How Software Architects Focus Their Attention

Eoin Woods*, [0000-0001-7858-1552] Rabih Bashroush [0000-0002-9610-0027]

* - corresponding author
University of East London, University Way, London, E16 2RD, UK
{eoin.woods, rabih.bashroush}@uel.ac.uk

Abstract. As part of our software architecture research and practice we have found that a common difficulty for new architects is knowing where to focus their attention to maximise their effectiveness. This led us to wonder whether successful experienced architects have any common techniques or heuristics that they use to help them achieve this. In an earlier study where, having interviewed experienced architects, we found that in fact there were some common heuristics that they use, we created a simple model based on an analysis of their advice. In this paper we explain how we validated that model with a wider survey of experienced enterprise and software architects and, from the findings of that study, extended the model with an additional dimension. This resulted in our model having four primary guidelines, which are: focus on stakeholder needs and priorities, prioritise time according to risks, delegate as much as possible, and ensure team effectiveness.

Keywords: software architecture, software architecture decision making, software architect effectiveness.

1 Introduction

In our research and practice in the field of software architecture, we have noticed and experienced that it is difficult for software architects to focus their attention. The software architect's responsibilities are broad and in principle they can be involved in almost any technical aspect of a project from requirements to operational concerns.

However, we also observe that successful software architects appear to be very good at focusing their attention effectively, which led us to wonder how they achieve this. They may use time management techniques (like [2]) but we wondered whether there are common role-specific heuristics which could be taught to new architects.

In a previous study [15], we decided to investigate this via a questionnaire-based study of a group of experienced architects. We discovered that there are common heuristics which experienced architects use to focus their attention and we created a model to capture and relate them.

In this paper we explain how, in a second study, we then validated the model with a much wider group of software and enterprise architects, via an online questionnaire, and refined the model based on their input.

In the next section of this paper we present the refined model and, in the rest of the paper, we briefly summarise the earlier work and then explain how we went about validating the original model and identifying that an extension to it was required.

The contribution of our work is not the heuristics in our model, as most of them are quite familiar to experienced practitioners. Our contribution is to capture them clearly in a simple, coherent model, and to validate their value to experienced practitioners. Our work makes the implicit knowledge held in the heads of experienced architects explicit and accessible. We believe that this makes the model a useful reminder for experienced practitioners and an effective teaching aid for new architects.

2 A Model for Focusing Architectural Attention

Our experience-based model to guide architects where to focus their attention in order to maximise their effectiveness is shown in Fig. 1.

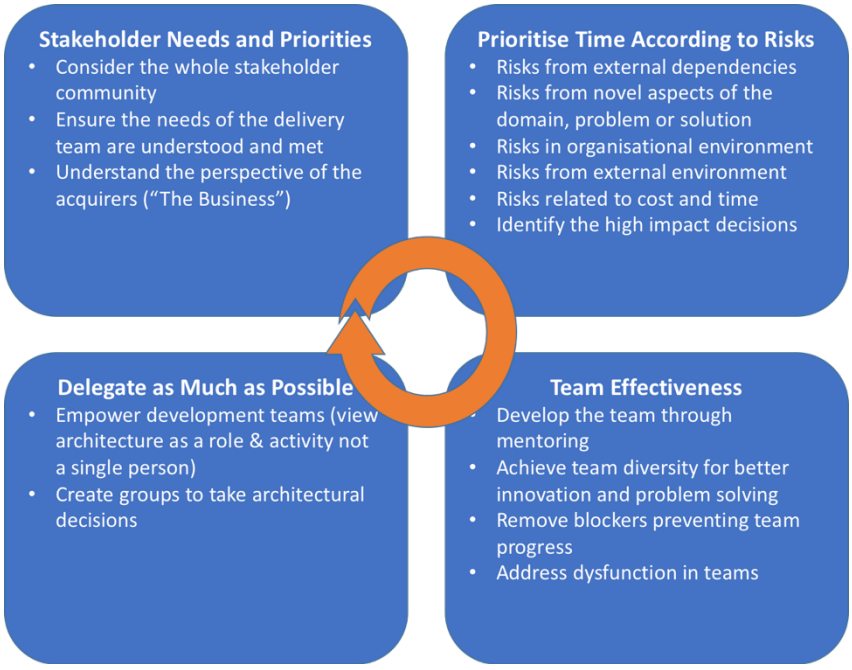

Fig. 1. Refined Model for Focusing Architectural Attention

The model is comprised of 4 aspects: Stakeholder Needs and Priorities, Prioritise Time According to Risks, Delegate as Much as Possible and Team Effectiveness. It

provides a guide, or an aide memoir, on how to prioritise architectural work to maximise its effectiveness. Each of these aspects is explained in the subsections below.

