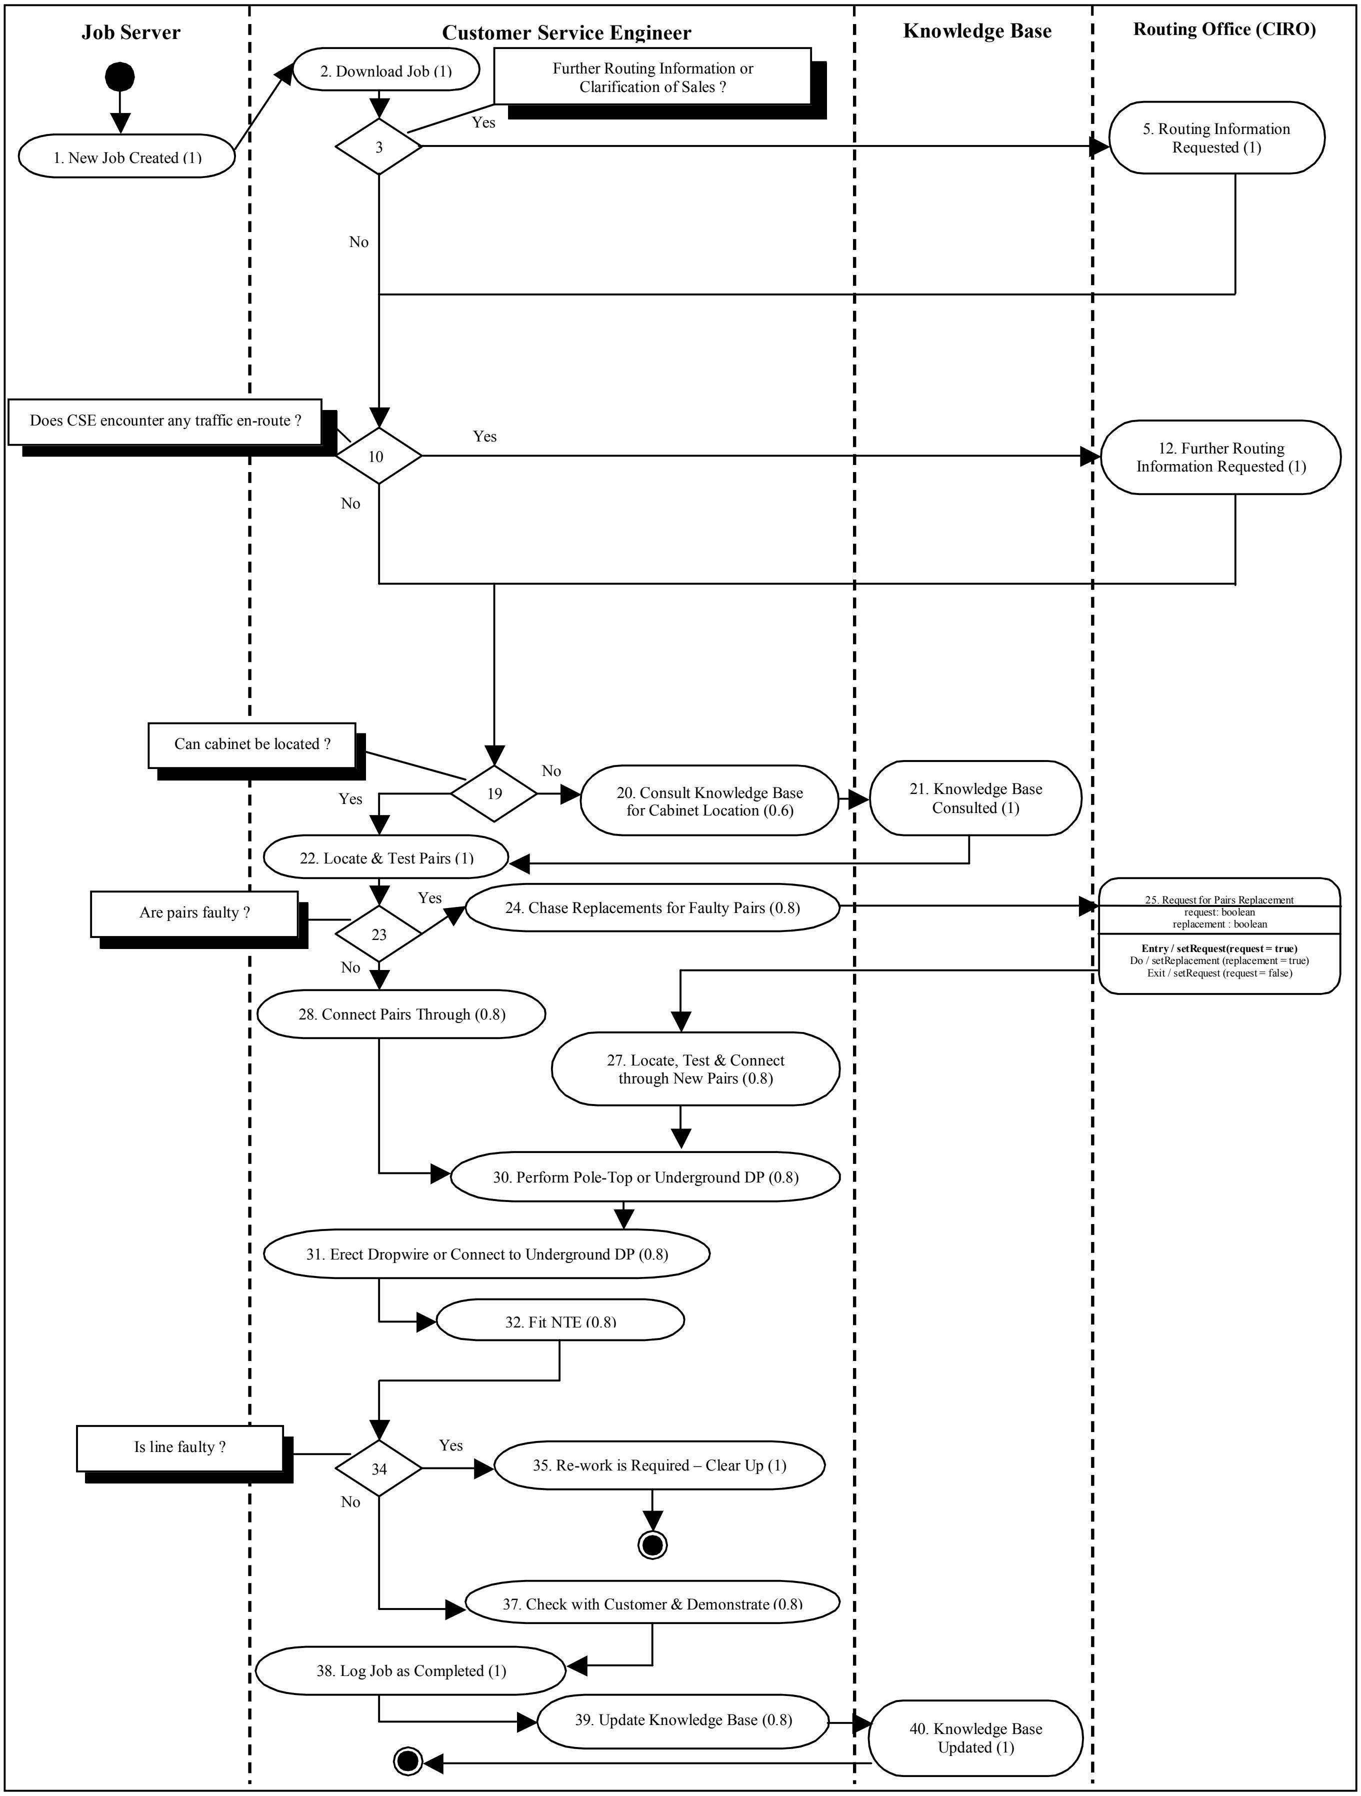

Figure 3: Workflow Process Model Stage for a Communications Worker – User Expertise 0.8

