

Chapter 2

Systems Architecture

Abstract This introductory chapter discusses the important system architectural aspects of a transmission electron microscope and how these are influenced by new market demands. It starts by discussing system-wide architectural considerations and the challenges imposed by current and future industrial markets. Then it looks at how new market demands (e.g. automation, repeatability and ease-of-use) change the architectural key drivers significantly, and therefore lead to *architectural stress* in the current system. Finally, architectural patterns and new design concepts for a next system architecture for the transmission electron microscope are discussed.

Keywords Systems architecture • Electron microscope • Key drivers • System qualities

2.1 Architectural Considerations

Richard Doornbos

Embedded Systems Institute, The Netherlands

Firstly, it is important to understand how the current state-of-the-art electron microscopes can be characterised as a system. In general, one can say that they are precision-critical: overall system performance—in particular resolution—is highly determined by the accuracy and precision of its components. As current TEMs are capable of extreme magnification, e.g. for visualising single atoms, it should be clear that the quality requirements of all critical components are at the cutting edge of what is technically possible.

Electron microscopes are long-lived systems, as main components like column, mechanics and electronics remain operational for at least 15 years. Only peripheral components, such as computer hardware and system software, are typically updated on a shorter time scale. A main characteristic of these long-lived systems is that they generally suffer from obsolescence: that components needing to be replaced are already out of production.

Traditionally, electron microscopes are intended for basic and exploratory research in the fields of material sciences, performed in an academic context. To successfully operate these systems takes an expert in physics, usually dedicated for years to a particular system. They will be both knowledgeable about the internals of the microscope, to optimally control its operation, and understand the physics of the electron-specimen interaction and the microscopic application to be investigated.

It is clear from the above that current electron microscope system designs focus on key qualities like resolution and image quality.

Trends

As the resolution of the best modern TEMs (e.g. Titan from FEI Company, introduced in 2005) is in the order of 50–80 pm ($1 \text{ pm} = 10^{-12} \text{ m}$), it is often said that the resolution race is over. There is no strong need to significantly increase the resolution further as the size of the smallest atom lies in the order of 120 pm. Therefore this probably marks the end of current challenges.

Looking at current trends in microscopy applications, we generally see a significant move from human observation towards measurement and quantification. There is a growing need for statistical evidence, requiring repetitive accurate measurements and routine applications.

It is hard to explore these trends in detail, as a very large variety of applications exist or are being developed. The market is also rather fragmented. Typical research applications aim for ultimate and unique results that will lead to scientific recognition. However, industrial applications aim for entirely reliable results in the area of (troubleshooting in) production process development and process quality control.

New *types* of microscopy applications are being developed, enabling the measurement of the smallest signals as computational support increases. The power of modelling and reconstruction techniques in this area is impressive.

In electron microscopes, the pace of development has traditionally been governed by the speed of technological and engineering innovation regarding higher resolution. With the resolution race coming to a halt, things may change towards a more customer driven innovation model.

A New Direction

To understand what the new markets are and what these different customers want from the electron microscope, it is very important to thoroughly analyse the new




Fig. 2.1 Hierarchy of an (example) industrial usage of electron microscopy

needs and requirements. What is exactly meant by ‘routine measurements’ and ‘automated nanometrology’? The answer will probably be different for each customer, but a common denominator should result.

For a company producing systems in this field it is also important to search for new ‘key differentiators’ and ‘distinctive features’, e.g. strong automation, high stability and system predictability, excellent ease-of-use, etc.

Of course, a change of direction like this is likely to have more profound business consequences. Not only may markets change, but different business models may become more profitable. In particular, the need for more automated measurements suggests the role of a *solution provider*. This entails involvement in the measurement problem, provisioning of measurement solutions and reporting support. Activities such as keeping the system operational will often also be part of the solution provider’s responsibilities.

An important technical consequence is the embedding of the system into the customer’s workflow. This involves the specimen flow, such as sampling, preparation, loading and unloading, disposal, together with the integration of the measurement system into the customer’s computational and IT infrastructure.