2.1 Stakeholder Needs and Priorities

The first theme which emerged strongly in our study was to focus on the needs and priorities of the stakeholders involved in the situation. The principle that architecture work involves working closely with stakeholders is widely agreed [3, 13] and this theme reinforces that.

Architects need to focus significant effort to make sure that stakeholder needs and priorities are understood, in order to maximise focus on the critical success factors for a project and maximise the chances of its success. Three specific heuristics to achieve this which emerged from the study are:

- *Consider the whole stakeholder community*. Spend time understanding the different groups in the stakeholder community and avoid the mistake of just considering obvious stakeholder groups like end-users, acquirers and the development team. As the architecture methods referenced above note, ignoring important stakeholders (like operational staff or auditors) can prevent the project meeting its goals and cause significant problems on the path to production operation.
- *Ensure that the needs of the delivery team are understood and met*. Spend sufficient time to ensure that the delivery team can be effective. What is the team good at? What does it know? What does it not know? What skill and knowledge gaps does it have? These areas need attention early in the project so that architecture work avoids risks caused by the capabilities of the team and that time is taken to support and develop the team to address significant weaknesses.
- *Understand the perspective and perceptions of the acquirers of the system*. Acquirers are a key stakeholder group who judge its success and usually have strategic and budgetary control, so can halt the project if they are unhappy. Addressing this group's needs, perceptions and concerns emerged as an important factor for experienced architects in our study. Acquirers are often senior managers and so may be distant from the day-to-day reality of a project and need regular, targeted, clear communication to understand their concerns and ensure that they have a realistic view of the project.

2.2 Prioritise Time According to Risks (Driven by Impact x Probability)

During a project, an effective approach to prioritising architectural attention is to use a risk driven approach to identify the most important tasks. If the significant risks are understood and mitigated, then enough architecture work has probably been completed. If significant risks are unmitigated, then more architecture work is needed. The specific heuristics to consider for risk assessment are:

- *Risks from external dependencies*. Understand your external dependencies because you have little control over them, and they need architectural attention early in the project and whenever things change.
- *Risks from novel aspects of the domain, problem, or solution*. Another useful heuristic, from the experience of our study participants, is to focus on novelty in your project. What is unfamiliar? What problems have you not solved before? Which technology is unproven? The answers to these questions highlight risks and the participants in our study used them to direct their effort to the most important risks to address.
- *Risks in the organisational environment*. Each organisation is different and there are nearly always risks specific to an environment such as the internal political situation, what is possible in the organisational culture, and the maturity of the organisation with respect to architecture, change and risk. Different organisations have different cultures and capabilities for change, which can create risks. The speed which different sorts of risk can be addressed can also be affected by organisational factors and so may cause you to change where you focus attention. Participants in our study noted the importance of "situational awareness" [14] to allow risks specific to the organisational environment to be identified and addressed.
- *Risks from the external environment*. Nearly all organisations exist in a complex ecosystem of interacting partners, customers, regulators, competitors and other actors and they can be a source of risk for many systems. So can general trends and changes in the industry that the organisation exists within (such as a changing regulatory environment, or industry wide pressures such as reducing margins on products or services).
- *Risks related to cost and time*. Most architects will report that they are often expected to achieve challenging goals in unrealistic timescales or with unrealistic cost estimates. Many of our study participants reported that they needed to focus significant attention on risks resulting from cost and time.
- *Identify the high impact decisions*. Prioritise architecture work that will help to mitigate risks where many people would be affected by a problem (e.g. problems with the development environment or problems that will prevent effective operation) or where the risk could endanger the programme (e.g. missing regulatory constraints).

2.3 Delegate as Much as Possible

Delegation was an unexpected theme that emerged from our study. The architects who mentioned this theme viewed themselves as a potential bottleneck in a project and focusing attention on the delegation and empowerment of others was a way to minimize this. Delegation was also seen as a way of freeing the architect to focus on the most important aspects of the project.