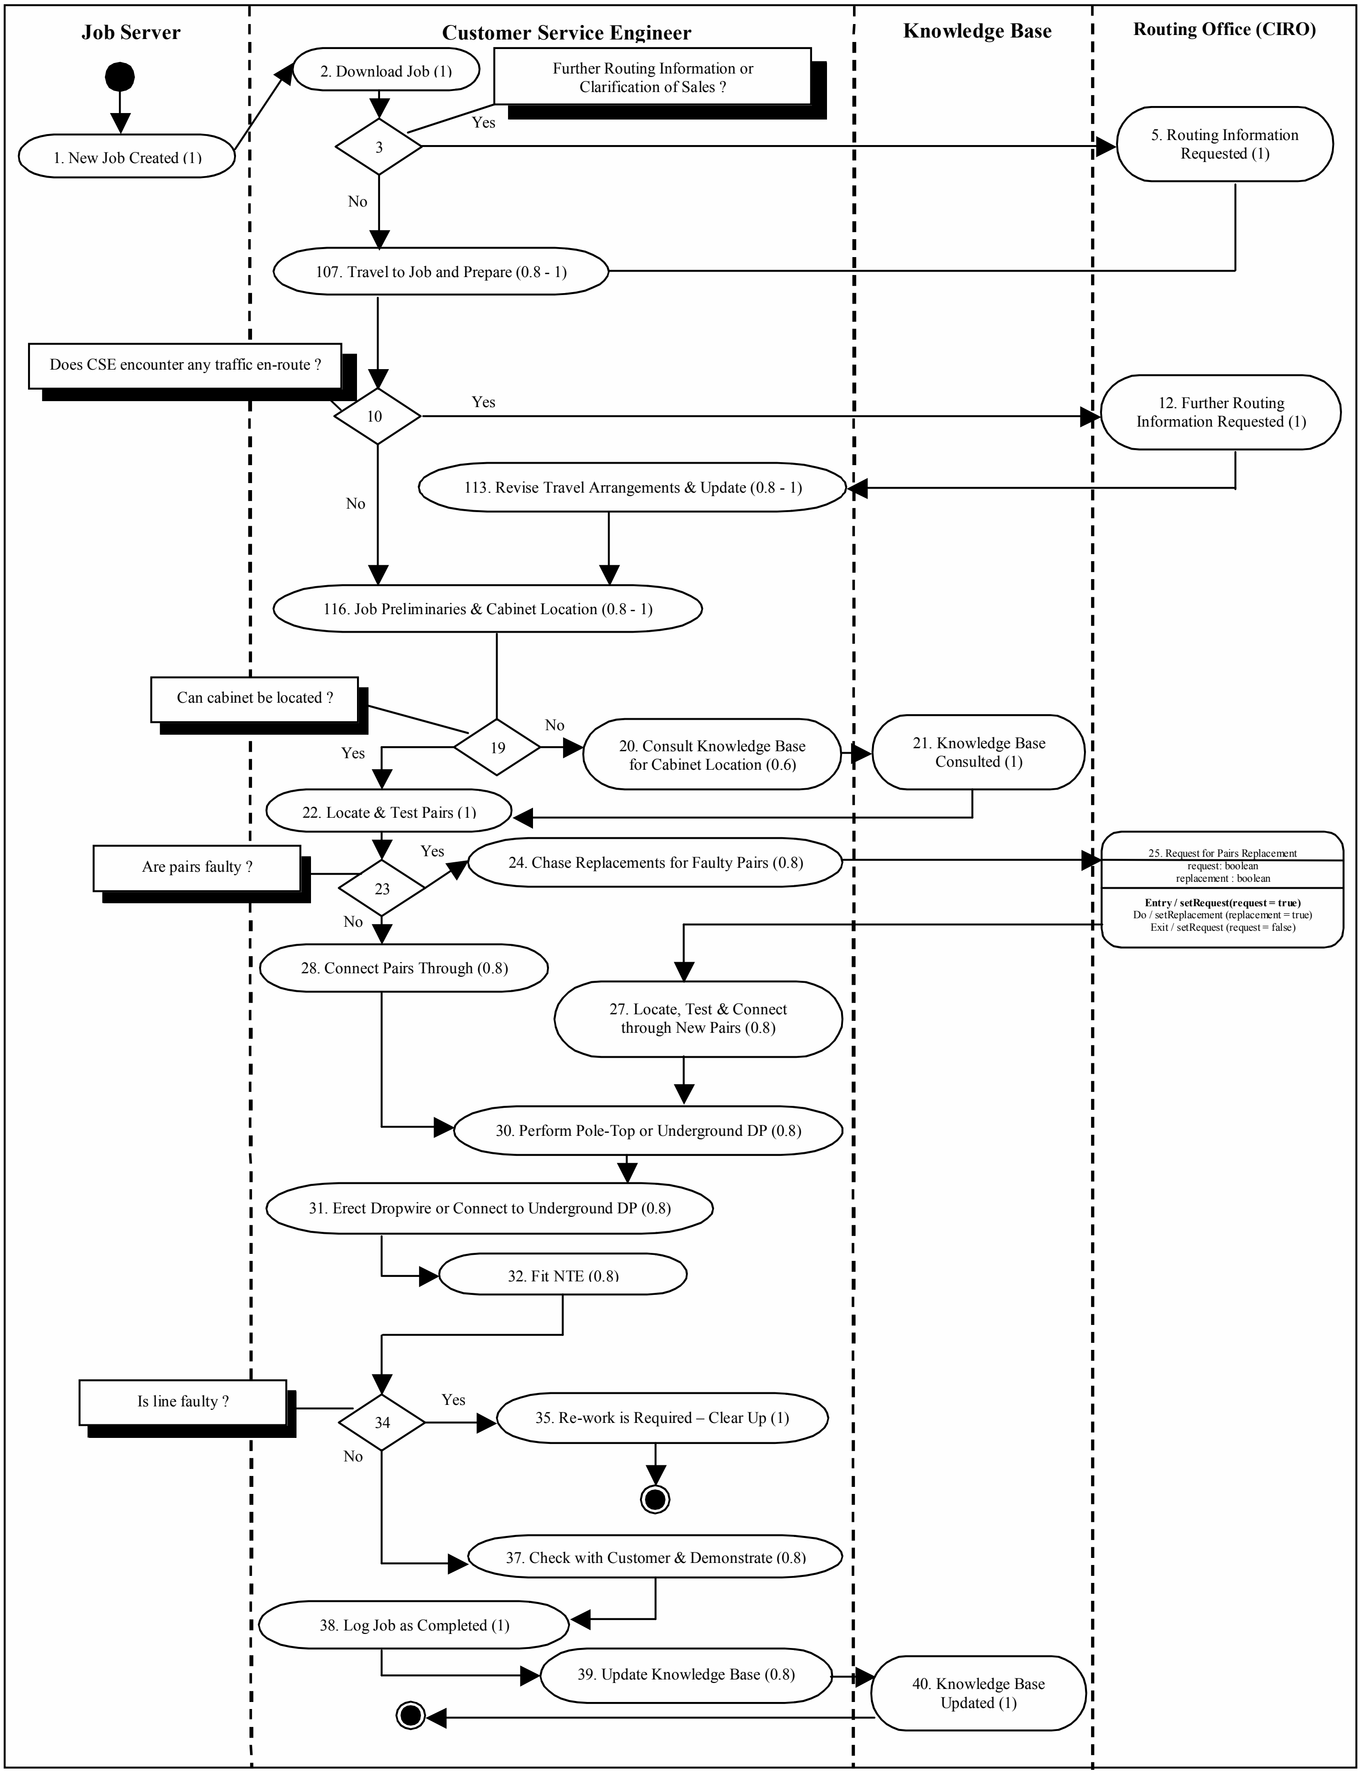

Figure 4: Workflow Process Model Stage for a Communications Worker with Advanced Tasks – User Expertise 0.8

It should be immediately clear what the problem with this is because the diagram has become too sparse, especially up near the top end of the diagram where the task expertises were set low. The result of this is that the meaning of the workflow will be lost. Although the more expert user doesn't need to see all the tasks, simply removing the tasks that are of an expertise lower than his/her expertise just results in a 'disjointed story of events'. To compensate for that, it is necessary to fill the gaps that are left with something to keep the workflow meaning (figure 4). For this diagram, it is proposed advanced tasks should be substituted for about every 3 or 4 'elementary' tasks. So for example task number 116. 'Job Preliminaries & Cabinet Location' has replaced 16. 'Find Premises & Van', 17. 'Introduce Self, Discuss & Plan Job' and 18. 'Travel to Cabinet' (N.B. the task numbers are arbitrary).

So a simple comparison of user expertise & task expertise needs to be refined to include a minimum task expertise and a maximum task expertise. So …

if (Task Expertise Min <= User Expertise <= Task Expertise Max)
{ Deliver Task }

The insertion of advanced tasks for a few elementary tasks is also consistent with working practices, in that sometimes expert users tend to skip tasks in a workflow process because they feel that they know the job, but in doing so follow incorrect procedures through over-confidence.

KM

There is some applicability here to the field of knowledge management (KM) [10][11]. Nonaka [12] draws on the work of Polyani in his distinction between tacit and explicit knowledge. Tacit knowledge is personal & context-specific and can be difficult for an individual to express, whereas explicit knowledge is that which is codified, i.e. written down in some form or other, e.g. in books, documents, manuals, computer programs etc.

He talks of four modes of knowledge conversion ([12], p62). Tacit to tacit knowledge is classified as socialization, and is exemplified by a master craftsman passing on his skills to an apprentice. Externalisation is the process of turning an individual's tacit knowledge into something that is written down for others to draw upon. Combination is the combination of small concepts or ideas into grander concepts, or maybe one could make a case for a set of software components combined into a software system. Finally internalization is the act of taking explicit, codified knowledge back on board as an individual's tacit knowledge.

So there is a kind of cycle from tacit knowledge through explicit knowledge back to tacit knowledge. However, this is only at the level of an individual within the so-called epistemological dimension. Individual knowledge must be dissipated throughout an organisation in order for new knowledge to be created within the organisation as a whole. So it is possible to envisage two dimensions of knowledge creation, the epistemological dimension, which we have already looked at and the ontological dimension, which is the level at which the knowledge resides within the organisation ([12], p57).