Investigation of the customer’s workflow provides new perspectives on the role and importance of the electron microscope in the entire process. There are new insights in process optimisation with respect to throughput. They show that speed increase in the electron microscopic measurements sometimes has very little effect on overall throughput (see the example in Fig. 2.1).

The new types of applications and their industrial context strongly influence the design of a new electron microscope systems architecture. The architecture should provide more flexibility and adaptability with respect to the customer’s way of working. Generalizing over the entire field introduced so far, the new *key drivers* [1] are: cost effective analysis, accurate and precise measurements, and ease-of-use.

Note that these drivers are significantly different from the traditional key drivers of high resolution and image quality. Therefore, addressing the new market and applications using the current systems architecture is very difficult as the inherent limitations hamper, or even block, this process. This problem is called '*stress in the architecture*'.

With a new direction, new system functions and operational modes are required. At the workflow and workflow element level (see Fig. 2.1), we can distinguish two operation modes needed by an anticipated future automated microscope. First, there must be a batch-mode routine analysis mode, especially suited to trouble shooting and improvement of e.g. industrial production process steps. The second operational mode is an in-line, real-time analysis of production process samples, particularly for industrial production process control.

The most important microscopy functions at application-level are:

- Determination of nano-particles' sizes
- Determination of the three dimensional shape of nano-particles or aggregates, and
- Determination of the chemical composition of nano-particles.

To be able to perform these measurements cost-effectively, while eliminating repetitive labour (e.g. manually analysing thousands of particles), automation is essential. We can distinguish several major elements in the measurement process that are eligible for automation: specimen preparation, specimen loading and unloading, imaging by the electron microscope, and image measurement and data interpretation (this may include off-line image processing).

At the electron microscopic measurement level we need new basic system functions. These include automated acquisition of an optimum image (e.g. 'sharp', minimum distortions, sufficient signal-to-noise ratio), automated acquisition of a specified area of the specimen, etc.

The *key system qualities* [1] which should follow from the key drivers can be related to the levels indicated in Fig. 2.1. At the highest level, availability and reliability (note that the electron microscope system is part of the workflow!), and adaptability and flexibility with respect to customer needs are the most important system qualities, as they will lead to cost-effective analyses.

At the microscopy application level, throughput or short time-to-result, and understandability are the most important system qualities, as they lead to cost-effective analyses and address ease-of-use.

At the electron microscopic measurement level, system stability, predictability, precision, reproducibility, and robustness are the most important system qualities, as they lead to result correctness and address ease-of-use.

The consequences for the operator will be significant, as the built-in automation eliminates the need to deeply understand the system to a large extent. Therefore less dedicated education is needed and 'non-PhD operators' can use the system. However, specialised knowledge is still needed to fully exploit the new possibilities, especially on automation scripts, automated application, and experiment

design. This signifies the split between a microscope operator and an experiment designer/application specialist. The user-system interaction will also be significantly different: there is less user control of the system's operation, particularly during automated measurement runs. We suggest adding a clear user notification mechanism about the so-called *under the hood* operation (to know what is going on), and about the effects of automation. Effects can include a certain amount of extra sample damage, the extra time needed for system control, and the possibility of irrecoverable errors.

Fortunately, less time will be spent by operators on tuning the system, making more time available for the actual measurement task. Associated with this change we expect that more off-line and remote analysis work will become the norm. The user's trust in the generated data should be maintained. We suggest introducing sufficient feedback and checkpoints in e.g. automated image interpretation, by using partial/user-assisted instead of fully automated image interpretation.

For a system architect responsible for the new family of electron microscopes, this change leads to many complex questions and discussions, and opens up various investigations. Besides participating in finding out what functionality is required, the system architect's role is to provide answers on how to realise these new systems given the design of the existing systems, the current development organisation, and time and cost constraints.

The main technical questions are about which new techniques, methods, or even system architectures are suitable. The system architect is asked to give a concrete answer to the question: how can we adapt (add-on, redesign) current (precision-critical and obsolescent) systems to achieve the new requirements (typically still in vague terms, or in terms of physical components) within given constraints (cost, time, and organisation)?