The general message of this theme is to delegate as much architecture work as possible to the person or group best suited to perform it. This prevents individuals becoming project bottlenecks, allowed architects to spend more time on risk identifica-

tion and mitigation, and spreads architectural knowledge through the organisation. The heuristics that were identified to help achieve this are:

- *Empower the development teams*. To allow delegation and work sharing, architects need to empower (and trust) the teams that they work with. This allows governance to become a shared responsibility and architecture to be viewed as an activity rather than something that is only performed by one person or a small group. This causes architectural knowledge, effort, and accountability to be spread across the organisation, creates shared ownership, reduces the load on any one individual and prevents a single individual from delaying progress.
- *Create groups to take architectural responsibilities*. A related heuristic is to formalise delegation and create groups of people to be accountable for specific aspects of architectural work. For example, in a large development programme, an architecture review board can be created to review and approve significant architectural decisions. Such a group can involve a wide range of expertise from across the programme and beyond, so freeing a lead architect from much of the effort involved in gathering and understanding the details of key decisions, while maintaining effective oversight to allow risks to be controlled and technical coherence maintained. Similarly, a specific group of individuals could be responsible for resilience and disaster recovery for a large programme, allowing them to specialise and focus on this complex area, and allowing a lead architect to confidently delegate to them, knowing that they will have the focus and expertise to address this aspect of the architecture.

2.4 Team Effectiveness

A theme that emerged when we validated our initial model with a wider group was the need to spend time making sure that the development team was as effective as possible. The participants who highlighted this factor were concerned with developing the individuals in the team and ensuring that the team was as diverse as possible, to provide it with a range of skills and perspectives.

Other aspects of this theme were the importance of architecture work being used to quickly unblock the team when it hit difficulties and the importance of technical leaders, like the architect, to step in when needed to make sure that the team was functioning well and to address any dysfunctional behaviour observed.

The heuristics identified as being important for achieving team effectiveness were:

- *Develop the team through mentoring*. Every team should be on a collective journey towards improvement and hopefully every individual in a team is on a similar personal journey to be the best that they can be. People doing architecture work tend to be experienced, so a valuable area to focus attention is developing the individuals and the team as a whole, through thoughtful, intentional mentoring.
- *Achieve team diversity for better innovation and problem solving*. To innovate and identify good solutions to problems, it is valuable to have a range of experience,

perspectives and skills in the team. Our study participants indicated that a valuable use of time is building diverse teams that can achieve this.

- *Remove blockers preventing team progress*. Development and support teams often end up blocked by technical or organisational factors, so spending time resolving these problems is a valuable focus for many architects.
- *Address dysfunction in teams*. Sometimes teams don't work well, and it requires someone who is close to the team, and respected by them, but outside the team structure, to identify the problem and suggest solutions. People doing architecture work are often close to the teams but outside their structure, and have the respect, soft-skills and experience to resolve team problems. This use of architectural time can have huge benefits when dysfunctional behaviour is observed in teams.

2.5 Summary

This model provides a simple guide to focusing architectural attention during a project. It is comprised of 4 aspects: Stakeholder Needs and Priorities, Prioritise Time According to Risks, Delegate as Much as Possible and Team Effectiveness. We believe that it can be an effective guide or reminder on the best ways to focus architecture work during a project.

In common with any set of heuristics, the model is only a starting point and must be considered, interpreted, and applied in a context specific way by the architects and teams who use it. However, as we explain later in the paper, it has validated well against a reasonably broad survey of experienced, practicing architects and so we believe that it is a useful guide upon which to build a personal approach for prioritisation.

3 Related Work

When we started investigating this topic, we were primarily interested in how practitioners really worked. However, we also performed a literature search to find related work from the research community.

We did not find any studies investigating the specific topic we are interested in, but an architectural method which helps architects to focus their attention is Risk and Cost Driven Architecture (RCDA) [11]. This method transforms the architect's approach from defining architectural structures early in a project, to providing a stream of decisions throughout it, prioritising their work using the risk and cost of open decisions. This guides the architect to focus on the important architectural decisions to work on at any point in time but does not provide any guidance beyond that aspect of their work. So, while valuable, it is quite narrow in this regard. Also, while a recognised approach, it isn't very widely used in the industry, so we were interested how the practitioners who don't use RCDA prioritise their attention.

We also found some very specific advice from a very experienced architect and researcher [9] that architects should spend 50% of their time on architecting, 25% on

inbound communication and 25% on outbound communication. However, this is anecdotal advice based on personal experience, so we don't know how many (if any) practitioners follow this advice.

In the research domain we found a research on the prioritisation of requirements [4, 6] and a literature review of this area as of 2014 [1]. Prioritising requirements is related to focusing architectural attention, but it is only one factor from a large possible set, so this research was not very relevant to our investigation.

Finally, there is a large amount of mainstream business literature on time management (such as [1, 8]) however we were interested in providing more specific advice for software architects rather than this sort of more general advice.