So if you combine the four modes of knowledge conversion that occur within a single epistemological dimension with the ontological dimension, then what occurs is a knowledge spiral.

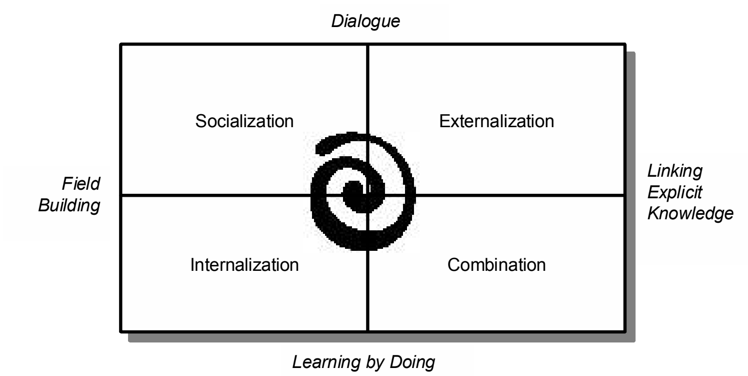

Figure 5: The Knowledge Spiral ([12], p71)

The knowledge spiral is a cycle around the four modes of knowledge conversion (socialisation, externalisation, combination & internalisation), and is therefore known alternatively as the SECI cycle. This knowledge then dissipates upwards throughout larger and larger groups within an organisation until eventually knowledge is held inter-organisationally. Tacit knowledge must still be held within individual's heads by definition, but it becomes commonly held beliefs and knowledge throughout the organisation.

Regarding a mobile worker, it is an individual under consideration, rather than an organisation as a whole. Therefore this is at the single, base ontological level, i.e. that of the individual. What is it that would stimulate the SECI cycle ? Well an obvious candidate is a knowledge base of the mobile worker's ideas, hints and tips. This also has applicability to the second question posed at the beginning of this section, i.e. what to do if the user deviates from the set of tasks.

Firstly, one needs to ask why a user would deviate from a set of tasks. The most likely scenario would be that the user, particularly if he/she is relatively experienced, feels that the tasks, as laid out, don't fit the situation that he/she is experiencing. Clearly, in this case, there is a definite need for the knowledge base. If the tasks don't fit the situation then the workflow process model is insufficient to cover all eventualities, and as far as the user is concerned is incorrect, which would likely result in deviation. The workflow process model would then need to be adjusted in order to accommodate the new situation highlighted by the user. So as the system gets used more and more often, the novel situation should come about less often. This is also having the effect of stimulating movement around the SECI cycle and up the ontological levels within the organisation. Users' tacit knowledge is codified initially in the loose form of hints and tips and then more systematically into updates to the given workflow process model. This knowledge can then be internalised into the tacit knowledge of less experienced users. So knowledge is spread upwards throughout the organisation from the less experienced up to the more expert.

Expressing Logic

It should be apparent that the stage presented in figures 2-4 isn't just made up of tasks. There are also decision points, which result in multiple paths. There could also be parallel processing points, or in UML terms, synchronisation bars at which paths split and join again. The workflow process model diagrams are actually just UML activity diagrams, so the components that make up those diagrams are the ones that are applicable to them. Therefore the system devised to deliver the primary inputs to the algorithm must be capable of expressing the logic of the workflow process model or the UML activity diagram, as posed in question 3 above.

There could probably be several different ways of expressing the logic of the diagrams, but the one chosen for the project just flattens out the decisions and pathways into a set of records. So for example the top part of the stage from the start point down to the decision point "Does CSE encounter any traffic en route ?" can be represented as a set of logic sequences …

								Path
0	*1*	*2*	*3D*					*0*
3D	*7*	*8*	*9*	*10D*				*0*
3D	*4*	*5*	*6*	*7*	*8*	*9*	*10D*	*1*

Table 1 : Logic Sequences

The numbers are the task numbers that although arbitrary, are just designated roughly in sequence for convenience. The 'D' indicates a decision point rather than an ordinary task and a sequence runs from a start point, decision point or parallel processing point to an end point or another decision / parallel processing point. The path parameter indicates the number of paths that lead out of a decision / parallel processing point. So if there is only one path, it will just be designated zero. However if there is more than one path, they will be numbered accordingly from zero upwards. In the typical case of yes/no answers or true/false, path 0 indicates the false path and path 1 indicates the true path. Given the information above, it is possible to reconstruct the workflow process model.