Concepts and Possible Solutions

In finding answers to the questions above it is important to analyse the current system for bottlenecks and system tensions that arise due to the new requirements. A simple back-of-the-envelope throughput calculation using the current TEM shows that only low numbers of samples per hour can be analysed. This is, for example, mainly due to slow loading and unloading of specimens. Another analysis reveals that automated imaging of a large area (needing multiple images) requires improved stability and the correction of side-effects. These issues are addressed in detail in Part III of this book ([Chaps. 5–7](#)).

Finding conceptual solutions is a core task of the system architect. There are few formal approaches or tools to support this. However, by examining the architectural patterns, creativity can be stimulated and steered in the right direction. In the next section, the architectural patterns discovered in current TEMs are discussed and new system concepts for future electron microscopes are suggested.

2.2 Design Concepts for Global Electron Microscope Control

Sjir van Loo

Embedded Systems Institute, The Netherlands

Traditionally, an electron microscope operator often just wants to take a picture in order to visualise certain aspects of the sample at hand. And, as with a photo camera, the operator first has to set the microscope into an operational mode to ensure images fulfil certain quality parameters. The operator may control a relatively large number of settings, many of them while observing the generated image. The visual feedback loop is used to optimise the image until the operator decides that the image is ‘good enough’. Tasks the operator performs include inserting the sample into the microscope, setting the magnification, bringing the sample into focus, removing astigmatism from the image, finding a feature of interest on the sample, and so on.¹

After the operator has set the microscope into a desired setting, without touching any of the controls image quality will slowly degenerate over time. There are ‘aging’ effects of the sample, caused by bombarding the sample continuously with high energy electrons, and contamination effects. Besides these, most of the effects observed are caused by internal and external physical phenomena. They include relaxation effects of the electrical and mechanical parts, temperature changes in these parts, external pressure changes, and other phenomena that influence the microscope’s settings. So the operator has to change the settings continuously to maintain the desired image quality. Essential for this discussion is that on one hand, an operator may actively change the image by adjusting settings like magnification, sample position or electron beam related parameters like energy and spot size, while on the other hand the machine will slowly drift away from its set point by internal and external influences.

Regarding automating microscope usage, the goal is to change the architecture so that the system shows autonomous behaviour. This involves replacing many of the actions normally performed by the operator. At first this implies that a measurement is automated and consequently the user has to be able to write an application program, performing the desired measurements. As an example see the application described in [Sect. 5.2](#). The application programmer’s most important task is to program the analysis of the images acquired. As such it is very helpful if the images acquired always have the quality desired. That is, they are ‘sufficiently’ in focus and astigmatism free, and not, or barely, drifting away. For simplicity, we will assume that ‘sufficiently’ can be defined and is fixed for all applications. See

¹ In doing so the user builds up a mental model of the sample, how it is structured, which parts are elevated, where the interesting parts are, etc. The actual image is often just a reminder to trigger the mental model.

also [Sect. 6.3](#). The system has to perform a number of actions *under the hood*, i.e. invisible to the application programmer.

For the sake of this discussion we will divide the control actions the system can perform into three different groupings:

1. Electron microscope settings

These are the global application choices for a number of controls that define the way the microscope is used. They include, among others, the mode (e.g. camera or scanning acquisition mode), the magnification, the electron beam parameters like spot size, and many others. Typically these do not change during image acquisition. Settings are assumed to be constant.

2. Position controls

Several means exist to position the sample in 3D space with respect to the position of the beam. See [Chap. 7](#).

3. Image controls

These are used to control important image quality parameters such as focus and absence of astigmatism. They cannot be ‘set’, but have to be optimised based on observation of the images generated. We will limit ourselves to focus and astigmatism.

If changes are made to the electron microscope settings or the position controls, the image controls may have to be adjusted as the image quality might degrade, e.g. get out of focus. The impact of some changes will be very limited, while for other changes the precise effect is unknown. This suggests a model in which changes in settings and positions can affect the image controls according to a set of simple rules. These could be “after changing the spot size we have no knowledge about the state of focus and astigmatism”, “one step change in magnification has a limited effect on focus and no effect on astigmatism”, and “a stage move where (move distance $\times$ magnification $<$ value) has a limited effect on focus and no effect on astigmatism”.