4 Research Method

When planning this research, we selected a qualitative research approach because we wanted to explore the "lived-experiences" of expert practitioners by asking them questions to encourage reflection and insight [12] rather than assessing performance or alignment with specific practices via quantitative means.

The process was organised into four distinct stages.

- Stage 1: gathering primary data using semi-structured interviews with practitioners.
- Stage 2: analysis of the primary data and creation of a preliminary model.
- Stage 3: validation of the preliminary model via a structured online questionnaire, completed by practitioners in relevant architecture roles (primarily software, solution, and enterprise architects).
- Stage 4: analysis of the validation data and refinement of the preliminary model into a final, validated model.

The first two stages were reported in [15] but we will briefly explain the whole process here for the sake of clarity.

We chose to gather our primary data using semi-structured interviews, providing interviewees with a written introduction to the question we wanted to answer and some questions to start them thinking. The content of the interviews was analysed through iterative thematic coding and, as suggested in [12], the process of collection and analysis was iterative and exploratory rather than a rigid linear one.

This exercise produced a set of heuristics that the architects use with themes to classify them. A heuristic had to be mentioned by at least three of the participants (a third of them) for us to consider it significant enough to be included in the model. We combined the themes and heuristics to form a simple model (the "preliminary model") of how experienced architects go about prioritizing their effort.

Once we had the preliminary model we published it at a research conference [15] and via a LinkedIn post[1] and created an online questionnaire to allow architecture practitioners to evaluate and comment on the usefulness of the model. We publicised the survey via LinkedIn, Twitter and email to our network of architects.

[1] https://www.linkedin.com/pulse/focusing-software-architects-attention-eoin-woods

We received 84 responses to the survey that answered our closed-ended questions, of which 50 contained answers to the open-ended questions. We used the closed-ended questions to evaluate the usefulness of the model and analysed the open-ended responses to identify themes missing from the model.

The model was validated strongly across respondents from different locations, with varying amounts of experience, and from different architectural specialisations. A small number of suggestions for improvement emerged from the answers to the open-ended questions. These suggestions were used to revise and extend the model, creating an improved final version, that reflected the input from the respondents.

A description of the four stages of the research method is presented in the following sections of the paper.

5 Stages 1 and 2: The Initial Study

Our primary data gathering was performed using a semi-structured, face-to-face survey of 8 experienced software architecture practitioners in 4 countries. As Stages 1 and 2 were previously reported in [15] we just summarise the work here.

We found the participants by approaching suitable individuals from our professional networks. We were looking for practitioners who had a minimum of 10 years' professional experience and who worked as architects in the information systems domain (rather than architects from – for example – embedded systems).

We focused on the information systems domain because we know from experience that working practices differ between professional domains like information systems and embedded systems. Hence, we thought it more likely that we could create a useful model if we limited ourselves to one broad domain, at least initially.

Our preliminary model for focusing architectural attention is shown in Fig. 2.

.

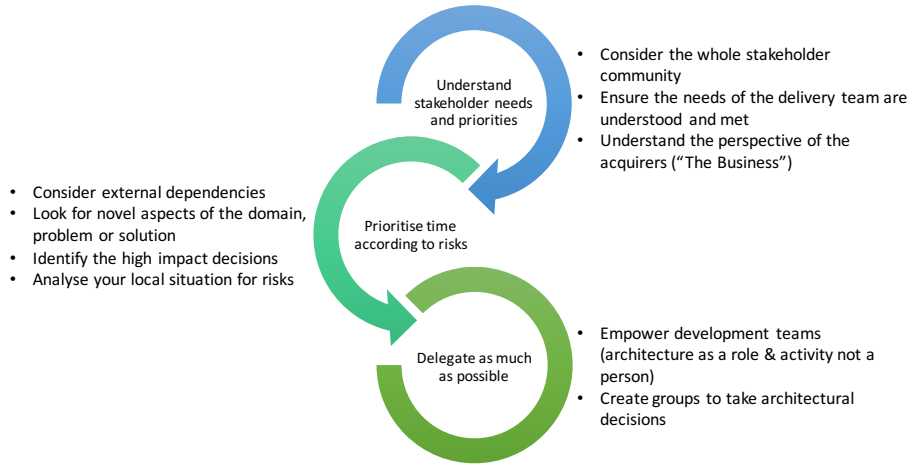

Fig. 2. Preliminary Model for Focusing Architectural Attention

The three categories of heuristic that the study revealed were: the need to focus on stakeholder needs, the importance of considering risks when deciding on where to focus attention, and finally the importance of spending time to achieve effective delegation of responsibilities. These categories form the structure of our model and remind the architect of the general ways in which they should focus their attention. The categories and heuristics are explained in [15] and Section 2 of this paper.