In order for a system to reconstruct the model within software, it would need the records to be written in a consistent way. There would also need to be a software component that understands how to reconstruct the model in the way that a human would on paper. A table of logic records would need something like the following fields …

logic_id :	Primary key field
task_id :	Matches the task number
decision_flag :	Indicates whether this task is a decision point
parallel_flag :	Indicates whether this task is a parallel processing point
decision_path :	Matches path parameter
stage_id :	Foreign key field to the stage

So the first line of the logic sequences above would result in 4 records.

logic_id	*task_id*	*decision_flag*	*parallel_flag*	*decision_path*	*stage_id*
1	*0*	*False*	*False*	*0*	*3*
2	*1*	*False*	*False*	*0*	*3*
3	*2*	*False*	*False*	*0*	*3*
4	*3*	*True*	*False*	*0*	*3*

Table 2 : Records for First Logic Sequence

Task_id, numbered 0 indicates a start point of which there is only one in any given stage. There could be more than one stop point, so these could be numbered within a range, 990-999 say. The next line in the logic sequences, the false path out of decision point 3 would result in 8 more records etc.

As the logic sequences start and end at decision points or parallel processing points, it should be clear that these should be treated as tasks that must always be delivered, regardless of the user's expertise. It has already been demonstrated that workflow meaning is lost as the user's expertise increases, leading to the need for additional advanced tasks. If the decision points were not always delivered, the meaning would be lost altogether because the logic sequences would not be clear and the system would not know what order to (potentially) deliver tasks to the user.

Note that the logic records do not actually tell the system whether to actually deliver the tasks, that decision is taken by the algorithm based on the user & task expertise which are set against the respective records for the user & task. So what are actually missing from the logic sequences above are the advanced tasks. In figure 4, advanced tasks have been substituted for some of the more basic tasks, which would not get delivered to the more expert user (user expertise 0.8 in this case).

So in the logic sequence from the start point to decision point 10, an advanced task, numbered 107, has been inserted in place of tasks 7-9. So this would actually have to be in the logic sequence, after tasks 7-9 say, as below…

									Path
0	*1*	*2*	*3D*						*0*
3D	*7*	*8*	*9*	*107*	*10D*				*0*
3D	*4*	*5*	*6*	*7*	*8*	*9*	*107*	*10D*	*1*

Table 3 : Logic Sequences with Advanced Task

This would result in an extra couple of logic records inserted into the false and true paths out of decision point 3. There is nothing in those records to indicate whether that task should be delivered. That information is held on the task record as the task expertise min (minimum) and the task expertise max (maximum), which would be set at 0.8 & 1 in this case. Consequently the algorithm would determine that this task only gets delivered to users with expertise above 0.8.

Further Insights

Turning to question 4 of the set of questions posed at the beginning of this section. Are there elements of user context that will be revealed through new case studies that are not yet considered important ? Well clearly this question can't be answered yet, but the research project is currently working with an organisation in the area of health. This will present a different kind of case study because it won't be one that is entirely based on field workers. The project has also been permitted access to this organisation to interview members of staff in order to determine workflow processes from the top down and bottom up. This will allow the project to get in at the beginning of the documentation of a workflow process model, and also to steer the work somewhat towards the needs of the project. So hopefully, there will be area of the study that will reveal new insights that are applicable to the project, specifically elements of user context that have not been considered and how knowledge is cycled throughout the organisation.

Question 5 refers to an issue concerned with the algorithm itself, and the delivery of attached data. The algorithm may determine the most bandwidth appropriate choice to deliver to the user according to a snapshot taken of the network availability at the given point of delivery. However, what happens if the user gets sent a lesser choice than he/she desires because the network availability is low ? It would seem that the most appropriate thing to do in this situation is for the user to be asked if they are happy with the data choice delivered or whether they want to wait the appropriate length of time to deliver their desired choice. Of course the irony with this is that by the time it comes to actually delivering their desired choice, the network availability may have increased sufficiently to actually accommodate it within the designated download time.

A Matrix of Data Choices

As an answer to question 6), the further question that could be posed is what happens if the network availability would permit delivery of the highest data choice, a 20Mb video file say, but the user is of sufficient expertise that he/she would rather receive a more succinct version, e.g. an advanced text file or perhaps a smaller video file ? The diagram below illustrates the point..