For our architecture these rules are important. Given that we want to provide the user with images of sufficient quality, we need a way of determining if the system needs to perform an action, like focusing the image. So, if an application changes the microscope settings or repositions the sample, the system has to intercept and invoke an appropriate action as defined by the given rules. The system must maintain an invariant, informally defined as providing the application with images of sufficient quality only.

Invariant

For the following it is useful to introduce a set of discrete states for the key characteristics. For focus we come to three states: OK|CLOSE|UNKNOWN. For absence of astigmatism, we have two states: OK|UNKNOWN. It turns out that when combining these only 3 states can result from the above rules: (OK, OK), (CLOSE, OK), (UNKNOWN, UNKNOWN).




Fig. 2.2 State diagram for the invariant

To be able to acquire a useful image we define that an invariant expression on image controls should hold. The invariant expression is straightforward: $(\text{Focus} == \text{OK}) \text{ AND } (\text{Absence-of-astigmatism} == \text{OK})$, denoted as $\text{Inv}(F, A) == (\text{OK}, \text{OK})$. As initially the status is unknown, we assume the invariant not to hold and we first have to restore it, starting with status $(\text{UNKNOWN}, \text{UNKNOWN})$ (Fig. 2.2).

In automated applications the system has to maintain the invariant, i.e. restore the invariant after it has been disturbed. As the disturbances are described in the rules mentioned above, the system can determine when a restore operation has to be invoked. It does this by intercepting the operations the application performs on microscope settings and position controls. As these operations in the current architecture are already executed by the microscope system on request of the application, adding interception is easy and straightforward. The knowledge that the system has about the effect of the disturbance, i.e. state $(\text{CLOSE}, \text{OK})$ or $(\text{UNKNOWN}, \text{UNKNOWN})$, is used to determine which restore operation to use. Of course the restore operation used for $(\text{UNKNOWN}, \text{UNKNOWN})$ is always usable. However, it turns out that a restore operation that uses the knowledge that only the focus has to be restored (when the system is close to focus) is much faster. See also Sects. 3.1 and 6.3 where the two restore operations are described.

Note that relaxation effects and external influences might also cause the system to drift away from the optimal image controls and hence disturb the invariant. However, within the range of magnifications used in our experiments (see Sect. 3.1) and the time frames of the type of automated applications we are interested in, it shows that focus and astigmatism drift is negligible. A time-out value on the time the invariant holds might possibly be required in other applications, but was not relevant for this project. Position drift, however, is another relevant disturbance (see also Sect. 7.3). The correcting position drift might cause a disturbance of the invariant, as during executing this correction the rules on invariant disturbance also apply!

Sensor and Payload Data

From the above it is clear that to be able to acquire an image with suitable quality for the microscopy application we first have to acquire a number of images to restore the invariant. We will refer to the image we want to acquire for the application as payload data (image), and images to restore the invariant as sensor data (image).

Of course it would be very nice if we had independent sensors to measure the image quality directly. This would allow a control loop to be built so that the system is always ready to acquire payload images. However, such a sensor is not available.² In current electron microscopy the only sensor available with sufficient sensitivity is the instrument itself. Consequently we have to use the instrument both for acquiring payload data and sensor data. As well as being required for restoring the invariant, sensor data is also needed to allow for drift correction when acquiring payload data over a long period of time. This is the case in EDX, X-Ray data collection, or tomography.

Sensor data serves a different purpose than payload data and therefore has different requirements with respect to properties such as image content, number of pixels, and acquisition time. For instance, in order to restore the invariant it is not necessary to acquire sensor images of the same size as the payload image. We could use a smaller size, which is faster in acquisition time and speeds up the overall process. However, there should be sufficient information in the image for robust operation of the invariant restore operation. Consequently, selecting an area from the full image to be used for this purpose might not be straightforward. Furthermore, changing the settings and position controls in order to switch between sensor data and payload data should not trigger any change in the invariant. This means the freedom to define sensor data acquisition is bound by the intended acquisition of payload data.