6 Stage 3: Validating the Preliminary Model

6.1 The Questionnaire

Once we had a preliminary model, we wanted to validate its usefulness with a much larger group of experienced practitioners using a structured online questionnaire.

The questionnaire asked the respondents to read the model and then comment on its credibility and usefulness. We asked both closed questions, that asked respondents to rate the model on 5-point scales, and open-ended questions that allowed the respondents to consider whether there were aspects of focusing attention that we had missed and to collect general comments on the model. Finally, we asked some closed classification questions to allow us to understand who had completed the survey, while preserving their anonymity if desired.

We asked three closed-ended questions to find out whether the respondent thought that the model was credible and useful. These questions and possible responses were:

- Q1. "Is this model similar to how you focus architectural attention in your work already?" Not at all similar / Not Very Similar / Somewhat Similar / Quite Similar / Very Similar
- Q2. "Would you find this model helpful in guiding architectural attention for maximum benefit?" Definitely Not / Probably Not / Possibly / Probably Yes / Definitely Yes
- Q3. "Are the areas of risk mentioned in the "Prioritise time according to risks" activity valuable?" Definitely Not / Probably Not / Somewhat / Probably Yes / Definitely Yes

The open-ended questions that we asked were:

- Q4. "Are there other general areas of risk that should be added to "Prioritise time according to risks" that would be applicable to most (information) systems and environments? If so please list and briefly explain them."
- Q5. "Are there any significant factors missing from the model which you use to focus your architectural work?"
- Q6. "Do you have any other comments on the model or the survey"

The closed-ended questions we asked to allow us to classify the respondents and their possible answers were:

- Q7. "What environment do you work in?" Industry / Industrial Research / Academic / Other (please specify)
- Q8. How many years of post-graduation experience do you have? 1-5 years / 5-10 years / 10-15 years / 15-20 years / More than 20 years
- Q9. What is your main job role? Software Architect / Enterprise Architect / Software Designer / Researcher / Other (please specify)
- Q10. Where in the world are you based? North America / South America / Europe (inc. UK) / Middle-East and Africa / Asia-Pacific / Other (please specify)

Having trialled the questionnaire ourselves, and with two other individuals, we expected most respondents to take 10 – 15 minutes to complete it.

6.2 The Respondents

To use the questionnaire to validate the model, we needed to find a suitable set of architects who could read it and complete the survey for us. We found our initial respondents via a LinkedIn post (https://www.linkedin.com/pulse/focusing-software-architects-attention-eoin-woods/) that appeared in the LinkedIn news feed of practitioners, which resulted in 23 people completing the survey successfully.

To gain more responses to the survey, we sent a targeted email to practicing software, solution and enterprise architects in our professional network, which resulted in 61 more responses to the survey, making a total of 84 completed surveys.

About a third of the respondents identified themselves as software architects, about a quarter as enterprise architects, about 12% as software designers, 10% as solution architects, and 5% as technical architects. Four respondents didn't complete this answer and four had other job titles (a risk assessor, a technical manager and systems engineer, a project manager and a strategy consultant).

We asked the respondents to classify their work environment as Industry, Industrial Research or Academic and a few respondents self-identified as working in the public sector. 78 respondents (~90%) were from industry or building systems in the public sector (several of whom identified as both "Industry" and "Industry Research"), one was from an academic work environment and 5 (~6%) did not answer this question.

We then asked respondents where they worked geographically, and 55% of respondents identified themselves as from Europe, 30% from the Americas and only 7% from Asia-Pacific and a single correspondent from the Middle East and Africa.

We discuss the possible impact of geographical location when we consider threats to validity, but we think that we achieved good cross-geographic participation, but still ended up with a strong bias to Western Europe and North America.

The final classification we asked our respondents for was the number of years of experience that they had. Over half of them (55%) had at least 20 years of post-graduation experience, 17% had 15-20 years of experience, 15% had 10-15 years of experience 7% had 5-10 years if experience and only one respondent had less than 5 years of experience. Four of our correspondents did not answer this question.

6.3 The Closed-Ended Responses

As mentioned earlier, we structured the questionnaire into two parts, the closed-ended questions that asked people to rate the usefulness of the model and the open-ended questions that asked whether we had missed anything important from it. In this section, we review and analyse the responses for the closed-ended questions.

