ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOCPR01

GRADUATE SOFTWARE ENGINEER DEVELOPMENT PROGRAM AT DIAMOND LIGHT SOURCE

A. A. Wilson*, T. M. Cobb, U. K. Pedersen, Diamond Light Source, Oxfordshire, United Kingdom

Abstract

Diamond Light Source is the UK's synchrotron facility. The support and development of the beamlines and accelerators at Diamond requires a significant quantity of specific knowledge and skills; the opportunity to acquire these beforehand is not available to many early in their career. This limits the field of candidates who can begin working independently at the level of software systems engineer. The graduate software engineer development program was started in 2015 to provide a route for engineers who are recent graduates or new to the field to develop the required skills and experience. Over the course of two years it comprises a series of projects in different groups, mentored on-the-job training and organized training courses. The program has recently been expanded to cover all groups in the Scientific Software, Controls and Computation department at Diamond, with an intake of four new engineers per year. This paper presents the structure and development of the program and invites discussion with other organizations to share knowledge and experience.

INTRODUCTION

The field of control systems for particle accelerators is highly specialised. To work independently in the field, a significant amount of specific knowledge and skills are required. It can therefore be difficult for organizations to find and recruit people with the requisite experience to start working independently in this function. Meanwhile there are few opportunities for early career professionals to gain the necessary experience in advance.

In 2015, the Controls Group at Diamond implemented a Graduate Training Program in order to create a route to hire talented but inexperienced people who are either recent university graduates or new to the field.

This program was later expanded to include the adjacent groups in the Scientific Software, Controls and Computation Department (SSCC). The SSCC department covers the work on the "full stack" of software and computing at Diamond, encompassing computing, systems administration, electronics, controls, data acquisition and analysis, and business applications.

This paper outlines the structure of the program, its current status and our plans for future developments. The purpose of the paper is to invite discussion: to share experience with the community and to gain some ideas to improve the program further.

* andrew.wilson@diamond.ac.uk

PROGRAM STRUCTURE

The program consists of two years of training. Participants begin in a graduate role, and move to the experienced engineer role following completion of the program.

Recruitment

The recruitment process takes place in November and December for a start date in the Autumn of the following year. This follows the pattern of many graduate training programs in industry, targeting students in their final year of university; in this way we aim to be competitive with such programs.

We aim to attract university leavers with a background in a STEM subject area (Science, Technology, Engineering and Mathematics), not only those with a background in Computer Science. We look particularly for an interest and some experience in software or computing, for example, a personal interest or an undergraduate project.

First Year

In the first year, participants undertake four projects, each lasting three months, in different groups within SSCC. The projects are generally stand-alone, although sometimes consecutive projects will build on each other, approaching different aspects of the same topic. As far as possible they are aligned to the interests of the individual and the needs of the business. Figure 1 shows the number of projects undertaken to date per group.

The purpose of these projects is for the participant to gain experience of working in different areas of the business; to sample the different technologies and working environments; and to build a network of connections which will be useful in their later role. The host groups benefit from the extra effort to complete a necessary piece of work, and the new ideas and different approaches that the participants bring. The projects are often interdisciplinary and therefore promote collaboration between the different groups.

In addition to the participant, the projects involve several people whose roles are as follows:

Supervisor Defines the scope and goals of the project; provides guidance and assistance to the participant throughout the project; involves the participant in some day-to-day work of the host team; arranges project meetings.

Customer As the target customer, defines the requirements for the project and reviews the results. This role is often performed by the Supervisor.

Technical support Additional people who may be involved if the supervisor does not have all the technical information needed to support the project. For example, the technical support helps plan the architecture and review the results.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOCPR01

We emphasize the importance of documenting the work produced. Each project has a wiki page which is used throughout the project cycle: by the supervisor and others to record the background, requirements, milestones and rough timeline; and by the participant to record the progress and outcomes of the project.

Second Year

Participants spend the second year in one team, where they are assigned a supervisor. The purpose is to learn "on the job" under their supervision and guidance.

In the course of the year, the participant picks up more responsibility. For example, in Beamline Controls, they take on support for one or more of Diamond's 35 photon beamlines throughout the year. The purpose is to progress towards independent, unsupervised work in the team, at the level of an experienced Engineer.

After the end of the second year, having achieved independent working and responsibility, the goal is for the participants to successfully complete the company's "grade review" process, thereby achieving a promotion to the role of experienced Engineer. They normally continue to work in the same team where they spent the second year, but may transition to other areas of SSCC depending on individual career aspirations and the needs of the organization.