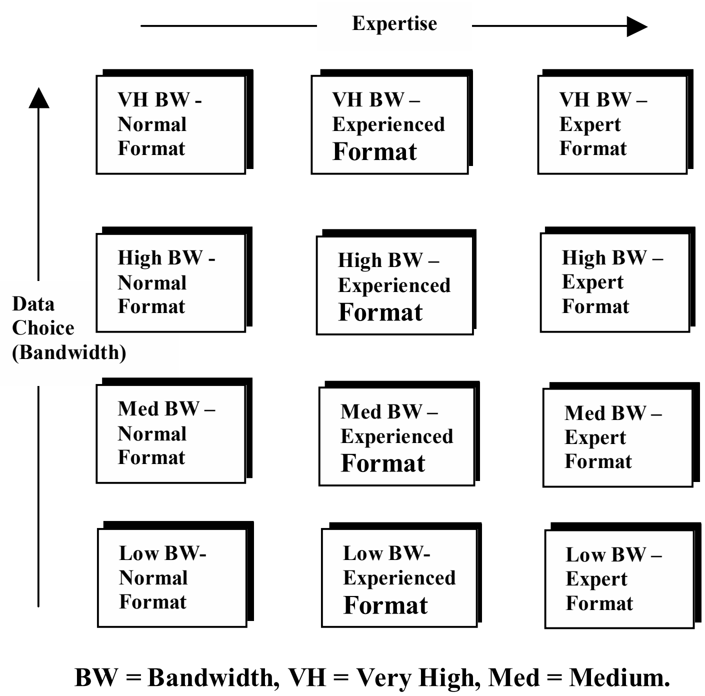

Figure 6: Data Choices Matrix

Here the normal formats correspond to attached data that is delivered if only the network availability is taken into account. User expertise is taken into consideration on the delivery of tasks but the maximum data choice permissible is delivered regardless. In order to deliver a data format more closely aligned to expertise, a matrix of data choices would need to be constructed as shown in figure 6. Although it isn't immediately apparent from the diagram, there is a degradation in file size and therefore required bandwidth down both axes, i.e. down the data choice axis and across the expertise axis as expertise increases. If this diagram was constructed on a strictly Cartesian graph, the lines across would not be parallel but curve down to the right in correspondence to the reduction in required bandwidth. This argument illustrates that maybe an expertise value should also be placed on the attached data as well as the tasks. This will be something that will be explored in future work, as will the topic of the next and final question.

Location

Question 7 asks how location actually has any bearing on the system. The question seems somewhat odd at first because it would seem obvious that location ought to be a vital piece of user context, as deemed by the plethora of location based services at the planning stage in the mobile industry right now [13]. In actual fact, other than delivering the appropriate data choice back to the correct user location, its bearing on the algorithm isn't obvious.

Where location would have a bearing is on pieces of equipment out in the field that the mobile user might have to maintain or fix. A workflow process model might indicate that a user ought to engage with a piece of equipment in some way, but that piece of equipment referred to in the model is representative of all pieces of equipment of that type. However, if there are many instances of that equipment, they might be in different states of repair. Therefore their location suddenly becomes important as far as the system is concerned, in order to differentiate between instances of that equipment type. What this would entail is an inventory of all pieces of equipment for the given company & their current state. The given workflow process model would need a relationship to the equipment type that is applicable at that stage.

A possible solution for this actually feeds back into question 4, which enquires whether new case studies will reveal anything new about user context. It is clear that equipment state is not always applicable to a given worker's role. For example, in one branch of the health sector, delivery of patients to a hospital might be the primary goal. Therefore, it could be suggested that equipment state is a piece of context, specific to mobile field workers. This would fit within the generic-specific workflow process model concept outlined earlier. Perhaps it is wise to bear in mind here another of the key concepts drawn from knowledge management [14], that a researcher ought to be reflexive ('reflective'). Reflexivity is the recognition by the researcher that his/her perspective, embedded in language, culture, (scientific) community etc, has an influence on the effects observed So it is important to stress that this could be a case of fitting the concepts to the model, rather than forming the model from them. However, Dey and Abowd's general definition of user context is fully accommodated, and it is clear that 1 non-specific model could ever fit all possible variations in workflow process models.

A set of criteria that could be passed within a specific context for mobile field engineers, as regards equipment state, could be equipment type, equipment & faults where equipment is of a given equipment type and faults are specific to that equipment type. At a given location there would be any number of a particular equipment type, and each of these pieces of equipment could have a different set of faults, i.e. equipment state of a particular equipment type could vary. So this could be passed to the workflow process model and acted upon. However, this would only be the case if the specific workflow process model was written in such a way as to exploit it. For example, a decision point might lead to different paths based on given faults with a particular equipment type. This emphasises the specific nature of the applicability of location to equipment state. It is passed as extra information in the form of a specific context, and is acted upon providing the specific set of workflow process tasks are set up to exploit any differences in equipment state.