The first question we asked was to find out if the model was similar to how experienced architects already focused their attention, to assess the basic credibility of the model for experienced architects. 75% of respondents indicated that it was "very similar" or "quite similar" to their existing approach for focusing their attention, 20% said it was "somewhat similar", 5% said it was "not very similar" to how they worked, and no respondents replied that it was "not at all similar". These responses suggest that the model validates strongly against the participants' existing practice.

The second question attempted to establish, whether the respondents thought that model would be useful in practice. 27% responded that it was "definitely useful", 43% that it was "probably useful", 26% said "possibly useful" and 3 respondents (4%) said "probably not". These responses suggest that most of the participants see probable value in the model (i.e., 70% see it as definitely or probably useful).

Finally, we wanted to check that the areas of risk we had identified as important within the "prioritise time according to risks" heuristic were valuable to a practicing architect. 43% of respondents indicated "definitely yes", 37% indicated "probably yes", 15% responded "somewhat", 4% as "probably not" and a single respondent indicated "definitely not". The single individual who indicated "definitely not" was an enterprise architect in the 10 – 15 years of experience group, who commented in the open-ended questions that he did not believe that it was possible to define general software development risks in a useful way.

From this response, 80% of respondents believe that the areas of risk were "definitely" or "probably" valuable, suggesting that this aspect of the model should be of value to many practitioners.

In summary, having analysed the answers to the closed-ended answers in our survey, we conclude that our model is likely to be credible and useful for the architects who responded to our survey and broadly aligns with the prioritization approach used by many experienced architects.

We interpret these results as a successful validation of the model, but we were also interested in how the model could be improved and so we used the responses to the open-ended questions in the survey to find themes that we might have missed.

6.4 The Open-Ended Responses

As explained earlier, we asked two open-ended, questions, Q4, to identify missing risk factors from the "prioritise time according to risks" heuristic ("*are there other general areas of risk that should be added to "prioritise time according to risks" that would be applicable to most (information) systems and environments?*") and Q5, to ask whether we had missed any aspects of the model ("*are there any significant factors missing from the model which you use to focus your architectural work?*"). We

had 44 responses to Q4, about missing risk factors, and 51 responses to Q5, about missing areas of risk.

Given the nature of these responses, we again used a simple thematic coding analysis to analyse them, coding each one initially using straightforward, descriptive labels, reflecting the language in the response, then refining this with further coding steps, to identify higher-level categories to group the responses into.

For the first question, Q4, we initially coded the responses to 37 distinct categories, plus two null categories for the initial coding of "None" and "General Comment". The responses suggested a diverse range of possible risk areas, and when we refined the coding to find common concepts, this resulted in 24 higher level categories.

We attempted to refine this further but did not find further meaningful refinements as we tried further rounds of coding and ended up with a very long "tail" of risk areas with only a single mention in the responses. We ended up with 5 categories that had 4 responses or more: Organisational Environment (11 occurrences), Stakeholders (6 occurrences), Cost (6 occurrences), Time (4 occurrences) and External Environment (4 occurrences). We chose to focus on categories with at least 4 occurrences as this represents approximately 5% of the total respondents to the survey and we judged this to be high enough to include as risk areas for the "prioritising time according to risks" element, in the refined version of the model, presented in Section 2.

For the second open-ended question, Q5, on missing aspects of the model, we initially coded the responses into 43 distinct categories and continued with the process of refining the coding further, ending up with 26 higher level categories. As with the responses to Q4, many of the categories were only mentioned once and only four were mentioned 4 times or more: Team Effectiveness (10), Benefits (7), Stakeholders (6) and Time (5). Of these factors, "Stakeholders" are already a significant factor in the model and the comments provided in these cases were suggesting a particular emphasis on certain stakeholders or method of dealing with stakeholders, suggesting that a new element was not needed in the model.

Adding a completely new aspect to the model is a significant step and so we only wanted to consider this for aspects which had been identified as important by a significant number of respondents to the survey. Hence we decided to add a new element to the model to reflect the "Team Effectiveness" theme as it was the only additional aspect that at least 10% of the respondents had identified as important.

Finally, we also received 51 general comments in the open-ended questions which we thematically coded into 23 groups, most of which had one or two comments in them. However, there were 14 "Positive Comments", 6 about "How the Architect Should Work" and 5 on the "Presentation of the Model".

These comments were interesting but only the five comments on the presentation of the model suggested the need to change the model. These comments consistently suggested that our graphical presentation indicated a linear process, whereas we actually meant to communicate a continuous process throughout the project lifecycle, so this was an indicator that we needed a better graphical representation for the mode.