TRAINING

In addition to the on-the-job training, participants attend formal courses on a range of technical topics, depending on the role:

Programming skills C/C++, Java, Python
Tools & frameworks EPICS control system framework [1]
Hardware e.g. Delta Tau motion controller configuration

Each participant is encouraged to attend an international collaboration meeting or conference to learn about collaborations and developments in the wider field.

There is a regular program of informal meetings, alternating talks by participants about their projects and informal question-and-answer sessions with invited speakers from elsewhere in the organization.

CARE AND WELLBEING

In addition to the general management of the scheme, there are two part-time roles relating to keeping things running smoothly:

- A graduate program coordinator to manage the allocation of projects, supervisors and seating; the coordinator attends the kick-off and wrap-up meetings for each project.
- A second person in charge of "pastoral support" to talk to people; collect feedback on the projects; identify any potential problems; organize the "extra-curricular" activities and make sure that people are generally happy.

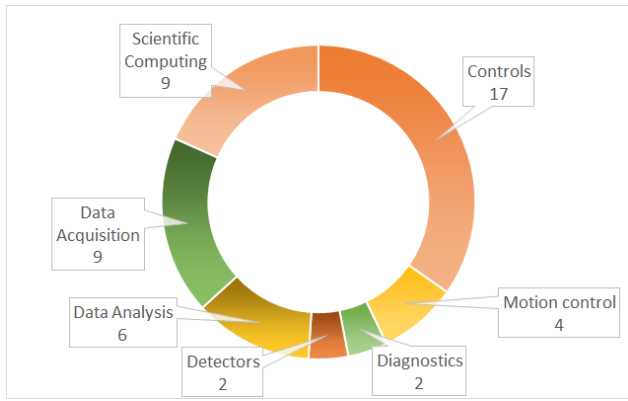

Figure 1: Number of projects completed per group in SSCC.

Due to projects being distributed across different groups, the graduates generally will not work directly with one another during the first year, and may not have the opportunity to spend time getting to know each other. We recognize that a strong graduate community is beneficial, as more established SSCC graduates can provide both technical and non-technical support from the perspective of a peer. They can help new graduates to navigate the complex infrastructure of the organization, find the right people to ask for help and work more effectively on their projects. This in turn alleviates the load on supervisors who are often more senior employees. To further this aim, the person in charge of pastoral support organizes various activities, such as coffee meetings and exchanges of tours with graduates from nearby facilities. As a new initiative in 2019, we have held an SSCC Graduate Away Day, shortly after the arrival of several new starters, with team-building and activities to create opportunities for the graduates to build these connections in a relaxed environment.

EXAMPLE PROJECTS

We describe here two example projects to give an impression of the scope and structure.

Dashboards for Detector Status

Task Develop a web-based dashboard system to monitor the status of detector systems at photon beamlines

Combines data from multiple sources in order to monitor the health of the system and understand the root causes of faults. Data sources include:

- Values over time of EPICS Process Variables (PVs) from the Archiver Appliance;
- Server metrics collected in a Graphite database;
- Counts of application log messages of different categories, collected in a Graylog database.

A Grafana [2] dashboard was used for the web interface. The system allows concurrent display of important parameters such as software state & errors, hardware state, server load and network load.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOCPR01

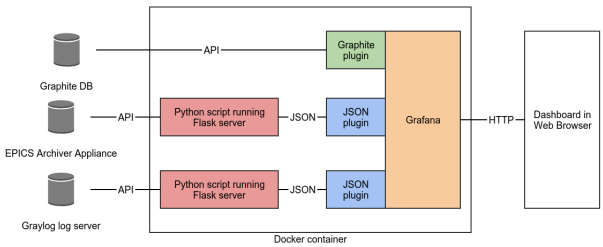

Figure 2: Simplified block diagram showing data sources for Detector Status dashboard.

The project consisted of liaising with target customers to identify the key information to display; understanding how and from where to gather this information; writing scripts to gather and format the information; packaging the scripts to run as a maintainable service in a container; and designing an intuitive layout for effective display of the information on the dashboard. Figure 2 shows an overview of the architecture.

The product has now been rolled out for Eiger [3] and Excalibur [4] detectors at six Diamond beamlines.

Groups Project supervised in Scientific Computing, with customers in Controls.