Sensor and payload data acquisitions sometimes have to be interleaved. The interleaving patterns are determined by the simultaneous execution of multiple invariant restoring operations. Some sensor data acquisitions, such as for optimizing focus or absence of astigmatism, have to be finished before the payload acquisition. Other corrections, such as drift compensation, need sensor data acquisitions inherently interleaved with the payload data acquisition. One example is the repeated switching to sensor data acquisition *during* payload acquisition: after a few scanned lines of the payload image, each time a small sensor image is acquired. This sensor image is used to estimate the drift. The influence of the interleaving pattern may play an important role when dynamics are concerned. Luckily most microscopy applications investigate static samples. The interleaving pattern is also important for the control system as the sensor images will not be available at a regular pace. Therefore the control system should be able to cope with this variability.

² Although in [Sect. 6.1](#) a sensor that is very helpful is introduced.




Fig. 2.3 The value of a key system parameter over time, showing the various phases

A Transient, Stationary or Stable System

Another characterisation of the system state is useful for our purposes. A change of system settings takes time as physical changes have to take place. For instance, a change in the lens current to re-establish focus is not instantaneous because of the lens dynamics. The time between applying the new current to the lens and the moment the lens is stable again is referred to as a transient. During transients all data acquired is considered to be invalid, for both payload and sensor data. This corresponds with our requirement that acquired data is only valid when the invariant holds.

Figure 2.3 shows the behaviour of a key system parameter, such as focus, over time. The large change is called transient, leading to a stationary state. In this state small changes are allowed to either restore or maintain the invariant, therefore stable phases are interchanged with short, unstable phases. Transients are intentionally applied to the system and usually explicit, whereas the small changes in the stationary phase are implicit and performed autonomously by the system. Note that more invariant restoring operations may occur in a stationary phase.

Using the above-defined concepts, the system controls its behaviour according to the model of Fig. 2.2. However, we need signals from the system to determine the beginning and end of transients and states. Dedicated sensors (e.g. magnetic field sensors) may provide input to models that describe key system parameters from which signalling could be derived, providing a signal at the start and end of a transient. These signals can contribute to driving the state model.

The combination of design concepts described in this section will help in creating a new systems architecture for an automated electron microscope. The concepts have to be extended, refined and worked out in more detail to finally realise a predictable and well-behaved system. And more concepts will have to be developed. However, with these concepts we were able to run the application of automatic particle sizing, as described in Sect. 5.2, by implementing the concepts on the experimentation platform in the Concept Car (Sect. 3.1).

Roughly the system behaves as follows:

The system maintains the invariant state, initially (UNKNOWN, UNKNOWN), by intercepting all microscope setting commands and all positioning commands invoked

by the application. When the application invokes an acquire image command, the system checks the invariant state. Then, if required, it invokes the procedure (see [Sect. 6.3](#)) to restore the invariant, unnoticeable to the application. This procedure acquires sensor images as desired. After restoring the invariant the system acquires the requested payload data and returns this data to the application. During payload acquisition the system corrects for drift by interleaving payload data and sensor data acquisitions.³ After providing the payload to the application, the application proceeds, performing the desired analysis and then invoking new commands intercepted by the system. This possibly results in a disturbance of the invariant again, etc.

Independent of the procedure above, all invocations of commands that change the lens settings, be it from the application or internal in the system, e.g. to correct for focus, wait until the system is stable before proceeding to a subsequent acquisition of sensor or payload data. The implementation of this may be simple, based on a worst case time out, or very advanced, using a feedback signal from the control unit. However, the concept remains the same. Note that in some use cases images may need to be acquired at high frequency, like in full interactive control by the operator. At these times the system should provide direct visual feedback to the operator. Also, payload acquisition will run continuously, overriding the aforementioned under the hood procedures. This provides backwards compatibility with respect to user behaviour.

Acknowledgments The author likes to acknowledge Bernard van Vlimmeren and Hugo van Leeuwen of FEI for their valuable discussions.

Reference

1. G. Muller, *Systems Architecting-A Business Perspective* (CRC Press, Boca Raton, 2012)

³ This is only possible in scanning acquisition mode, not in camera acquisition mode.