7 Stage 4: The Refined Model

We took the results of the open-ended question analysis described in Section 6.4 and used them to add missing features to the model, improve the list of risks to suggest for time prioritisation and improve the model using the advice provided in the general comment responses to the survey. The result of this work is the model that was presented in Section 2.

As can be seen, if you compare the final model in Section 2 with the preliminary model in Section 5, three significant changes have been made:

1. An additional feature, "Team Effectiveness", has been added to the final model, because this theme was noted as an important missing feature of the model by more than 10% of the respondents to the survey.
2. The list of risks to use to guide time prioritisation has been extended and refined based on common suggestions made from question 4 in the survey.
3. In response to comments in the survey' open-ended questions, the graphical presentation of the model has been altered to try to emphasise that it is not a linear "process" but a set of activities to be performed throughout the project lifecycle.

The result is a model that guides an architect where to focus their attention during a project, focusing on four themes: Stakeholder Needs and Priorities, Prioritise Time According to Risks, Delegate as Much as Possible and Team Effectiveness.

In common with any set of heuristics, the model must be considered, interpreted, and applied in a context specific way, but it validated well against a reasonably broad survey of experienced, practicing architects and so we believe that it should be a useful guide for many practitioners.

We did not ask participants in the study whether they had a particular architectural domain specialisation (such as web-based systems, data analytics systems, embedded systems and so on) however we didn't target any specific group as we did in the preliminary study (where we limited ourselves to information systems architects) and we didn't get any comments about applicability to specific domains in the open-ended answers to the survey. This leads us to have reasonable confidence that the lessons captured in the model are quite widely applicable.

8 Threats to Validity

Specific steps we took to ensure the integrity of this work included focusing on the practitioner community (the intended users of the model), focusing on experienced respondents who have the experience to evaluate the model, finding a reasonably large, geographically distributed group to validate it for us, structuring the questionnaire to allow disagreement as well as confirmation, and analysing the results in a careful, structured manner to allow the data to lead us to the conclusions, to avoid unconscious bias. However, we acknowledge that there are potential limitations to any qualitative study, which could threaten our study's validity.

There are four main types of threat to the validity of a study like this, namely construct, internal, external and conclusion validity as defined in [10].

Construct validity is concerned with the relationship between theory and observation. Common threats when using questionnaires are the phrasing of the questions and using too many closed-ended questions. We kept the questions brief and refined the questionnaire wording after testing it. We provided open-ended questions for the participants to explain, expand or clarify their answers.

Internal validity is concerned with the validity of the causality relationship between the observations and the outcomes of the study. We addressed this by using very straightforward analysis so the threats to the correctness of the analysis we performed are minor. We also reviewed each respondent's responses for coherence.

External validity is concerned with the generalisability of the results of the study. In our case the key risk is an unrepresentative respondent population. We mitigated this risk through a geographically distributed, relatively large respondent population. However, a residual risk is the lack of representation from Asia. We mitigated concerns about experience and competence by targeting experienced architects. We know a significant percentage of the respondents at least slightly and have confidence in their ability to validate the model. This leaves us with a residual risk that our participants may share more common opinions than a random sample, but anecdotally we believe that they are similar to most practitioners we have met over the years.

Conclusion validity is concerned with the validity of the relationship between the data obtained in the study and the conclusions that have been drawn from it. We mitigated the possibility of asking the wrong questions by using a semi-structured interview in the first stage and providing extensive opportunity for open-ended responses in the third stage. We mitigated risks of analytical mistakes by reviewing and cross checking our work and using a simple, repeatable process. We avoided unconscious bias by using a structured coding process for open-ended question analysis, to allow us to be led by the data.

In summary, we designed and executed the study carefully but acknowledge that there are some threats to its validity which could threaten the generalizability of our results. The most significant is the lack of Asian participation, however a model useful in Europe and America would still be a valuable outcome.

9 Future Work

The refined model is now ready for dissemination to the practitioner community to see if it proves as useful in practice as our survey of the preliminary model suggests. We have already published the refined model in a less formal style via LinkedIn[2], resulting in a number of positive comments. We will also try to publish a summary of it in practitioner-oriented publications and publicise it through practitioner conferences, if it proves to be of interest to programme selection committees.

[2] https://www.linkedin.com/pulse/revisiting-how-people-prioritise-software-work-eoin-woods

We could also run a further study to establish if there are useful elements missing from the model, such as those related to design, implementation, modelling and other more general and technical aspects of architecture work. These may not have been mentioned by the experienced practitioners because they have largely been internalised due to their level of expertise. Or it is possible that they are not all that useful for prioritisation of architectural attention.