3. Conclusion

A model has been developed based around a J2EE design for a distributed system to deliver a user's context to a previously documented workflow process model. In this way, the user context, the workflow process model, and in addition, the network availability form three inputs to an optimised information flow algorithm. The algorithm determines whether to deliver back workflow tasks to the user based on the user's expertise and authority against the equivalent values on the task. Data may be attached to an individual task and exist as a set of data choices. The algorithm in its more complex form determines which of these data choices to return to the user, based on the most bandwidth appropriate choice deliverable within an appropriate cut off time given the network availability at the point of delivery.

During the design of the J2EE model, a set of key questions was asked of the model to determine whether it could withstand the situations to which they refer. It can be argued that it is this elucidation of ideas that justifies the adoption of a design approach for the early stages of the work. Development towards a full J2EE application may seem preferable, and indeed is not ruled out for future work, but the application is just the delivery system around the algorithm, and it is the algorithm that is the main focus of our work. Research time is at a premium and in the early and middle stages it was felt that more could be gained from focusing in on the algorithm. Further work is ongoing on the example scenarios or case studies, allowing the system further examples of workflow process models around which to base the system. As mentioned in one of the key questions, it is possible that further insights could be revealed on the nature of context itself.

Other than the case studies, future work points to the development of a full J2EE system, or more likely the development of a small prototype based around the algorithm. Inputs to that algorithm have been identified by working through the J2EE design and the posing of the set of key questions. This has allowed the algorithm to be isolated and tested individually, but it is unlikely that this could have been done without prior rigorous examination of the full system, part of which has been set out in this paper.

4. Acknowledgement

The authors wish to acknowledge the Engineering and Physical Sciences Research Council (EPSRC) who originally funded this research work by grant GR/L 05129 and who are continuing to support the research through grant GR/N 35823.

References

[1] M. Hope, T. Chrisp, N. Linge, Improving working practices within the utility industry through mobile context aware geographic information systems: *8th International Symposium of ACMGIS*, Washington D.C., USA, November 10th-11th 2000.

[2] A.K. Dey and G.D. Abowd, Towards a better understanding of context and context-awareness: *2000 Conference on Human Factors in Computing Systems (CHI 2000)*, The Hague, The Netherlands, April 3, 2000.

[3] D.R. Bates, N. Linge, M. Hope, Context Aware Information Services for Mobile Workers, *3rd Annual Postgraduate Symposium on the Convergence of Telecommunications, Networking and Broadcasting, (PGNet 2002)*, Liverpool John Moores University, UK, 2002, 390-395.

[4] D. Hollingsworth, *The Workflow Reference Model*, http://www.wfmc.org/standards/docs/tc003v11.pdf, (Workflow Management Coalition, Jan 1995).

[5] Workflow Management Coalition, *XML-Based Workflow and Process Management Standards: XPDL, Wf-XML*, http://xml.coverpages.org/wf-xml.html, (Workflow Management Coalition, May 2002).

[6] S. Bennett, J. Skelton and K. Lunn, *Schaum's Outlines of UML* (International: McGraw-Hill, 2001).

[7] J. Holt, *UML for systems engineering – watching the wheels* (London, UK: IEE Professional Applications of Computing Series 2, 2001).

[8] M. Hope and N. Linge, Determining the Propagation of IEEE 802.11 Radio LANs for Outdoor Applications, *Proceedings of the 24th Conference on Local Computer Networks (LCN 99)*, Lowell, Massachusetts, 18-20 October 1999, 49-50.

[9] Sun Microsystems Inc., *The Java™ 2 Enterprise Edition Developer's Guide, Version 1.2.1,* http://java.sun.com/j2ee/sdk_1.2.1/devguide1_2_1.pdf, (Sun Microsystems Inc, May 2000).

[10] C.W. Choo, *The Knowing Organization – How Organizations Use Information to Construct Meaning, Create Knowledge and Make Decisions* (New York, NY, USA: Oxford University Press, 1998).

[11] R. McAdam and S. McCreedy, A Critical Review of Knowledge Management Models, *The Learning Organization, MCB University Press*, *6*(3), 1999, 91-100.

[12] I. Nonaka and H. Takeuchi, *The Knowledge-Creating Company – How Japanese Companies Create the Dynamics of Innovation* (Oxford, UK: Oxford University Press, 1995).

[13] *IEE European Workshop on Location Based Services 2002*, London, UK, 16th September 2002.

[14] P. Johnson and J. Duberley, *Understanding Management Research – An Introduction to Epistemology* (London, UK., Thousand Oaks, CA., New Delhi, India: Sage Publications, 2000).