Training Rigs for Hardware Triggered Scanning

Task Diamond has developed a standard hardware and software stack for hardware-triggered continuous scanning at photon beamlines [5]. In order to support and roll out this stack more widely, there was a need for a set of demonstration rigs to be used in the training of Controls and Data Acquisition staff. The purpose of this project was to develop a training rig, consisting of several positioning stages with a motion controller, a PandA [6] for triggering and position capture, and a PC to run the control and acquisition software. The training rig can collect data using a visible light source and a Gigabit Ethernet camera. A simplified block diagram of a rig is shown in Fig. 3. The project involved:

- Configuring and commissioning the motion system;
- Configuring the PC;
- Deploying EPICS control processes, Malcolm [7] for scanning control, GDA [8] for experiment control and DAWN [9] for data analysis & visualization;
- Commissioning the stack to acquire and analyze data.

Four rigs have been produced and can demonstrate the acquisition and analysis of Tomography data. They have been used in a pilot training course held in Summer 2019.

Groups Project supervised in Motion Control, involving collaboration with Scientific Computing and a customer in Controls.

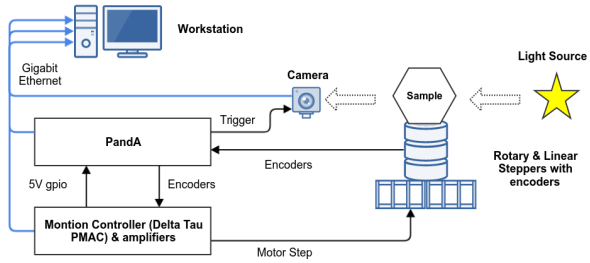

Figure 3: Simplified block diagram for Hardware Triggered Scanning training rig.

CHALLENGES & FUTURE DEVELOPMENTS

The relatively small intake can make some organizational aspects inefficient. For example, critical numbers of attendees are required before it is feasible to hold each training course. A new starter may therefore have to wait for some time before receiving some of the training, which could mean they do not receive it at the optimum time in their development. We are piloting online learning to supplement the formal courses, so that if some specific knowledge is required for a project, it can be quickly gained. On the other hand, it is worth noting that the small intake allows a flexible and personal approach to arranging the placements.

Historically, start dates for new graduates have been distributed across the year; the reason for moving towards aligning start dates each year in the Autumn is that aspects of administration can be made more efficient and the reach of recruitment can be increased.

Having a high ratio of staff in training means there is a limited supply of mentors to spend the required time in the second year. A certain amount of this is inevitable, but some ideas to alleviate it are:

- Some of the more routine aspects of the training provided by the second-year supervisor, for example on the tools, processes and development environment in Controls, could be provided centrally. This would reduce the burden on the supervisor.
- In order to attract more volunteers to be supervisors, we could increase the internal promotion of the scheme.

The possibility of collaboration with other institutions could provide some interesting opportunities and we would like to explore this further.

CONCLUSION

The program is established and working well; engineers who complete it have a strong foundation of experience and move easily into working independently on beamlines, supporting the scientists to operate the instruments. We continue to see the benefits of the program spread across the groups in SSCC. The statistics in Table 1 illustrate the current scale and status.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOCPR01

Table 1: Current Statistics

New starters per year	4
Total number of participants	15
Total completing the program and grade review	4
Projects completed	49
Talks on projects and other topics	15

ACKNOWLEDGEMENTS

For additional contributions to this paper, the authors wish to thank Austen Rose, Mark Heron, Emma Dixon, Bryan Tester, Niamh Dougan and Will Rogers.

REFERENCES

[1] EPICS: Experimental Physics and Industrial Control System, https://epics-controls.org/

[2] Grafana, https://grafana.com/

[3] G. D. Yendell *et al.*, "Implementing Odin as a Control and Data Acquisition Framework for Eiger Detectors", presented at the 17th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'19), New York, NY, USA, Oct. 2019, paper THBPP05, this conference.

[4] G. D. Yendell *et al.*, "Odin - a Control and Data Acquisition Framework for Excalibur 1M and 3M Detectors", in *Proc. 16th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 966–969. doi:10.18429/JACoW-ICALEPCS2017-TUPHA212

[5] R. D. Walton *et al.*, "Mapping Developments at Diamond", in *Proc. 15th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015, pp. 1111–1114. doi:10.18429/JACoW-ICALEPCS2015-THHB3001

[6] S. Zhang *et al.*, "PandABox: A Multipurpose Platform for Multi-technique Scanning and Feedback Applications", in *Proc. 16th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 143–150. doi:10.18429/JACoW-ICALEPCS2017-TUAPL05

[7] T. M. Cobb *et al.*, "Malcolm: A Middlelayer Framework for Generic Continuous Scanning", in *Proc. 16th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 780–784. doi:10.18429/JACoW-ICALEPCS2017-TUPHA159

[8] GDA: Generic Data Acquisition, http://www.opengda.org/

[9] DAWN: Data Analysis Workbench, https://dawnsci.org/