Finally, another possible area of study is the validation of Philippe Kruchten's insight [9] that they should spend 50% of their time on architecting, 25% on inbound communication and 25% on outbound communication.

10 Conclusion

Our experience and informal discussion with architects over many years suggested that they find it difficult to decide how to focus their attention to maximise their effectiveness. We were interested in how experienced practitioners solved this problem and whether there were commonly used heuristics. To investigate this, we used a four-step process of investigation.

We started with a semi-structured interview process with eight experienced practitioners and concluded that there are some shared heuristics which practitioners use, but that practicing architects are not aware that the heuristics are common and shared. We found that the heuristics clustered into three groups: focus the architects attention on stakeholders, use their time to address specific risks and delegate as much as possible, in order to give them as much time for architecture work as possible.

We then created a simple structured model to capture and explain the heuristics that emerged from the initial study and we published this via social media channels. In the next step, we asked practitioners to complete a survey to comment on the usefulness of the model and whether anything had been missed. 84 responses were received to the survey, mainly from European and North American software, solution and enterprise architects with over 10 years of professional experience.

When we analysed the survey responses we found that the model validated well, as 70% of the respondents think it would probably or definitely be useful, but we found that we had missed several important risk factors which are commonly used for prioritisation and we had missed an element of focusing attention, which is the need to ensure overall team effectiveness. We added these missing elements to the model.

These findings are not completely unexpected and many of the heuristics in the model are familiar. However, neither the participants or ourselves knew that these were the important, shared heuristics before we undertook the study, so we believe that the model that we have created will have value as a teaching aid and as an aide memoir for experienced practitioners.

We have started to publicise the model via social media and plan to continue this by incorporating it into practitioner-oriented articles and conference talks. If the model gains some acceptance over a period of time, there would be value in a future survey of its users to review the model's usefulness after experience of using it.

11 References

1. Achimugu, P., Selamat, A., Ibrahim, R., Mahrin, M.: A systematic literature review of software requirements prioritization research, Information and Software Technology, Volume 56, Issue 6, June 2014, Pages 568-585.
2. Allen, D.: Getting Things Done: The Art of Stress-free Productivity. 2nd edn. Piatkus (2015).
3. Bass L., Clements, P., Kazman, R.: Software architecture in practice. 3rd edn. Addison Wesley, Upper Saddle River, NJ (2012).
4. Berander, P., Andrews, A.: Requirements prioritization. In: Aurum, A., Wohlin, C. (eds) Engineering and managing software requirements., pp. 69-94. Springer, Heidelberg (2005).
5. Harindran A., Chandra V.: Research Methodology, Pearson, (2017).
6. Herrmann, A., Daneva, M.: Requirements prioritization based on benefit and cost prediction: an agenda for future research. In: Tetsuo, T. (ed) International Requirements Engineering, 2008. RE'08. 16th IEEE. IEEE, 2008.
7. Karlsson, J., Ryan K.: A cost-value approach for prioritizing requirements. IEEE Software 14(5) 67-74 (1997).
8. Koch, K.: The 80/20 Principle: The Secret of Achieving More with Less. Nicholas Brearley Publishing (2007).
9. Kruchten, P.: What do software architects really do?, The Journal of Systems and Software, 81(12), 2413–2416 (2008).
10. Matt, G.E., Cook, T.D.: Threats to the validity of research synthesis. Cooper, H., Hedges, L.V. (eds) The Handbook of Research Synthesis, pp. 503-520. Russell Sage Foundation, New York. 503-520 (1994)
11. Poort, E.R., van Vliet H.: RCDA: Architecting as a risk-and cost management discipline. Journal of Systems and Software 85(9) 1995-2013 (2012).
12. Reimer, F.J., Quartaroli, M.T., Lapan, S.D.: Qualitative Research: An Introduction to Methods and Designs. Wiley, London (2012).
13. Rozanski, N., Woods, E.: Software systems architecture, working with stakeholders using viewpoints and perspectives, 2nd edn. Addison Wesley, Upper Saddle River, NJ (2011).
14. Wikipedia, Situational Awareness, https://en.wikipedia.org/wiki/Situation_awareness, last accessed 2017/04/10.
15. Woods, E., Bashroush, R.: A Model for Prioritization of Software Architecture Effort. In: Lopes A., de Lemos R. (eds) Software Architecture. ECSA 2017. Lecture Notes in Computer Science, vol 10475, August 2017. Springer.